L

Learning in Games

JOHN NACHBAR
Department of Economics, Washington University, St. Louis, USA

Article Outline

Glossary

Bayesian learning In repeated games, a model in which each player best responds to her prior, which is a probability distribution over her opponent's behavior strategies.

Behavior strategy In a repeated game, a behavior strategy for player i gives, for each possible date and each possible history of play in the game up to that date, a probability distribution over i's actions next period. This includes the possibility that player i may play some action for certain.

Belief learning In repeated games, a model in which players best respond to prediction rules.

Prediction rule In a two-player repeated game, a *deterministic prediction rule* gives a probability distribution over the opponent's actions next period as a deterministic function of the history of the game. In a *stochastic* prediction rule, the distribution over the opponent's actions can depend on history probabilistically. A deterministic prediction rule that gives player 2's forecast about player 1's actions is formally equivalent to a behavior strategy for player 1.

Repeated games A repeated game is an extensive form game representing a repeated strategic interaction. The interaction being repeated is called the *stage game*. In a *discounted* repeated game, payoffs in the repeated game are a geometrically weighted sum of the payoffs each period from the stage game. The weight on period t payoffs is δ^t, where $\delta \in (0, 1)$ is the *discount factor*. A player who is patient has a discount factor close to 1.

Definition of the Subject

In the context of this article, *learning* refers to a particular class of dynamic game theoretic models. In models in this class, players are "rational" but not necessarily in equilibrium: players forecast, possibly inaccurately, the future behavior of their opponents and optimize, or ε optimize, with respect to their forecasts. Two objectives are to model out-of-equilibrium behavior by sophisticated players and to understand when, or whether, play might converge to equilibrium.

Learning models are a branch of a larger literature on out-of-equilibrium behavior in dynamic games. In other branches of the literature, players are modeled as "adaptive". Players do not forecast and they do not optimize. Rather, they follow some other form of behavioral rule, such as imitation, regret minimization, or reinforcement. Learning models, especially those that attempt to capture sophisticated behavior, are most appropriate in settings where players have a good understanding of their strategic environment and where the stakes are high enough to make forecasting and optimization worthwhile. For surveys of evolutionary/adaptive models, see ► Evolutionary Game Theory and [11,73,74].

Introduction

This essay surveys research on learning in games. The subject is as old as game theory itself, dating to the work of Cournot in the early nineteenth century. But, with some important exceptions, research on learning in games was largely dormant until the late 1980s.

I have divided the survey into two parts. The first part covers *deterministic* learning, in which each player optimizes with respect to a prediction rule that gives a probability distribution over the opponent's next period actions as a deterministic function of the history of the game to date. Cournot's original dynamic model falls into this category. The second part briefly surveys some of the literature on *stochastic* learning, in which each player optimizes with respect to a prediction rule that is a probabilistic function of history. This survey expands and updates my earlier survey [58]. For a complementary survey, with a different emphasis and more detail on some topics, see [28].

One can formulate learning within the context of very general dynamic games. For this survey, however, I focus on the most studied case: discounted infinitely repeated games with perfect monitoring (at each stage, players know the sequence of action profiles chosen at previous stages). In Subsect. "Convergence to What Sort of Equilibrium?", I discuss an important variant: repeated extensive form games in which the perfect monitoring assumption is violated. Also for simplicity, I focus mostly on two-player games.

Deterministic Learning

Classical Learning Models

I begin by introducing two "classical" learning models, Cournot's best response dynamics and fictitious play. In both models, players are unsophisticated: they cannot learn to forecast even simple patterns in opponent behavior, such as deterministic alternation between two stage games actions. But Cournot best response and fictitious play are historically important, easy to describe, relatively tractable, heavily studied, often used as building blocks in more elaborate dynamic models (e. g. [65]), and can be used to illustrate some of the major themes in the literature.

Cournot's Best Response Dynamics Following Cournot [12], consider two mineral spring owners, A and B, who are each trying to make money by selling water. Suppose that the profit to firm A is $(13 - Q^A - Q^B)\, Q^A - Q^A$, where Q^i is the output in one period by firm i, $13 - Q^A - Q^B$ gives the market price, and output costs 1 per unit, hence total cost to firm A is Q^A. The quantity Q^A that maximizes A's profit depends on Q^B, and so in order to choose Q^A, firm A must forecast Q^B. In a model of sophisticated behavior, firm A would want to forecast not only Q^B next period but also how B might respond in subsequent periods. But, for the moment, assume that A is myopic, meaning that A chooses Q^A to maximize next period's expected profit. Then A should set,

$$Q^A = \frac{12 - \mathbb{E}[Q^B]}{2},$$

where $\mathbb{E}[Q^B]$ is A's expectation of Q^B.

What should $\mathbb{E}[Q^B]$ be? In Cournot's story, A puts probability one on B continuing to produce next period whatever B produced this period, hence $\mathbb{E}[Q^B_{t+1}] = Q^B_t$. Cournot argued that if both firms adjust in this way, then eventually Q^A and Q^B converge to the same number, Q^*, which must satisfy,

$$Q^* = \frac{12 - Q^*}{2},$$

hence

$$Q^* = 4\,.$$

In modern terminology, the symmetric output profile (4, 4) is the Nash equilibrium of the single period version of this game. A generalized version of Cournot's mineral springs game has become the benchmark model in the economic field of Industrial Organization. Cournot not only identified the Nash equilibrium of this game, but he also provided a dynamic rationale for why the Nash equilibrium might arise. Cournot, in his exposition, was actually more interested in the dynamic story than in the equilibrium.

Fictitious Play Fictitious play and its variants, first introduced by Brown [10] in the 1950s, is the single most heavily studied learning model. The term "fictitious play" refers to the idea that players would run through the fictitious play dynamic in their minds, prior to the start of actual play, as a way of forecasting their opponent's behavior once play begins. In recent years, however, fictitious play has usually been interpreted as a model of actual behavior in real time.

In its modern form, fictitious play dictates that each player predicts that the probability that her opponent will play an action, say L, next period is a weighted sum of an initial probability on L and the frequency with which L has been chosen to date. The weight on the frequency is $t/(t + k)$, where t is the number of periods thus far and $k > 0$ is a fixed parameter. The larger is k, the longer the initial probability on L dominates the player's forecast about L.

In fictitious play, players are *as if* myopic: for any k and any discount factor, each player effectively believes he cannot influence his opponent's future behavior (see the discussion in Subsect. "Belief Learning and Bayesian Learning"), and so it is optimal to maximize current pe-

riod payoff, ignoring the future. As this observation suggests, forecasting under fictitious play is extremely naive. Fictitious play is, in fact, mathematically very close to regret minimization, an adaptive learning model; see [35] and [36]. To justify fictitious play for sophisticated players, [21] suggests thinking of environments in which there is a large population of players randomly matched (an environment normally associated with evolutionary models), rather than a standard repeated game; see also [29].

Belief Learning and Bayesian Learning

In general, *belief learning* refers to any learning model in which players best respond (or ε-best respond) to a deterministic prediction rule. Both the Cournot dynamic and fictitious play are examples of belief learning. An alternative approach to modeling learning is *Bayesian*. Recall that in a repeated game, a *behavior strategy* gives, for every history, a probability over the player's stage game actions next period. In a Bayesian learning model, each player chooses a behavior strategy that best responds (or ε-best responds) to a *prior*, a probability distribution over the opponent's behavior strategies. A basic observation in the theory of deterministic learning is that belief learning and Bayesian learning are mathematically equivalent.

Explicitly, player 1's prediction rule about player 2 is mathematically identical to a behavior strategy for player 2. Thus, any belief learning model is equivalent to a Bayesian model in which each player optimizes with respect to a prior that places probability one on her prediction rule, now reinterpreted as the opponent's behavior strategy.

Conversely, any Bayesian model is equivalent to a belief learning model. Explicitly, for any prior over player 2's behavior strategies there is a degenerate prior, assigning probability one to a particular behavior strategy, that is equivalent in the sense that both priors induce the same distributions over play in the game, no matter what behavior strategy player 1 herself adopts. This is a form of Kuhn's theorem [50]; for a version applicable to repeated games, see [4]. I refer to the behavior strategy used in the degenerate prior as a *reduced form* of the original prior. Thus, any Bayesian model is equivalent to a Bayesian model in which each player's prior places probability one on the reduced form, and any such Bayesian model is equivalent to a belief learning model.

To illustrate, consider fictitious play. For simplicity, consider a stage game with just two actions, L and R. By an i. i. d. strategy for player 2, I mean a behavior strategy in which player 2 plays L with probability q, independent of history. Thus, if $q = 1/2$ then player 2 always randomizes 50:50 between L and R. Fictitious play is equivalent to a degenerate Bayesian model in which each player places probability one on the fictitious play prediction rule, and one can show that this is equivalent in turn to a non-degenerate Bayesian model in which the belief is represented as a beta distribution over q; see, for example, [25]. The uniform distribution over q, for example, corresponds to taking the initial probability of L to be 1/2 and the parameter k to be 2. Thus, fictitious play can be thought of as a Bayesian model in which each player is certain that the other is playing an i. i. d. strategy, but isn't sure which i. i. d. strategy. Note that since each player believes the other is i. i. d., each believes she has no influence on the others future behavior. Hence, as noted earlier, myopic optimization is optimal for any discount factor.

Likewise, the Cournot prediction rule is Bayesian. In particular, one can form the degenerate prior from the Cournot prediction rule. To my knowledge, the Cournot rule has no compelling non-degenerate Bayesian interpretation.

The fact that both the Cournot dynamic and fictitious play can be interpreted as Bayesian underscores the fact that there is no presumption that Bayesian players are sophisticated. It is possible for a Bayesian to have an "idiotic" prior. I take up the issue of sophisticated learning below, especially in Subsect. "Sophisticated Learning".

What Should "Convergence" Mean?

The primary research question in the learning literature is whether play converges to equilibrium. One of the characteristics of the literature is a running dialog about what "converges to equilibrium" ought to mean.

Convergence in What Game? For concreteness, suppose that we are looking for convergence to Nash equilibrium; similar comments apply to convergence to other forms of equilibrium, such as correlated equilibrium.

In the models that are the focus of this essay, learning takes place within the context of a repeated game. Therefore, it is natural to look for convergence to Nash equilibria of the repeated game. Repeated play of a stage game Nash equilibrium always constitutes a Nash equilibrium of the repeated game, but typically there are also other types of repeated game equilibria.

First, even if players do play stage game Nash equilibria, they need not play the *same* stage game Nash equilibrium in every period. As an example, consider the repeated Battle of the Sexes, a stage game of which is given in Table 1. Alternation between (a, a) in odd periods and (b, b) in even periods, both of which are stage game

Learning in Games, Table 1
Battle of the Sexes

	a	*b*
a	8, 10	0, 0
b	0, 0	10, 8

Learning in Games, Table 2
Matching Pennies

	a	*b*
a	1, −1	−1, 1
b	−1, 1	1, −1

Nash equilibria, is a perfectly reasonable repeated game Nash equilibrium.

Second, if the discount factor is close enough to 1, the folk theorem implies that there will be many repeated game Nash equilibrium, and some of these may involve play of actions that, in the stage game, are not part of any stage game equilibrium (for more on the folk theorem, see [30]). The standard example is the repeated prisoner's dilemma, in which the repeated game can have Nash equilibria that sustain cooperation, even though cooperation is strictly dominated in the stage game.

More generally, learning stories take place in the context of a larger dynamic game. Learning may lead to Nash equilibrium play in the continuation of this larger game, but Nash equilibrium in the continuation of the larger game need not correspond to Nash equilibrium-like play in a component or projection of that game.

Convergence in What Sense? Until the 1990s, most of the work on learning focused on whether the empirical marginal frequencies of realized play converged to a Nash equilibrium of the stage game.

As an illustration, consider repeated Matching Pennies, the stage game for which is given in Table 2. The Nash equilibrium of repeated Matching Pennies, for any discount factor, calls for players to randomize 50:50 in every period, regardless of history. Under convergence of the empirical marginal frequencies, play is said to converge to Nash equilibrium if each player plays a half of the time. This criterion is satisfied if play alternates deterministically between (a, a) and (b, b). In contrast, a Nash equilibrium play path looks random; in particular, with high probability, over any finite set of dates, the four pure outcomes (a, a), (a, b), (b, a), and (b, b) occur about equally often. Convergence of the empirical marginal frequencies is thus a very weak convergence criterion.

A somewhat tougher criterion is convergence of the empirical *joint* frequencies. In the matching pennies example, this would require that each of the four pure action profiles gets played 1/4 of the time. This convergence criterion eliminates the previous example but it is still extremely weak. It would, for example, allow convergence to the deterministic sequence (a, a), (a, b), (b, a), (b, b), (a, a), … to count as convergence to a Nash equilibrium. In contrast, in the Nash equilibrium, by the strong law of large numbers, (a, a) gets played at dates $t = 1, 5, 9, \ldots$ only 1/4 of the time. These difficulties with defining convergence in terms of empirical frequencies were first emphasized by Fudenberg and Kreps; see [21] and also the related paper by Jordan [43].

Choice of convergence standard is obviously a matter of balance. Choose too weak a standard, as in the examples above, and convergence is arguably not meaningful. Choose too strong a standard and one gets impossibility results, even though positive results are still available for other, weaker but arguably still satisfactory, forms of convergence. For a stark example, consider the Battle of the Sexes game of Table 1. If players start in the repeated game Nash equilibrium in which they play (a, a) in odd periods and (b, b) in even periods then play converges (trivially) to Nash equilibrium play in any continuation game. But note that, strictly speaking, we get a *different* Nash equilibrium depending on whether the starting date of the continuation game is odd or even. So a convergence standard that requires that play be close to the *same* repeated game Nash equilibrium in every continuation game yields non-convergence, even in this trivial example. For a more subtle example illustrating the tradeoffs in choosing the right convergence standard, see Subsect. "Payoff Uncertainty".

Loosely, the strongest form of convergence that one can generally hope for is that, to an outside observer, play over finite continuation histories looks asymptotically like play of some repeated game Nash equilibrium or ε-Nash equilibrium. Again, one must allow for the possibility of convergence to the play of different Nash equilibria in different continuation games, even along the same play path.

Convergence to What Sort of Equilibrium? Although my focus here is primarily on Nash equilibrium, other solution concepts are often of interest. Obvious alternatives are rationalizable strategy profiles and correlated equilibria and both have received attention (for example [7,16], and [61]).

Two other variants of Nash equilibrium, less common in standard game theory texts, have also proved important in the literature. First, it is often natural to assume that

players ε optimize rather than exactly optimize. If players only ε optimize then convergence will be to an ε-Nash equilibrium rather than to an exact Nash equilibrium. It is common in the literature to model myopic players as using a logit selection from the stage game's ε-best response correspondence. In this case, ε-Nash equilibrium in the stage game takes the form of a McKelvey and Palfrey *quantile response equilibrium* (QRE) [52], a solution concept that has become important in the experimental game theory literature (see, for example, [32]). While any Nash equilibrium is an ε-Nash equilibrium, ε-Nash equilibria are typically not Nash equilibria, although the difference is small if ε is small.

A second important variant arises in repeated games in which the stage game has a non-trivial extensive form. In such settings, the perfect monitoring assumption may be untenable. In particular, it may make sense to assume instead that while players observe the outcome of the stage game, they do not observe the full strategy profile for the stage game.

To take a concrete example, consider the repeated ultimatum bargaining game, in which, in the stage game, player A makes an offer in the form of an integer $x \in [0, 100]$ and player B either rejects, yielding a payoff profile of $(0, 0)$, or accepts, yielding a payoff profile of $(x, 100 - x)$. It may make sense to assume that while player A can observe player B's action (accept or reject) in response to the actual offer, she cannot observe player B's entire stage game strategy, since this would mean observing how B would have responded to every other possible offer.

For this setting, Fudenberg and Levine [22] propose *self-confirming equilibrium* (SCE) as an alternative to Nash equilibrium. There are actually several forms of SCE, depending on the setting and the information of the players. The basic idea is that in a SCE of the stage game, players are optimizing with respect to predictions about their opponents that are correct along the stage game play path. Predictions about how the opponent would behave off the stage game play path, however, may be wrong. SCE bears a family resemblance to Hahn's conjectural equilibrium [33] and Kalai and Lehrer's subjective equilibrium [46].

For players to learn to make correct predictions about stage game behavior off the play path, they must experiment with alternative play. Fudenberg and Levine show, in the context of a large population, steady-state learning model [23], that if players are sufficiently patient then they have enough incentive to experiment, and thus will learn to play a stage game Nash equilibrium. In other settings, or if players are not sufficiently patient, it may be plausible to expect players to end up at SCE that are far from any Nash equilibrium. Fudenberg and Levine suggest that this may explain some real world behaviors that may otherwise seem at variance with standard game theory predictions. A provocative example can be found in [27].

Learning in Games, Table 3
The Shapley game

	a	b	c
a	1, 0	0, 0	0, 1
b	0, 1	1, 0	0, 0
c	0, 0	0, 1	1, 0

Finally, some other solution concepts have received attention, notably Samet and Kalai's persistent retracts [48], Basu and Weibull's CURB (closed under rational behavior) sets [5], and Voorneveld's minimal prep sets [70] and [71]. These solution concepts tend to arise most readily in stochastic learning models, which I discuss in Sect. "Stochastic Learning".

Convergence in Classical Learning Models

Neither Cournot best response dynamics nor fictitious play exhibit universal convergence, by which I mean convergence for all stage games. It is easy to construct non-convergence examples for Cournot best response, even for games with pure strategy equilibria. For fictitious play, there are classes of games that cause convergence problems even when one considers only weak forms of convergence, such as convergence of the empirical marginal distributions. The classic example of a problem game, due to Shapley [69], is given in Table 3. In this game, if the prediction rules concentrate initial probability on, say, (b, a), then play cycles, starting at (a, a) then moving to (a, c), then to (c, c) and so on. Moreover the cycle becomes slower and slower, so that not even the empirical marginal frequencies converge to the unique Nash equilibrium of the stage game, which has each player randomize (1/3, 1/3, 1/3). See also [43]. Hart and Mas-Colell [37] prove a kind of generalization of the Shapley example: for a large class of dynamics that include a continuous time version of fictitious play as a special case, one can always find a stage game and initial conditions for which convergence to stage game equilibrium play fails.

On a more positive note, Cournot best response dynamics do converge to pure strategy stage game equilibrium if the stage game is solvable by the iterated deletion of strictly dominated strategies. See [7,55], and, for a more general class of models [53].

As for fictitious play, there are ε optimizing variants of fictitious play that yield convergence to approximate stage game equilibrium play (possibly mixed) for all zero sum

games, all games with an interior ESS (evolutionarily stable strategy), and all common interest games, in addition to all games that are strict dominance solvable, with the approximation closer the smaller is ε. Somewhat weaker convergence results are available for supermodular games. These claims follow from results in [39], which builds on [6] and [21], and which also provides additional references to the large literature on fictitious play.

The use of ε optimization, or something analogous (such as small, privately observed, payoff shocks in each period), is necessary. In a standard repeated game setting, for almost every stage game, standard fictitious play, with exact optimization, yields pure continuation play paths (see [21]). Thus, in games like repeated matching pennies, standard fictitious play cannot give convergence to an equilibrium play path. This is an instance of a recurring modeling theme in the learning literature: it is often easier to get positive convergence results for *vp* optimization than for exact optimization.

Kalai–Lehrer Learning

Kalai and Lehrer [45] (hereafter KL) take a Bayesian perspective and ask what conditions on priors are sufficient to give convergence to equilibrium, or approximate equilibrium, play. I find it helpful to characterize KL, and related papers, in the following way. Say that a prior profile (giving a prior for each player) has the *learnable best response property* (LBR) if there is a profile of best response (or ε-best response) strategies (LBR strategies) such that, if the LBR strategies are played, then each player learns to predict the play path.

A player *learns to predict the play path* if her prediction of next period's play is asymptotically as good as if she knew her opponent's behavior strategy. If the behavior strategies call for randomization then players accurately predict the distribution over next period's play rather than the realization of next period's play. For example, consider a 2×2 game in which player 1 has stage game actions T and B and player 2 has stage game actions L and R. If player 2 is randomizing 50:50 every period and player 1 learns to predict the play path, then for every ε there is a time, which depends on the realization of player 2's strategy, after which player 1's next period forecast puts the probability of L within ε of 1/2. (This statement applies to a set of play paths that arises with probability one with respect to the underlying probability model; I gloss over this sort of complication both here and below.)

If LBR holds, and players are using their LBR strategies, then, asymptotically, the continuation play path is an approximate equilibrium play path of the continuation repeated game. The exact sense in which play converges to equilibrium play depends on the strength of learning and of optimization. See KL and also [66] (both of which focus on exact optimization) and [60] (which considers ε optimization).

KL, building on work in the probability literature on the merging of measures, notably Blackwell–Dubins [8], show that a strong form of LBR holds if beliefs satisfy an absolute continuity condition: each player assigns positive probability to any (measurable) set of play paths that has positive probability given the players' actual strategies. A strong sufficient condition for this is that each player assigns positive, even if extremely low, prior probability to her opponent's actual strategy, a condition that KL call *grain of truth*.

Universal Convergence

The work described thus far leaves open whether one can write down *any* deterministic learning model, even an implausible one, that exhibits universal convergence, that is, a learning model that, for a given stage game form (giving stage game actions but not payoff functions), gives convergence to Nash equilibrium play in the repeated game for all (or at least for generic) discount factors and stage game payoff functions.

Implicit in most work on this question is an assumption that the learning model is "uncoupled" in the sense of [37]: player i's prior over her opponent's repeated game strategies does not depend on the specification of her opponent's stage game payoffs. Cournot best response and fictitious play, as typically employed, are both examples of uncoupled learning models. One can get convergence in a "coupled" model simply by having the players play a Nash equilibrium of the repeated game. So, some degree of uncoupling is needed to avoid triviality. On the other hand, full uncoupling is a strong assumption from the perspective of sophisticated learning, since it effectively rules out introspective reasoning about ones opponent. And full uncoupling presumably makes convergence more difficult to achieve. Indeed, the main result of [37] is an impossibility theorem on convergence for certain classes of adaptive learning models.

For belief learning models, there are two competing intuitions on (uncoupled) universal convergence. One intuition is that universal convergence, even to ε equilibrium, is impossible because universal learnability is impossible: for any prior, there are opposing strategies that a player will fail to forecast, even approximately. One can show this via a diagonalization argument along the lines of [63] and it also follows as a corollary of results discussed in Sub-

sect. "Sophisticated Learning". In fact, for any given prior, the set of strategies one can learn to forecast is small, in a sense that can be made precise. A classic reference on the difficulty of prediction, in the context of general stochastic processes, is [20]. For a useful up to date survey, see the literature review in [1].

The competing intuition is that universal learnability is not necessary for convergence. The fact that both players are engaged in a game provides structure that conceivably, could force posteriors to be correct, at least along the path of play, even if priors are fundamentally wrong. Something like this happens in fictitious play. In the standard Bayesian interpretation of fictitious play (see Subsect. "Belief Learning and Bayesian Learning"), each player is certain that her opponent is playing an i. i. d. strategy even though neither is, because an i. i. d. strategy is not optimal. Yet, in many games, fictitious play generates behavior that looks asymptotically i. i. d. and play does converge to that of a Nash equilibrium.

Based on the first intuition and also on other negative results like those in [19] and [37], there has been a consensus that universal convergence is impossible for deterministic learning. Recently, however, Noguchi has reported a positive result [60]. Noguchi argues that, starting from any pair of priors, one can construct augmented priors with the property that (a) if a player could have learned to predict a strategy under the old prior, then she can also learn to predict that strategy under the new prior and (b), for ε optimization, a form of LBR holds, hence convergence to ε-equilibrium play obtains, where ε can be made arbitrarily small. One could take the original priors to be the fictitious play priors. Or one could be more ambitious and build various forms of pattern recognition into the original priors.

Sophisticated Learning

A number of papers investigate classes of prediction rules that exhibit desirable properties, such as the ability to detect certain kinds of patterns in opponent behavior. Important examples include [3,24,26,67].

In [57], I consider the issue of sophistication from a Bayesian perspective. For simplicity, focus on 2 player games. Fix a profile of priors and a subset of behavior strategies for each player and consider the following criteria for these strategy subsets.

Learnability For any strategy profile drawn from the strategy subsets, both players learn to prediction the play path.

Richness Informally, richness requires that if a behavior strategy is included in one of the strategy subsets then certain variations on that strategy must be included as well. Richness, called CSP in [57], is satisfied automatically if the strategy subsets consist of all strategies satisfying a standard complexity bound, the same bound for both players. Thus richness holds if the subsets consist of all strategies with k-period memory, or all strategies that are automaton implementable, or all strategies that are Turing implementable, and so on.

Consistency Each player's subset contains a best response to her belief.

The motivating idea is that, for priors to be considered sophisticated, a necessary (but not sufficient) condition is that the priors can be interpreted as probability distributions over strategy subsets satisfying these criteria. [57] studies whether any such priors exist.

Consider, for example, the Bayesian interpretation of fictitious play in which priors are probability distributions over the i. i. d. strategies. The set of i. i. d. strategies satisfies learnability and richness. But for any stage game in which neither player has a weakly dominant action, the i. i. d. strategies violate consistency: any player who is optimizing will not be playing i. i. d.

As shown in [57], this feature of Bayesian fictitious play extends to all Bayesian learning models. For large classes of repeated games, for *any* profile of priors and for *any* strategy subsets satisfying richness, if learnability holds then consistency fails.

Let me make a few remarks. First, since the set of all strategies satisfies richness and consistency, it follows that for any profile of priors there is a strategy profile that the players will not learn to predict. This can also be shown directly by a diagonalization argument along the lines of [13] and [63]. The impossibility result of [57] can be viewed as a game theoretic version of [13]. For a description of what subsets *are* learnable, see [59].

Second, suppose that the strategy subsets are generated by some standard definition of complexity, the same for both players. Then, as noted above, richness holds. Suppose further that there are priors for which learnability holds. This will be the case, for example, for Turing implementable strategies, since the set of such strategies is countable. Then, for such priors, for large classes of repeated games, consistency fails: best responses must violate the complexity bound.

Third, if one constructs a Bayesian learning model satisfying learnability and consistency then LBR (see Subsect. "Kalai–Lehrer Learning") holds and, if players play their LBR strategies, play converges to equilibrium play. This identifies a potentially attractive class of Bayesian

models in which convergence obtains. The impossibility result says, however, that if learnability and consistency hold, then player beliefs must be partially equilibrated in the sense of, in effect, excluding some of the strategies required by richness.

Fourth, the theorem is robust. It holds under ε optimization, for ε small. It holds even for fairly weak definitions of prediction (definitions that allow occasional but persistent forecasting errors, for example). And while learnability, as defined, requires prediction for any strategy profile drawn from the strategy subsets, this is only for ease of exposition. A variation of the result says that if the subsets are rich then for any strategy profile for which prediction holds, either that profile is not (ε) optimal or it is not contained in the strategy subsets.

Last, consistency is not *necessary* for convergence. See Subsect. "Universal Convergence". The impossibility result is a statement about the ability to construct Bayesian models with certain properties; it is not a statement about convergence to equilibrium per se.

Payoff Uncertainty

Suppose that, at the start of the repeated game, each player is privately informed of his or her stage game payoff function, which remains fixed throughout the course of the repeated game. Refer to player i's stage game payoff function as her *payoff type*. Assume that the joint distribution over payoff functions is independent (to avoid correlation issues that are not central to my discussion) and commonly known.

Each player can condition her behavior strategy in the repeated game on her realized payoff type. A mathematically correct way of representing this conditioning is via distributional strategies; see [54].

For any prior about player 2, now a probability distribution over player 2's distributional strategies, and given the probability distribution over player 2's payoff types, there is a behavior strategy for player 2 in the repeated game that is equivalent in the sense that it generates the same distribution over play paths. Again, this is essentially Kuhn's theorem.

Say that a player *learns to predict the play path* if her forecast of next period's play is asymptotically as good as if she knew the reduced form of her opponent's distributional strategy. This definition specializes to the previous one if the distribution over types is degenerate. If distributional strategies are in Nash equilibrium (also known in this context as a Bayesian Nash equilibrium), then, in effect, each player is optimizing with respect to a degenerate belief that puts probability one on her opponent's actual distributional strategy, and in this case players trivially learn to predict the path of play.

One can define LBR (see Subsect. "Kalai–Lehrer Learning") for distributional strategies and, as in the payoff certainty case, one can show that LBR implies convergence to Nash equilibrium play in the repeated game with payoff types.

More interestingly, Nash equilibrium play in the repeated game with payoff types implies convergence to Nash equilibrium play of the *realized* repeated game – the repeated game determined by the realized type profile. This line of research was initiated by Jordan in [42]. Other important papers include KL [45] (see Subsect. "Kalai–Lehrer Learning"), [41,44], and [62] (which studies recurring rather than repeated games).

Suppose first that the realized type profile has positive probability. In this case, if a player learns to predict the play path then, as shown by KL, her forecast is asymptotically as good as if she knew both her opponent's distributional strategy *and* her opponent's realized type. LBR then implies that actual play, meaning the distribution over play paths generated by the realized behavior strategies, converges to equilibrium play of the realized repeated game. For example, suppose that the type profile for matching pennies gets positive probability. In the unique equilibrium of repeated matching pennies, players randomize 50:50 in every period. Therefore, LBR implies that if the matching pennies type profile is realized then each player's behavior strategy in the realized repeated game involves 50:50 randomization asymptotically.

If the distribution over types admits a continuous density, so that no type profile receives positive probability, then convergence is more complicated. Suppose that players are myopic and that the realized stage game is like matching pennies, with a unique and fully mixed equilibrium. Given myopia, the unique equilibrium of the realized repeated game calls for repeated play of the stage game equilibrium. In particular, it calls for players to randomize. It is not hard to show, however, that in a type space game with a continuous density, exact optimization calls for each player to play a pure strategy for almost every realized type (this is a generalization of a point made in the context of fictitious play in Subsect. "Convergence in Classical Learning Models"). Thus, for almost every realized type profile in a neighborhood of a game like matching pennies, actual play (again meaning the distribution over play paths generated by the realized behavior strategies) cannot converge to Nash equilibrium play of the realized repeated game, *even if the distributional strategies are in Nash equilibrium*. Foster and Young [18] provide a generalization for non-myopic players.

This impossibility result is not robust to weakening optimization to ε optimization; the positive convergence results for ε-optimizing fictitious play provide one illustration, see Subsect. "Convergence in Classical Learning Models". A more subtle point, however, is that, even for exact optimization, a form of convergence obtains that, while weaker than convergence of actual play, is still very strong.

For simplicity, assume that each player knows the others distributional strategy and that these strategies form a (Bayesian) Nash equilibrium. Then to an outsider, for almost any type profile, observed play looks asymptotically like Nash equilibrium play in the realized repeated game (this follows from the main theorem in [62]). In particular, in a neighborhood of a game like matching pennies, for almost any type profile, observed play looks random. Since, by [18], actual play in this setting cannot converge to equilibrium, and in particular cannot be random, the implication is that convergence to equilibrium involves a form of purification in the sense of [34], a point that has been emphasized by [41] and [62]. To a player in the game, opponent behavior likewise looks random because, even though she knows her opponent's distributional strategy, she does not know her opponent's type. As play proceeds, each player in effect learns more about her opponent's type, but never enough to zero in on her opponent's realized, pure, behavior strategy.

Finally, the difficulties with characterizing sophistication in Bayesian learning models extend to models with payoff uncertainty, with learnability, richness, and consistency redefined in terms of distributional strategies; see [56].

Stochastic Learning

For a concrete example of a stochastic prediction rule, consider Young's variation on fictitious play [72]. In standard fictitious play, a player gets to observe the entire history of the game to date. In a bounded memory version of fictitious play, a player gets to observe only the last, say, s periods. In Young's variation on bounded memory fictitious play, a player observes only an $\ell < s$-period sample of the last s periods, with the sample drawn according to a probability distribution that assigns positive probability to every possible sample of length ℓ. If the player's forecast about period $t+1$ depends on her information in period t, then her prediction rule will be stochastic: her forecast will depend on the realization of the date t sample. With a stochastic prediction rule, a player at the beginning of the game cannot compute her own future forecast as a function of history.

A range of stochastic learning models have been studied, nearly all for myopic optimization. In some cases (e. g. [19]), the stochastic prediction rules have a quasi-Bayesian interpretation: most of the time, players optimize with respect to fixed prediction rules, as in a Bayesian model, but occasionally players switch to new prediction rules, implicitly abandoning their priors.

In this last part of the survey, I discuss two of the main themes in the stochastic learning literature. One is familiar: convergence to equilibrium. The other is new: calibration.

Convergence in Stochastic Learning Models

Although details can differ substantially from model to model, convergence results for stochastic learning models often have a similar feel. The stochastic nature of the prediction rule introduces an element of randomness into behavior. This randomness can cause players to stumble upon close-to-equilibrium play in the stage game and players, once near equilibrium, tend to stay near equilibrium, if not forever then at least for relatively long periods of time.

One thread of the literature on convergence in stochastic learning models poses specific models and investigates what sort of solution concepts emerge. In the case of Young's stochastic fictitious play model [72], the stochastic dynamics serve as a selection mechanism: learning tends to concentrate on some stage game Nash equilibrium but not others, even distinguishing between multiple strict Nash equilibrium, a phenomenon that also appears in many stochastic evolutionary models. In other stochastic learning models, convergence is sometimes not to Nash equilibrium but instead to sets containing Nash equilibria, such as CURB sets or minimal prep sets; see [40,49], and [64].

A second thread of the literature asks whether it is possible to construct *any* stochastic learning model that yields universal convergence to Nash equilibrium. The focus on stochastic prediction rules was motivated in part by the failure to find universal convergence results using deterministic prediction rules; see also Subsect. "Universal Convergence". The seminal paper on global convergence with stochastic prediction rules is [19]. For more recent work, see [31] and [38].

Calibration

In [17], Foster and Vohra showed that if one allows stochastic prediction rules then one can get approximate calibration on *every* realized history. That is, no matter what strategy the opponent plays, the prediction rule will approximately satisfy a statistical condition called calibra-

tion – and one can make the approximation to exact calibration as close as one likes. Calibration, in turn, implies a weak form of convergence to correlated equilibrium; see [16].

Calibration is a weak condition, so weak that it is possible for a prediction rule to be calibrated without it being accurate in any reasonable sense; see [47]. Nevertheless, calibration is sufficiently strong that a simple diagonalization argument shows that no deterministic prediction rule can satisfy it even approximately [63]. Foster and Vohra's discovery that one can get approximate calibration using stochastic prediction was startling.

Foster and Vohra's calibration result has spawned another literature, distinct from its application to learning, on whether it is possible for an observer to distinguish between an expert (someone who knows the true form of some stochastic process) and a charlatan (someone who uses the Foster and Vohra prediction rule, or one of its more elaborate cousins, to pass the observer's tests). Papers in this currently very active area of current research include [2,14,15,51,68].

Future Directions

Directions for future work, both for the learning literature and, more broadly, for the general literature on out-of-equilibrium dynamics in games, include the following.

Benchmark learning One of the goals of the literature is to establish a benchmark model of sophisticated learning. [57] sheds doubt on whether this is possible, but the desiderata considered in [57] are not based on axiomatic, decision theoretic criteria. The question of whether it is possible to build an axiomatic theory of sophisticated learning is unexplored.

Environments This survey has focused on repeated two-player games, but many other learning environments and patterns of interaction are possible. One can, for example, study environments in which players are linked through a network as in [9], or environments in which players are behaviorally heterogeneous. The possibilities are endless. The task is to enrich the scope of research without simply generating an ever expanding catalog of models, which leads to . . . .

Empirical testing One objective of the learning literature is to understand how real people behave in real games. Work tying learning models to actual behavior is underway [11] but still in its infancy. The appropriateness of a model may depend on the environment and in at least in some cases behavioral heterogeneity may be relevant.

Application An understanding of how players behave outside of equilibrium will enrich the application of game theoretic models. One example concerns institutional design, an instance of which is the choice of game for selling goods (a particular type of auction, for instance). Different institutions may have similar Nash equilibria but very different dynamic properties, and the dynamic properties could therefore play a role in the choice of institution.

Bibliography

1. Al-Najjar N (2007) Decision makers as statisticians. Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston
2. Al-Najjar N, Weinstein J (2008) Comparative testing of experts. Econometrica 76(3):541–559
3. Aoyagi M (1996) Evolution of beliefs and the Nash equilibrium of normal form games. J Econ Theory 70:444–469
4. Aumann R (1964) Mixed and behaviour strategies in infinite extensive games. In: Dresher M, Shapley LS, Tucker AW (eds) Advances in game theory. Annals of Mathematics Studies, vol 52. Princeton University Press, Princeton, pp 627–650
5. Basu K, Weibull AW (1991) Strategy subsets closed under rational behavior. Econ Lett 36:141–146
6. Benaim M, Hirsch M (1999) Mixed equilibria arising from fictitious play in perturbed games. Games Econom Behav 29: 36–72
7. Bernheim BD (1984) Rationalizable strategic behavior. Econometrica 52(4):1007–1028
8. Blackwell D, Dubins L (1962) Merging of opinions with increasing information. Ann Math Stat 38:882–886
9. Blume LE (1993) The statistical mechanics of strategic interaction. Game Econ Behav 5(3):387–424
10. Brown GW (1951) Iterative solutions of games by fictitious play. In: Koopmans TJ (ed) Activity analysis of production and allocation. John Wiley, New York, pp 374–376
11. Camerer C (2008) Behavioral game theory. In: Blume LE, Durlauf SN (eds) The new Palgrave dictionary of eocnomics, 2nd edn. MacMillan, New York
12. Cournot A (1838) Researches into the mathematical principles of the theory of wealth. Kelley, New York; translation from the French: Bacon NT (1960)
13. Dawid AP (1985) The impossibility of inductive inference. J Am Stat Assoc 80(390):340–341
14. Dekel E, Feinberg Y (2006) Non-Bayesian testing of an expert. Rev Econ Stud 73:893–906
15. Feinberg Y, Stewart C (2008) Testing multiple forecasters. Econometrica 76(3):561–582
16. Foster D, Vohra R (1997) Calibrated learning and correlated equilibrium. Game Econ Behav 21:40–55
17. Foster D, Vohra R (1998) Asymptotic calibration. Biometrika 85:379–390
18. Foster D, Young P (2001) On the impossibility of predicting the behavior of rational agents. Proc Natl Acad Sci 98: 12848–12853
19. Foster D, Young P (2003) Learning, hypothesis testing, and Nash equilibrium. Game Econ Behav 45:73–96

20. Freedman D (1965) On the asymptotic behavior of Bayes estimates in the discrete case ii. Ann Math Stat 36:454–456
21. Fudenberg D, Kreps D (1993) Learning mixed equilibria. Game Econ Behav 5(3):320–367
22. Fudenberg D, Levine DK (1993) Self-confirming equilibrium. Econometrica 61(3):523–545
23. Fudenberg D, Levine DK (1993) Steady state learning and Nash equilibrium. Econometrica 61(3):547–574
24. Fudenberg D, Levine DK (1995) Universal consistency and cautious fictitious play. J Econ Dyn Control 19:1065–1089
25. Fudenberg D, Levine DK (1998) Theory of learning in games. MIT Press, Cambridge
26. Fudenberg D, Levine DK (1999) Conditional universal consistency. Game Econ Behav 29:104–130
27. Fudenberg D, Levine DK (2006) Superstition and rational learning. Am Econ Rev 96:630–651
28. Fudenberg D, Levine DK (2008) Learning and equilibrium. Harvard University, Cambridge
29. Fudenberg D, Takahashi S (2007) Heterogeneous beliefs and local information in stochastic fictitious play. Harvard University, Cambridge
30. Fudenberg D, Tirole J (1991) Game theory. MIT Press, Cambridge
31. Germano F, Lugosi G (2007) Global Nash convergence of Foster and Young's regret testing. Game Econ Behav 60:135–154
32. Goeree JK, Holt CA (2001) Ten little treasures of game theory and ten intuitive contradictions. Am Econ Rev 91:1402–1422
33. Hahn F (1977) Exercises in conjectural equilibrium. Scand J Econ 79:210–226
34. Harsanyi J (1973) Games with randomly disturbed payoffs: A new rationale for mixed-strategy eequilibrium points. Int J Game Theory 2:1–23
35. Hart S, Mas-Colell A (2000) A simple adaptive procedure leading to correlated equilibrium. Econometrica 68(5):1127–1150
36. Hart S, Mas-Colell A (2001) A general class of adaptive strategies. J Econ Theory 98:26–54
37. Hart S, Mas-Colell A (2003) Uncoupled dynamics do not lead to Nash equilibrium. Am Econ Rev 93:1830–1836
38. Hart S, Mas-Colell A (2006) Stochastic uncoupled dynamics and Nash equilibrium. Game Econ Behav 57:286–303
39. Hofbauer J, Sandholm W (2002) On the global convergence of stochastic fictitious play. Econometrica 70(6):2265–2294
40. Hurkens S (1995) Learning by forgetful players. Game Econ Behav 11:304–329
41. Jackson M, Kalai E (1999) False reputation in a society of players. J Econ Theory 88(1):40–59
42. Jordan JS (1991) Bayesian learning in normal form games. Game Econ Behav 3:60–81
43. Jordan JS (1993) Three problems in learning mixed-strategy Nash equilibria. Game Econ Behav 5(3):368–386
44. Jordan JS (1995) Bayesian learning in repeated games. Game Econ Behav 9:8–20
45. Kalai E, Lehrer E (1993) Rational learning leads to Nash equilibrium. Econometrica 61(5):1019–1045
46. Kalai E, Lehrer E (1993) Subjective equilibrium in repeated games. Econometrica 61(5):1231–1240
47. Kalai E, Lehrer E, Smorodinsky R (1999) Calibrated forecasting and merging. Game Econ Behav 18(1–2):151–169
48. Kalai E, Samet D (1984) Persistent equilibria in strategy form games. Int J Game Theory 13:129–144
49. Kets W, Voorneveld M (2005) Learning to be prepared. Tilburg University, Tilburg
50. Kuhn HW (1964) Extensive games and the problem of information. In: Dresher M, Shapley LS, Tucker AW (eds) Contributions to the theory of games, vol II. Annals of Mathematics Studies, vol 28. Princeton University Press, Princeton, pp 193–216
51. Lehrer E (2001) Any inspection is manipulable. Econometrica 69(5):1333–1347
52. McKelvey R, Palfrey T (1995) Quantal response equilibrium for normal form games. Game Econ Behav 10:6–38
53. Milgrom P, Roberts J (1991) Adaptive and sophisticated learning in repeated normal form games. Game Econ Behav 3: 82–100
54. Milgrom P, Weber R (1985) Distributional strategies for games with incomplete information. Math Oper Res 10:619–632
55. Moulin H (1984) Dominance-solvability and cournot stability. Math Soc Sci 7(1):83–102
56. Nachbar JH (2001) Bayesian learning in repeated games of incomplete information. Soc Choice Welf 18(2):303–326
57. Nachbar JH (2005) Beliefs in repeated games. Econometrica 73:459–480
58. Nachbar JH (2008) Learning and evolution in games: Belief learning. In: Blume LE, Durlauf SN (eds) The new Palgrave dictionary of eocnomics, 2nd edn. Palgrave MacMillan, New York
59. Noguchi Y (2005) Merging with a set of probability measures: A characterization. Kanto Gakuin University, Tokyo
60. Noguchi Y (2007) Bayesian learning with bounded rationality: Convergence to ε-Nash equilibrium. Kanto Gakuin University, Tokyo
61. Nyarko Y (1994) Bayesian learning leads to correlated equilibria in normal form games. Econ Theory 4:821–841
62. Nyarko Y (1998) Bayesian learning and convergence to Nash equilibria without common priors. Econ Theory 11(3):643–655
63. Oakes D (1985) Self-calibrating priors do not exist. J Am Stat Assoc 80(390):339
64. Sanchirico CW (1996) A probabilisitc model of learning in games. Econometrica 64(6):1375–1393
65. Sandholm W (2007) Evolution in Bayesian games ii: Stability of purified equilibrium. J Econ Theory 136:641–667
66. Sandroni A (1998) Necessary and sufficient conditions for convergence to Nash equilibrium: The almost absolute continuity hypothesis. Game Econ Behav 22:121–147
67. Sandroni A (2000) Reciprocity and cooperation in repeated coordination games: The principled-player approach. Game Econ Behav 32(2):157–182
68. Sandroni A (2003) The reproducible properties of correct forecasts. Int J Game Theory 32(1):151–159
69. Shapley L (1962) On the nonconvergence of fictitious play. Tech Rep RM-3026 RAND
70. Voorneveld M (2004) Persistent retracts and preparation. Game Econ Behav 51:228–232
71. Voorneveld M (2004) Preparation. Game Econ Behav 48: 403–414
72. Young PH (1993) The evolution of conventions. Econometrica 61(1):57–84
73. Young PH (2008) Adaptive heuristics. In: Blume LE, Durlauf SN (eds) The new Palgrave dictionary of eocnomics, 2nd edn. Palgrave MacMillan, New York
74. Young PH (2008) Stochastic adaptive dynamics. In: Blume LE, Durlauf SN (eds) The new Palgrave dictionary of eocnomics, 2nd edn. Palgrave MacMillan, New York

Learning and Planning (Intelligent Systems)

Ugur Kuter
Institute of Advanced Computer Studies and Department of Computer Science, University of Maryland, College Park, USA

Article Outline

Glossary

Automated planning Automated Planning is the problem of generating a sequence of actions for an initial configuration of the world that, when executed, produces a final configuration that satisfies a specified set of goal conditions.

Planning knowledge Almost all of successful AI planning systems developed in recent years use some sort of *search-control knowledge* for effective planning. Sometimes the planning knowledge is specified in terms of *domain-independent heuristics*, i. e., heuristics intended for use in many different planning domains, and sometimes they are specified in terms of *domain-specific control knowledge* (i. e., tailored to a specific problem domain).

Learning for problem-solving and planning Learning for planning is the process of acquiring auxiliary knowledge related to a planning problem, which can be used by a planning system to understand better the underlying planning domain and to control its search for generating plans in that domain.

Definition of the Subject

The ability to produce knowledge about past experiences and exploit that knowledge in an operational context in later problem-solving and planning sessions is an important attribute of any intelligent system, human- or AI-based alike. Automated Planning and Learning is the research paradigm that focuses on the development of intelligent systems and technologies that combine the ability to make decisions and generate courses of actions (i. e., plans) with the capability to reason and produce knowledge about past experiences, future problems that the system needs to tackle, and strategies about how to tackle them.

Probably the first work that laid a formal treatment for this combination is the early planning system STRIPS [17], developed in early 1970s. The STRIPS planning system was an evidence that planning and learning are usually two pieces of an intelligent system, where the knowledge acquired via learning is used to enhance the problem-solving and planning experience later. Following its vision, many approaches have been developed in order to bring the techniques from the machine-learning community into planning. Most notably, *explanation-based learning (EBL)* and *inductive learning* techniques have been the most popular to contribute to AI planning.

The common objective of using learning for planning is to develop and use automated techniques to produce some sort of planning knowledge that is used to improve the performance of a planner. Sometimes the learned knowledge describes search strategies for "how to plan" (i. e., "what a planner should do in order to generate solutions"), and thus, guides the planner to focus only on the relevant parts of the planning problem and generate correct plans effectively. Sometimes the learned knowledge specifies the potential failures a planner could encounter during its search and their explanations. The planner uses such failure knowledge in order not to repeat similar mistakes over and over again. Other times, the planning knowledge is about the dynamics of the underlying environment; e. g., knowledge about the possible unexpected effects of the actions when executed in different circumstances. In all cases, the learned knowledge serves as a guidance to a planner in generating solutions to the planning problems.

Perhaps the biggest obstacle for developing planning and learning systems has been the so-called *utility problem*; that is, the trade-off between the time and resources spent for learning a sound and complete body of knowledge for planning and the cost of exploiting that knowledge during planning, which sometimes is shown to overshadow the effort spent in planning itself. Because of this problem, regrettably, the most important and remarkable developments in automated planning in recent years did not involve any learning component; instead, recent planning approaches include very efficient planning techniques that use controlled search with domain-specific

and/or domain-independent heuristics, constraint-satisfaction techniques for reasoning with time and resources, and model-checking based planning algorithms.

A key challenge for most of the recent planning approaches is that they require a domain expert to provide some sort of planning knowledge to the system, thus, replacing the functionality of a learning component by expert-level knowledge given as input to the planner. In many realistic planning problems, however, such planning knowledge may not be completely available; this is partly because it is very hard to compile such knowledge due to the complexities in the domains, e. g., evacuation and rescue operations, and it is partly because there is no expert to provide it, e. g., space operations. As the practicality of the planning, problem-solving and decision-making systems has been expanded to such complex, interesting, real-world challenges over the years, however, the fellowship of planning and learning has become once again important for the development of such systems.

Introduction

Traditional Artificial Intelligence (AI) planning, also known as *classical planning*, is the problem of generating a sequence of actions for an initial configuration of the world that, when executed, produces a final configuration that satisfies a specified set of goal conditions. The mathematical formulation of classical planning assumes that the environment must contain finitely many objects, and configurations of those objects describe the states of the environment. A classical planner always knows what is true and what is false in a state of the world. Planner's actions have deterministic outcomes; i. e., when executed, an action has a single and an instantaneous effect on the state of the world. Furthermore, the planner's actions are the only cause of change in the world. Thus, the world evolves in discrete and deterministic time steps when a plan (i. e., a sequence of deterministic actions) is executed.

It has been both theoretically and experimentally demonstrated that classical planning is a hard problem [14]. Over the years, some of the most impressive recent advances in classical AI planning are based on the use of "planning knowledge" for organizing the search space and guiding the search during planning. Sometimes the planning knowledge is specified in terms of *domain-independent heuristics*, i. e., heuristics intended for use in many different planning domains, and sometimes they are specified in terms of *domain-specific control knowledge* (i. e., tailored to a specific problem domain). In any case, this is usually in the form of auxiliary knowledge about the dynamics of the planning domain (i. e., semantic descriptions of possible actions and changes in the world), certain properties of the objectives of possible planning problems in that domain, and how to guide the planning process to solve those planning problems.

A key challenge for most of the successful planners is the requirement of a domain expert to provide the sort of background planning knowledge mentioned above. In many realistic planning domains, however, such planning knowledge may not be completely available; this is partly because it is very hard for the experts to compile such knowledge due to the complexities in the domains and it is partly because there is limited access to an expert to provide it. Here are several examples of such real-world planning domains and the currently-deployed planning techniques/systems in those domains:

- **Space Operations** Probably the most important success story of AI and AI planning technologies is their use in space operations. MAPGEN, the Mixed-initiative Activity Plan Generation tool, of NASA is a crucial part of the Mars Exploration Rover (MER) ground planning systems. The motivation for the MAPGEN system is to better support users that need to rapidly build activity plans that have to satisfy complex rules and meet the resource limits [2]. MAPGEN's planning component is based on the paradigm of constraint-based planning, where activities and states are described by predicate statements that hold over temporal intervals. The interval time-points and the predicate parameters are represented by variables connected by constraints. This approach supports a variety of complex planning constructs, including: activities with temporal durations, states that expire, exogenous events, complex constraints on parameters, temporal constraints linking activities and states, and subgoaling rules with conditions and disjunctions.
  Currently, MAPGEN uses mixed-initiative techniques (i. e., human-in-the-loop during the planning process) and automated reasoning methods in order to acquire its constraint-based planning knowledge. For the times the human expert is not available, learning such safety and operational constraints would be important and useful complement of the existing planning technologies in MAPGEN, since a planning system that can learn such knowledge to develop ways on how to safely and correctly operate in the world holds great promise for success.
- **Rescue and Evacuation Operations** Rescue and Evacuation operations [6,37,38] is the task of planning how to evacuate groups of people who may be in danger. Generally in these situations, there will be an incom-

plete description of the environment, in the form of standard requirements and operating procedures. Formulating an evacuation plan can be quite complex: typically there will be hundreds of tasks to be carried out. These tasks will depend on a wide range of factors: sources of danger, available resources, geography, weather predictions, political issues, and so forth. Complete information about the current state will never be available; the planning must include dynamic information gathering, and plans must be formulated with an incomplete world state. For such a problem, the planning must be done by a human expert or under the supervision of a human expert. It is unrealistic to expect that a planning system could produce good plans by itself, and awed evacuation plans could yield dire consequences.

Rescue and evacuation plans often require detailed analysis and knowledge about previous experiences. Learning lessons from past experiences of an evacuation planning system and from the information gathered from external knowledge sources will be crucial, especially when the human experts are not accessible and when the evacuation problem is too complex to be solved and managed by humans alone.

- **Vehicle Movement Planning** Planning the movements of vehicles, such as ships in naval operations and Unmanned Aerial Vehicles (UAVs) in surveillance missions, is important both commercially and militarily [11,19]. The state space and action space are effectively infinite: states include positions and velocities of vehicles, and actions correspond to movements of the vehicles along various routes. Since movements of different vehicles may occur concurrently, it is important to make sure they do not interfere with each other. The outcomes and durations of the actions cannot be known with certainty, because of factors such as weather, currents, and the behavior of the vehicles' operators. Elaborate simulation tools are available to aid in planning ship movements but the planning is still mostly done manually [1].

All of the domains above require some sort of planning knowledge to be used during the planning and execution process, to understand and exploit the underlying dynamics and properties of the environment in which the planning is taking place, and the ability to use past experiences in order to guide future decisions in solving problems. Over the years, several different machine-learning paradigms have been used to support automated planning. The common objective of using learning for planning is to develop and use automated techniques to learn some knowledge that is used to improve the performance of a planner. Many different techniques have been developed with this objective, including *learning macro-operators*, e. g., [5,36], *learning search control knowledge* [16,32,34], *learning of task hierarchies* [8,42,44] and *learning plan abstraction* [3,26].

Using explanation-based learning techniques have been investigated in mid-90s [25] in order to speed up the planning process where there is no planning knowledge available to give to the planner. These techniques mainly focused on observing the failures that occur during planning and generalize them into conditions that the planner should avoid during its search. These generalized conditions are then used for search-space pruning in the later stages of the planning process. Existing works on learning macro-operators (e. g., [5,36]) also falls in the category of speed-up learning, as do work on learning search control knowledge (e. g., [16,32,35]).

Recently, several architectures have been proposed to learn hierarchical planning knowledge from a collection of plan traces and from a given action model [8,42,44]. [42]'s X-Learn, for example, uses inductive generalization to learn task decomposition constructs, which relate subgoals and conditions for decomposing those goals into their subgoals. By grouping goals in this way, task models are learned that lead to significant speed-ups in problem-solving. As another example, [8,40] achieves the same objective via problem-solving techniques similar to explanation-based learning. Two other recent studies [23,51] proposed algorithms to learn the applicability conditions of task-decomposition constructs, given the hierarchical relationships between the tasks.

This paper describes some basic concepts in automated planning systems where machine learning has been of assistance in order to produce planning knowledge. We start with a summary of traditional automated planning paradigm of AI and an overview of analytical machine-learning methods that have been deemed suitable for integration with planning systems. Then, we present an analysis of four different forms of approaches that have integrated planning and learning techniques to achieve different objectives. We conclude with our remarks on the lessons learned in the Planning and Learning research so far and with some possible future directions.

Artificial Intelligence Planning, Traditionally

Traditional Artificial Intelligence (AI) planning, also known as *classical planning*, is the problem of generating a sequence of actions that, when executed in an initial configuration of the world, produces a final configu-

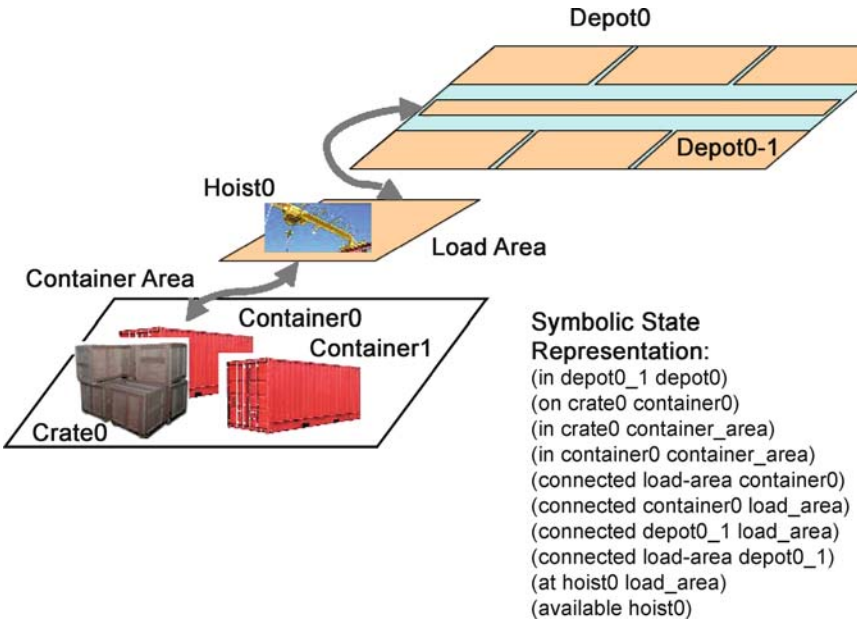

Learning and Planning (Intelligent Systems), Figure 1
A state description in the Simple-Storage domain

ration that satisfies a specified set of goal conditions. The mathematical formulation of classical planning assumes that the environment must contain finitely many objects, and configurations of those objects describe the states of the environment. More formally, a *state* is a collection of logical ground atoms. Classical planners usually make the *closed-world assumption*: the set of logical ground atoms in a state describe the facts that are true about the environment, whereas any logical atom that does not appear in the state is assumed to be false. Thus, a classical planner always knows what is true and what is false in a state of the world.

As an example, consider a simple transportation planning domain, a simple variant of a benchmark planning domain from the International Planning Competition (IPC) at 2006 [18], that we are going to use for illustration purposes here and throughout this paper. The original domain at the IPC is called *Storage*, and we will call our version as *Simple-Storage* in this paper. This domain involves moving a certain number of crates from some containers to some depots by a hoist. Each crate has a weight known to the planner. Each depot area can contain a number of crates in stacks, provided that in a stack, a heavier crate must always be placed lower than a lighter one. Inside a depot, each hoist can move according to a specified spatial map connecting different areas of the depot. Once the hoist is holding a crate for delivery and moved into a depot, it cannot move outside of the depot without dropping the crate.

Figure 1 illustrates a snapshot state of the world in this domain. This state specifies a depot and several storage areas in it, a hoist, a crate, two containers, and the spatial map of the depot areas. In particular, the state describes what is true in the world at this instance; e. g., the *crate*0 is in the *container*0, there is an empty slot in *depot*0 at area *depot*0–1, the *hoist* is empty and available for use in the loading area, and so on. It also specifies information about what is not true in the world; e. g., we know that *crate*0 is not in *container*1, we know that the hoist is not in any depot or container area, and so forth.

A classical planner's actions have always deterministic outcomes; i. e., when executed, an action has a single and an instantaneous effect on the state of the world. Actions are specified via an *action schema*, or equivalently, through a set of *planning operators*. A planning operator is an expression of the following form (*head, preconditions, effects*), where *head* is an expression of the form $(name, arg_1, \ldots, arg_k)$ such that *name* is a symbol denoting the name of the operator and $arg_1, \ldots, arg_k$ are the *arguments* of the planning operator: each argument a_i is either a variable or a constant symbol. The *preconditions* and *effects* of the operator are specified as logical formulas that can contain both positive and negative literals. A negative literal in the effects of the operator specifies that the literal will become false in the state of the world after the operator is applied (i. e., the atom will be deleted from the state). A positive literal in the effects becomes true in the world after the operator is applied.

For example in our *Simple-Storage* domain, we will specify five different planning operators for this planning domain: a pick-up operator for lifting a crate by a hoist, a put-down operator for dropping a crate by a hoist, a move operator for moving a hoist inside and outside a depot, and a locate operator for moving a hoist from one area of a depot to another one. The following is an abstract description of the pick-up operator in a Lisp-like notation:

operator for **pick-up**
:parameters (?c – crate ?a1 – area ?a2 – area ?p – place)
:precondition (and (connected ?a1 ?a2) (at ?h ?a2)
(available ?h)(on ?c ?a1) (in ?a1 ?p))
:effect (and (not (on ?c ?a1)) (clear ?a1)
(not (available ?h))(holding ?h ?c) (not (in ?c ?p))) .

An *action* is a ground instance of a planning operator, where all of the variable symbols that appear in the operator definition are replaced by constant symbols. We will denote the preconditions and effects of an action a by Pre(a) and Eff(a), respectively.

An action a is applicable in a state s, if the preconditions of a are satisfied in s; i. e., if $s \models \text{pre}(a)$. The *result* of applying an action a in a state s, result(s, a), is the state s' that is produced by removing every negative literal in Eff(a) from s and by adding every positive literal in Eff(a). A *plan* is a sequence (i. e., totally-ordered set) of actions. The *result* of applying a plan $\langle a_1, \ldots, a_k \rangle$ in a state s is the state $s' = \text{result}(\text{result}(\ldots, \text{result}(\text{result}(s, a_1), a_2), \ldots), a_k)$. We denote this by $s' = \text{result}(s, \langle a_1, \ldots, a_k \rangle)$ for simplicity of the notation.

As the reader may have noticed already, the planner's actions are the only cause of change in the world dynamics of classical planning. Also, it follows from the above definitions that classical planning assumes that the world evolves only in discrete and deterministic time steps when a plan (i. e., a sequence of deterministic actions) is executed.

A *classical planning problem* is defined as a tuple (s_0, G, O) where s_0 is the initial state (a symbolic representation of the state of the world at the time that the plan executor will begin executing its plan), G is a set of goal states (a symbolic representation of the desired state or states the executor will end executing the plan), and O is the set of planning operators. A *solution* to a classical planning problem (s_0, G, O) is a plan $\langle a_1, \ldots, a_k \rangle$ such that $\text{result}(s_0, \langle a_1, \ldots, a_k \rangle) \in G$.

Solution methods for solving classical planning problems are categorized as *state-space search* and *plan-space search* techniques:

State-Space Planning State-space search techniques perform a search over the state space. If the search is done in a forward-chaining fashion, then the planner starts from the initial state of the planning problem, and it successively generates new states by applying actions, until a goal state is reached. In backward-chaining state-space search, the planner starts from a goal state and applies successively the inverse of the actions until it reaches to the initial state.

Plan-Space Planning In plan-space planning, the planner starts with the initial plan, and successively adds new actions and fixes the flaws in the plan, until no flaw is left and the plan is a solution to the current planning problem. The initial plan consists of only two special actions, start and finish. The start action does not have any preconditions and its effects encode the initial state of the planning problem. The finish action does not have any effects and its preconditions specify a logical condition that encodes the goals of the planning problem. Any precondition that is not established in a partial plan is called a flaw, and the planner successively enhances the initial plan by adding new actions whose effects establish the preconditions of the existing actions in the plan. This process continues until no flaws left in the plan, which then encodes the solution of the input planning problem.

Several variants of state-space and plan-space search paradigms have been developed in automated planning over the years. A key issue with any planning algorithm that uses either of the search paradigms is that the amount for the search that the planner performs could be huge, and therefore, the planner could take a lot of time before it returns a solution plan. Some of the most impressive advances in the recent years are based on various techniques for organizing the search space and guiding the search. Sometimes, these techniques are *domain independent*, i. e., they can work in any planning domain, and sometimes they are *domain-specific*, i. e., they exploit characteristics of the underlying planning domain. When *domain-specific planning knowledge* is used for planning, sometimes the planners themselves are domain-specific, i. e., they work only in a particular domain such as process planning [21] or bridge [46]. In other cases, the planner is *domain-configurable*, i. e., it consists of a domain-independent planning engine, plus a language for writing domain-specific problem-solving information.

For modeling and learning structured planning knowledge about a problem domain, one of the best-known formalisms is *Hierarchical Task Networks (HTNs)*. An HTN planner formulates a plan by decomposing tasks (i. e., symbolic representations of activities to be performed) into smaller and smaller subtasks until tasks are reached that can be performed directly. The basic idea was devel-

oped in the mid-70s [45,47], and the formal underpinnings were developed in the mid-90s [13].

An *HTN planning problem* description consists of the initial state and the goal task network (a set of tasks to be performed, along with some constraints over those tasks that must be satisfied). A solution to an HTN planning problem is a *plan*, as in classical planning; i. e., a sequence of actions that, when executed in the initial state, perform the desired tasks.

In order to generate solutions for the planning problems, an HTN planner uses an *HTN domain description* that contains mainly two kinds of knowledge artifacts: *methods* and *operators*.[1] The operators are like the planning operators used in any classical planner. The action symbols of these operators are designated as *primitive tasks* (i. e., tasks that we know how to perform directly). Any task that does not correspond to an operator name is a *non-primitive* task.

Each method is a prescription for how to accomplish a non-primitive task by decomposing it into subtasks (which may be either primitive or non-primitive tasks). A method consists of four elements: (1) the task that the method can be used to accomplish, (2) the set of preconditions which must be satisfied for the method to be applicable, (3) the subtasks to accomplish, (4) any constraints on the bindings of objects to variables in the above and on the orderings of the subtasks to be accomplished.

For example, consider the task of moving a collection of crates from their containers to their designated depot areas. Here is an informal description of a method we might use for the task move_crate in the Simple-Storage domain:

method move_crate_by_hoist for the task move_crate:
- *applicability conditions:* the hoist is in working order and is present at the location of the container, the crate does not have any other crate on top of it
- *subtasks:* (1) pick up the crate, (2) move the hoist to the designated depot, (3) locate the depot area, and (4) put down the crate
- *constraints:* do subtask (1) before subtask (2), do subtask (2) before subtask (3), and do subtask (3) before subtask (4).

Machine Learning Methods in Planning

The most commonly investigated machine-learning paradigms that were used in automated planning systems are *explanation-based learning* and *inductive learning*. We will overview these paradigms below; the subsequent sections will discuss some examples of their integration with planning and describe why these paradigms have been the most popular among researchers so far.

Explanation-Based Learning (EBL) EBL is an *analytical learning* method, where an explanation-based learner uses prior, common-sense background knowledge about the problem-solving domain and deductive reasoning over that background knowledge to augment the information originated from a certain number of input training examples. The background knowledge in EBL helps the learning process in several ways, described below.

First, the deductive reasoning in EBL over the background knowledge and the training examples specifies those features of the hypothesis space that are relevant to the underlying learning problem. The objective of a learning problem is to derive an explanation (i. e., a logical proof) of a given target concept. For example, a target concept in our Simple-Storage planning problem described earlier could be a simple or complex logical sentence such as "it is safe to place the crate x on top of the crate y in the depot", or in a space-operations domain, it could be "any day without any sandstorms on Mars is good for the rover Opportunity to explore the Victoria's Crater".

Given the target concept (i. e., the hypothesis concept), an EBL algorithm attempts to generate a logical proof of that concept using the training examples and the background theory. The background domain theory specifies any additional, general and common-sense background knowledge about the underlying problem. For example, the background theory for stacking containers above might include rules as follows:

$$SafeCheck(x, y) \leftarrow Crate(x), Crate(y), Stackable(x, y).$$

$$Stackable(x, y) \leftarrow Weight(x, wx), Weight(y, wy), wx < wy.$$

A training example would most naturally be modeled as the values of certain attributes of domain and the target concepts; for example, in the Simple-Storage domain, a training example would contain information about the weights of the crates such as $Weight(Crate\,0,10)$ and $Weight(Crate\,1,20)$. From this information, the EBL system deduces that $SafetyCheck(Crate\,0, Crate\,1)$ is true.

[1] HTN planners may have other kinds of knowledge artifacts as well. For example, the SHOP2 planner [39] has axioms that can be used to infer conditions about the current state.

If there is no logical proof of the target concept generated by deductive reasoning over the background theory and the training examples, then the algorithm discards the hypothesis concept. Otherwise, each concept that appears in the proof is relevant to the hypothesis and each concept that does not appear there is not relevant. Note that, with this deductive proof-generation capability, an EBL algorithm can discover relevant concepts to the target that cannot be generated from the training examples alone, which is an important advantage of EBL over other learning paradigms as this allows EBL to generate more general hypotheses to learning problems.

Deductive reasoning over the input background knowledge in EBL also allows a learning algorithm to focus on only the relevant portions of the hypothesis space, ignoring the rest. As a result, the complexity of the hypothesis space and the learning process is usually much less than that in other learning paradigms that cannot exploit such background knowledge.

A key drawback of EBL that surfaces in some cases is that the background knowledge must be complete and perfect in order for EBL to work and generate correct hypotheses as described above. The reason for this is that the deductive proof-generation in EBL cannot differentiate between perfect and imperfect knowledge, and even if it could, the deductive proof-generation in EBL would not be able to handle it. Thus, if the input knowledge contains errors or it is incomplete, then there is no guarantee that the hypotheses found by EBL algorithms will be correct solutions to the learning problems. There has been some attempts to alleviate this issue with EBL, and we are going to touch those works below when we are discussing *Inductive Learning* methods.

Inductive Learning Inductive learning focuses on the problem of producing (i. e., inducing) general hypotheses from specific training examples. Inductive learners generalize from the observed examples by identifying features that empirically distinguish positive examples from negative ones. Version-space learning, decision-tree learning, inductive logic programming, neural networks, and genetic algorithms are all examples of the inductive learning paradigm.

The primary difference between inductive learning and explanation-based learning is the use of background knowledge in the latter. In inductive learning, a learner selects a hypothesis from the given hypothesis space that fits best to the input training examples $(\langle x_1, f(x_1)\rangle, \ldots, \langle x_n, f(x_n)\rangle)$ in order to model an unknown target concept f. For example, consider our target concept *safe stacking* from above. In addition to this target concept, we must also specify all of the relevant attributes in the planning domain for an inductive learner to learn this concept. In this case, let's suppose the following is the set of those relevant attributes:

- *Holding*: whether the hoist is holding a crate or not;
- *WeightTop*: whether the crate to be placed on top has light, medium or heavy weight;
- *WeightBottom*: whether the crate to be placed at the bottom has light, medium or heavy weight;

Some training examples for inductively learning our safety-stacking concept from the above attributes will be as follows:

Holding = *True*, *WeightTop* = *Light*,
WeightBottom = *Heavy*, *SafeStack* = *True*
Holding = *True*, *WeightTop* = *Light*,
WeightBottom = *Medium*, *SafeStack* = *True*
Holding = *True*, *WeightTop* = *Heavy*,
WeightBottom = *Medium*, *SafeStack* = *False*
Holding = *True*, *WeightTop* = *Heavy*,
WeightBottom = *Light*, *SafeStack* = *False*
Holding = *True*, *WeightTop* = *Heavy*,
WeightBottom = *Heavy*, *SafeStack* = *True*

From these examples, an inductive learner will produce the following definition of the SafetyStack concept:

Holding = *True* AND
IF *WeightTop* = *Light* THEN *SafeStack* = *True*; ELSE
IF *WeightTop* = *Heavy* AND *WeightBottom* = *Heavy*
THEN *SafeStack* = *True*; ELSE
IF *WeightTop* = *Heavy* THEN *SafeStack* = *False*;

The reasoning behind this learned concept is as follows. First of all, in every training example the learner has seen *Holding* is *true*, and therefore, it appears as such in the definition of the target concept. The first and second examples specify the case in which $WeightTop = Light$ for two different values for *WeightBottom*. Thus, the inductive learner generalizes from the value of the latter attribute for the case *WeightTop* = *Light* and *SafetyStack* = *True*. The third and the fourth examples causes the same type of generalization for the case we have *WeightTop* = *Heavy* and *SafetyStack* = *False*. Finally, the remaining if-condition in the above learned concept is originated from the last example.

As learning is based on only the input training examples, the learned hypothesis can only be an approximation of the target concept in inductive learning. The

learning algorithms rely on training examples that are sufficiently many and rich in content, and lacking any further information or rich training samples, they assume that the best hypothesis regarding the unseen examples is the hypothesis that best fits the observed examples. This strong assumption about the training examples is, in some cases, a key practical limit on inductive learning algorithms.

In addition to the above machine-learning approaches, there are two other approaches that have been incorporated with automated planning. Case-based learning is an *instance-based learning* method, where given some training examples, a CBL algorithm simply stores the training examples as *cases* in order to be retrieved, adapted, and used in later problem-solving sessions. CBL uses complex, rich, and relational symbolic representations to represent its cases regarding the training examples. The retrieval and combination of cases to solve to the current training example may rely on knowledge-based reasoning and search-intensive problem-solving methods. Given such complex logical descriptions of already-seen examples, a CBL algorithm attempts to map those descriptions to a newly-seen example and generate an explanation/solution for that example. Note that the methodology of CBL is in contrast to explanation-based and inductive learning, which produce a general approximation of a target concept/hypothesis from training examples and sometimes background knowledge. Generalizing beyond the already-seen examples are postponed in CBL until new ones are received.

Reinforcement Learning (RL) emerged as another paradigm where the learning and the planning processes are modeled under the same framework (rather than a combination of the two in some fashion), with a focus on non-classical planning domains where actions may have multiple, non-deterministic outcomes and the objective is to produce a mapping from states of the world to actions that optimizes some utility function. We will come back to RL later in this paper in Sect. "Reinforcement Learning" and discuss it in more detail as a separate topic, as it is a quite different planning and learning approach than the way above learning paradigms have been usually investigated in the context of automated planning.

Learning Domain Models

An important assumption of traditional automated planning is that the planning algorithms require complete and correct descriptions of all of the possible planning operators in a planning domain; otherwise, the plan they generate does not guarantee to be a correct solution to planning problems. In many practical problems, however, it may not be possible to compile the complete models of the actions due to the uncertainties in the world, incomplete information about the semantics of the actions, and other practical reasons. Thus, it is important to have the ability to learn action models in order to produce complete semantics for them and plan with them.

The primary approach for learning action models is by observation. Intuitively, a learner is demonstrated information about the semantics of the actions (i. e., about the preconditions and effects), allowed for generalizing from that information, and practice with the generalized knowledge to discover mistakes and correct hypothesis. This process continues until the learned action models provide failure-free planning for the planning problems in a domain.

An observation for an action model can be understood in terms of the state-transitions induced by that action: let s be a state in which the action a is applicable and suppose applying a in s produces s'. Then, the tuple (s, a, s') is an observation of the action model; i. e., the states s and s' give information about the preconditions of the action and its effects. Note that we are not assuming anything about the action here; we do not necessarily know a's preconditions and effects; we only observed that a was applied in s and produced s'.

The observations over the action models could be given in terms of solution plans as training examples, or in terms of single-action executions in the world. The first time an observation is obtained on the action a, the learner makes the entire state s as the preconditions of a and the entire state s' as its effects. If the learner is provided with a domain theory that specifies information about what is not true in those states, it is also possible for the learner to exploit that theory to hypothesize negated preconditions and effects for a as well. As the learning procedure sees more observations for the action a, it generalizes from those initial preconditions and effects.

At any point before the learning finishes, the acquired planning operators are not necessarily complete or correct in the following ways. First, the learned preconditions for a planning operator could be *over-specific* or it could be *over-general*. In this case, the preconditions of the planning operator will have redundant and perhaps incorrect conditions and a planner that uses this operator will have to achieve all of those unnecessary preconditions during planning. In some cases, i. e., if the planner a plan-space planner, this does force the planning system to do unnecessary search by attempting to insert actions to establish those preconditions and so on. The entire process is deviated from the actual objectives of the planning problem in hand. In a state-space planner, the unnecessary

conditions may make the operator inapplicable in certain states, which again, induces unnecessary search effort. Furthermore, unnecessary preconditions can make many solvable problems unsolvable because it may not be possible to achieve some of them.

If the preconditions are over-general, then this increases the branching factor of the search space of the planner (since a planning operator may become applicable in a state that it should not), degrading the performance of the planner significantly.

When some effects of an operator are not learned, then a planner generates incorrect states during its search for a plan. This makes the later decisions on the applicability of the operator even more problematic, since, for example, an operator that should not be applicable in a state may become applicable and incorrect states with over-general preconditions of operators makes the branching factor even larger.

As a result of the aforementioned factors, a learning algorithm that produces action models by observation usually require several examples of such models in order to refine its learned knowledge.

Learning Search-Control Knowledge

As we discussed in Sect. "Artificial Intelligence Planning, Traditionally", a planning algorithm is essentially a search procedure over a space of planning constructs, such as world states in state-space planning and partially-ordered plans in plan-space planning. At each point during its search, the planner must make some decisions and choose among many alternative ways of how to expand its search. Sometimes those choices leads to failures and the planner backtracks in order to try other alternatives. Other times, the choices lead to solution plans; sometimes, those choices produce good quality plans very efficiently, sometimes they are not efficient, and sometimes they result in poor solutions.

Each decision that the planner makes among alternative choices during the search is a learning opportunity to enhance the planner search-control functionality. If a planner is able to understand and learn good choices against bad ones and its mistakes that result in failures, then the planner does not necessarily repeat those mistakes in later problems, if equipped with sufficient learning and planning capabilities. In the following, we are going to discuss such learning and planning capabilities for controlled search during planning.

Explanation-Based Methods

Probably the most common approach to learning control knowledge is to use analytical learning techniques where the systems learns by problem solving, and the most popular analytical method is Explanation-Based Learning (EBL). As we described in Sect. "Machine Learning Methods in Planning", an EBL system learns by analyzing explanations of problem-solving behaviors; in particular, the system builds an explanation for why an example is a member of some target concept. When used to learn planning control knowledge, the example usually corresponds to a decision made by the planner during its search and the target concept is explaining why the decision was a good one or a bad one. The explanation built by EBL can be used to construct a control rule which can then be used to help make future planning decisions that were similar to the original example.

During planning, a planner's problem-solver goes through several critical decision points. In a plan-space planner, for example, these decision points correspond to the choice of which causal links should be established and how, and threat selection and resolution (including promotion, demotion, and separation). In a forward-chaining state-space planner, those decisions correspond to the choice of an action applicable to the current state among all of the applicable ones.

The most common learning opportunity during planning is when the system reaches an impasse; i. e., when the planner reaches to a decision point, but fails to make one. This may happen when there are no applicable actions in the current state during the search or there is no way to resolve a threat on a causal link in a partial plan. At such an impasse, the system that exploits an EBL component can attempt to generate an explanation of the failure that produced the impasse. The failures depend on the type of planning being done in the system. When a plan-space planner reaches an impasse, the failure(s) can be explained in terms of the inconsistency of the ordering and binding constraints associated with the partial plan, or in terms of an unsatisfiable open condition. In a state-space planner, the failures could be understood in terms of the unsatisfiable preconditions of the inapplicable action(s).

As an illustration of how search failures are used as explanations, consider the search tree of a forward – chaining state – space planner as shown in Fig. 2. The search node that is marked as a red box in this tree is the point of impasse, since there are no actions applicable in that state. In our Simple-Storage planning problem, for example, the red box could represent a state that specifies a crate that is in a depot area, another crate that is held by the hoist, and the hoist is in that particular depot area. In this configuration, suppose the crate on the hoist, say crate2, is heavier than the one in the depot area . Therefore, the action (put-

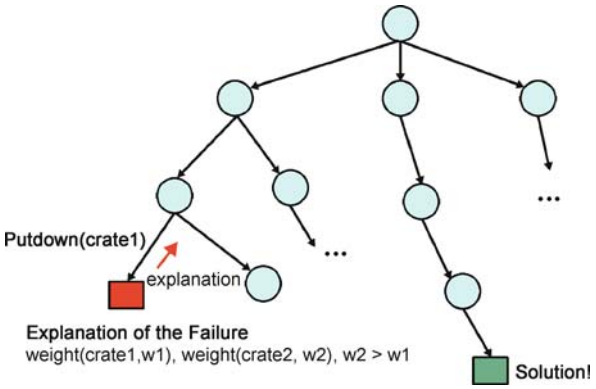

Learning and Planning (Intelligent Systems), Figure 2
The example of a search tree generated by a forward-chaining state-space planner

down crate2) is not applicable. The explanation for this failure can be understood by

$$weight(crate\,1, w\,1), weight(crate\,2, w\,2), w\,2 > w\,1\ .$$

While backtracking, the planner propagates back this explanation to the parent node of the failed one and through that node, the rest of of the search node to be expanded. When this explanation is generated and propagated, the system uses this explanation as a search-control rule and it never attempts to move the first crate on top of the second in the later search process.

Learning from failures in the way described above is the simplest form of using EBL for producing planning control knowledge. Another possible source for learning during planning is the successful search traces that a planner generates. For example, consider again the search tree shown in Fig. 2. The decision points in the search trace that leads to an optimal solution to the input planning problem, as marked in the figure, could be used for learning control rules that prefers these decisions or their generalizations over the other search traces in the figure. In order for an EBL component to learn somewhat more effective control rules than the sort described above, the system must be provided a domain theory that includes rules for deriving relevant features of the planning problems and the search process, as well as templates for the control rules to be learned by the system.

As an example, suppose the domain theory for the Simple-Storage domain includes the following control-rule template:

$$\textbf{select action}\ (\text{move}, ?hoist, ?from, ?to) \leftarrow COND_1, \ldots, COND_k\ ,$$

where each $COND_i$, $i = 1, \ldots, k$, specifies a template for any logical formula that can be evaluated given a state and any combination of concepts that cannot directly inferred from the state, but could be learned using the domain theory. For example, suppose the planner successfully applied the action move(hoist0, load-area, depot0-1-1) in a state consists of the following:

$$\begin{aligned}
&in(depot\ 0_1, depot\ 0), in(crate\ 1, depot\ 0_1),\\
&\quad holding(crate\ 0)\\
&in(container\ 0, container_area),\\
&\quad connected(load_area, container_area)\\
&connected(container_area, load_area),\\
&\quad connected(depot\ 0_1, load_area)\\
&connected(load_area, depot\ 0_1), at(hoist\ 0, load_area)
\end{aligned}$$

Then, with the above domain theory, an EBL component of the system will learn a control rule that is similar to the following:

$$\begin{aligned}
&\textbf{select action}\ (\text{move}, hoist\ 0, load_area, depot\ 0_1)\\
&\quad \leftarrow holding(crate\ 0), at(hoist\ 0, load_area),\\
&\quad\quad at(crate\ 1, depot\ 0_1),\\
&\quad\quad SafeStack(crate\ 0, crate\ 1)\ ,
\end{aligned}$$

where the definition of the concept *SafeStack*(*crate* 0, *crate* 1) is learned via the domain theory given in Sect. "Machine Learning Methods in Planning".

The search-control rules using an EBL approach as above can be learned either during planning or after planning. In the former case, the planner propagates its explanations of failures while backtracking to the upper levels of the search tree and forms the control rules based on the domain theory during this propagation. The control rules generated in this fashion could then be used in the later branches of the search tree and in the other planning problems in the same planning domain. In the latter case, the planner generates the entire search tree that includes both the failed and successful search traces and post-processes the generated tree in order to derived the search-control rules. These rules are then used in other planning problems. Learning via backtracking has the advantage of not requiring to wait until the entire search tree is generated: the learner could produce rules in an incremental fashion and use those rules as they are generated during planning. Learning after planning, on the other hand, cannot produce and use any control rules until the search tree is constructed for an input problem. However, as the search tree contains more information in terms of comparisons between failed and successful decisions and optimal and sub-optimal search traces, which is not usually available during a backtracking search, this approach can learn more expressive and powerful control rules.

Unfortunately due to its reliance on an human expert specified domain theory and only a few examples, standard EBL can often produce complex, sometimes overly-specific and sometimes overly-general control rules that cannot be utilized very well in new planning problems [4,32]. This situation is commonly known as the utility problem, where even though the learned rules are correct, the cost of testing their applicability to new planning situations often outweighs their savings. Another problem is that EBL methods are difficult to apply in domains where it is hard to construct a complete and tractable domain theory. These drawbacks usually make it difficult to successfully apply EBL methods to real-world problems.

Inductive Methods

Inductive learning techniques can acquire planning control rules by looking at examples of positive and negative planning decisions. These examples are usually found by analyzing a solution to a planning problem and by solving a set of training problems and extracting examples based on good and bad planning decisions. Positive learning examples for inductive methods are those planning decisions that a hypothesized search-control rule correctly classifies – i. e., the control rule should not eliminate the planning decisions that yield to solutions and guides the planner to take that decisions. Negative examples, on the other hand, are those planning decisions that does not yield to a solution but enforced by the control rule on the planner. In the former case, if the control rule eliminates a good decision, then this means that the rule is over-specific: in that case, an inductive learner generalizes the rule in order to classify all positive examples. In the latter, the case the negative example describes a case where the rule is overly-general since it accepts a decision it is not supposed to. In this case, the learning component specialized this rule by adding extra constraints while correctly classifying all positive and negative examples.

The difference between inductive methods and EBL is that the latter is a deduction based method: EBL methods use a expert-provided domain theory in order to produce search-control knowledge as explained above. Inductive methods, on the other hand, build control rules by performing an inductive search through the space of possible rules, where the objective is to produce search-control knowledge that cover some positive planning decisions but do not cover any negative planning decisions.

The inductive search is usually guided by a domain independent bias that helps the learning component to choose between several search-control rules that applies to the same or very similar planning decisions. For example, both of the following rules describe the situations to the planner in which it should prefer a particular crate over another:

1. **select action** (move, *hoist* 0, *load_area*, *depot* 0_1)
   $\leftarrow$ *holding*(*crate* 0), *at*(*hoist* 0, *load_area*),
   at(*crate* 1, *depot* 0_1),
   SafeStack(*crate* 0, *crate* 1).
2. **select action** (move, *hoist* 0, *load_area*, *depot* 0_1)
   $\leftarrow$ *holding*(*crate* 0), *at*(*hoist* 0, *load_area*),
   at(*crate* 1, *depot* 0_1),
   in(*container* 0, *container_area*),
   connected(*load_area*, *container* 0),
   connected(*container* 0, *load_area*),
   connected(*depot* 0_1, *load_area*),
   connected(*load_area*, *depot* 0_1),
   SafeStack(*crate* 0, *crate* 1) .

The first rule is much more simpler than the second one since the latter includes some redundant features of the decision on the crate, which can also be inferred from the first rule. One very-commonly used bias is Occam's razor which prefers simple rules such as the first one above over more complex ones.

A main advantage of inductive learning techniques over other learning methods such as EBL is that they tend to learn very general control knowledge since they examine a number of different examples. This quality is important for building useful general control knowledge, which can be more effectively applied to new planning problems. Unfortunately, inductive techniques have several disadvantages as well. Perhaps, most importantly, inductive methods usually require large numbers of examples to acquire effective control information. Also, these methods can quickly become computationally intractable since they often search through large amounts of information when building control rules. These disadvantages often prevent inductive methods from being successfully used for planning problems.

However, inductive learning of search-control knowledge can be combined with EBL methods to develop systems that could address the issues with the both approaches when exploited independently. One way to combine the two approaches is to use EBL methods to generate a first draft of search-control knowledge. As we discussed above, EBL methods usually produce very specific search-control knowledge given the expert-provided domain theory. Furthermore, the learning mostly depends on this input domain theory, which might be incorrect and in-

complete. In such cases, EBL methods could produce overgeneral knowledge, or even incorrect control rules. However, this knowledge can then be refined that knowledge via inductive generalization and specialization over the set of training examples.

Learning Operational Planning Procedures

The techniques described in the previous section enhance the performance of a planner by learning search-control knowledge about what to do and what not to do during a problem-solving session and using that knowledge on later problems in order to organize the search space and not to do similar mistakes all over again. Despite the effectiveness of these techniques, in many domains, the performance of a planner can be improved even more greatly by inferring and exploiting operational knowledge about the domain structure that is not explicitly encoded in the formulation input planning problems (i. e., it is not encoded in terms of the planning operators, initial and the goal state).

As many human skills are procedural and hierarchical in nature, with complex procedures defined in terms of simpler ones, it is usually preferable to describe the structural domain planning knowledge in a form that conforms to a human's understanding of the planning problems and domains. In the simplest form, such knowledge is described in terms of *macro-operators*, i. e., sequences of planning operators (i. e., actions) that are grouped together since all or most solution plans in the planning domain require those actions to be executed in that sequence. More expressive planning knowledge is described in multiple hierarchies and usually in terms of HTN methods discussed in Sect. "Artificial Intelligence Planning, Traditionally". In the rest of this section, we are going to present these two forms of of domain planning knowledge and describe basic methods for how to learn such knowledge and use in planning.

Learning Macro-Operators

Informally, a *macro-operator* is perhaps the simplest form of planning knowledge about how a planning domain works and what characteristics that the planning problems in that domain exhibit in terms of the actions they contain. A macro-operator specifies, therefore, a grouping of the standard planning operators such that the operators (or more precisely, the actions that are the ground instantiations of those planning operators) are executed together in all or most of the solution plans in the domain.

Let O be a set of planning operators in a classical planning domain. Then, formally, a *macro-operator* is a tuple of the form $(O', \prec, \sigma)$ where $O' \subseteq O$ is a multi-set of classical planning operators, $\prec$ is a partial-ordering relation over the operators in O', and σ is a variable-binding substitution over the variables that appear in the operators of O'. Note that the set of planning operators O' may contain more than one instance of the same planning operator.

A planning process that involves macro-operators has usually two components: learning a set of macro-operators for a domain, and afterwards, using them during planning. The solution methods for learning macro-operators first generate solutions to a number of planning problems in the underlying domain, and then, process the structural properties of those solutions in order to infer correct and useful operator groupings as macros. The simplest way of realizing this learning approach is to identify patterns of actions in those solution plans. For example, consider our example Simple-Storage planning domain and suppose we have the following plan:

$$\langle move(hoist\,0, depot\,0\_1, load\_area), move(hoist\,0, load\_area, container\,0), pickup(crate\,0, container\,0, container\_area), move(hoist\,0, container\,0, load\_area), move(hoist\,0, load\_area, depot\,0\_1), putdown(crate\,0, depot\,0\_1, depot\,0)\rangle .$$

In two places in this plan, there are sequences of move operations for moving the hoist between the load area, depot, and the container. This pattern is a good candidate as a macro-operator because we could expect that any solution plan in this domain would involve a hoist performing similar move operations between a load area, a container, and a depot.

A more complicated and better approach would be to use structural semantic information in those example plans in order to learn the macro-operators. The basic structural semantic property in all planning domains is the causal dependencies between the actions in the solutions. A causal dependency is the relationship between two actions in the plan, where an effect of one of the actions establish a precondition of the other. A causal dependency is *positive* when an action a produces an effect p and the action a' has p as one of its preconditions. A causal dependency is *negative* if a produces $\neg p$ and p is a precondition of a'. Note that the causal dependencies above are very similar to the causal links that we described for plan-space planning previously in Sect. "Artificial Intelligence Planning, Traditionally".

Figure 3 shows the graph of causal dependencies within this plan. This graph reveals certain patterns in the solution that may not be trivially extracted from a pattern-based learning approach as the one above. For example, the causal dependency between the actions pickup(*crate* 0,

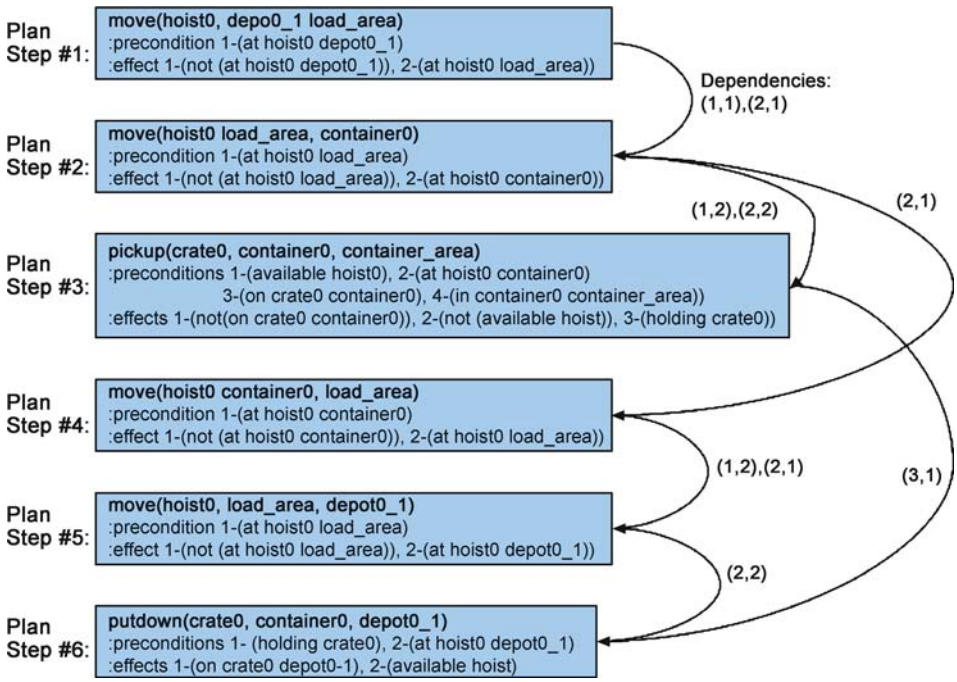

Learning and Planning (Intelligent Systems), Figure 3

A simple example showing the causal dependencies between the preconditions and effects of the actions in a solution plan. Each precondition and effect of an action (i.e., plan step above) is labeled by positive numbers 1, 2, ... A dependency between two actions a_1 and a_2 is shown as a pair of numbers (x, y) where x is the label of a effect of a_1 and y is the label of the precondition of a_2

container 0, *container_area*) and move(*hoist* 0, *container* 0, *load_area*) due to fact that both actions operate the same hoist is not necessarily clear from a syntactic analysis of the above plan. Similar is true for the dependency between the actions move(*hoist* 0, *load_area*, *depot* 0_1) and putdown(*crate* 0, *depot* 0_1, *depot* 0). Such dependencies help a macro-operator learner to produce better groupings of actions such as the following one:

$\langle pickup(?crate, ?container, ?area), move(?hoist, ?container, ?load\_area), move(?hoist, ?load\_area, ?depot\_area), putdown(?crate, ?depot\_area, ?depot)\rangle$.

Typical solution methods for learning such macro-operators in our Simple-Storage domain involve searching through the subgraphs of the causal-dependency graph of a solution plan and identifying patterns in those subgraphs. Then, the learning algorithm compares those patterns with those generated from other solution plans in the domain and generalizes into a set of macro-operators for the domain. Those macro-operators are then added into the set of planning operators O originally specified from the domain. This way, the macro-operators are treated by a planning algorithm as regular operators, ensuring the soundness and completeness of the planning algorithm.

The learned macro-operators are used for speeding up the search in a planning algorithm by causing jump-aheads in the search space, reducing the amount of work done by the search and the amount of backtracking. It has been experimentally demonstrated in the International Planning Competitions that learning and using macros in classical planners increases the performances of the planners significantly, allowing them to solve larger planning problems faster [5,12]. One important issue with using macro-operators, however, is that the search for macros over the possible subgraphs of solution dependencies could generate many macro-operators, most of which are redundant and some of which are not useful for planning. Thus, typical learning methods for macro-operators include heuristic selection and filtering capabilities to prune the set of generated macro-operators and use a small, but heuristically most useful ones in later planning.

Learning HTN Methods

Macro-operator learning is the simplest form of learning procedural domain planning knowledge that is described by observed groupings of actions that an executor performs in the plans. Although macro-operators can improve the performance of a planner as described above,

their expressivity in terms of the planning knowledge they specify is limited. For example, it is not possible to say and learn "if there is a crate that needs to be moved to a depot and there is a hoist at the crate's current location, move the crate with that hoist; otherwise, use a hoist that is the nearest one to the crate's location" in terms of a macro operator.

More specifically, macro-operators can only specify domain structures that is observable in the solution plans; they cannot include complex operational planning procedures for the domain that are not clearly explicit in the observations. Furthermore, macro-operators do not include any preconditions that does not appear in the original planning operators already; the conditions under which a macro-operator is applicable in a state is usually determined by the causal dependencies and the preconditions of the operators involved in that macro. This means that the applicability conditions of a macro-operator are sound, but do not specify any additional strategical knowledge about when and how to apply that macro operator.

Learning the kinds of strategical planning knowledge mentioned above requires more expressive languages in which such knowledge could be formally written. One of the most successful approaches that address such expressive has been languages that allows for specifying procedural and hierarchical planning domain structures, such as Hierarchical Task Networks, and their variants. There has been several recent techniques have been developed in order to learn HTNs or certain portions of HTN knowledge; we are going to present the basic ideas of this paradigm below.

As an illustration of the kinds of structured knowledge we are aiming to learn, recall our example of an HTN method for moving a crate from its container to its designated depot. We copy the example from Sect. "Artificial Intelligence Planning, Traditionally" here:

method move_crate_by_hoist for the task move_crate:
applicability conditions: the hoist is in working order and is present at the location of the container, the crate does not have any other crate on top of it
subtasks: (1) pick up the crate, (2) move the hoist to the designated depot, (3) locate the depot area, and (4) put down the crate
constraints: do subtask (1) before subtask (2), do subtask (2) before subtask (3), and do subtask (3) before subtask (4).

There are several issues that a learner must address while learning the above components of a method. Here is a summary:

- A learner must produce/hypothesize the structural (i. e., hierarchical and procedural) organization of knowledge. This involves, in an HTN framework, determining the structure of the HTN methods, i. e., the task a method can decompose and into which subtasks it does so.
- The applicability conditions under which a method should be invoked.
- Any binding and ordering constraints regarding the subtasks and the variable symbols in the method.

Typical solution techniques for learning such HTN methods involve a sort of explanation-based learning (EBL) methodology. As we described in Sect. "Machine Learning Methods in Planning", any EBL-based learning method requires a domain theory to be provided to the learning for deductive proof generation process. For HTN-method learning, the input domain theory also needs to augmented with some *operational information* regarding a planning domain. For example, consider the following rules for our *Simple-Storage* problem:

Move_Crate (?*crate*, ?*from_loc*, ?*to_depot*)
← *Crate*(?*crate* 2), *in*(?*crate* 2, ?*to_depot*),
SafeStack(?*crate*, ?*crate* 2),
PickedUp(?*crate*, ?*from_loc*).

PickedUp (?*crate*, ?*from_loc*)
← pickup(?crate, ?container, ?area),
move(?*hoist*, ?*container*, ?*load_area*),
move(?*hoist*, ?*load_area*, ?*depot_area*),
putdown(?*crate*, ?*depot_area*, ?*depot*).

Note that the above rules contain facts that can be evaluated directly in a state such as *in*(?*crate* 2, ?*to_depot*), high-level domain concepts that could be learned by EBL such as *SafeStack*, and descriptions of actions and higher-level activities that could be performed in the domain. For example, the first rule above specifies a higher-level activity called Move_Crate. It further specifies some conditions under which this activity could be accomplished in the body of that rules such as the learned concept *SafeStack*. The second rule specifies a concept that corresponds to a grouping of primitive actions that can be executed in the world. Note that the body of this rule is essentially a specification of a macro-operator that could also be learned by the techniques described in the earlier. The *PickedUp* concept is the crucial piece that allows to learn a hierarchy, as it appears both in the head of the second rule and in the body of the first one.

An operational activity specification allows an EBL-based learning method to accomplish several things. The rules in the specification provide a general common-sense description of associating the activities in the domain with the factual world state predicates. These predicates could be used to specify the goals that an activity could be expected to accomplished when performed in the world, or other side effects associated with the activity.

Second, as the proof tree of a deduction engine in an EBL-based learner will use and include information from task specification rules, the proof itself gives significant and strong clues about the structure of an HTN method and hierarchy of the methods. In other words, a learning algorithm for HTNs uses the proof tree to generate the relationships between tasks and their subtasks, as well as the applicability conditions associated with those relationships using the proof trees generated by the EBL mechanisms. These proof trees are further refined to produce the HTN methods by identifying the tasks in those trees and their subtasks, and organizing the subtrees to generate the applicability conditions and the variable-bindings regarding them.

Note that the above learning schema only produces simple HTN structures, where the only constraints over the orderings of the subtasks of a task is mostly totally-ordered, as dictated by the proof tree. Furthermore, generally HTN methods could encode complex groupings of subtasks (that is, more formally, HTNs are capable of modeling context-free languages) and complex recursion over those tasks. Depending on the problems, the practical implementations of the learning schema described above could learn some simple forms of recursion and methods that encode a bit more complex than regular languages, but not structural knowledge as expressive as standard HTNs. Extending this learning schema remains an important research direction in automated planning and learning.

Reinforcement Learning

Reinforcement Learning (RL) focuses on the problem of learning to choose actions for possible situations that an intelligent system could encounter while acting and interacting with the environment. The objective of a reinforcement learning is to use rewards and punishments (i. e., reinforcements) it receives from the environment while acting, in order to produce a *policy* (i. e., a mapping from situations to actions) that maximizes some utility function defined in terms of the reinforcements. In this respect, reinforcement learning naturally diverges from the other machine-learning paradigms, esp., the ones we have mentioned in this paper, since RL does not learn from any problem knowledge either in the form of training examples or in the form of background theory. Instead, the emphasis in RL is to learn from solely interactions with the world in order to achieve an explicitly stated goal.

When seen as a form of interleaved planning and execution paradigm, RL is a very effective form of generating plans (i. e., policies) in non-classical domains. In non-classical domains, actions may have multiple possible outcomes and the planner usually does not know which outcome will occur in the world when the action is executed. The planning process is usually characterized as trial-and-error search and the interactive learning built into the core of RL provides a way for reasoning about such multiple outcomes of the actions, since the learner executes and observes the effects of the actions immediately and weights that execution with other actions that might have been executed in terms of the rewards and punishments it receives.

Furthermore, interleaving planning and execution does not require that a RL algorithm know the underlying domain model, i. e., the preconditions and the effects of the actions. Whether an action is successfully applied in a situation is given through the reward mechanisms and the RL algorithm learns the effects of the actions via its observations of the world after the execution of the action. Thus, in this respect, RL can also be seen as a learning paradigm domain understanding since the outcome policies specify knowledge about how the domain works. Due to this property, most researchers believe that reinforcement learning is very suitable for robotic applications, systems control, games, and other applications.

Typical solution methods for RL problems, such as Q-Learning [50], are based on dynamic programming and simulation techniques, and they usually require exploring all or most of the state space in order to generate optimal or near-optimal policies. In complex problems, the sizes of the state spaces are usually prohibitive, and as a result, RL algorithm usually have efficiency problems in such settings. One way to alleviate this problem is to incorporate the aforementioned learning methods for planning into RL algorithms in order to learn within RL some sort of planning knowledge about what to do and what not to do by using background knowledge. An example of this approach is the integration of explanation-based learning techniques with RL, as described in [10].

Summary and Future Directions

The ability to produce knowledge about past experiences and exploit that knowledge in an operational context in

later problem-solving and planning sessions is an important attribute of any intelligent system, human- or AI-based alike. Automated Planning and Learning is the research paradigm that focuses on the development of intelligent systems and technologies that combine the ability to make decisions and generate courses of actions (i. e., plans) with the capability to reason and produce knowledge about past experiences, future problems that the system needs to tackle, and strategies about how to tackle them.

We have described the basic concepts in automated planning systems where machine learning has been of assistance in order to produce planning knowledge. We presented first a summary of traditional automated planning paradigm of AI and an overview of analytical machine-learning methods that have been deemed suitable for integration with planning systems. Then, we presented an analysis of four different forms of approaches that have integrated planning and learning techniques to achieve different objectives.

Learning domain models from observation has been investigated in several works such as [28,30,31,48,49,52,53]. Our formulation for learning domain models in this paper is partly based on the early work described in [48] and partly an analysis of all the other works on this topic. Among those that are more advanced than our formulation here, Levine and DeJong [30] describes a different approach to learn domain models that is based on explanation-based learning techniques. The work in [28] describes a way not only to learn action models, but also some temporal characteristics of the actions such as the duration it takes an action when executed. It is important to note that all these recent developments in learning domain models go beyond the expressivity of the earlier approaches, and this holds great promise for their applicability in realistic applications.

Explanation-based learning has been probably the most popular machine-learning method applied in automated planning, as evidenced by the works reported in [4,7,9,20,24,25,27,32,33] and others. Inductive learning, on the other hand, has been less investigated in planning, probably due to the correctness and reliability of the produced knowledge when there are not many training examples, compared to EBL's deductive reasoning guarantees. Examples of systems that used inductive learning for planning include [4,15,22,31,43] and others. Many of these inductive approaches, however, include an EBL component in the learning process in one way or another to reduce the uncertainty with the training examples. Thus, induction is usually used to overcome the limitations of EBL to produce over-specific or over-general knowledge, and to produce knowledge that is reliable based on a domain theory even if there are not sufficiently many training examples.

Existing works on learning macro-operators (e. g., [5,36]) are sometimes categorized as speed-up learning mentioned above, as do work on learning search control knowledge (e. g., [16,32,35]). Recently, several architectures have been proposed to learn hierarchical planning knowledge from a collection of plan traces and from a given action model [8,42,44]. Our formulations of learning macro-operators and operational procedures (i. e., HTNs) described in this paper are based on the previous works [5,8,36,40] on the topic.

Although our description of learning macro-operators and planning procedures is mostly based on EBL methods, rather than inductive learning, there has also been some advances on the latter approach for learning operational planning procedures. For example, [42]'s X-Learn, for example, uses inductive generalization to learn task decomposition constructs, which relate subgoals and conditions for decomposing those goals into their subgoals. By grouping goals in this way, task models are learned that lead to significant speed-ups in problem-solving. As another example, [8,40] achieves the same objective via problem-solving techniques similar to explanation-based learning. Two other recent studies [23,51] proposed algorithms to learn the applicability conditions of task-decomposition constructs, given the hierarchical relationships between the tasks.

Despite all the approaches described in this paper that have been very successful in learning for planning, we believe there are still interesting planning problems that machine-learning has not yet or just started to contribute. Here are several examples of such problems:

Learning for Multi-Planner Systems

There has been a growing recognition that sophisticated, *multi*-planner systems are needed for full automation in real-world applications involving complex, uncertain, and time-directed situations. In those situations, extensive planning knowledge is required but is difficult to obtain. This is partly because of the complexities in the environments, e. g., rescue and evacuation operations, and it is partly because there is often no expert to provide it, e. g., space operations. In such complex domains, a planning system that can learn such knowledge to develop ways on how to safely and correctly operate in the world holds great promise for success.

A key challenge of automated planning, when *multiple heterogeneous planners*, is "safety learning"; i. e., the problem of automated extraction of some safety con-

straints to identify and ensure the correct behavior of the entire system. A general-purpose use of the learned safety constraints further motivates a hybrid approach, which involves combinations of the learning methods described in this paper, such as inductive and explanation-based learning. It also requires the ability to maximize the information from the world, as such information is not available and reliable in most real life situations. For this latter requirement, Bayesian Learning approaches could be used in automated planning, although examples of Bayesian Learning for planning are rare.

Planning and Learning Under Uncertainty

All of the learning methods and opportunities we described in the paper for automated planning involve traditional planning problems, where the actions are deterministic and the state information is always complete and correct. In complex planning domains, however, there is usually a huge amount of uncertainty where an action may have a large number of possible outcomes, or even we may not know how many different possible things is going to happen when we execute an action.

We described Reinforcement Learning (RL) as a learning method for planning in planning domains that involve such uncertainties. Although RL can learn mappings from states to actions in non-traditional planning domains, these mappings are usually solutions to planning problems, rather than planning knowledge about a domain. Domain-specific knowledge has been used to search in reinforcement learning to address efficiency problems [10,41]. These approaches are based on hierarchical abstraction techniques that are somewhat similar to HTN planning. Given a domain, the hierarchical abstraction of that domain is analogous to an instance of the decomposition tree that an HTN planner might generate. However, the abstractions must be supplied in advance by the user, rather than being generated on-the-fly by the HTN planner. An important future direction is therefore to investigate learning methodologies for HTNs in the context of reinforcement learning in non-traditional planning domains.

Another approach for handling uncertainty in planning domains is to interleave the planning and learning process with the execution of some the planned actions in order to obtain key information about the world. This information can be then used to determine how to proceed with the planning and the learning processes. The interleaving of planning, learning, and execution has been investigated for mostly domains without uncertainties as in the works of [29,40]. Our description of learning HTNs for planning in this paper were based on these works. However, learning HTNs as in these works are very simple and not expressive enough to deal with uncertainty, and thus, it is still an open research area.

Acknowledgments

This work was supported by DARPA's Transfer Learning and Integrated Learning programs. The opinions in this paper are those of the author and do not necessarily reflect the opinions of the funders.

Bibliography

1. Aha DW (2002) Plan deconfliction, repair, and authoring in EDSS. Technical report, Progress report, Naval Research Laboratory
2. Ai-Chang M, Bresina J, Charest L, Hsu J, Jnsson AK, Kanefsky B, Maldague P, Morris P, Rajan K, Yglesias J (2003) Mapgen planner: Mixed-initiative activity planning for the mars exploration rover mission. In: Printed Notes of ICAPS 03 System demos
3. Bergmann R, Wilke W (1996) On the role of abstraction in case-based reasoning. In: European Workshop on Case-Based Reasoning (EWCBR-96), pp 28–43
4. Borrajo D, Veloso M (1997) Lazy incremental learning of control knowledge for efficiently obtaining quality plans. Artif Intell Rev 11(1–5):371–405
5. Botea A, Müller M, Schaeffer J (2005) Learning partial-order macros from solutions. In: International Conference on Automated Planning and scheduling (ICAPS-05). AAAI Press, Menlo Park, pp 231–240
6. Burstein M, Brinn M, Cox M, Hussain T, Laddaga R, McDermott D, McDonald D, Tomlinson R (2007) An architecture and language for the integrated learning of demonstrations. In:Burstein M, Hendler J (eds) AAAI Workshop Acquiring Planning Knowledge via Demonstration. AAAI Press, Menlo Park, pp 6–11
7. Chien SA (1989) Using and refining simplifications: Explanation-based learning of plans in intractable domains. In: International Joint Conference on Artificial Intelligence (IJCAI-89), Detroit. Morgan Kaufmann, San Francisco, pp 590–595
8. Choi D, Langley P (2005) Learning teleoreactive logic programs from problem solving. In: International Conference Inductive Logic Programming (ILP-05), Bonn. Springer, New York, pp 51–68
9. DeJong GF, Mooney R (1986) Explanation-based learning: An alternative view. Mach Learn 1(2):145–176
10. Dietterich TG (2000) Hierarchical reinforcement learning with the MAXQ value function decomposition. JAIR 13:227–303
11. DOT (1999) An assessment of the US marine transportation system, a report to congress. Technical report, US Department of Transportation, p 103
12. Edelkamp S, Hoffmann J (2004) International planning competition. http://ipc.icaps-conference.org
13. Erol K, Hendler J, Nau DS (1996) Complexity results for hierarchical task-network planning. AMAI 18:69–93
14. Erol K, Nau DS, Subrahmanian VS (1995) Complexity, decidability and undecidability results for domain-independent planning. Artif Intell 76(1–2):75–88

15. Estlin TA, Mooney RJ (1997) Learning to improve both efficiency and quality of planning. In: IJCAI, pp 1227–1233, Nagoya, Japan. Morgan Kaufmann, San Francisco
16. Fern A, Yoon S, Givan R (2004) Learning domain-specific control knowledge from random walks. In: ICAPS, Whistler. AAAI Press, Menlo Park, pp 191–199
17. Fikes RE, Nilsson NJ (1971) STRIPS: A new approach to the application of theorem proving to problem solving. Artif Intell 2:189–208
18. Gerevini A, Dimopoulos Y, Haslum P, Saetti A (2006) International planning competition. http://zeus.ing.unibs.it/ipc-5/
19. Goldman R (2004) Adapting research planners for applications. In: ICAPS workshop on Connecting Planning Theory with Practice
20. Gratch J, DeJong G (1992) Composer: A probabilistic solution to the utility problem in speed-up learning. In: AAAI, San Jose. AAAI Press, Menlo Park, pp 235–240
21. Hebbar K, Smith SJJ, Minis I, Nau DS (1996) Plan-based evaluation of designs for microwave modules. In: Proc ASME Design Technical Conference, Irvine, August 1996
22. Huang Y, Kautz H, Selman B (2000) Learning declarative control rules for constraint-based planning. In: International Conference on Machine Learning (ICML-00). Morgan Kaufmann, San Francisco, pp 415–422
23. Ilghami O, Nau DS, Muñoz-Avila H, Aha DW (2005) Learning preconditions for planning from plan traces and HTN structure. Comput Intell 21(4):388–413
24. Kambhampati S (2000) Planning graph as (dynamic) CSP: Exploiting EBL, DDB and other CSP techniques in Graphplan. JAIR 12:1–34
25. Kambhampati S, Katukam S, Qu Y (1996) Failure driven dynamic search control for partial order planners: An explanation based approach. Artif Intell 88(1–2):253–315
26. Knoblock CA (1993) Generating abstraction hierarchies: An automated approach to reducing search in planning. Kluwer, Norwell
27. Laird J, Newell A, Rosenbloom P (1987) SOAR: An architecture for general intelligence. Artif Intell 33(1):1–67
28. Lanchas J, Jimenez S, Fernandez F, Borrajo D (2007) Learning action durations from executions. In: Proceedings of the ICAPS-07 Workshop on AI Planning and Learning, Providence, 22 Sep 2007
29. Langley P, Choi D (2006) Learning recursive control programs from problem solving. J Mach Learn Res 7:493–518
30. Levine G, DeJong GF (2006) Explanation-based acquisition of planning operators. In: International Conference on Automated Planning and Scheduling (ICAPS-06), The English Iave District, 6–10 June 2006. AAAI Press, Menlo Park, pp 152–161
31. Matthew D, Oates T, Cohen PR (2000) Learning planning operators in real-world, partially observable environments. In: AIPS, Breckenridge, 14–17 April 2000. AAAI Press, Menlo Park
32. Minton S (1988) Learning effective search control knowledge: An explanation-based approach. Technical Report TR CMU-CS-88-133, School of Computer Science, Carnegie Mellon University
33. Mitchell T, Keller R, Kedar-Ceballi S (1986) Explanation-based generalization: A unifying view. Mach Learn 1(1):47–80
34. Mitchell TM (1977) Version spaces: A candidate elimination approach to rule learning. In: IJCAI. AAAI Press, Cambridge, pp 305–310
35. Mitchell TM (1997) Machine learning. McGraw-Hill, New York
36. Mooney RJ (1988) Generalizing the order of operators in macro-operators. In: Machine Learning, International Conference on Machine Learning (ML-88), Ann Arbor, 12–14 June 1988. Morgan Kaufmann, San Francisco, pp 270–283
37. Muñoz-Avila H, Aha DW, Breslow L, Nau DS (1999) HICAP: An interactive case-based planning architecture and its application to non-combatant evacuation operations. In: AAAI/IAAI Proceedings, Orlando, 18–22 July 1999. AAAI Press, Menlo Park, pp 870–875
38. Muñoz-Avila H, Breslow LA, Aha DW, Nau DS (1998) Description and functionality of NEODocTA. Technical Report AIC-96-005, Naval Research Laboratory, Navy Center for Applied Research in Artificial Intelligence
39. Nau DS, Cao Y, Lotem A, Muñoz-Avila H (1999) SHOP: Simple hierarchical ordered planner. In: International Joint Conference on Artificial Intelligence (IJCAI-99), Stockholm, 31 July–6 Aug 1999. Morgan Kaufmann, San Francisco, pp 968–973
40. Nejati N, Langley P, Konik T (2006) Learning hierarchical task networks by observation. In: Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, 25–29 June 2006. ACM Press, New York
41. Parr R (1998) Hierarchical control and learning for Markov decision processes. University of California, Berkeley
42. Reddy C, Tadepalli P (1997) Learning goal-decomposition rules using exercises. In: International Conference on Machine Learning (ICML-97), Nashville, 8–12 July 1997. Morgan Kaufmann, San Francisco, pp 278–286
43. Reddy C, Tadepalli P (1999) Learning horn definitions: Theory and application to planning. New Gener Comput 17(1): 77–98
44. Ruby D, Kibler DF (1991) SteppingStone: An empirical and analytic evaluation. In: AAAI, Anaheim, July 1991. Morgan Kaufmann, San Francisco, pp 527–531
45. Sacerdoti E (1975) The non-linear nature of plans. In: International Joint Conference on Artificial Intelligence (IJCAI-75), Tiblisi, pp 206–214
46. Smith SJJ, Nau DS, Throop T (1998) Computer bridge: A big win for AI planning. AI Mag 19(2):93–105
47. Tate A (1977) Generating project networks. In: International Joint Cinference on Artificial Intelligence (IJCAI-77), Cambridge, 22–25 Aug 1977, pp 888–893
48. Wang X (1994) Learning by observation and practice: A framework for automatic acquisition of planning operators. In: AAAI
49. Wang X (1994) Learning planning operators by observation and practice. In: International Conference on AI Planning Systems (AIPS-94), Chicago, June 1994. AAAI Press, Mento Park, pp 335–340
50. Watkins CJCH, Dayan P (1992) Q-learning. Mach Learn 8(3–4):279–292
51. Xu K, Munoz-Avila H (2005) A domain-independent system for case-based task decomposition without domain theories. In: AAAI, Pittsburgh, July 2005. AAAI Press, Menlo Park, pp 234–240
52. Yang Q, Wu K, Jiang Y (2005) Learning actions models from plan examples with incomplete knowledge. In: International Conference on Automated Planning and Scheduling (ICAPS-05), Monterey, June 2005. AAAI Press Menlo Park, pp 241–250
53. Yang Q, Wu K, Jiang Y (2007) Learning action models from plan examples using weighted max-sat. Artif Intell 171(2–3): 107–143

Learning, System Identification, and Complexity

M. VIDYASAGAR
Tata Consultancy Services, Software Units Layout, Hyderabad, India

Article Outline

Glossary

Accuracy A significant parameter in statistical learning theory that describes the error between the current estimate and the truth.

Algorithmic complexity The computational complexity of computing a hypothesis in a learning problem, on the basis of observations.

ARMA Auto-Regressive Moving Average. It is a specific kind of model for a dynamical system.

Confidence An estimate of the confidence with which we can make a particular statement, that is, a bound on the probability that the statement is false.

PAC Probably Approximately Correct. It is a widely used model for statistical learning.

Sample complexity The number of samples required to learn to a specific accuracy and confidence.

Vapnik–Chervonenkis dimension A combinatorial parameter that captures the "richness" of a set of concepts to be learnt.

Definition of the Subject

Definitions

$A \Delta B$ Symmetric difference between the sets A and B.
A_m Algorithm to be applied to m labeled inputs.
C A concept class, that is, a collection of subsets.
$d_P(T, H_m)$ Generalization error with the target set T and the hypothesis H_m.
H_m Hypothesis generated by the algorithm A_m.
$H_m(T; \mathbf{x})$ Hypothesis generated when the target set is T and the sample is $\mathbf{x}$.
$I_T(\cdot)$ Indicator function of the set T.
$J(\theta)$ Cost function (often, though not always, the sum of squares).
$M(\epsilon, C, d_P)$ ϵ-packing number of the concept class C with the pseudometric d_P.
$N(\epsilon, C, d_P)$ ϵ-covering number of the concept class C with the pseudometric d_P.
$r(m, \epsilon)$ Learning rate function; depends on the number of inputs m and the accuracy ϵ.
θ Parameter that characterizes the model used in identification.
$\hat{\theta}_t$ Estimate of θ at time t.
Θ Set of possible parameters θ.
$\mathbf{u}_t$ Input to the system to be identified at time t.
$\mathbf{y}_t$ Output of the system to be identified at time t.

System identification refers to the problem of fitting a model, from within a given set of models, to a series of observations (or measurements). By tradition, system identification is *recursive*, that is, the data is given one observation at a time, and ideally the next model is obtained as a function of the previously fit model and the current observation. One consequence of having a recursive identification procedure is that, once we have used the latest observation to update the model, the observation can be discarded, because the latest model embodies the effect of all the past observations.

Statistical learning is a closely related problem in which it is assumed that the data at hand is generated by a "true but unknown" target concept (or function), and the objective is to fit a model to the data at hand. Unlike in system identification, statistical learning is generally not recursive. In other words, at any given point in time, *the entire set of past observations* is often used to do the model fitting.

Perhaps the major difference between system identification theory and statistical theory is that the latter is heavily focused on so-called "finite time estimates." In other words, by tradition most theorems in system identification theory are *asymptotic in nature*, and tell us what happens as the number of observations approaches infinity. In contrast, by its very nature, statistical learning theory is devoted to the derivation of explicit upper bounds on just far the current model is from the unknown target. These bounds are a function of the number of observations, and the desired accuracy and confidence parameters (these are defined precisely later on).

Introduction

One of the earliest known examples of system identification is "curve fitting," which is the problem of fitting

a curve from within a given family to set of points in the plane (or some higher dimensional space). The problem of "curve fitting" goes back to antiquity, with some of the major contributions being made two centuries ago by Gauss and Laplace. System identification is a natural extension of curve fitting to the situation where the models are not "static" functions, but rather dynamical systems. The dynamical nature of the models introduces additional complexities into the theory.

In a learning problem, the issue of "complexity" is central, whether it is "sample complexity" or "algorithmic complexity." The former refers to the number of samples needed to achieve a desired level of accuracy and confidence, and the latter refers to the difficulty of computing a model on the basis of the given set of observations. Both types of complexity are studied in this article.

The difference between system identification and statistical learning is one of emphasis more than anything else. Thus it is not surprising that it is possible to view system identification as a problem in statistical learning. In this article, this connection is brought out explicitly.

We begin the article by reviewing the problem of curve fitting using a least squares optimization criterion, and then show how the basic ideas can be extended to the problem of system identification. Then we introduce the problem of PAC (probably approximately correct) learning. Then we discuss the issues of complexity. Then we show how identification can be viewed as a kind of learning problem. We conclude with some directions for future research.

Least Squares Curve Fitting – A Precursor

Many of the issues that underlie system identification can be readily understood in the context of the much simpler problem of "curve fitting" using a least squares error criterion.

Suppose we are given a series of pairs of vectors $(\mathbf{x}_i, \mathbf{y}_i)$, $i = 1, \ldots, l$, where $\mathbf{x}_i \in \mathbb{R}^n, \mathbf{y}_i \in \mathbb{R}^m \forall i$. The objective is to fit the data with a linear relationship of the form $\mathbf{y} = A\mathbf{x}$, where $A \in \mathbb{R}^{m \times n}$, such that the least squares error criterion

$$J(A) = \sum_{i=1}^{l} \| \mathbf{y}_i - A\mathbf{x}_i \|_2^2$$

is minimized, where $\| \cdot \|_2$ denotes the Euclidean norm

$$\| \mathbf{v} \|_2 = (\mathbf{v}^t \mathbf{v})^{1/2} .$$

Using the notation

$$Y := [\mathbf{y}_1 \ldots \mathbf{y}_l] \in \mathbb{R}^{m \times l}, \; X := [\mathbf{x}_1 \ldots \mathbf{x}_l] \in \mathbb{R}^{n \times l} ,$$

we can easily establish that the optimal choice of the matrix A is given by

$$A^* = CD^{-1} , \tag{1}$$

where

$$C = YX^t , \quad D = XX^t .$$

This formula presupposes that the matrix D is nonsingular and therefore invertible. Since

$$D = XX^t = \sum_{i=1}^{l} \mathbf{x}_i \mathbf{x}_i^t ,$$

it is easy to see that D is nonsingular if and only if the set of vectors $\{\mathbf{x}_1, \ldots, \mathbf{x}_l\}$ contains n linearly independent vectors. This leads to the first observation, namely: Unless there is sufficient data, the curve-fitting problem does not have a unique solution.

Now suppose that the data is in fact generated by a "true but unknown" linear relationship, where the data is corrupted by noise. In other words, suppose that

$$\mathbf{y}_i = A_0 \mathbf{x}_i + \mathbf{v}_i ,$$

where A_0 is the "true" matrix, and $\{\mathbf{v}_i\}$ is a sequence of random vectors that are pairwise independent. For the moment, suppose that the vectors $\mathbf{x}_i$ are deterministic and not random. Then the measurements are described by

$$Y = AX + V ,$$
$$\text{where } V = [\mathbf{v}_1 \ldots \mathbf{v}_l] .$$

Now from this formula and (1), the following facts follow readily.

1. Since the measurement noise vectors are random, the optimal choice of A after l measurements, call it A_l, is also random.
2. If the measurement noise vectors $\mathbf{v}_i$ are all zero mean, then the optimal choice A_l is *unbiased*, that is, the expected value of A_l equals the true matrix A_0 for each value of l.
3. If the measurement noise vectors are normally distributed around zero, then A_l is normally distributed around the true value A_0.
4. Under very broad conditions, for example if the noise vectors are a sequence of i.i.d. (independent, identically distributed) vectors with finite variance, then the estimate A_l converges almost surely to the true value A_0 as $l \to \infty$.

We shall see that many of the same features are present in the system identification problem as well.

System Identification

System identification refers to the problem of fitting a dynamical system model to a given sequence of temporal (i. e., time-spaced) measurements. Thus one is given a set of input-output pairs $\{(u_t, y_t), t \geq 0\}$, as well as a family of models $\{\boldsymbol{h}(\theta), \theta \in \Theta\}$. The idea is to choose the "best possible" model from within this family to represent the data.

System identification is a vast and well-developed subject. In a brief essay such as this one, it is difficult to do more than to scratch the surface of the topic. The reader is referred to the book [9] for a very authoritative and comprehensive treatment.

Problem Formulation

Suppose one is given a sequence of input-output measurements $\{(\mathbf{u}_t, \mathbf{y}_t), t \geq 0\}$, where $\mathbf{u}_t$ is the input to the system at time t, and $\mathbf{y}_t$ is the output of the system at time t. The family of models $\{h(\theta), \theta \in \Theta\}$ can consist of linear recursive relationships, or nonlinear mappings. Perhaps the simplest class of models is the so-called ARMA (Auto-Regressive Moving Average) model, in the form

$$\mathbf{y}_t = \sum_{i=1}^{k} A_i(\theta)\mathbf{y}_{t-i} + B_i(\theta)\mathbf{u}_{t-i} \,, \tag{2}$$

where $A_i(\theta), B_i(\theta)$ are matrices of appropriate dimension that depend on the parameter $\theta \in \Theta$, where Θ is a subset of some Euclidean space. Note that the input $\mathbf{u}_t$ is not necessarily "white" noise, but could be a combination of a deterministic input and "colored" noise. The relationship (2) can be rewritten in the form

$$[I - A(\theta) - B(\theta)]\phi_t = \mathbf{0} \,, \tag{3}$$

where

$$A(\theta) = [A_1(\theta) \ldots A_k(\theta)],$$

$$B(\theta) = [B_1(\theta) \ldots B_k(\theta)], \phi_t = \begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \vdots \\ \mathbf{y}_{t-k} \\ \mathbf{u}_{t-1} \\ \vdots \\ \mathbf{u}_{t-k} \end{bmatrix} .$$

To estimate the parameter θ_t at time t, it is common to choose it so as to minimize the accumulated least-squares error. Since we wish to fit a model of the form (3) to the data, the vector

$$\begin{aligned} \mathbf{e}_l(\theta) &= [I - A(\theta) - B(\theta)]\phi_l \\ &= \mathbf{y}_l - \sum_{i=1}^{k} A_i(\theta)\mathbf{y}_{l-i} + B_i(\theta)\mathbf{u}_{l-i} \end{aligned}$$

gives a goodness of fit of the model to the data, when the parameter is chosen as θ. Therefore we can define the cost function at time t as

$$\begin{aligned} J_t(\theta) &= \frac{1}{t}\sum_{l=k}^{t} \| \mathbf{e}_t \|_2^2 \\ &= \frac{1}{t}\sum_{l=k}^{t} \left\| \left[\mathbf{y}_l - \sum_{i=1}^{k} A_i(\theta)\mathbf{y}_{l-i} + B_i(\theta)\mathbf{u}_{l-i} \right] \right\|_2^2 \end{aligned}$$

as a scalar measure of how good a model parameter θ is. It is easy to see that the normalizing factor $1/t$ is inserted to ensure that the quantity J_t remains well behaved as $t \to \infty$. Thus it is logical to choose the estimate $\hat{\theta}_t$ such that

$$\hat{\theta}_t = \arg\min_{\theta \in \Theta} J_t(\theta) \,.$$

Note that $\hat{\theta}$ is our estimate of the parameter θ, based on the observations up to time t.

The above description applies when the model is a linear ARMA model. However, the approach itself is far more general. Define the vector ϕ_t as above, and suppose that the family of models consists of a set of relationships of the form

$$\mathbf{f}(\theta, \phi_t) = \mathbf{0} \,.$$

If $\mathbf{f}(\theta, \cdot)$ is a linear relationship of the form (3) for each choice of the parameter $\theta \in \Theta$, then we get the familiar parametrized ARMA model. As before, given a sequence of observations, the error at time t is defined as

$$\mathbf{e}_t(\theta) = \mathbf{f}(\theta, \phi_t) \,,$$

while the accumulated least squares error criterion is defined as before, as

$$J_t(\theta) = \frac{1}{t}\sum_{l=k}^{t} \| \mathbf{e}_t \|_2^2 \,.$$

In case the models are linear, choosing the least-squares error criterion leads to a very elegant solution of the optimization problem. But in case the models are not linear, there is no particular advantage in choosing the least-squares error criterion. Instead, one can choose some "loss" function $\ell : \mathbb{R}_+ \to \mathbb{R}_+$ such that $\ell(\cdot)$ is monoton-

ically increasing, and choose

$$J_t(\theta) = \frac{1}{t}\sum_{l=k}^{t} \ell(\| \mathbf{e}_t \|_2)$$

as the cost function at time t. Once again, the estimate $\hat{\theta}_t$ is chosen as the minimizer of $J_t(\theta)$ as θ varies over Θ.

Suppose now that the data comes from a "true but unknown" system whose output measurements are corrupted by measurement noise. Specifically, suppose it is the case that

$$\mathbf{y}_t = \sum_{i=1}^{k} A_i(\theta_{\text{true}})\mathbf{y}_{t-i} + B_i(\theta_{\text{true}})\mathbf{u}_{t-i} + \mathbf{v}_t \,,$$

where the measurement noise $\mathbf{v}_t$ is an i.i.d. zero-mean sequence. For convenience, let us define

$$\mathbf{y}_t^0 = \sum_{i=1}^{k} A_i(\theta_{\text{true}})\mathbf{y}_{t-i} + B_i(\theta_{\text{true}})\mathbf{u}_{t-i}$$

to be the corresponding output sequence of the same system, but without any measurement noise. Because the current output $\mathbf{y}_t$ depends not only on the current and past inputs $\mathbf{u}_{t-1}$ but also the past outputs $\mathbf{y}_{t-i}$, the difference $\mathbf{y}_t - \mathbf{y}_t^0$ is a function of not just the current measurement noise $\mathbf{v}_t$, but also of the past measurement noise vectors $\mathbf{v}_{t-i}$. As a result, even if the measurement noise $\mathbf{v}_t$ consists of independent random variables, the dynamical nature of the system ensures that the difference between the actual measured output $\mathbf{y}_t$ and the ideal, noiseless measurement $\mathbf{y}_t^0$ is a sequence of *dependent* random variables. In other words, the dynamical nature of the system ensures that we cannot any longer make an analogy with the "static" least squares curve-fitting problem of the previous section.

Behavior of Identification Algorithms

We have seen that a widely used approach to system identification is to formulate the error criterion $J_t(\theta)$, and at each instant of time, choose the model parameter $\hat{\theta}_t$ so as to minimize $J_t(\theta)$. As time goes on, this approach generates a sequence of estimated parameters $\{\hat{\theta}_t\}$. Now we can ask three successively stronger questions, as follows:

First, as more and more data is provided to the identification algorithm, does the estimation error between the outputs of the identified model and the actual time series approach the minimum possible estimation error achievable by any model within the given model class? To clarify this question, let us define the quantity $J(\theta)$ as the limit as $t \to \infty$ of the cost function $J_t(\theta)$. It is easy to see that the limit exists under very mild conditions. Now define

$$J^* = \min_{\theta \in \Theta} J(\theta)$$

to be the best possible error measure. The question that is being asked is: Does the quantity $J(\hat{\theta}_t)$ approach this minimum value? In other words, is the performance of the identification algorithm asymptotically optimal?

Second, does the identified model converge to "the set of best possible models" within the given model class? In other words, suppose that the minimum possible estimation error is achievable by one or more "best possible models," and let Θ^* the set of all θ such that $J(\theta) = J^*$. The question is: does the distance between the estimated vector $\hat{\theta}_t$ and the set Θ^* approach zero at $t \to \infty$? Note that the question, as posed, permits the sequence $\{\hat{\theta}_t\}$ not to have a limit in the conventional sense, but to keep bouncing around nearer and nearer the set Θ^*. The difference between the first question and the second question is this. The first question requires only that the scalar $J(\hat{\theta}_t)$ should converge to the optimal value J^*, while the second question requires that the estimated parameter vector $\hat{\theta}_t$ should converge to the set of minimizers Θ^*.

Third, assuming that the data is generated by a "true" system whose output is corrupted by measurement noise, does the identified model converge to the "true" system? In other words, if both the true system and the family of models are parametrized by a vector of parameters, does the estimated parameter vector converge to the true parameter vector? The difference between the second question and the third question is the following: In the second question, it is *not* assumed that the data is generated by a "true" system corresponding to a "true" parameter θ_{true} belonging to the set Θ. Accordingly, in the second question, all that we ask is that the sequence of estimates $\{\hat{\theta}_t\}$ should converge to the set of minimizers of the function J. In contrast, in the third question, we add the assumption that the data to which the identification algorithm is applied has a specific structure, namely that it is generated by a true but unknown parameter vector θ_{true}. The requirement in the third question is that the sequence of estimates $\{\hat{\theta}_t\}$ should converge to θ_{true} as $t \to \infty$.

From a technical standpoint, Questions 2 and 3 are easier to answer than Question 1. Since identification is often carried out recursively, the output of the identification algorithm is a sequence of estimates $\{\theta_t\}_{t \geq 1}$. Suppose that we are able to show that Question 1 has an affirmative answer, i. e., that $J(\theta_t) \to J^*$, where J^* denotes the minimum possible estimation error. In such a case, with some additional assumptions it is possible to answer both Questions 2 and 3 in the affirmative.

Traditionally a positive answer to Question 2 is assured by assuming that Θ is a *compact* set, which in turn ensures that the sequence $\{\theta_t\}$ contains at least one convergent subsequence. For convenience, let us relabel this subsequence again as $\{\theta_t\}$. If the answer to Question 1 is "yes," if θ^* is any limit point of the sequence, and if $J(\theta)$ is continuous (or at worst, lower semi-continuous) with respect to θ, then it readily follows that $J(\theta_t) \to J^*$. In other words, the model $h(\theta^*)$ is an "optimal" fit to the data among the family $\{h(\theta), \theta \in \Theta\}$. Coming now to Question 3, suppose θ_{true} is the parameter of the "true" system, and let h_{true} denote the "true" system. In order for Question 3 to have an affirmative answer, the true system h_{true} must belong to the model family $\{h(\theta), \theta \in \Theta\}$; otherwise we cannot hope that $h(\theta_t)$ will converge to h_{true}. The traditional way to ensure that $\theta_{\text{true}} = \theta^*$ is to assume that the input to the true system is "persistingly exciting" or "sufficiently rich," so that the only way for $h(\theta^*)$ to match the performance of h_{true} is to have $\theta^* = \theta_{\text{true}}$.

Because of the arguments presented above, the main emphasis in system identification theory has been to study conditions to ensure that Question 1 has an affirmative answer, i. e., that the identification algorithm is asymptotically optimal. This question can be readily addressed using the following "universal" approach. We have defined $J(\theta)$ as the limit as $t \to \infty$ of the quantity $J_t(\theta)$. As remarked earlier, this limit exists under very mild conditions, *for each fixed* θ. Thus it can be said without much fuss that $J_t(\theta) \to J(\theta)$ as $t \to \infty$, again *for each fixed* θ. However, if a stronger property holds, namely that $J_t(\theta) \to J(\theta)$ as $t \to \infty$, *uniformly with respect* θ, then Question 1 has an affirmative answer. This is the approach adopted in [8]. In many ways, this approach to answering Question 1 is the common thread that binds system identification theory to learning theory. An alternate approach to Question 1 based on ergodic theory can be found in [5,6].

The fourth question that can be asked is: Suppose that the estimated parameter sequence $\{\hat{\theta}_t\}$ converges to the true parameter θ_{true}. As we have already seen, since the data that acts as the input to the identification algorithm is generated at random, the resulting estimates $\hat{\theta}_t$ are also random. Thus we may wish to know, not merely that the random variable $\hat{\theta}_t$ converges to θ_{true} in some probabilistic sense (say in probability, or almost surely), but also, the asymptotic distribution of the estimate $\hat{\theta}_t$ around the true value θ_{true}. To make this point clear, let us return to the linear least-squares curve-fitting problem. We know that if the measurement noise consists of i.i.d. gaussian variables, then the optimal estimate also has a gaussian distribution centered around the true value, and that the variance around the true value goes to zero at a rate that can be quantified fairly precisely. When we address this question in the context of dynamical systems however, the situation is considerably more complex, and one is obliged to use very sophisticated mathematics in order to analyze the asymptotic behavior of the estimated parameter sequence. Nevertheless, it is possible to say something about the asymptotic behavior of the estimates. The reader is invited to consult Chap. 9 of [9] for details.

To keep the exposition simple, we have discussed only one possible approach to system identification. Minimizing some kind of generalized least squares error criterion is only one way to choose an estimated parameter $\hat{\theta}_t$. There are other popular methods, such as the "instrument variables" method. In the setting of linear systems especially, it is possible to evolve quite different approaches based on frequency domain (instead of time domain) methods, based on correlating the input and output variables. Here again, the book [9] is an excellent reference. Another excellent reference is [10].

Statistical Learning Theory

Problem Formulation

A parallel theme to the problem of system identification is statistical learning. This kind of learning theory comes in many varieties, and the type of learning studied here is known as Probably Approximately Correct (PAC) learning theory. Even within PAC learning theory, there are two aspects, namely: statistical and computational. The present section discusses the statistical aspects of PAC learning, while the next section discusses the computational aspects.

Suppose one is trying to identify an unknown set on the basis of observation. The notion of PAC (probably approximately correct) learnability gives mathematical formalization of this intuitive idea. Suppose X is a set, and that C is a collection of subsets of X. Suppose also that P is a probability measure on the set X, which may or may not be known ahead of time to the learner. The learning problem is formulated as follows: There is a fixed but unknown concept $T \in C$, called the **target** concept. The objective is to "learn" the target concept on the basis of observation, consisting of i.i.d. samples $x_1, \ldots, x_m \in X$ drawn in accordance with P. For each sample x_j, an "oracle" tells the learner whether or not $x_j \in T$; equivalently, the oracle returns the value of $I_T(x_j)$, where $I_T(\cdot)$ is the indicator function of T. Thus, after m samples have been drawn, the information available to the learner consists of the "labeled multisample"

$$[(x_1, I_T(x_1)), \ldots, (x_m, I_T(x_m))] \in [X \times \{0,1\}]^m .$$

The objective is to construct a suitable approximation to the unknown target concept T on the basis of the labeled sample, using an appropriate algorithm. For the purposes of the present discussion, an "algorithm" is merely an indexed family of maps $\{A_m\}_{m\geq 1}$, where

$$A_m : [X \times \{0,1\}]^m \to C .$$

Thus an algorithm merely maps sets of labeled multisamples into elements of the concept class C. Suppose m i.i.d. samples have been drawn, and define

$$H_m(T;\mathbf{x}) = A_m[(x_1, I_T(x_1)), \ldots, (x_m, I_T(x_m))] .$$

The set $H_m(T;\mathbf{x})$ is referred to as the **hypothesis** generated by the algorithm. When there is no danger of confusion, one can abbreviate $H_m(T;\mathbf{x})$ by H_m.

In order to analyze how close the hypothesis H_m is to the unknown target set T, we need a *quantitative* measure of the disparity. Given two sets $A, B \subseteq X$, define their symmetric difference $A\Delta B$ by

$$A\Delta B = (A \cap B^c) \cup (A^c \cap B) = (A \cup B) - (A \cap B) .$$

Thus $A\Delta B$ consists of those points that belong to precisely one of the sets (and not the other). Now define the "generalization error"

$$d_P(T, H_m) = P(T\Delta H_m) .$$

Then it is easy to see that $d_P(T, H_m)$ is precisely the probability that a randomly chosen test sample is *misclassified* by the hypothesis H_m.

Roughly speaking, the algorithm $\{A_m\}$ can be said to "learn" the target concept T if the distance $d_P(T, H_m)$ approaches zero as $m \to \infty$. However, since the samples $\mathbf{x}_i$ are random, the resulting hypothesis H_m is also "random" and the generalization error $d_P(T, H_m)$ is also random. Define

$$r(m,\epsilon) = \sup_{T\in C} P^m\{\mathbf{x} \in X^m : d_P[T, H_m(T;\mathbf{x})] > \epsilon\} . \quad (4)$$

The quantity $r(m,\epsilon)$ is called the "learning rate" function.

Definition 1 The algorithm $\{A_m\}$ is said to be **probably approximately correct (PAC)** if $r(m,\epsilon) \to 0$ as $m \to \infty$ for each $\epsilon > 0$. The concept class C is said to be **PAC learnable** if there exists an algorithm that is PAC.

Thus the algorithm $\{A_m\}$ is PAC if, for each target concept T, the distance $d_P(T, H_m)$ converges to zero in probability as $m \to \infty$, and moreover, the rate of convergence is *uniform* with respect to the unknown target concept T.

Note that, if an algorithm is PAC, then for each pair of numbers $\epsilon, \delta > 0$, there exists a corresponding integer $m_0 = m_0(\epsilon, \delta)$ such that

$$P^m\{\mathbf{x} \in X^m : d_P[T, H_m(T;\mathbf{x})] > \epsilon\} \leq \delta , \quad \forall m \geq m_0 , \ \forall T \in C . \quad (5)$$

In other words, if at least m_0 i.i.d. samples are drawn and the algorithm is applied to the resulting labeled multisample, it can be said with confidence at least $1 - \delta$ that the hypothesis H_m produced by the algorithm is within a distance ϵ of the true but unknown target concept T. In the above expression, the quantity ϵ is referred to as the "accuracy" of the algorithm, while δ is referred to as the "confidence." The integer $m_0(\epsilon, \delta)$ is called the "sample complexity" of the algorithm. In all but the simplest cases, one can only obtain upper bounds for $m_0(\epsilon, \delta)$, but this is sufficient.

The above statement provides some motivation for the nomenclature "PAC." If the algorithm $\{A_m\}$ is PAC, the hypothesis H_m produced by the algorithm need not *exactly* equal T; rather, H_m is only an *approximately* correct version of T, in the sense that $d_P(H_m, T) \leq \epsilon$. Even this statement is only *probably* true, with a confidence of at least $1 - \delta$.

The notion of PAC learnability is usually credited to Valiant (see [12]), who showed that a certain collection of Boolean formulas is PAC learnable; however, the term "PAC learnability" seems to be due to [1]. It should be mentioned that Valiant had some additional requirements in his formulation beyond the condition that $r(m,\epsilon) \to 0$ as $m \to \infty$. In particular, he required that the mapping A_m be computable in polynomial time. These additional requirements are, in some sense, what distinguish computational (PAC) learning theory from statistical (PAC) learning theory. The references on computational learning theory in the bibliography go into computational issues in greater depth.

PAC learning is closely related to an older topic in probability theory known as empirical process theory, which is devoted to a study of the behavior of empirical means as the number of samples becomes large. One of the standard references in the field, namely [13], is actually about empirical process theory. In the book [15] PAC learning theory is treated side by side with the uniform convergence problem, and it is shown that each has something to offer to the other. However, in the present write-up we do not discuss the uniform convergence problem. Finally, the book [16] removes the assumption that the learning samples must be independent, and replaces it with the weaker assumption that the learning samples form a so-called mixing process.

Conditions for PAC Learnability

In this section we summarize the main results that provide necessary and sufficient conditions for a collection of sets to be PAC learnable. A central role is played by a combinatorial parameter known as the Vapnik–Chervonenkis (VC)-dimension. This notion is implicitly contained in the paper [14].

Definition 2 Let $\mathcal{A}$ be a collection of subsets of X. A set $S = \{x_1, \ldots, x_n\} \subseteq X$ is said to be **shattered** by $\mathcal{A}$ if, for every subset $B \subseteq S$, there exists a set $A \in \mathcal{A}$ such that $S \cap A = B$. The **Vapnik–Chervonenkis dimension** of $\mathcal{A}$, denoted by VC-dim($\mathcal{A}$), equals the largest integer n such that there exists a set of cardinality n that is shattered by $\mathcal{A}$.

One can think of the VC-dimension of $\mathcal{A}$ as a measure of the "richness" of the collection $\mathcal{A}$. A set S is "shattered" by $\mathcal{A}$ if $\mathcal{A}$ is rich enough to distinguish between all possible subsets of S. Note that S is shattered by $\mathcal{A}$ if (and only if) one can "pick off" every possible subset of S by intersecting S with an appropriately chosen set $A \in \mathcal{A}$. In this respect, perhaps "completely distinguished" or "completely discriminated" would be a better term than "shattered." However, the latter term has by now become standard in the literature. Note that, if $\mathcal{A}$ has finite VC-dimension, say d, then it is not rich enough to distinguish all subsets of *any* set containing $d + 1$ elements or more; but it is rich enough to distinguish all subsets of *some* set containing d elements (but not necessarily *all* sets of cardinality d).

From the definition it is clear that, in order to shatter a set S of cardinality n, the concept class $\mathcal{A}$ must contain at least 2^n distinct concepts. Thus an immediate corollary of the definition is that every finite collection of sets has finite VC-dimension, and that

$$\text{VC-dim}(\mathcal{A}) \leq \lg(|\mathcal{A}|) .$$

The importance of the VC-dimension in PAC learning theory arises from the following theorem, which is due to Blumer et al.; see [3]. In the statement of the theorem, we use the notion of a "consistent" algorithm. An algorithm is said to be "consistent" if it perfectly reproduces the training data. In other words, an algorithm is consistent if, for every labeled multisample $((x_1, I_T(x-1), \ldots, (x_m, I_T(x_m))$, we have that

$$I_{H_m}(x_i) = I_T(x_i) \forall i ,$$

where H_m is the hypothesis produced by the algorithm.

Theorem 1 *Suppose the concept class C has finite VC-dimension; then C is PAC-learnable for every probability measure over X. Let d denote the VC-dimension of C, and let $\{A_m\}$ be any consistent algorithm. Then the algorithm is PAC; moreover, the quantity $r(m, \epsilon)$ defined in (4) is bounded by*

$$r(m, \epsilon) \leq 2 \left(\frac{2em}{d} \right)^d 2^{-m\epsilon/2} , \tag{6}$$

for every probability measure over X. Any consistent algorithm learns to accuracy ϵ and confidence δ provided at least

$$m \geq \max \left\{ \frac{8d}{\epsilon} \lg \frac{8e}{\epsilon}, \frac{4}{\epsilon} \lg \frac{2}{\delta} \right\} \tag{7}$$

samples are used.

Thus Theorem 1 shows that the finiteness of the VC-dimension of the concept class C is a sufficient condition for so-called "distribution-free" PAC learnability. Moreover, (7) gives an estimate of the sample complexity, that is, the number of samples that suffice to achieve an accuracy of ϵ and a confidence of δ. Note that the sample complexity estimate is $O(1/\epsilon \log(1/\epsilon))$ for fixed δ, and $O(\log(1/\delta))$ for fixed ϵ. This is often summarized by the statement "confidence is cheaper than accuracy."

There is a converse to Theorem 1. It can be shown that, in order for a concept class C to be PAC learnable for *every* probability measure on the set X, and for the learning rate to be uniformly bounded with respect to the underlying probability measure, the concept class C must have finite VC-dimension.

Theorem 1 provides not only a condition for PAC-learnability but also relates the learning rate to the VC-dimension of the concept class. Thus it is clearly worthwhile to determine upper bounds for the VC-dimension of various concept classes. Over the years, several researchers have studied this problem, using a wide variety of approaches, ranging from simple counting arguments to extremely sophisticated mathematical reasoning based on advanced ideas such as algebraic topology, model theory of real numbers, and so on. Since the list of relevant references would be too long, the reader is referred instead to Chap. 10 of [15,16] which contain a derivation of many of the important known results together with citations of the original sources.

The Complexity of Learning

In this section, we discuss briefly the complexity of a learning problem. Complexity can arise in two ways: sample complexity, and algorithmic complexity. Sample complexity refers to the rate at which the number $m_0(\epsilon, \delta)$ grows with respect to $1/\epsilon$ and $1/\delta$, where $m_0(\epsilon, \delta)$ is the smallest value of m_0 such that (5) holds. Algorithmic complexity refers to the complexity of computing the hypothesis H_m

as a function of the number of samples of m, Each kind of complexity has its own interesting features.

Sample Complexity

Let us begin by discussing the notions of covering numbers, and packing numbers. Given a concept class C consisting of subsets of a given set X and a probability measure P on X, let us observe that the quantity d_P defines a pseudometric on the collection of (measurable) subsets of X, as follows: Given two subsets $A, B \subseteq X$, define

$$d_P(A, B) = P(A \Delta B) ,$$

where $A \Delta B$ is the symmetric difference between the sets A and B. Note that d_P is only a pseudometric because $d_P(A, B) = 0$ whenever $A \Delta B$ is a set of measure zero – it need not be the empty set (meaning that $A = B$). However, this is a minor technicality, and we can still refer to d_P as defining a "distance" between subsets of X.

Given the concept class C, the probability measure P (which in turn defines the pseudometric d_P), and a number ϵ, we say that a finite collection $\{C_1, \ldots, C_n\}$ where each $C_i \in C$ is an **ϵ-cover** of C if, for every $A \in C$, we have $d_P(A, C_i) \leq \epsilon$ for at least one index i. To put it in other words, the collection $\{C_1, \ldots, C_n\}$ is an ϵ-cover of C under the pseudometric d_P if every set $A \in C$ is within a distance of ϵ from at least one of the sets C_i. The **ϵ-covering number** $N(\epsilon, C, d_P)$ is defined to be the smallest number n such that there exists an ϵ-cover of cardinality n. Note that, in some cases, an ϵ-cover may not exist, in which case we take the ϵ-covering number to be infinity.

The next, and related, notion is the ϵ-packing number of a concept class. A finite collection $\{C_1, \ldots, C_m\}$ where each $C_i \in C$ is said to be **ϵ-separated** if $d_P(C_i, C_j) > \epsilon \; \forall i \neq j$. In other words, the finite collection $\{C_1, \ldots, C_m\}$ is ϵ-separated if the concepts are pairwise at a distance of more than ϵ from each other. The **ϵ-packing number** $M(\epsilon, C, d_P)$ is defined to be the largest number m such that there exists an ϵ-separated collection of cardinality m. It is not difficult to show that

$$\begin{aligned} M(2\epsilon, C, d_P) &\leq N(\epsilon, C, d_P) \\ &\leq M(\epsilon, C, d_P) \quad \forall \epsilon > 0 . \end{aligned}$$

See for instance Lemma 2.2 of [16].

The next two theorems, due to Benedek and Itai [2], are fundamental results on sample complexity.

Theorem 2 *Suppose C is a given concept class, and let $\epsilon > 0$ be specified. Then C is PAC-learnable to accuracy provided the covering number $N(\epsilon/2, C, d_P)$ is finite. Moreover, there exists an algorithm that is PAC with sample complexity*

$$m_0(\epsilon, \delta) \leq \frac{32}{\epsilon} \ln \frac{N(\epsilon/2, C, d_P)}{\delta} .$$

Thus a concept class with a finite covering number $N(\epsilon/2, C, d_P)$ is learnable to accuracy ϵ. The next theorem provides a kind of converse.

Theorem 3 *Suppose C is a given concept class, and let $\epsilon > 0$ be specified. Then any algorithm that is PAC to accuracy ϵ requires at least $\lg M(2\epsilon, C, d_P)$ samples, where $M(2\epsilon, C, d_P)$ denotes the 2ϵ-packing number of the concept class C with respect to the pseudometric d_P.*

One consequence of the theorem is that a concept class C is learnable to accuracy ϵ only if $M(2\epsilon, C, d_P)$ is finite. This fact is used to construct examples of "non-learnable" concept classes; see for instance Example 6.10 in [16]. However, the question of interest here is whether there exists a concept class that is learnable, but for which the sample complexity $m_0(\epsilon, \delta)$ grows faster than any polynomial in $1/\epsilon$.

In Sect. 3.3 in [7], the author constructs a concept class with the property that

$$M(1/n, C, d_P) \geq 2^{2^{c \cdot n^{1/4}}} ,$$

where c is some constant. Thus it follows from Theorem 3 that the sample complexity satisfies

$$m_0(\epsilon/2, \delta) \geq 2^{c \cdot (1/\epsilon)^{1/4}} , \quad \forall \epsilon .$$

So it is after all possible to have a situation where a concept class is learnable, but the sample complexity grows faster than any polynomial in the reciprocal accuracy $1/\epsilon$.

Algorithmic Complexity

We have seen from Theorem 1 that if a concept class C has finite VC-dimension, then every *consistent* algorithm is PAC. In the original paper [12], one studies not a single concept class C, but rather an indexed family of concept classes $\{C_n\}$, for every integer n. The objective is to determine whether there exists an algorithm that PAC-learns each concept class, such that the computational complexity of learning C_n is polynomial in n. Note that we are *not* referring here to the sample complexity. In most problems in computer science, each concept class C_n is finite and therefore automatically has finite VC-dimension. Rather, we are referring to the difficulty of constructing a consistent hypothesis, and the computational complexity of this task as a function of n. In the original paper [12], it is shown that the problem of learning 3-CNF (3 conjunctive

normal forms) is PAC learnable in this more restrictive sense. Note that 3-CNF in n Boolean variables consists of all conjunctions of clauses in which each clause is a disjunction of no more than three Boolean variables or their negations. For instance, if $n = 5$, then

$$(x_1 \vee \neg x_2 \vee x_5) \wedge (x_2 \vee x_3 \vee \neg x_4) \wedge (\neg x_1 \vee \neg x_3) \wedge (\neg x_2 \vee \neg x_3 \vee \neg x_5)$$

belongs to 3-CNF. In this instance, Valiant shows that it is possible to construct a consistent hypothesis in polynomial time in n.

Shortly after the publication of this paper, a very simple example was produced of a learning problem in which it is NP-hard to construct a consistent hypothesis. Specifically, consider the family of all conjunctions of at most k clauses, where each clause is a disjunction of various literals or their negations (without limit). For instance, if $n = 5$, then

$$(x_1 \vee \neg x_2 \vee \neg x_3 \vee x_5) \wedge (\neg x_1 \vee x_3 \vee \neg x_4 \vee \neg x_5) \wedge (\neg x_1 \vee x_2 \vee x_4 \vee \neg x_5)$$

is an example of this kind of formula. In this case, it can be shown that, for each fixed k and n, the VC-dimension of the concept class is polynomial in k and n. Thus it follows from Theorem 1 that the *sample complexity* of learning is polynomial. However, it is shown in [11] that the *computational complexity* of finding a consistent hypothesis is intractable unless $P = NP$. In contrast, if each such Boolean formula is expressed as a 3-DNF, then it is possible to construct a consistent hypothesis in polynomial time with respect to n.

System Identification as a Learning Problem

The Need for a Quantitative Identification Theory

By tradition, identification theory is *asymptotic* in nature. In other words, the theory analyzes what happens as the number of samples approaches infinity. However, there are situations where it is desirable to have *finite time*, or nonasymptotic, estimates of the rate at which the output of the identification process converges to the best possible model. For instance, in "indirect" adaptive control, one first carries out an identification of an unknown system, and after a finite amount of time has elapsed, designs a controller based on the current model of the unknown system. If one can quantify just how close the identified model is to the true but unknown system *after a finite number of samples*, then these estimates can be combined with robust control theory to ensure that indirect adaptive control leads to closed-loop stability.

Among the first papers to state that the derivation of finite-time estimates is a desirable property in itself are Weyer, Williamson and Mareels; see [20,21]. In those papers, it is assumed that the time series to which the system identification algorithm is applied consists of so-called "finitely-dependent" data; in other words, it is assumed that the time series can be divided into blocks that are independent. In a later paper [19], the author replaced the assumption of finite dependence by the assumption that the time series is β-mixing. However, he did not derive conditions under which a time series is β-mixing. In a still later paper [4], the authors study the case where the data is generated by general linear systems (Box–Jenkins models) driven by i.i.d. Gaussian noise, and the model family also belongs to the same class. In this paper, the authors say that they are motivated by the observation that "signals generated by dynamical systems are not β-mixing in general". This is why their analysis is limited to an extremely restricted class. However, their statement that dynamical systems do not generate β-mixing sequences is simply incorrect. Specifically, the all the systems studied in [4] are themselves β-mixing! In another paper [18], the present author has shown that *practically any exponentially stable system driven by i.i.d. noise with bounded variance* is β-mixing. Hence, far from being restrictive, the results of [19] have broad applicability, though he did not show it at the time.

Problem Formulation

System identification viewed as a learning problem can be stated as follows: One is given a time series $\{(u_t, y_t)\}$, where u_t denotes the input to the unknown system at time t, and y_t denotes the output at time t. We have that $u_t \in U \subseteq \mathbb{R}^l$, $y_t \in Y \subseteq \mathbb{R}^k$ for all t, where for simplicity it is assumed that both U and Y are compact sets. One is also given a family of models $\{h(\theta), \theta \in \Theta\}$ parametrized by a parameter vector θ belonging to a set Θ. Usually Θ is a subset of a finite-dimensional Euclidean space. At time t, the data available to the modeler consists of all the past measurements until time $t-1$. Based on these measurements, the modeler chooses a parameter θ_t, with the objective of making the best possible prediction of the next measured output y_t. The method of choosing θ_t is called the identification algorithm.

To make this formulation a little more precise, let us introduce some notation. Define $\mathcal{U} := \prod_{-\infty}^{\infty} U$, and define $\mathcal{Y}$ analogously. Note that the input sequence $\{u_t\}$ belongs to the set $\mathcal{U}$, while the output sequence belongs to $\mathcal{Y}$. Let $\mathcal{U}_{-\infty}^{0}$ denote the one-sided infinite cartesian product $\mathcal{U}_{-\infty}^{0} := \prod_{-\infty}^{0} U$, and for a given two-sided infinite se-

quence $\mathbf{u} \in \mathcal{U}$, define

$$\mathbf{u}_t := (u_{t-1}, u_{t-2}, u_{t-3}, \dots) \in \mathcal{U}_{-\infty}^0 .$$

Thus the symbol $\mathbf{u}_t$ denotes the infinite past of the input signal $\mathbf{u}$ at time t. The family of models $\{h(\theta), \theta \in \Theta\}$ consists of a collection of maps $h(\theta), \theta \in \Theta$, where each $h(\theta)$ maps $\mathcal{U}_{-\infty}^0$ into Y. Thus, at time t, the quantity $h(\theta) \cdot \mathbf{u}_t =: \hat{y}_t(\theta)$ is the "predicted" output if the model parameter is chosen as θ. The quality of this prediction is measured by a "loss function" $\ell \colon Y \times Y \to [0, 1]$. Thus $\ell(y_t, h(\theta) \cdot \mathbf{u}_t)$ is the loss we incur if we use the model $h(\theta)$ to predict the output at time t, and the actual output is y_t.

Since we are dealing with a time series, all quantities are *random*. Hence, to assess the quality of the prediction made using the model $h(\theta)$, we should take the *expected value* of the loss function $\ell(y_t, h(\theta) \cdot \mathbf{u}_t)$ with respect to the law of the time series $\{(u_t, y_t)\}$, which we denote by $\tilde{P}_{\mathbf{u},\mathbf{y}}$. Thus we define the objective function

$$J(\theta) := E\left[\ell(y_t, h(\theta) \cdot \mathbf{u}_t), \tilde{P}_{\mathbf{u},\mathbf{y}}\right] . \tag{8}$$

Note that $J(\theta)$ depends solely on θ and nothing else.

A key observation at this stage is that the probability measure $\tilde{P}_{\mathbf{u},\mathbf{y}}$ is *unknown*. This is because, if the statistics of the time series are known ahead of time, then there is nothing to identify!

Definition 3 System Identification Problem: Given the time series $\{(u_t, y_t)\}$ with *unknown* law $\tilde{P}_{\mathbf{u},\mathbf{y}}$, construct if possible an iterative algorithm for choosing θ_t as a function of t, in such a way that

$$J(\theta_t) \to \inf_{\theta \in \Theta} J(\theta) .$$

A General Result

Since the law $\tilde{P}_{\mathbf{u},\mathbf{y}}$ of the time series is not known, the objective function $J(\theta)$ cannot be computed exactly. To circumvent this difficulty, one replaces the "true" objective function $J(\cdot)$ by an "empirical approximation," as defined next. For each $t \geq 1$ and each $\theta \in \Theta$, define the empirical error

$$\hat{J}_t(\theta) := \frac{1}{t} \sum_{i=1}^{t} \ell[y_i, h(\theta) \cdot \mathbf{u}_i] . \tag{9}$$

If we could choose $\ell(y, z) = \| y - z \|^2$, then

$$\hat{J}_t(\theta) = \frac{1}{t} \sum_{i=1}^{t} \| y_i - \hat{y}_i(\theta) \|^2$$

would be the average cumulative mean-squared error between the actual output y_i and the predicted error $\hat{y}_i(\theta)$, from time 1 to time t. Note that the least squares error is in general unbounded, whereas the loss function ℓ is supposed to map into the bounded interval $[0, 1]$. To circumvent this technical difficulty, we can choose

$$\hat{J}_t(\theta) = \frac{1}{t} \sum_{i=1}^{t} \eta\big(\| y_i - \hat{y}_i(\theta) \|^2 \big) ,$$

where $\eta(\cdot)$ is any kind of saturating nonlinear function, for instance $\eta(s) = \tanh(s)$. This corresponds to the choice $\ell(y, z) = \tanh(\| y - z \|^2)$.

Note that, unlike the quantity $J(\theta)$, the function $\hat{J}(\theta)$ *can* be computed on the basis of the available data. Hence, in principle at least, it is possible to choose θ_t so as to minimize the "approximate" (but computable) objective function $\hat{J}(\theta)$ in the hope that, by doing so, we will somehow minimize the "true" (but uncomputable) objective function $J(\theta)$.

The next theorem gives some sufficient conditions for this approach to work. Specifically, if the estimates $\hat{J}_t(\theta)$ converge *uniformly* to the correct values $J(\theta)$ as $t \to \infty$, where the uniformity is with respect to $\theta \in \Theta$, then the above approach works. The importance of this uniform convergence property was recognized very early on in identification theory. In particular, Caines [5,6] uses ergodic theory to establish this property, whereas Ljung [8] takes a more direct approach.

Theorem 4 *At time t, choose θ_t^* so as to minimize $\hat{J}_t(\theta)$; that is,*

$$\theta_t^* = \arg\min_{\theta \in \Theta} \hat{J}_t(\theta) .$$

Let

$$J^* := \inf_{\theta \in \Theta} J(\theta) .$$

Define the quantity

$$q(t, \epsilon) := \tilde{P}_{\boldsymbol{u},\boldsymbol{y}}\{ \sup_{\theta \in \Theta} |\hat{J}_t(\theta) - J(\theta)| > \epsilon \} . \tag{10}$$

Suppose it is the case that $q(t, \epsilon) \to 0$ as $t \to \infty, \forall \epsilon > 0$. Then

$$\tilde{P}_{\boldsymbol{u},\boldsymbol{y}}\{J(\theta_t^*) > J^* + \epsilon\} \to 0 \quad \text{as } t \to \infty, \forall \epsilon > 0 . \tag{11}$$

In other words, the quantity $J(\theta_t^*)$ converges to the optimal value J^* in probability, with respect to the measure $\tilde{P}_{\boldsymbol{u},\boldsymbol{y}}$.

Corollary 1 *Suppose that $q(t, \epsilon) \to 0$ as $t \to \infty, \forall \epsilon > 0$. Given $\epsilon, \delta > 0$, choose $t_0 = t_0(\epsilon, \delta)$ such that*

$$q(t, \epsilon) < \delta \forall t \geq t_0(\epsilon, \delta) . \tag{12}$$

Then

$$\tilde{P}_{u,y}\{J(\theta_t^*) > J^* + \epsilon\} < \delta \forall t \geq t_0(\epsilon/3, \delta)\,. \tag{13}$$

Proofs of the theorem and corollary can be found in [17].

A Result on the Uniform Convergence of Empirical Means

The property of Theorem 4 does indeed hold in the commonly studied case where y_t is the output of a "true" system corrupted by additive noise, and the loss function ℓ is the squared error, provided the model family and the true system satisfy some fairly reasonable conditions, namely: each system is BIBO (bounded input, bounded output) stable, all systems in the model family have exponentially decaying memory, and finally, an associated collection of input-output maps has finite P-dimension. Note that the P-dimension is a generalization of the VC-dimension to the case of maps assuming values in a bounded interval $[0, 1]$. By identifying a set with its support function, we can think of the VC-dimension as a combinatorial parameter associated with binary-valued maps. The P-dimension is defined in Chap. 4 of [15,16]

Define the collection of functions $\mathcal{H}$ mapping $\mathcal{U}$ into $\mathbb{R}$ as follows:

$$g(\theta) := \mathbf{u} \mapsto \| (f - h(\theta)) \cdot \mathbf{u}_0 \|^2 \colon \mathcal{U} \to \mathbb{R}\,,$$

$$\mathcal{G} := \{g(\theta) \colon \theta \in \Theta\}\,.$$

Now the various assumptions are listed.

A1. There exists a constant M such that

$$|g(\theta) \cdot \mathbf{u}_0| \leq M, \forall \theta \in \Theta, \mathbf{u} \in \mathcal{U}\,.$$

This assumption can be satisfied, for example, by assuming that the true system and each system in the family $\{h(\theta), \theta \in \Theta\}$ is BIBO stable (with an upper bound on the gain, independent of θ), and that the set U is bounded (so that $\{u_t\}$ is a bounded stochastic process).

A2. For each integer $k \geq 1$, define

$$g_k(\theta) \cdot \mathbf{u}_t := g(\theta) \cdot (u_{t-1}, u_{t-2}, \ldots, u_{t-k}, 0, 0, \ldots)\,. \tag{14}$$

With this notation, define

$$\mu_k := \sup_{\mathbf{u} \in \mathcal{U}} \sup_{\theta \in \Theta} |(g(\theta) - g_k(\theta)) \cdot u_0|\,.$$

Then the assumption is that μ_k is finite for each k and approaches zero as $k \to \infty$. This assumption essentially means that each of the systems in the model family has decaying memory (in the sense that the effect of the values of the input at the distant past on the current output becomes negligibly small).

A3. Consider the collection of maps $\mathcal{G}_k = \{g_k(\theta) \colon \theta \in \Theta\}$, viewed as maps from U^k into $\mathbb{R}$. For each k, this family $\mathcal{G}_k$ has finite P-dimension, denoted by $d(k)$.

Now we can state the main theorem.

Theorem 5 *Define the quantity $q(t, \epsilon)$ as in (10) and suppose Assumptions A1 through A3 are satisfied. Given an $\epsilon > 0$, choose $k(\epsilon)$ large enough that $\mu_k \leq \epsilon/4$ for all $k \geq k(\epsilon)$. Then for all $t \geq k(\epsilon)$ we have*

$$q(t, \epsilon) \leq 8k(\epsilon) \left(\frac{32e}{\epsilon} \ln \frac{32e}{\epsilon} \right)^{d(k(\epsilon))} \cdot \exp(-\lfloor t/k(\epsilon) \rfloor \epsilon^2/512M^2)\,, \tag{15}$$

where $\lfloor t/k(\epsilon) \rfloor$ denotes the largest integer part of $t/k(\epsilon)$.

Theorem 6 *Let all notation be as in Theorem 5. Then, in order to ensure that the current estimate θ_t satisfies the inequality $J(\theta_t) \leq J^* + \epsilon$ (i. e., is ϵ-optimal) with confidence $1 - \delta$, it is enough to choose the number of samples t large enough that*

$$\lfloor t/k(\epsilon) \rfloor \geq \frac{512M^2}{\epsilon^2} \left[\ln \left(\frac{24k(\epsilon)}{\delta} \right) + d(k(\epsilon)) \ln \left(\frac{32e}{\epsilon} \right) + d(k(\epsilon)) \ln \ln \left(\frac{32e}{\epsilon} \right) \right]. \tag{16}$$

Bounds on the P-Dimension

In order for the estimate in Theorem 5 to be useful, it is necessary for us to derive an estimate for the P-dimension of the family of functions defined by

$$\mathcal{G}_k := \{g_k(\theta) \colon \theta \in \Theta\}\,, \tag{17}$$

where $g_k(\theta) \colon U^k \to \mathbb{R}$ is defined by

$$g_k(\theta)(\mathbf{u}) := \| (f - h(\theta)) \cdot \mathbf{u}_k \|^2\,,$$

where

$$\mathbf{u}_k := (\ldots, 0, u_k, u_{k-1}, \ldots, u_1, 0, 0, \ldots)\,.$$

Note that, in the interests of convenience, we have denoted the infinite sequence with only k nonzero elements as $u_k, \ldots, u_1$ rather than $u_0, \ldots, u_{1-k}$ as done earlier. Clearly this makes no difference. In this section, we state and prove such an estimate for the commonly occurring case where each system model $h(\theta)$ is an ARMA model where the parameter θ enters linearly. Specifically, it is

supposed that the model $h(\theta)$ is described by

$$x_{t+1} = \sum_{i=1}^{l} \theta_i \, \phi_i(x_t, u_t), \;\; y_t = x_t \,, \tag{18}$$

where $\theta = (\theta_1, \ldots, \theta_l) \in \Theta \subseteq \mathbb{R}^l$, and each $\phi_i(\cdot,\cdot)$ is a polynomial of degree no larger than r in the components of x_t, u_t.

Theorem 7 *With the above assumptions, we have that*

$$\begin{aligned} P\text{-}dim(\mathcal{G}_k) &\leq 9l + 2l \lg[2(r^{k+1} - 1)/(r-1)] \\ &\approx 9l + 2lk \lg(2r) \quad \textit{if} \quad r > 1 \,. \end{aligned} \tag{19}$$

In case $r = 1$ so that each system is linear, the above bound can be simplified to

$$P\text{-}dim(\mathcal{G}_k) \leq 9l + 2l \lg(2k) \,. \tag{20}$$

It is interesting to note that the above estimate is *linear* in both the number of parameters l and the duration k of the input sequence **u**, but is only logarithmic in the degree of the polynomials ϕ_i. In the practically important case of linear ARMA models, even k appears inside the logarithm.

Future Directions

At present, system identification theory is a mature, widely used subject. Statistical learning theory is newer and still needs to find its feet. In particular, the various estimates for sample complexity given here are deemed to be "too conservative" to be of practical use. Thus considerable research needs to be carried out on the issue of sample complexity. The interplay between system identification theory and statistical learning theory perhaps offers the greatest scope for interesting new work. At present, statistical learning theory allows us to derive *finite time estimates* of the disparity between the current estimates of a system, and the true but unknown system. However, these estimates do not always translate readily into a metric "distance" between the true but unknown system and the current estimate. In order to be useful in the context of robust controller design, it is necessary to derive estimates for some kind of metric distance. This is an important topic for future research.

Acknowledgment

The author thanks Erik Weyer for his careful reading and comments on an earlier draft of this article.

Bibliography

Primary Literature

1. Angluin D (1987) Queries and concept learning. Mach Learn 2:319–342
2. Benedek GM, Itai A (1988) Learnability by fixed distributions. First Workshop on Computational Learning Theory. Morgan-Kaufmann, San Mateo, pp 80–90
3. Blumer A, Ehrenfeucht A, Haussler D, Warmuth M (1989) Learnability and the Vapnik–Chervonenkis dimension. J ACM 36(4):929–965
4. Campi MC, Weyer E (2002) Finite sample properties of system identification methods. IEEE Trans Auto Control 47:1329–1334
5. Caines PE (1976) Prediction error identification methods for stationary stochastic processes. IEEE Trans Auto Control AC-21:500–505
6. Caines PE (1978) Stationary linear and nonlinear system identification and predictor set completeness. IEEE Trans Auto Control AC-23:583–594
7. Hammer B (1999) Learning recursive data, Math of Control. Signals Syst 12(1):62–79
8. Ljung L (1978) Convergence analysis of parametric identification methods. IEEE Transa Auto Control AC-23:77–783
9. Ljung L (1999) System Identification: Theory for the User. Prentice-Hall, Englewood Cliffs
10. Pintelon R, Schoukens J (2001) System Identification: A frequency domain approach. IEEE press, Piscataway
11. Pitt L, Valiant L (1988) Computational limits on learning from examples. J ACM 35(4):965–984
12. Valiant LG (1984) A theory of the learnable. Comm ACM 27(11):1134–1142
13. Vapnik VN (1982) Estimation of Dependences Based on Empirical Data. Springer, Heidelberg
14. Vapnik VN, Chervonenkis AY (1971) On the uniform convergence of relative frequencies to their probabilities. Theory Prob Appl 16(2):264–280
15. Vidyasagar M (1997) A Theory of Learning and Generalization. Springer, London
16. Vidyasagar M (2003) Learning and Generalization: With Applications to Neural Networks and Control Systems. Springer, London
17. Vidyasagar M, Karandikar RL (2001) System identification: A learning theory approach. Proc IEEE Conf on Decision and Control, Orlando, pp 2001–2006
18. Vidyasagar M, Karandikar RL (2001) A learning theory approach to system identification and stochastic adaptive control. In: Calafiore G, Dabbene F (eds) Probabilistic and Randomized Algorithms for Design Under Uncertainty. Springer, Heidelberg, pp 265–302
19. Weyer E (2000) Finite sample properties of system identification of ARX models under mixing conditions. Automatica 36(9):1291–1299
20. Weyer E, Williamson RC, Mareels I (1996) Sample complexity of least squares identification of FIR models. Proc 13th World Congress of IFAC, San Francisco, pp 239–244
21. Weyer E, Williamson RC, Mareels I (1999) Finite sample properties of linear model identification. IEEE Trans Auto Control 44(7):1370–1383

Books and Reviews

Anthony M, Biggs N (1992) Computational Learning Theory. Cambridge University Press, Cambridge
Devroye L, Györfi L, Lugosi G (1996) A probabilistic theory of pattern recognition. Springer, New York
Haussler D (1992) Decision theoretic generalizations of the PAC

model for neural net and other learning applications. Info Comput 100:78–150
Kearns M, Vazirani U (1994) Introduction to Computational Learning Theory. MIT Press, Cambridge
Kulkarni SR, Vidyasagar M (2001) Learning decision rules under a family of probability measures. IEEE Trans Info Thy
Natarajan BK (1991) Machine Learning: A Theoretical Approach. Morgan-Kaufmann, San Mateo
van der Vaart AW, Wallner JA (1996) Weak convergence and Empirical Processes. Springer, New York
Vapnik AN, Chervonenkis AY (1981) Necessary and sufficient conditions for the uniform convergence of means to their expectations. Theory Prob Appl 26(3):532–553
Vapnik VN (1995) The Nature of Statistical Learning Theory. Springer, New York
Vapnik VN (1998) Statistical Learning Theory. Wiley, New York
Weyer E, Campi MC (2002) Non-asymptotic confidence ellipsoids for the least squares error estimate. Automatica 38:1529–1547

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion

Ralf Metzler[1], Aleksei V. Chechkin[2], Joseph Klafter[3,4]

[1] Physics Department, Technical University of Munich, Garching, Germany
[2] Institute for Theoretical Physics NSC KIPT, Kharkov, Ukraine
[3] School of Chemistry, Tel Aviv University, Tel Aviv, Israel
[4] Freiburg Institute for Advanced Studies (FRIAS), University of Freiburg, Freiburg, Germany

Article Outline

Glossary

Anomalous diffusion Deviations of the linear time dependence $\langle x^2(t)\rangle = 2Kt$, of the mean squared displacement in absence of an external bias, in the form of a power-law: $\langle x^2(t)\rangle = 2K_\zeta t^\zeta$. Here, K_ζ is the anomalous diffusion constant of dimension $\mathrm{cm}^2/\mathrm{sec}^\zeta$. In the range $0 < \zeta < 1$, we deal with subdiffusion, whereas $\zeta > 1$ describes superdiffusion. Lévy flights have a diverging mean squared displacement.

Continuous time random walk A theory that describes a random motion by assigning each jump a jump length x and a waiting time t elapsing between two successive jumps, drawn from the two probability densities $\lambda(x)$ and $\psi(t)$, respectively. The two densities λ and ψ fully specify the probability density function $P(x,t)$ describing the random process. In Fourier–Laplace space, the propagator follows as $P(k,u) = u^{-1}\psi(u)/[1-\psi(u)\lambda(k)]$.

Fourier and Laplace transforms Integral transforms that aid in convenient solution of Linear partial differential equations. The definitions of the fractional operators used in the text correspond to Laplace and Fourier convolutions, such that these transformations also become useful there. The Laplace transform of a function $f(t)$ is defined as

$$f(u) \equiv \mathcal{L}\{f(t);u\} = \int_0^\infty f(t)\mathrm{e}^{-ut}\mathrm{d}t\,,$$

while the Fourier transform of $g(x)$ reads

$$g(k) \equiv \mathcal{F}\{g(x);k\} = \int_{-\infty}^\infty g(x)\mathrm{e}^{ikx}\mathrm{d}x\,.$$

Note that we denote the transform of a function by explicit dependence on the respective variable. For $u = 0$ and $k = 0$ the Laplace and Fourier transform are the averages of the function f or g, respectively. Tauberian theorems ascertain relations between the original function and its transform. For instance, the small u behavior $f(u) \sim 1-(u\tau)^\alpha$ implies long time scaling $f(t) \sim \tau^\alpha/t^{1+\alpha}$. See [35,46] for details.

Fractional differintegration The multiple derivative of an integer power is $\mathrm{d}^m x^n/\mathrm{d}x^m = n!/(n-m)!x^{n-m}$ for $m \le n$. The result is zero if $m > n$. Replacing the factorials by the Γ-function, one can generalize this relation to ${}_0D_x^q x^p = x^{p-q}\Gamma(1+p)/\Gamma(1+p-q)$. In particular, this includes the fractional differentiation of a constant, ${}_0D_x^q 1 = x^{-q}/\Gamma(1-q)$, that no longer vanishes. Fractional differentiation was first mentioned by Leibniz in a letter to de l'Hospital in 1695. The Riemann–Liouville and Riesz-Weyl fractional operators used in the following are a straightforward generalization of the Cauchy multiple integral, followed by regular differentiation.

Lévy flight A special type of continuous time random walk. A Lévy flight's waiting time distribution is narrow, for instance, Poissonian with $\psi(t) = \tau^{-1}\exp(-t/\tau)$, and the resulting dynamics therefore Markovian. The jump length distribution of a Lévy

flight is long-tailed: $\lambda(x) \sim \sigma^\mu |x|^{-1-\mu}$, with $0 < \mu < 2$, such that no second moment exists. The resulting PDF is a Lévy stable law with Fourier transform $P(k, t) = \exp(-K^{(\mu)} t|k|^\mu)$.

Lévy stable laws The central limit theorem states that the properly normalized sum of independent, identically distributed random variables with finite variance converges to a Gaussian limit distribution. A generalization of this theorem exists for the case with infinite variance, namely, the generalized central limit theorem. The related distributions are the Lévy stable laws, whose density in the simplest case have the characteristic function $p(k) = \int_{-\infty}^{\infty} p(x) \exp(ikx) \mathrm{d}x$ (Fourier transform) of the form $p(k) = \exp(-c|k|^\mu)$, with $0 < \mu < 2$. In direct space, this corresponds to a power-law asymptotic $p(x) \simeq c|x|^{-1-\mu}$. In the limit $\mu = 2$, the universal Gaussian distribution is recovered.

Lévy walks In contrast to Lévy flights, Lévy walks possess a finite mean squared displacement, albeit having a broad jump length distribution. This is possible by the introduction of a time penalty for long jumps through a coupling $\lambda(x)p(t|x)$ between waiting times and jump lengths, such that long jumps involve a longer waiting time. For instance, a δ-coupling of the form $\frac{1}{2}\lambda(x)\delta(|x| - vt)$ is often chosen, such that v plays the role of a velocity.

Strange kinetics Often, deviations from exponential relaxation patterns $\exp(-t/\tau)$ and regular Brownian diffusion are observed. Instead, non-exponential relaxation, for instance, of the stretched exponential form $\exp(-[t/\tau]^\alpha)$ $(0 < \alpha < 1)$ or of the inverse power-law form $t^{-\zeta}$, is observed, or anomalous diffusion behavior is found.

Weak ergodicity breaking In a system with a broadly distributed waiting time with diverging characteristic time scale, a particle can get stuck at a certain position for a long time. For instance, for a waiting time distribution of the form $\psi(t) \sim \tau^\alpha t^{-1-\alpha}$, the probability of not moving until time t scales as $\tau^\alpha t^{-\alpha}$, that is, the probability of not moving is therefore appreciable even for long times. A particle governed by such a $\psi(t)$ even at stationarity does not equally explore different domains of phase space.

Definition of the Subject

Classical Brownian motion, characterized by a mean squared displacement

$$\langle x^2(t) \rangle = 2Kt \,, \tag{1}$$

growing linearly in time in absence of an external bias, is the paradigm for random motion. Note that we restrict our discussion to one dimension. Brownian motion quantifies the jittery motion of coal dust particles observed by Dutchman Jan Ingenhousz in 1785 [49], the zigzagging of pollen grains in solution reported by Robert Brown in 1827 [16], and possibly the dance of dust particles in the beam of sunlight in a stairwell so beautifully immortalized in the famed poem by Lucretius [17]. In general, Brownian motion occurs in simple, sufficiently homogeneous systems such as simple liquids or gasses. In the continuum limit, Brownian motion is governed by the diffusion equation

$$\frac{\partial P(x,t)}{\partial t} = K \frac{\partial^2}{\partial x^2} P(x,t) \tag{2}$$

for the probability density function (PDF) $P(x, t)$ describing the probability $P(x, t)\mathrm{d}x$ of finding the particle at a position in the interval $x, \ldots, x + \mathrm{d}x$ at time t. For a point-like initial condition $P(x, 0) = \delta(x)$, the solution becomes the celebrated Gaussian PDF

$$P(x,t) = \frac{1}{\sqrt{4\pi K t}} \exp\left(-\frac{x^2}{4Kt}\right) . \tag{3}$$

The diffusion constant K fulfills the Einstein–Stokes relation $K = k_B T/(m\eta)$, where $k_B T = RT/N_A$ is the Boltzmann energy at temperature T, R is the gas constant, N_A the Avogadro number, m is the mass of the test particle, and η the friction coefficient. This relation between microscopic and macroscopic quantities was used for the determination of the Avogadro number by Perrin [90]. Examples of Perrin's recorded random walks and the jump length distribution constructed from the data for a time increment of 30 seconds are shown in Fig. 1. Figure 2 for two different ambient pressures displays time traces of the Brownian motion of a small mirror in air that led to an amazingly accurate determination of the Avogadro number by Kappler [51].

However, in many systems deviations from linear behavior (1) are observed [14,25,55,78,80,107,122]. These deviations can assume the power-law form

$$\langle x^2(t) \rangle = 2K_\zeta t^\zeta \,. \tag{4}$$

In such systems, for $0 < \zeta < 1$, we observe subdiffusion. Prominent examples of subdiffusion include charge carrier transport in amorphous semiconductors [92,101], diffusion of chemicals in subsurface aquifers [102], the motion of beads in actin gels [121], motion in chaotic maps [8,123], or the subdiffusion of biomacromolecules or granules in cells [41,110]; just to name a few (compare [80]

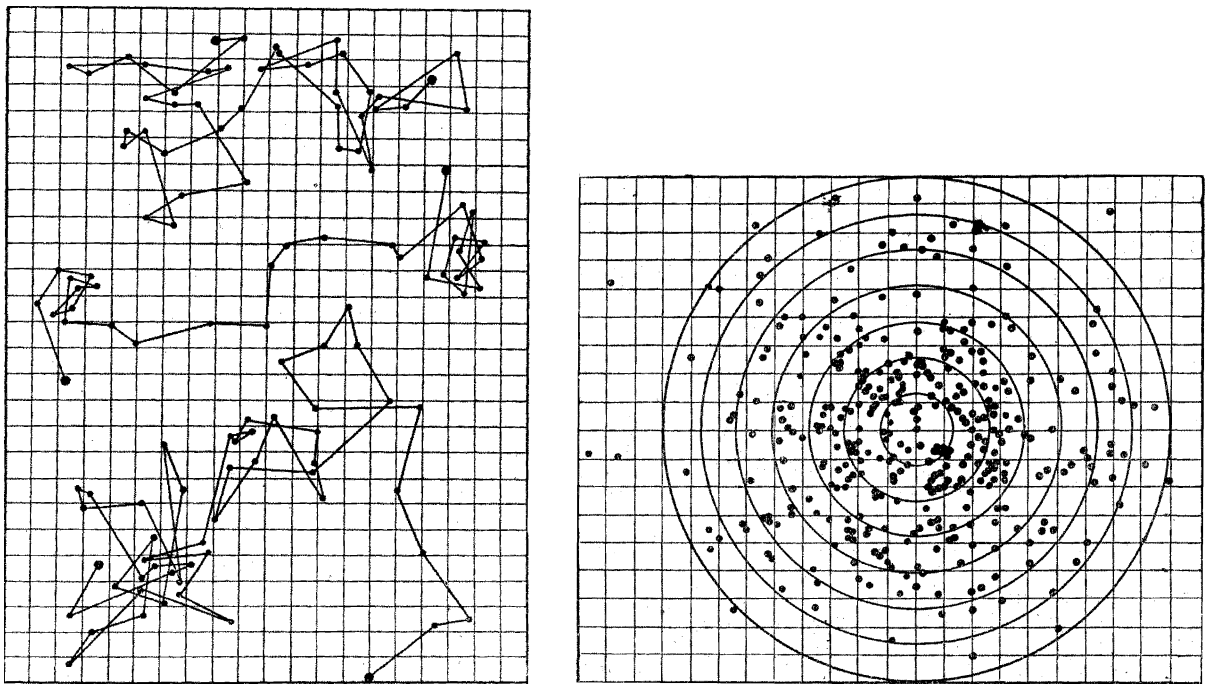

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 1
Random walk trajectories of putty particles in water recorded by Perrin [91]. *Left*: three designs obtained by tracing a small grain of putty at intervals of 30 sec. *Right*: the starting point of each motion event is shifted to the origin. The figure illustrates the continuum approach of the jump length distribution if only a large number of jumps is considered

for more details). Subdiffusion of this type is characterized by a long-tailed waiting time PDF $\psi(t) \simeq t^{-1-\alpha}$ $(0 < \alpha < 1)$, corresponding to the time-fractional diffusion equation

$$\frac{\partial P(x,t)}{\partial t} = K_\alpha \, {}_0D_t^{1-\alpha} \frac{\partial^2}{\partial x^2} P(x,t) \,, \tag{5}$$

with the fractional Riemann–Liouville operator [89,93,99]

$${}_0D_t^{1-\alpha} P(x,t) = \frac{1}{\Gamma(\alpha)} \frac{\partial}{\partial t} \int_0^t \frac{P(x,t')}{(t-t')^{1-\alpha}} \mathrm{d}t' \,. \tag{6}$$

From the latter definition, it becomes apparent that subdiffusion corresponds to a slowly decaying memory integral in the dynamical equation for $P(x,t)$.

The lack of a characteristic time scale (that is, the divergence of $\int_0^\infty t\psi(t)\mathrm{d}t$) of the waiting time PDF $\psi(t)$ $\simeq t^{-1-\alpha}$ no longer permits us to distinguish microscopic and macroscopic events. The result is that the diffusing particle can get stuck at a certain position for very long times, quantified by the sticking probability $\phi(t)$ $= \int_t^\infty \psi(t')\mathrm{d}t'$ of not moving. As detailed below the PDF $P(x,t)$ exhibits characteristic cusps at the location where the particle was initially released and, in the presence of an external drift, growing asymmetry. Moreover, a process whose measurement begins at $t = 0$ depends on the preparation time at some prior time, the so-called ageing effect [8,86]. Another effect in this context is that a particle no longer evenly distributes in a certain space, such that despite the existence of stationary states a weak ergodicity breaking occurs [11,13,45,65,66].

Contrasting with a subdiffusing particle are Lévy flights (LFs), that are based on a waiting time PDF with a finite characteristic time but a jump length distribution $\lambda(x) \simeq |x|^{-1-\mu}$ $(0 < \mu < 2)$ with diverging variance $\int_{-\infty}^{\infty} |x|^2 \lambda(x)\mathrm{d}x$. The diffusion equation for an LF becomes generalized to the space-fractional diffusion equation

$$\frac{\partial P(x,t)}{\partial t} = K^{(\mu)} \frac{\partial^\mu}{\partial |x|^\mu} P(x,t) \,, \tag{7}$$

where the fractional Riesz–Weyl operator is defined through its Fourier transform $\mathcal{F}\{\partial^\mu/\partial|x|^\mu \, P(x,t)\} = -|k|^\mu \, P(k,t)$. In Fourier space, one therefore recovers immediately the characteristic function $P(k,t) = \exp(-K^{(\mu)}|k|^\mu t)$ of a symmetric Lévy stable law. As can be seen from Eq. (7) LFs are Markovian processes. However, their trajectory now has a fractal dimension $d_f = \mu$ and

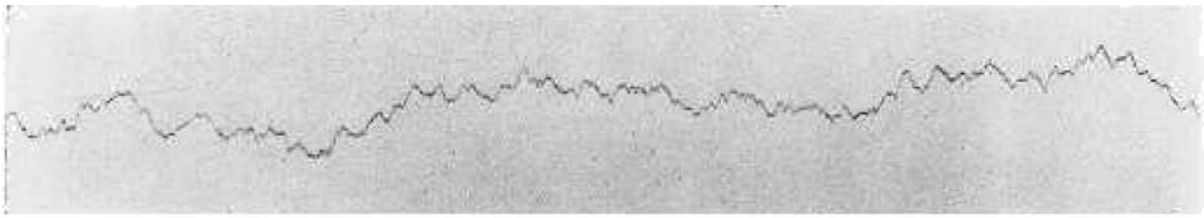

Registrieraufnahme der Brownschen Bewegung (natürlich Größe)
Direktionskraft 9,428 × 10^{-9} abs. Einh. Trägheitsmoment: 1×10^{-3} abs. Einh. Abstand Spiegel–Kamera: 72,1 cm.
Zeitmarke: 30 sec dx = 1 mm. a) Atmosphärendruck. Temperatur 13°C

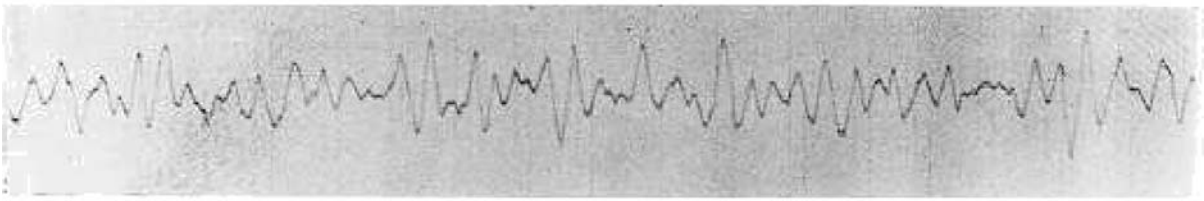

Registrieraufnahme der Brownschen Bewegung (natürlich Größe)
Direktionskraft 9,428 × 10^{-9} abs. Einh. Trägheitsmoment: 1×10^{-7} abs. Einh. Abstand Spiegel–Kamera: 72,1 cm.
Zeitmarke: 30 sec dx = 1 mm. b) 1×10^{-3}mm Hg. Temperatur 13°C

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 2
Erratic behavior of Brownian motion made visible in a high-precision measurement (including Kappler's original figure captions). Data from an Edelmann recorder, obtained by Kappler in 1931 [51]. Kappler monitored the Brownian motion of a small mirror (surface approx. 1mm^2), suspended from a fine quartz thread (several cm long and some μm thick). The mean squared of the torsional displacement, φ^2, follows the relation $D\varphi^2 = k_B T$, where D is the directional force of the suspension [51]. The facsimiles show two different realizations. From his data, Kappler obtained the Avogadro–Loschmidt number $N_L = 60.59 \cdot 10^{22} \pm 1\%$, to a remarkable accuracy

is characterized by local search changing with long excursions (see Fig. 3). This property has been shown to be a better search strategy than Brownian motion, as it oversamples less [64,116]. In fact, animals like albatross [115], spider monkeys [95], jackals [4], and even plankton [114] and bacteria [62] were claimed to follow Lévy search strategies. LFs also occur in diffusion in energy space [124], or in optical lattices [52]. Due to the clustering nature of their trajectory, LFs also exhibit a form of ergodicity breaking [68].

Introduction

For sums of independent, identically distributed random variables with proper normalization to the sample size, the generalized central limit theorem guarantees the convergence of the associated probability density to a Lévy stable density even though the variance of these random variables diverges [35,40,46,63,112]. Well-known examples of Lévy stable densities are the one-sided (defined for $x \geq 0$) Lévy–Smirnov distribution

$$f_{1/2,-1/2}(x) = \sqrt{\frac{1}{2\pi x^3}} \exp\left(-\frac{1}{2x}\right) , \tag{8}$$

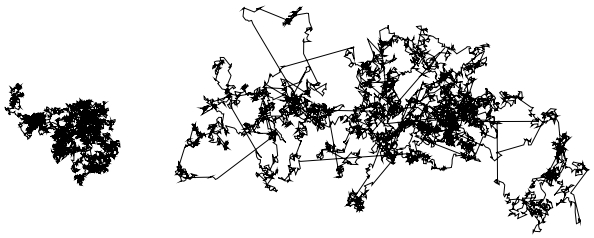

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 3
Trajectories of Brownian motion (*left*) and a Lévy flight of index $\mu = 1.5$ (*right*), both with the same number ($\simeq 7000$) of steps. The long sojourns and clustering appearance of the Lévy flight are distinct

related to the first passage time density of a Gaussian random walk process of passing the origin (see below), and the Cauchy (or Lorentz) distribution

$$f_{1,0}(x) = \frac{1}{\pi\left(1 + x^2\right)} . \tag{9}$$

In general, a Lévy stable density is defined through its

characteristic function of the PDF $f(x)$

$$\varphi(z) \equiv \mathcal{F}\{f(x)\} = \int_{-\infty}^{\infty} f_{\mu,\beta}(x) \mathrm{e}^{ikx} \mathrm{d}x \tag{10}$$

where

$$\log \varphi(z) = -|z|^{\mu} \exp\left\{ i \frac{\pi\beta}{2} \mathrm{sign}(z) \right\} , \tag{11}$$

for $\mu \neq 1$. Here, the skewness (or asymmetry) parameter β is restricted to the following region:

$$|\beta| \leq \begin{cases} \mu , & \text{if } 0 < \mu < 1 \\ 2 - \mu , & \text{if } 1 < \mu < 2 . \end{cases} \tag{12}$$

For $\beta = 0$, the corresponding Lévy stable density is symmetric around $x = 0$, while for $\beta = -\mu$ and $0 < \mu < 1$, it is one-sided. In general, a Lévy stable density follows the power-law asymptotic behavior

$$f_{\mu,\beta}(x) \sim \frac{A_{\mu,\beta}}{|x|^{1+\mu}} , \quad \mu < 2 , \tag{13}$$

with $A_{\mu,\beta}$ being a constant, such that for all Lévy stable densities with $\mu < 2$ the variance diverges

$$\langle x^2 \rangle = \infty . \tag{14}$$

Conversely, all fractional moments $\langle |x|^{\delta} \rangle < \infty$ for all $0 < \delta < \mu \leq 2$. From the above definitions it is obvious that the Lévy stable density $f_{2,0}$ corresponds to the Gaussian normal distribution

$$f_{2,0}(x) = \sqrt{\frac{1}{4\pi}} \exp\left(-\frac{1}{4} x^2 \right) \tag{15}$$

possessing finite moments of any order. In this limit, the generalized central limit theorem coincides with the more traditional, and universal, central limit theorem.

Brownian motion has traditionally been employed as the dominant model of choice for random noise in continuous-time systems, due to its remarkable statistical properties and its amenability to mathematical analysis. However, Brownian motion is just a single example of the Lévy family. Moreover, it is a very special and somewhat unrepresentative member of this family. Amongst the Lévy family, the Brownian member is the only motion with continuous sample-paths. All other motions have discontinuous trajectories, exhibiting jumps. Moreover, the Lévy family is characterized by selfsimilar motions. Brownian motion is the only selfsimilar Lévy motion possessing finite variance, while all other selfsimilar Lévy motions have an infinite variance.

Random processes whose spatial coordinates x or clock times t are distributed according to a Lévy stable density exhibit anomalies, that is, no longer follow the laws of Brownian motion. Consider a continuous time random walk process defined in terms of the jump length and waiting time distributions $\lambda(x)$ and $\psi(t)$ [87,101]. That is, each jump event of this random walk is characterized by a jump length x drawn from the distribution λ, and the time t between two jump events is distributed according to ψ. (Note that an individual jump is supposed to occur spontaneously.) In absence of an external bias, continuous time random walk theory connects $\lambda(x)$ and $\psi(t)$ with the probability distribution $P(x,t)\mathrm{d}x$ to find the random walker at a position in the interval $(x, x + \mathrm{d}x)$ at time t. In Fourier–Laplace space, $P(k,u) \equiv \mathcal{F}\{\mathcal{L}\{P(x,t); t \to u\}; x \to k\}$, this relation reads [54]

$$P(k,u) = \frac{1 - \psi(u)}{u} \frac{1}{1 - \lambda(k)\psi(u)} , \tag{16}$$

where $\mathcal{L}\{f(t)\} \equiv \int_0^{\infty} \exp(-ut) f(t) \mathrm{d}t$. We here neglect potential complications due to ageing effects. The following cases can be distinguished:

(i) $\lambda(x)$ is Gaussian with variance σ^2 and $\psi(t) = \delta(t - \tau)$. Then, to leading order in k^2 and u, respectively, one obtains $\lambda(k) \simeq 1 - \sigma^2 k^2$ and $\psi(u) \simeq 1 - u\tau$. From relation (16), one recovers the Gaussian probability density $P(x,t) = \sqrt{1/(4\pi K t)} \exp\{-x^2/(4Kt)\}$ with diffusion constant $K = \sigma^2/\tau$. The corresponding mean squared displacement grows linearly with time; see Eq. (1). This case corresponds to the continuum limit of regular Brownian motion. Note that here and in the following, we restrict the discussion to one dimension.

(ii) Assume $\lambda(x)$ still to be Gaussian, while for the waiting time distribution $\psi(t)$ we choose a one-sided Lévy stable density with stable index $0 < \alpha < 1$. Consequently, $\psi(u) \simeq 1 - (u\tau)^{\alpha}$, and the characteristic waiting time $\int_0^{\infty} t\psi(t)\mathrm{d}t$ diverges. Due to this lack of a time scale separating microscopic (single jump events) and macroscopic (on the level of $P(x,t)$) scales, $P(x,t)$ is no longer Gaussian, but is given by a more complex H-function [78,80,103]. In Fourier space, however, one finds the quite simple analytical form [78]

$$\begin{aligned} P(k,t) &= E_{\alpha}\left(-K_{\alpha} k^2 t^{\alpha}\right) \\ &= \sum_{0}^{\infty} \frac{\left(-K_{\alpha} k^2 t^{\alpha}\right)^n}{\Gamma(1 + \alpha n)} \end{aligned} \tag{17}$$

in terms of the Mittag–Leffler function. This generalized relaxation function of the Fourier modes turns

from an initial stretched exponential (KWW) behavior

$$P(k,t) \sim 1 - K_\alpha k^2 t^\alpha/\Gamma(1+\alpha) \sim \exp\left\{-K_\alpha k^2 t^\alpha/\Gamma(1+\alpha)\right\} \quad (18)$$

to a terminal power-law behavior [78]

$$P(k,t) \sim \left(K_\alpha k^2 t^\alpha \Gamma(1-\alpha)\right)^{-1} . \quad (19)$$

In the limit $\alpha \to 1$, it reduces to the traditional exponential $P(k,t) = \exp(-Kk^2 t)$ with a finite characteristic waiting time. Also the mean squared displacement changes from a linear to a power-law time dependence

$$\langle x^2(t)\rangle = 2K_\alpha t^\alpha/\Gamma(1+\alpha) , \quad (20)$$

with $K_\alpha = \sigma^2/\tau_0^\alpha$. This is the case of *subdiffusion*. We note that in x,t space the dynamical equation is the fractional diffusion equation [103]. In the presence of an external potential, it generalizes to the time-fractional Fokker–Planck equation [78,80,84], see also below.

(iii) Finally, take $\psi(t) = \delta(t-\tau)$ sharply peaked, but $\lambda(x)$ of Lévy stable form with index $0 < \mu < 2$. The resulting process is Markovian, but with diverging variance. It can be shown that the fractional moments scale as [82]

$$\langle |x(t)|^\delta\rangle \propto \left(K^{(\mu)} t\right)^{\delta/\mu} , \quad (21)$$

where $K^{(\mu)} = \sigma^\mu/\tau$. The upper index μ is chosen to distinguish $K^{(\mu)}$ from the subdiffusion constant K_μ. Note that the dimension of $K^{(\mu)}$ is $\mathrm{cm}^\mu/\mathrm{sec}$. From Eq. (16) one can immediately obtain the Fourier image of the associated probability density function,

$$P(k,t) = \exp\left\{-K^{(\mu)}|k|^\mu t\right\} . \quad (22)$$

From Eq. (11) this is simply a symmetric Lévy stable density with stable index μ, and this type of random walk process is most aptly coined a Lévy flight. A Lévy flight manifestly has regular exponential mode relaxation and is in fact Markovian. However, the modes in position space are no longer sharply localized like in the Gaussian or subdiffusive case. Instead, individual modes bear the hallmark of a Lévy stable density: diverging variance. We will see below how the presence of steeper-than-harmonic external potentials causes a finite variance of the Lévy flight, although a power-law form of the probability density remains.

In the remainder of this paper, we deal with the physical and mathematical properties of Lévy flights and subdiffusion. While we will be concerned mostly with the overdamped case, in the last section we will address the dynamics in velocity space in the presence of Lévy noise, in particular, the question of diverging variance of Lévy flights.

Lévy Flights

Underlying Random Walk Process

To derive the dynamic equation of a Lévy flight in the presence of an external force field $F(x) = -\mathrm{d}V(x)/\mathrm{d}x$, we pursue two different routes. One starts with a generalized version of the continuous time random walk (compare [85] for details; a slightly different derivation is presented in [10]).

To include the local asymmetry of the jump length distribution due to the force field $F(x)$ [72,85], we introduce the generalized transfer kernel $\Lambda(x,x') = \lambda(x-x')\,[A(x')\Theta(x-x') + B(x')\Theta(x'-x)]$ (and therefore $\Lambda(x,x') = \Lambda(x';x-x')$). As in standard random walk theory (compare [117]), the coefficients $A(x)$ and $B(x)$ define the local asymmetry for jumping left and right, depending on the value of $F(x)$. Here, $\Theta(x)$ is the Heaviside jump function. With the normalization $\int \Lambda(x',\Delta)\,\mathrm{d}\Delta = 1$, the fractional Fokker–Planck equation (FFPE) ensues [85]:

$$\frac{\partial}{\partial t}P(x,t) = \left(-\frac{\partial}{\partial x}\frac{F(x)}{m\eta} + K^{(\mu)}\frac{\partial^\mu}{\partial|x|^\mu}\right)P(x,t) . \quad (23)$$

Remarkably, the presence of the Lévy stable $\lambda(x)$ affects only the diffusion term, while the drift term remains unchanged and additive [36,37,85]. The fractional spatial derivative $\partial^\mu/\partial|x|^\mu$ of Riesz–Weyl type represents an integrodifferential operator defined through

$$\frac{\partial^\mu}{\partial|x|^\mu}P(x,t) = \frac{-1}{2\cos(\pi\mu/2)\Gamma(2-\mu)}\frac{\partial^2}{\partial x^2}\int_{-\infty}^{\infty}\frac{P(x',t)}{|x-x'|^{\mu-1}} , \quad (24)$$

for $1 < \mu < 2$, and a similar form for $0 < \mu < 1$ [22,93,99]. In Fourier space, for all $0 < \mu \le 2$ the simple relation

$$\mathcal{F}\left\{\frac{\partial^\mu}{\partial|x|^\mu}P(x,t)\right\} = -|k|^\mu P(k,t) \quad (25)$$

holds. In the Gaussian limit $\mu = 2$, all relations above reduce to the familiar second-order derivatives in x and thus the corresponding $P(x,t)$ is governed by the standard Fokker–Planck equation.

The FFPE (23) can also be derived from the Langevin equation [22,36,37,50]

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = -\frac{1}{m\gamma}\frac{\mathrm{d}V(x)}{\mathrm{d}x} + \xi_\mu(t)\,, \tag{26}$$

driven by white Lévy stable noise $\xi_\mu(t)$, defined through $L(\Delta t) = \int_t^{t+\Delta t} \xi_\mu\left(t'\right)\mathrm{d}t'$ being a symmetric Lévy stable density of index μ with characteristic function $p(k,\Delta t) = \exp\left(-K^{(\mu)}|k|^\mu \Delta t\right)$ for $0 < \mu \leq 2$. As with standard Langevin equations, $K^{(\mu)}$ denotes the noise strength, m is the mass of the diffusing (test) particle, and γ is the friction constant characterizing the dissipative interaction with the bath of surrounding particles.

A subtle point about the FFPE (23) is that it does not uniquely define the underlying trajectory [105]; however, starting from our definition of the process in terms of the stable jump length distribution $\lambda(x) \sim |x|^{-1-\mu}$, or its generalized pendant $\Lambda(x,x')$, the FFPE (23) truly represents a Lévy flight in the presence of the force $F(x)$. The strongly non-local character of Lévy flights poses certain difficulties when non-trivial boundary conditions are involved, as shown below.

Propagator and Symmetries

In absence of an external force, $F(x) = 0$, the exact solution of the FFPE is readily obtained as the Lévy stable density $P(k,t) = \exp\left(-K^{(\mu)}|k|^\mu t\right)$ in Fourier space. Back-transformed to position space, an analytical solution is given in terms of the Fox H-function [50,78,120]

$$P(x,t) = \frac{1}{\mu|x|}H^{1,1}_{2,2}\left[\frac{|x|}{\left(K^{(\mu)}t\right)^{1/\mu}}\,\middle|\,\begin{matrix}(1,1/\mu),(1,1/2)\\(1,1),(1,1/2)\end{matrix}\right], \tag{27}$$

from which the series expansion

$$P(x,t) = \frac{1}{\mu(K^{(\mu)}t)^{1/\mu}}\sum_{\nu=0}^{\infty}\frac{\Gamma([1+\nu]/\mu)}{\Gamma([1+\nu]/2)\Gamma(1-[1+\nu]/2)} \times \frac{(-1)^\nu}{\nu!}\left(\frac{|x|}{(K^{(\mu)}t)^{1/\mu}}\right)^\nu \tag{28}$$

derives. For $\mu = 1$, the propagator reduces to the Cauchy Lévy stable density

$$P(x,t) = \frac{1}{\pi\left(K^{(1)}t + x^2/[K^{(1)}t]\right)}\,. \tag{29}$$

We plot the time evolution of $P(x,t)$ for the Cauchy case $\mu = 1$ in Fig. 4 in comparison to the limiting Gaussian case $\mu = 2$.

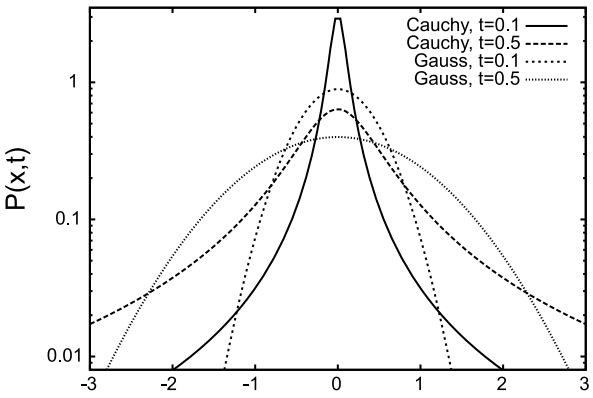

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 4

Cauchy distribution ($\mu = 1$) for two times in comparison to the Gaussian ($\mu = 2$). We chose $K^{(1)} = K = 1$. Note that the Cauchy distribution is narrower at the origin, and after crossing the Gaussian falls off in the much slower power-law fashion

Due to the point symmetry of the FFPE (23) for $F(x) = 0$, the propagator $P(x,t)$ is invariant under change of sign, and it is monomodal, that is, it has its global maximum at $x = 0$, the point where the initial distribution $P(x,0) = \delta(x)$ was launched at $t = 0$. Interestingly, the latter property is lost in the case of strongly confined Lévy flights discussed below. Due to their Markovian character, Lévy flights also possess a Galilei invariance [73,78]. Thus, under the influence of a constant force field $F(x) = F_0$, the solution of the FFPE can be expressed in terms of the force-free solution by introducing the wave variable $x - F_0 t$, to obtain

$$P_{F_0}(x,t) = P_0\left(x - \frac{F_0 t}{m\gamma}, t\right)\,. \tag{30}$$

This result follows from the FFPE (23), that in Fourier domain becomes [50]

$$\frac{\partial}{\partial t}P(k,t) = \left(-ik\frac{F_0}{m\gamma} - K^{(\mu)}|k|^\mu\right)P(k,t)\,, \tag{31}$$

with solution

$$P(k,t) = \exp\left(-\left[ik\frac{F_0}{m\gamma} + K^{(\mu)}|k|^\mu\right]t\right)\,. \tag{32}$$

By the translation theorem of the Fourier transform, Eq. (30) results. We show an example of the drift superimposed on the dispersional spreading of the propagator in Fig. 5.

Presence of External Potentials

Harmonic Potential In an harmonic potential $V(x) = \frac{1}{2}\lambda x^2$, an exact form for the characteristic function can

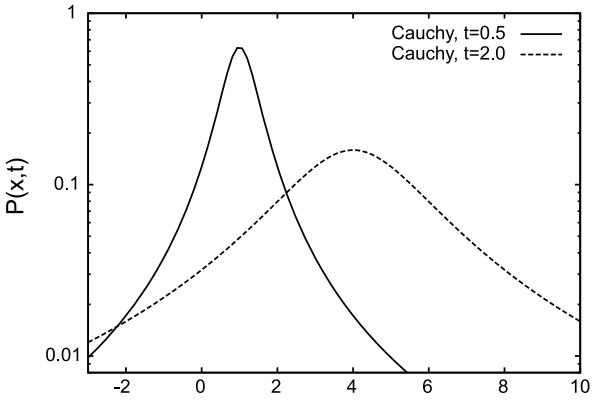

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 5
Cauchy distribution with $K^{(1)} = 1$ advected along a field $F_0/(m\gamma) = 2$, for different times

be found. Thus, from the corresponding FFPE in Fourier space,

$$\frac{\partial}{\partial t}P(k,t) = -\frac{\lambda}{m\gamma}k\frac{\partial}{\partial k}P(k,t) - K^{(\mu)}|k|^{\mu}P(k,t)\,, \quad (33)$$

by the method of characteristics one obtains

$$P(k,t) = \exp\left(-\frac{m\gamma K^{(\mu)}|k|^{\mu}}{\lambda\mu}\left[1-\mathrm{e}^{-\mu\lambda t/(m\gamma)}\right]\right) \quad (34)$$

for an initially central δ-peak, $P(x,0) = \delta(x)$ [50]. This is simply the characteristic function of a Lévy stable density with time-varying width. For short times, $1-\exp(-\mu\lambda t/[m\gamma]) \sim \mu\lambda t/[m\gamma]$ grows linearly in time, such that $P(k,t) \sim \exp\left(-K^{(\mu)}|k|^{\mu}t\right)$ as for a free Lévy flight. At long times, the stationary solution defined through

$$P_{\mathrm{st}}(k) = \exp\left(-\frac{m\gamma K^{(\mu)}|k|^{\mu}}{\lambda\mu}\right), \quad (35)$$

is reached. Notably, it has the same stable index μ as the driving Lévy noise. By separation of variables, a summation formula for $P(x,t)$ can be obtained similarly to the solution of the Ornstein–Uhlenbeck process in the presence of white Gaussian noise, but with the Hermite polynomials replaced by H-functions [50].

We note that in the Gaussian limit $\mu = 2$, the stationary solution, by necessity, must match the Boltzmann distribution corresponding to $\exp\left(-k_{\mathrm{B}}Tk^2/[2\lambda]\right)$. This requires that the Einstein–Stokes relation $K = k_{\mathrm{B}}T/(m\gamma)$ be fulfilled [113]. One might therefore speculate whether, for a system driven by external Lévy noise, a generalized Einstein–Stokes relation should hold, as was established for the subdiffusive case [78,84]. We will come back to this point below. As will be shown now, in steeper than harmonic external potentials, the stationary form of $P(x,t)$ even leaves the basin of attraction of Lévy stable densities.

Steeper than Harmonic Potentials To investigate the behavior of Lévy flights in potentials that are steeper than the harmonic case considered above, we introduce the non-linear oscillator potential

$$V(x) = \frac{a}{2}x^2 + \frac{b}{4}x^4\,, \quad (36)$$

that can be viewed as a next order approximation to a general confining, symmetric potential. It turns out that the resulting process differs from the above findings if a suitable choice of the ratio a/b is made. For simplicity, we introduce dimensionless variables through

$$x \to \frac{x}{x_0}\,; \quad t \to \frac{t}{t_0}\,; \quad a \to \frac{at_0}{m\gamma}\,, \quad (37)$$

where

$$x_0 = \left(\frac{m\gamma K^{(\mu)}}{b}\right)^{1/(2+\mu)}\,; \quad t_0 = \frac{x_0^{\mu}}{K^{(\mu)}}\,, \quad (38)$$

arriving at the FFPE

$$\frac{\partial}{\partial t}P(k,t) + |k|^{\mu}P(k,t) = \left(k\frac{\partial^3}{\partial k^3} - ak\frac{\partial}{\partial k}\right)P(k,t)\,. \quad (39)$$

Consider first the simplest case of a quartic oscillator with $a = 0$ in the presence of Cauchy noise ($\mu = 1$). In this limit, the stationary solution can be obtained exactly, yielding the expression

$$P_{\mathrm{st}}(x) = \frac{1}{\pi}\frac{1}{1-x^2+x^4} \quad (40)$$

plotted in Fig. 6. Two distinct new features are revealed in comparing the free Lévy flight with the Lévy flight in an harmonic potential: (1) Instead of the maximum at $x = 0$, one observes two maxima positioned at

$$x_m = \pm\sqrt{1/2}\,; \quad (41)$$

at $x = 0$, we find a local minimum. (2) There occurs a power-law asymptote

$$P_{\mathrm{st}}(x) \sim \frac{1}{\pi x^4} \quad (42)$$

for $x \gg 1$; consequently, this stationary solution no longer represents a Lévy stable density, and the associated mean squared displacement is finite, $\langle x^2\rangle < \infty$.

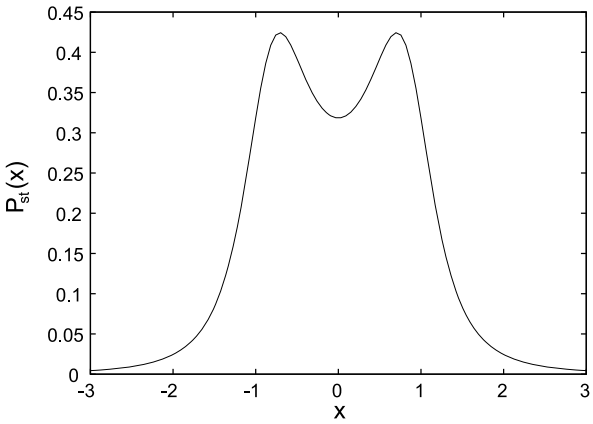

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 6
Bimodal stationary probability density $P_{st}(x)$ from Eq. (40). The maxima are at $\pm\sqrt{1/2}$

A more detailed analysis of Eq. (39) reveals [21,22], that (i) the bimodality of $P(x, t)$ occurs only if the amplitude of the harmonic term, a, is below a critical value a_c; (ii) for general μ, the asymptotic behavior is $P_{st}(x) \sim \pi^{-1} \sin(\pi\mu/2)\Gamma(\mu)|x|^{-\mu-3}$; (iii) and there exists a finite bifurcation time t_c at which the initially monomodal form of $P(x, t)$ acquires a zero curvature at $x = 0$, before settling in the terminal bimodal form.

Interestingly, in the more general power-law behavior

$$V(x) = \frac{|x|^c}{c} , \tag{43}$$

the turnover from monomodal to bimodal form of $P(x, t)$ occurs exactly when c becomes larger than 2. The harmonic potential is therefore a limiting case when the solution of the FFPE still belongs to the class of Lévy stable densities and follows the generalized central limit theorem. This is broken in a superharmonic (steeper than harmonic) potential. The corresponding bifurcation time t_c is finite for all $c > 2$ [22]. An additional effect appears when $c > 4$: there exists a transient trimodal state when the relaxing $\delta(x)$-peak overlaps with the forming humps at $x = \pm x_m$. At the same time, the variance is finite, if only $c > 4 - \mu$, following from the asymptotic stationary solution

$$P_{st}(x) \sim \frac{\sin(\pi\mu/2)\Gamma(\mu)}{\pi|x|^{\mu+c-1}} . \tag{44}$$

Details of the asymptotic behavior and the bifurcations can be found in [21,22]. From a reverse engineering point of view, Lévy flights in confining potentials are studied in [31].

First Passage and First Arrival of Lévy Flights

One might naively expect that a jump process of Lévy type, whose variance diverges (unless confined in a steep potential) may lead to ambiguities when boundary conditions are introduced, such as an absorbing boundary at finite x. Indeed, it is conceivable that for a jump process with extremely long jumps, it becomes ambiguous how to properly define the boundary condition: should the test particle be absorbed when it arrives exactly *at* the boundary, or when it crosses it *anyplace* during a non-local jump?

This question is trivial in the case of a narrow jump length distribution: all steps are small, and the particle cannot jump across a point (in the continuum limit considered herein). For such processes, one enforces a Cauchy boundary condition $P(0, t) = 0$ at the point $x = 0$ of the absorbing boundary, removing the particle once it hits the barrier after starting at x_0, where the dynamics are governed by Eq. (23) with $F(x) = 0$. Its solution can easily be obtained by standard methods, for instance, the method of images. This is completely equivalent to considering the *first arrival* to the point $x = 0$, expressed in terms of the diffusion equation with sink term:

$$\frac{\partial}{\partial t}\mathcal{P}(x, t) = K\frac{\partial^2}{\partial x^2}\mathcal{P}(x, t) - p_{fa}(t)\delta(x) , \tag{45}$$

defined such that $\mathcal{P}(0, t) = 0$. Note that the quantity $\mathcal{P}$ is no longer normalized, as probability decays to zero; for this reason, we use the notation $\mathcal{P}$. From Eq. (45) by integration we obtain the survival probability

$$\mathcal{S}(t) = \int \mathcal{P}(x, t)\mathrm{d}x \tag{46}$$

with $\mathcal{S}(0) = 1$ and $\lim_{t\to\infty} \mathcal{S}(t) = 0$. Then, the first arrival density becomes

$$p_{fa}(t) = -\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{S}(t) . \tag{47}$$

Equation (45) can be solved by standard methods (determining the homogeneous and inhomogeneous solutions). It is then possible to express $\mathcal{P}(x, t)$ in terms of the propagator $P(x, t)$, the solution of Eq. (23) with $F(x) = 0$ with the same initial condition, $P(x, 0) = \delta(x - x_0)$ and natural boundary conditions. One obtains

$$P(0, t) = \int_0^t p_{fa}(\tau)P(x_0, t - \tau)\,\mathrm{d}\tau , \tag{48}$$

such that the first arrival density corresponds to the waiting time distribution to jump from x_0 to 0 (or, vice versa, since the problem is symmetric). In Laplace space, this relation takes on the simple algebraic form $p_{fa}(u)$

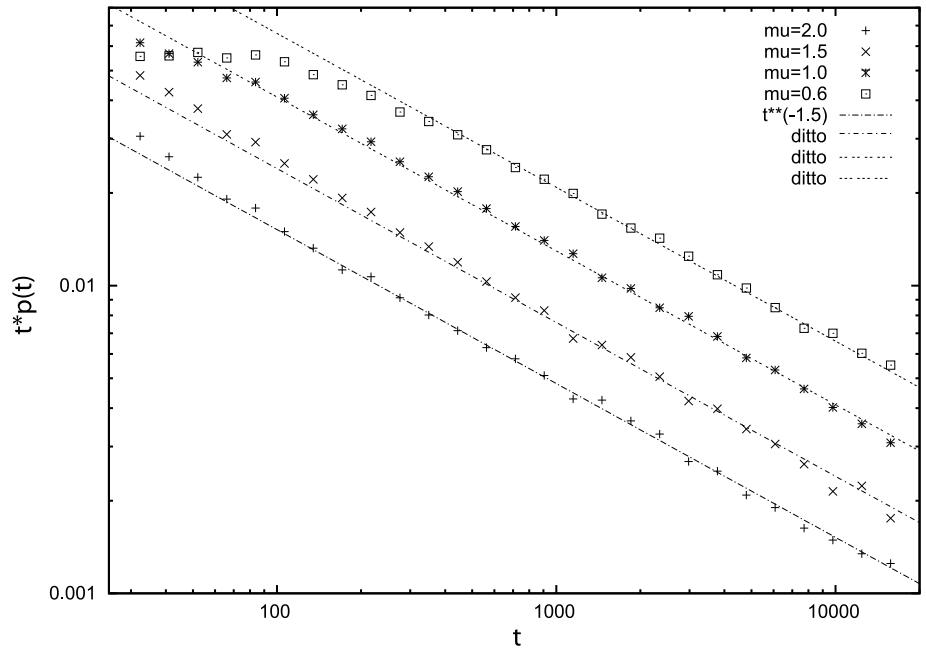

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 7
First passage density for various stable indices μ and initial position $x_0 = 10.0$ away from the absorbing boundary the universal $\sim t^{-3/2}$ scaling is distinct. Note that we plot $tp(t)$ in the y-axis

$= P(0,u)/P(x_0,u)$. Both methods, the explicit boundary value problem and the first arrival problem for Gaussian processes, produce the well-known first passage (or arrival) density of Lévy–Smirnov type (8),

$$p(t) = p_{\rm fa}(t) = \frac{x_0}{\sqrt{4\pi K t^3}} \exp\left(-\frac{x_0^2}{4Kt}\right) \sim \frac{x_0}{\sqrt{4\pi K t^3}}\,, \qquad (49)$$

with the asymptotic power-law decay $p(t) \sim t^{-3/2}$, such that no mean first passage time exists [46,97].

Long-tailed jump length distributions of Lévy stable form, however, endow the test particle with the possibility to jump across a certain point repeatedly. The first arrival necessarily becomes less efficient. Indeed, as shown in [20], the standard result (49) is generalized to

$$p_{\rm fa}(t) \sim C(\mu) \frac{x_0^{\mu-1}}{\left(K^{(\mu)}\right)^{1-1/\mu} t^{1-1/\mu}}\,, \quad \text{as } t \to \infty \qquad (50)$$

with $C(\mu) = \mu\Gamma(2-\mu)\Gamma(2-1/\mu)\sin(\pi[2-\mu]/2)\sin^2(\pi/\mu)/(\pi^2[\mu-1])$, and $1 < \mu \leq 2$ [20]. The long-time decay $\sim t^{-2+1/\mu}$ is slower than in (49).

One might naively assume that the first passage problem for Lévy flights (the particle is removed once it crosses the boundary) should be more efficient, that is, the first passage density $p(t)$ should decay more quickly, than for a narrow jump length distribution. However, as we have a symmetric jump length distribution $\lambda(x)$, the long outliers characteristic for these Lévy flights can occur both toward and away from the absorbing barrier. From this point of view it is not totally surprising to see the simulations result in Fig. 7, that clearly indicate a universal asymptotic decay $\sim t^{-3/2}$, exactly as for the Gaussian case.

In fact, for all Markovian processes with a *symmetric* jump length distribution, the Sparre Andersen theorem [35,96,108,109] proves, without knowing any details about $\lambda(x)$, the asymptotic behavior of the first passage time density universally follows $p(t) \sim t^{-3/2}$. The details of the specific form of $\lambda(x)$ enter only the prefactor and the pre-asymptotic behavior. A special case of the Sparre Andersen theorem was proved in [38] in which the particle is released at $x_0 = 0$ at time $t = 0$, and after the first jump an absorbing boundary is installed at $x = 0$. This latter case was simulated extensively in [125]. From a fractional diffusion equation point of view, it was shown in [20] that the fractional operator $\partial^\mu/\partial|x|^\mu$ needs to be modified to account for the fact that $\mathcal{P}(x,t) \equiv 0$ beyond the absorbing boundary. The fractional diffusion equation in the presence of the absorbing boundary therefore becomes [20]

$$\frac{\partial}{\partial t}\mathcal{P}(x,t) = \frac{K_\mu}{\kappa}\frac{\partial^2}{\partial x^2}\int_0^\infty \frac{\mathcal{P}(x',t)}{|x-x'|^{\mu-1}}{\rm d}x'\,, \qquad (51)$$

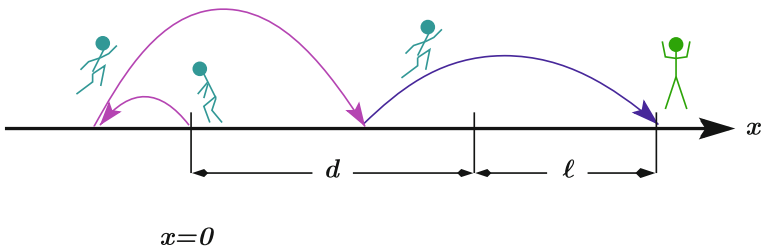

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 8
Schematic of the leapover problem: the random walker starts at $x = 0$ and after a number of jumps crosses the point $x = d$, overshooting it by a distance ℓ. For narrow jump length distributions, each jump is so small that crossing of the point d is equal to arriving at this point

where $\kappa = 2\Gamma(2-\mu)\,|\cos(\pi\mu/2)|$, such that the first term on the right hand side no longer represents a Fourier convolution. An approximate solution with Cauchy boundary condition reveals $p(u) \sim 1 - cu^{1/2}$, where c is a constant, indeed leading to the Sparre Andersen behavior $p(t) \sim t^{-3/2}$. The asymptotically exact result for the first passage time density of a Lévy flight on a semi-infinitive domain was derived in [58]:

$$p(t) \sim \frac{x_0^{\mu/2}}{\mu\sqrt{\pi K^{\mu}\Gamma(/mu/29)}} * t^{-3/2}\,.$$

This also demonstrates that the method of images no longer applies when Lévy flights are considered, for the images solution

$$\mathcal{P}_{\rm im}(x,t) = P(x-x_0,t) - P(x+x_0,t) \tag{52}$$

would be governed by the *full* fractional diffusion equation, and not Eq. (51), so the result for the first passage density, $p(t) \sim t^{-1-1/\mu}$ would decay faster than the Sparre Andersen universal behavior. A detailed discussion of the applicability of the method of images is given in terms of a subordination argument in [105]. We emphasize that this subtle failure of the method of images has been overlooked in literature previously [39,88], and care should therefore be taken when working with results based on such derivations. We also note that the method of images works in cases of subdiffusion, as the step length is narrow [75].

Leapover Properties of Lévy Flights

The statistics of first passage times is a classical concept for quantifying processes in which it is of interest when the dynamic variable crosses a certain threshold value for the first time. For processes with broad jump length distributions, another quantity is of interest, namely, the statistics of the first passage leapovers, that is, the distance ℓ by which the random walker overshoots the threshold value $x = d$ in a single jump (see Fig. 8). Surprisingly, for symmetric LFs with jump length distribution $\lambda(x) \sim |x|^{-1-\mu}$ $(0 < \mu < 2)$, the distribution of leapover lengths across x_0 is distributed like $p_l(\ell) \sim \ell^{-1-\mu/2}$, that is, it is much broader than the original jump length distribution. In contrast, for one-sided LFs jumps, the scaling of $p_l(\ell)$ bears the same index μ. Information on the leapover behavior is important to the understanding of how far search processes of animals for food or of proteins for their specific binding site along DNA overshoot their targets, or to better define stock market strategies determining when to buy or sell a certain relative to a given threshold price.

Using a general theorem for the first passage times and leapovers for homogeneous processes with independent increments, leapover properties for completely symmetric and fully asymmetric (one-sided) LFs are derived in [57]. The general case is considered in [59]. The basic results are as follows: For a completely symmetric LF with index μ, the distribution of first passage leapover lengths ℓ for a particle originally released a distance d away from the boundary reads (in scaled units)

$$p_l(\ell) = \frac{\sin(\pi\mu/2)}{\pi}\,\frac{d^{\mu/2}}{\ell^{\mu/2}(d+\ell)}\,. \tag{53}$$

The validity of this result is confirmed by extensive simulations – for more details refer to [57,58,59]. Note that p_l is normalized. In the limit $\mu \to 2$, p_l tends to zero if $\ell \neq 0$ and to infinity at $\ell = 0$, corresponding to the absence of leapovers in the Gaussian continuum limit. However, for $0 < \mu < 2$ the leapover PDF follows an asymptotic power-law with index $\mu/2$, and is thus broader than the original jump length PDF $\lambda(x)$ with index μ. This is a remarkable finding: while λ for $1 < \mu < 2$ has a finite characteristic length $\langle|x|\rangle$, this always diverges for $p_l(\ell)$ irrespective of μ.

In contrast, the result for completely asymmetric LFs has the form [31,57]

$$p_l(q) = \langle e^{-q\ell} \rangle = \frac{\sin(\pi\mu)}{\pi} \int_0^\infty d\ell e^{-q\ell} \frac{d^\mu}{\ell^\mu (d+\ell)} , \quad (54)$$

leading to the leapover PDF

$$p_l(\ell) = \frac{\sin(\pi\mu)}{\pi} \frac{d^\mu}{\ell^\mu (d+\ell)} , \quad (55)$$

which corresponds to the result obtained in [32,33] from a different method. Thus, for the one-sided LF, the scaling of the leapover is exactly the same as for the jump length distribution, namely, with exponent μ. This result compares favorably with simulations [57,59].

Kramers Problem for Lévy Flights

Many physical and chemical problems are related to the crossing of an energetic barrier driven by thermal fluctuations, including such problems as dissociation of molecules, nucleation processes, or the escape from an external, confining potential of finite height [44]. A particular example of barrier crossing in a double-well potential driven by Lévy noise was proposed for a long time series of paleoclimatic data [30]. Further cases where the crossing of a potential barrier driven by Lévy noise is of interest are in the theory of plasma devices [19], among others [80].

To investigate the detailed behavior of barrier crossing under the influence of external Lévy noise, we choose the rather generic double well shape

$$V(x) = -\frac{a}{2}x^2 + \frac{b}{4}x^4 . \quad (56)$$

Integrating the Langevin Equation (26) with white Lévy noise, we find an exponential decay of the survival density in the initial well:

$$p(t) = \frac{1}{T_c} \exp\left(-\frac{t}{T_c}\right) , \quad (57)$$

as demonstrated in Fig. 9. Lévy flight processes being Markovian, this is not surprising, since the mode relaxation is exponential [78,80]. More interesting is the question of how the mean escape time T_c behaves as a function of the characteristic noise parameters D and μ. In the regular Kramers problem with Gaussian driving noise the Arrhenius-type activation $T_c = C\exp(h/D)$ is followed, where h is the barrier height, and the prefactor C includes details of the potential. In the case of Lévy noise, a power-law form

$$T_c(\mu, D) = \frac{C(\mu)}{\left(K^{(\mu)}\right)^{\alpha(\mu)}} \quad (58)$$

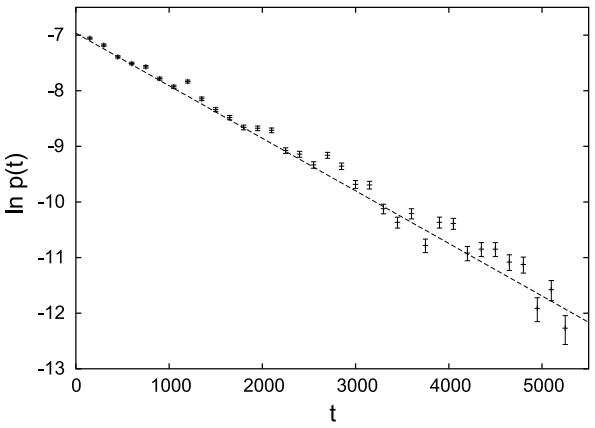

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 9
Probability density function $p(t)$ of barrier crossing times for $\mu = 1.0$ and $D = 10^{-2.5} \approx 0.00316$. The *dashed line* is a fit to Eq. (57) with mean escape time $T_c = 1057.8 \pm 17.7$

was assumed and corroborated by extensive simulations [23]. Detailed investigations [26] show that the scaling exponent $\alpha(\mu) = 1$ for all μ is strictly smaller than 2. As already proposed in [29] and derived in [48] in a somewhat different model, this means that, apart from a prefactor, the Lévy flight is insensitive to the external potential for barrier crossing, as confirmed by simulations [26]. Note that in comparison to [23], values of μ in the range $(0, 1)$ are also included. For large values of D, deviations from scaling are observed: eventually it will take only a single jump to cross the barrier when $D \to \infty$. Detailed studies show, indeed, that eventually the unit time step is reached, that is, $T_c \to 1$.

More on the "Pathology"

Despite their mathematical foundations in the generalized central limit theorem and their broad use in the sciences and beyond as description for statistical quantities, and despite the existence of systems (for instance, the diffusion on a polymer in chemical space mediated by jumps where the polymer loops back on itself [15,64,106]), the divergence of the fluctuations of Lévy processes is sometimes considered a pathology. This was already put forward by West and Seshadri [119], who pointed out that a Lévy flight in velocity space would be equivalent to a diverging kinetic energy. Here, we show that higher order dissipation effects lead to natural cutoffs in Lévy processes.

At higher velocities the friction experienced by a moving body starts to depend on the velocity itself [12]. Such non-linear friction is known from the classical Riccati equation $M dv(t)/dt = Mg - Kv(t)^2$ for the fall of

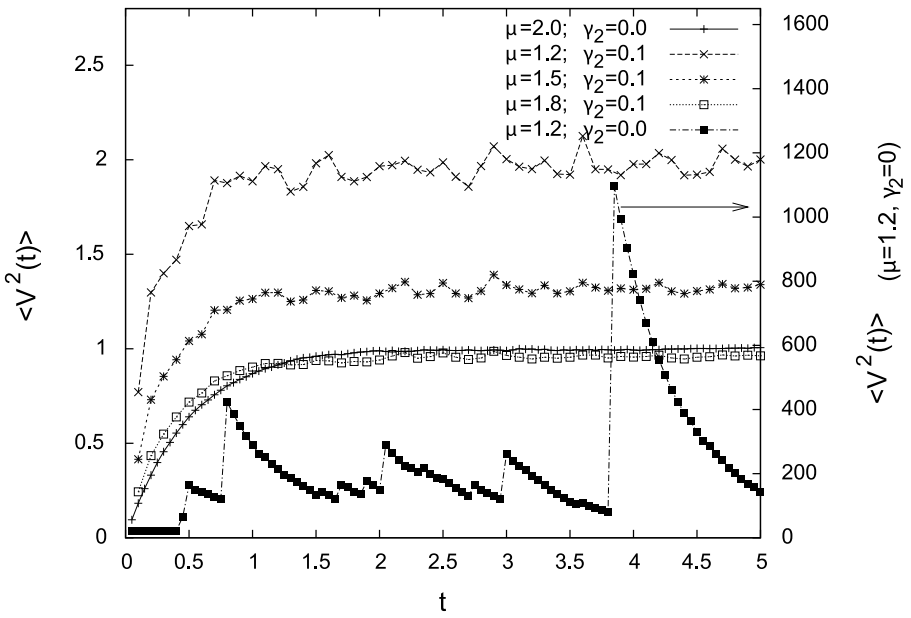

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 10
Variance $\langle V^2(t)\rangle$ as function of time t, with the quartic term set to zero, $\gamma_4 = 0$ and $\gamma_0 = 1.0$ for all cases. The variance is finite for the cases $\mu = 2.0, \gamma_2 = 0.0$; $\mu = 1.2, \mu = 1.5, \gamma_2 = 0.1$; and $\mu = 1.8, \gamma_2 = 0.1$. These correspond to the left ordinate. For the case $\mu = 1.2, \gamma_2 = 0.0$, the variance diverges, strong fluctuations are visible; note the large values of this curve corresponding to the right ordinate

a particle of mass M in a gravitational field with acceleration g [28], or autonomous oscillatory systems with a friction that is non-linear in velocity [3,12]. The occurrence of a non-constant friction coefficient $\gamma(V)$ leading to a non-linear dissipative force $-\gamma(V)V$ was highlighted in Klimontovich's theory of non-linear Brownian motion [56]. It is therefore natural that higher order, non-linear friction terms also occur in the case of Lévy processes.

We consider the velocity-dependent dissipative non-linear form (necessarily an even function) [24]

$$\gamma(V) = \gamma_0 + \gamma_2 V^2 + \ldots + \gamma_{2N} V^{2N} \quad \therefore \gamma_{2N} > 0 \quad (59)$$

for the friction coefficient of the Lévy flight in velocity space as governed by the Langevin equation

$$\mathrm{d}V(t) + \gamma(V)V(t)\mathrm{d}t = \mathrm{d}L(t) \quad (60)$$

with the constant friction $\gamma_0 = \gamma(0)$. $L(t)$ is again the μ-stable Lévy noise defined in terms of a characteristic function $p^*(\omega, t) = \mathcal{F}\{p(L,t)\} \equiv \int_{-\infty}^{\infty} p(L,t)\exp(i\omega L)\,\mathrm{d}L$ of the form $p^*(\omega, t) = \exp\left(-K^{(\mu)}|\omega|^\mu t\right)$ [63,100,112], where $K^{(\mu)}$ of dimension $\mathrm{cm}^\mu/\mathrm{sec}$ is the generalized diffusion constant. This is equivalent to the fractional Fokker–Planck equation [19,22,36,37,78,80]

$$\frac{\partial P(V,t)}{\partial t} = \frac{\partial}{\partial V}\Big(V\gamma(V)P\Big) + K^{(\mu)}\frac{\partial^\mu P}{\partial |V|^\mu}\,. \quad (61)$$

As we showed before by the example of Lévy flights in position space, the presence of the first higher order correction, $\gamma_2 V^2$ in the friction coefficient $\gamma(V)$ rectifies the Lévy motion such that the asymptotic power-law becomes steeper and the variance finite. When even higher order corrections are taken into consideration, higher order moments also become finite. Figure 10 shows an example for the second moment.

The effect on the velocity distribution of the process defined by Eqs. (60) and (61) for higher order corrections is demonstrated in Fig. 11 for the stationary limit, $P_{\mathrm{st}}(V) = \lim_{t\to\infty} P(V,t)$. While for smaller V, the character of the original Lévy stable behavior is preserved (the original power-law behavior, that is, persists to intermediately large V), for even larger V the corrections due to the dissipative non-linearity are visible in the transition(s) to steeper slope(s).

These dissipative non-linearities remove the divergence of the kinetic energy from the measurable subsystem of the random walker. In ideal mathematical language, the surrounding bath provides an infinite amount of energy

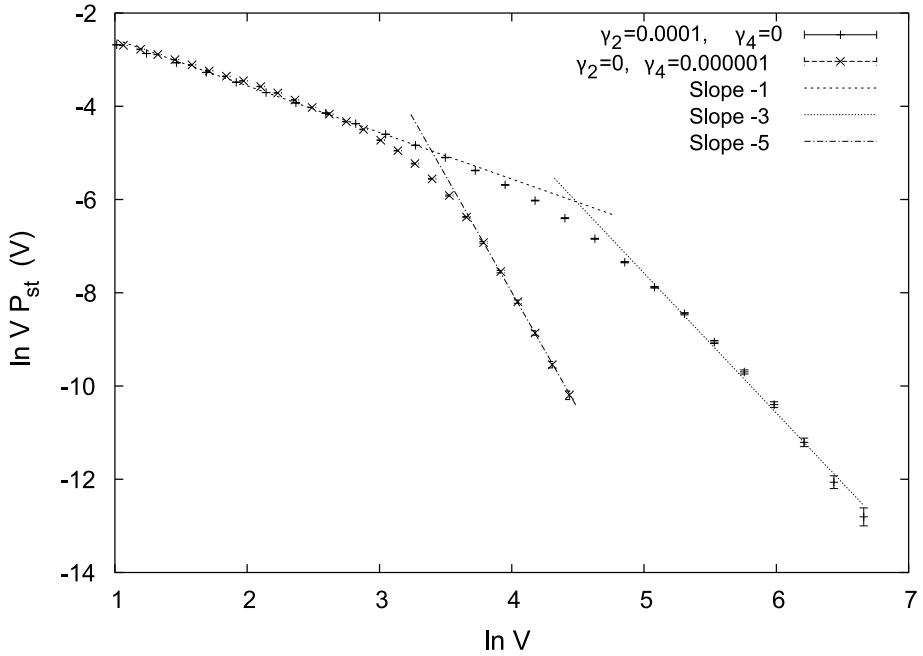

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 11
Stationary PDF $P_{st}(V)$ for $\gamma_0 = 1.0$ and (i) $\gamma_2 = 0.0001$ and $\gamma_4 = 0$; and (ii) $\gamma_2 = 0$ and $\gamma_4 = 0.000001$; with $\mu = 1.0$. The *lines* indicate the slopes − 1, − 3, and − 5

through the Lévy noise, and coupling via non-linear friction dissipates an infinite amount of energy into the bath, thereby introducing a natural cutoff in the kinetic energy distribution of the random walker subsystem. Physically, such divergencies are not expected, but correspond to the limiting procedure of large numbers in probability theory. We showed that both statements can be reconciled and that Lévy processes are, indeed, physical.

Also, Gaussian continuum diffusion exhibits non-physical features, possibly the most prominent being the infinite propagation speed inherent in the parabolic nature of the diffusion equation: even at very short times after system preparation in, say, a state $P(x, 0) = \delta(x)$, there has already arrived a finite portion of probability at large x. This problem can be corrected by changing from the diffusion to the Cattaneo (telegrapher's) equation. Still, for most purposes, the uncorrected diffusion equation is used. Or, one often uses natural boundary conditions even though the system under consideration is finite, since one might not be interested in the behavior at times when a significant portion of probability has reached the boundaries. In a similar sense, we showed that "somewhere out in the wings" Lévy flights are naturally cut off by dissipative non-linear effects. However, instead of introducing artificial cutoffs, knowing that for all purposes Lévy flights are a good quantitative description, and therefore meaningful, we use "pure" Lévy stable laws in physical models.

Bi-fractional Transport Equations

The coexistence of long-tailed forms for both jump length and waiting time PDFs was investigated within the CTRW approach in [126], discussing in detail the laminar-localized phases in chaotic dynamics. In the framework of fractional transport equations, combination of the waiting time PDF $\psi(t) \sim t^{-1-\alpha}$ $(0 < \alpha < 1)$ and jump length PDF $\lambda(x) \sim |x|^{-1-\mu}$ leads to a dynamical equation with fractional derivatives in respect to both time and space [67, 69,82,118]:

for $\alpha \leq 1$ and [82]

$$\frac{\partial}{\partial t} P(x,t) = K_\alpha^{(\mu)}\, {}_0D_t^{1-\alpha} \frac{\partial^\mu}{\partial |x|^\mu} P(x,t) \tag{62}$$

for $1 < \alpha \leq 2$.

As long as the condition $1 \leq \alpha \leq \mu \leq 2$ is met, both exponents can be chosen within the entire range [69]. For $\mu = 2$, in particular, this equation covers both sub- and superdiffusion up to ballistic motion, the latter corresponding to the wave equation $\alpha = \mu = 2$ [74]. A closed form solution of Eq. (62) can be found in terms of Fox's H-functions, [82], where some special cases permitting elementary solutions are also considered. Bi-fractional diffusion equations were also discussed in Refs. [5,7,42,47,98,111]. A bi-fractional Fokker–Planck equation with a power-law

dependence $\propto |x|^{-\theta}$ ($\theta \in \mathbb{R}$) of the diffusion coefficient was studied in [34,61].

Lévy Walks

Lévy walks correspond to the spatiotemporally coupled version of continuous time random walks. The waiting time and jump length PDFs, that is, are no longer decoupled but appear as conditional in the form $\psi(x,t) \equiv \lambda(x)p(t|x)$ (or $\psi(t)\tilde{p}(x|t)$) [54]. In particular, through the coupling $p(t|x) = \frac{1}{2}\delta(|x| - vt^{\nu})$, one introduces a generalized velocity v, which penalizes long jumps such that the overall process, the Lévy walk, attains a finite variance and a PDF with two spiky fronts successively exploring space [53,123]. Thus, Lévy walks have similar properties to generalized Cattaneo/telegraphers' equation-type models [27,73,81]. As we focus here on the properties of transport processes governed by the Langevin equation under Lévy noise, we only briefly introduce Lévy walks.

On the basis of fractional equations, formulations were obtained for the description of Lévy walks in the presence of non-trivial external force fields, with the same restriction to lower order moments in respect to a Lévy walk process [9,83]. Recently, however, a coupled fractional equation was reported [104], which describes a force-free Lévy walk exactly. Thus, it was shown that the fractional version of the material derivative $\partial/\partial t \pm \partial/\partial x$,

$$d_{\pm}^{\beta} P(x,t) \equiv {}_0D_t^{\beta} P(x \pm t, t)\,, \tag{63}$$

defined in Fourier–Laplace space through

$$\mathcal{F}\left\{\mathcal{L}\left\{d_{\pm}^{\beta} f(x,t); u\right\}; k\right\} \equiv (u \pm ik)^{\beta} f(k,u) \tag{64}$$

(where $\mathcal{F}$ acts on x and $\mathcal{L}$ on t) replaces the uncoupled fractional time operators – see also the detailed discussion of Lévy walk processes in [123]. Although one may argue for certain forms [104], there is so far no derivation for the incorporation of general external force fields in the coupled formalism. We note that a very similar fractional approach to Lévy walks was suggested in [70].

Subdiffusion and the Fractional Fokker–Planck Equation

Physical Foundation of Subdiffusion

The stochastic motion of a Brownian particle of mass m is described by the Langevin equation [18,60,113]

$$m\frac{\mathrm{d}^2x}{\mathrm{d}t^2} = -m\eta_1 v + F(x) + m\Gamma(t)\,, \quad v = \frac{\mathrm{d}x}{\mathrm{d}t}\,, \tag{65}$$

where $F(x)$ is an external force field, and η_1 is the friction coefficient. The erratic bombardment through the surrounding bath molecules is described by the fluctuating noise $\Gamma(t)$. To properly describe Brownian motion, $\Gamma(t)$ must be chosen δ-correlated (white) and Gaussian distributed. This is, the time averages of $\Gamma(t)$ are: $\overline{\Gamma(t)} = 0$, and $\overline{\Gamma(t)\Gamma(t')} = D\delta(t-t')$, where D is the noise strength. After averaging the fluctuations, the velocity moments become [18]

$$\langle \Delta v \rangle = \left(\eta v - \frac{F(x)}{m}\right)\Delta t\,,$$

$$\langle (\Delta v)^2 \rangle = \frac{2\eta k_{\mathrm{B}} T}{m}\Delta t + \mathcal{O}\left([\Delta t]^2\right)\,. \tag{66}$$

Both are proportional to Δt. For regular Brownian motion, these increments are used as expansion coefficients in the Chapman–Kolmogorov equation [18].

Subdiffusion now comes about through a so-called trapping scenario. Trapping describes the occasional immobilization of the test particle for a waiting time distributed according to the distribution $\psi(t)$. Here we assume that the particle leaves the trap with the same velocity it had prior to immobilization (this condition can be relaxed). In between trapping events, the particle is assumed to follow the regular Langevin equation such that each motion event on average lasts for a mean time τ^*. Choosing $\psi(t)$ with a finite characteristic waiting time one can show that this trapping scenario indeed preserves Brownian motion. However, once the characteristic waiting time diverges, it can be shown by application of the generalized central limit theorem that the occurrence of a large number of trapping events leads to subdiffusion and the fractional Klein–Kramers equation for the joint PDF $P(x,v,t)$ [71,76,77]. Integrating out the velocity coordinate, the fractional Fokker–Planck equation emerges for the PDF $P(x,t)$:

$$\frac{\partial P(x,t)}{\partial t} = {}_0D_t^{1-\alpha}\frac{\partial}{\partial x}\left(-\frac{F(x)}{m\eta_\alpha} + K_\alpha\frac{\partial}{\partial x}\right)P(x,t)\,. \tag{67}$$

where η_α is generalized friction coefficient of dimension $[\eta_\alpha] = \mathrm{sec}^{\alpha-2}$.

In mathematical terms the trapping scenario corresponds to a subordination principle: Trapping events cause a transformation of the internal step time of the random walk to a different observation time, corresponding to the relation

$$P(x,t) = \int_0^{\infty} \mathcal{E}_\alpha(s,t)P_1(x,s)\mathrm{d}s\,, \tag{68}$$

where $P_1(x,t)$ is the solution of the regular Brownian Fokker–Planck equation with $\alpha = 1$. Explicit forms of

the kernel $\mathcal{E}_\alpha(s,t)$ are known in terms of Fox' H-function [78]. In the Laplace domain, the kernel $\mathcal{E}_\alpha(s,t)$ has the comparatively simple form

$$\mathcal{E}_\alpha(s,u) = \frac{\eta_\alpha}{\eta_1 u^{1-\alpha}} \exp\left(-\frac{\eta_\alpha}{\eta_1} u^\alpha s\right) . \tag{69}$$

Note that this transformation guarantees the existence and positivity of the solution $P(x,t)$ of the fractional Fokker–Planck equation only if the corresponding Brownian problem possesses a well-defined solution. We note that recently an alternative approach to continuous time random walk subdiffusion in terms of δ-noise spike trains was presented [94].

Linear Response and Fluctuation-Dissipation Relation

From the fractional Fokker–Planck Equation (67) it is straightforward to prove that in the presence of a constant field F_0 the linear response relation

$$\langle x(t)\rangle_{F_0} = \frac{k_B T}{2} \langle x^2(t)\rangle_{F=0} \tag{70}$$

is fulfilled [78,84]. This is due to the fact that the waiting time distribution $\psi(t)$ is independent of force. Given the comparatively large traps required to create the long-tailed form of $\psi(t)$, this assumption appears reasonable for values of F_0 that are not too large. Experimentally, the linear response relation (70) was verified [43].

The stationary solution of the fractional Fokker–Planck Equation (67) is the standard Boltzmann–Gibbs equilibrium, $\lim_{t\to\infty} P(x,t) = \mathcal{N} \exp\{-V(x)/[k_B T]\}$, where $\mathcal{N}$ is a normalization constant [84]. This result is immediately obvious from the concept of subordination which changes the temporal spacing of events but preserves causal mode relaxation. As $F(x) = -V'(x)$, from the stationary solution of Eq. (67) by comparison with the Boltzmann–Gibbs form, one recovers the generalized Einstein–Stokes relation [84]

$$K_\alpha = \frac{k_B T}{m\eta_\alpha} , \tag{71}$$

reflecting the preservation of the fluctuation-dissipation relation for the subdiffusion process. An experimental verification of this relation was reported by [2].

Propagator

In the absence of an external force, the fractional diffusion equation is solved by Fox's H-function [78],

$$P(x,t) = \frac{1}{\sqrt{4K_\alpha t^\alpha}} H_{1,1}^{1,0}\left[\frac{|x|}{\sqrt{K_\alpha t^\alpha}} \,\middle|\, \begin{matrix}\left(1-\frac{\alpha}{2},\frac{\alpha}{2}\right)\\ (0,1)\end{matrix}\right], \tag{72}$$

an equivalent form to the original solution by [103]. Its series expansion reads [78]

$$P(x,t) = \frac{1}{\sqrt{4K_\alpha t^\alpha}} \sum_{n=0}^{\infty} \frac{(-1)^n}{n!\Gamma(1-\alpha[n+1]/2)} \left(\frac{x^2}{K_\alpha t^\alpha}\right)^{n/2} , \tag{73}$$

and asymptotically reaches the stretched Gaussian form

$$P(x,t) \sim \frac{1}{\sqrt{4\pi K_\alpha t^\alpha}} \sqrt{\frac{1}{2-\alpha}} \times \left(\frac{2}{\alpha}\right)^{(1-\alpha)/(2-\alpha)} \left(\frac{|x|}{\sqrt{K_\alpha t^\alpha}}\right)^{-(1-\alpha)/(2-\alpha)} \times \exp\left(-\frac{2-\alpha}{2}\left(\frac{\alpha}{2}\right)^{\alpha/(2-\alpha)} \left[\frac{|x|}{\sqrt{K_\alpha t^\alpha}}\right]^{1/(1-\alpha/2)}\right) , \tag{74}$$

valid for $|x| \gg \sqrt{K_\alpha t^\alpha}$. The latter corresponds to the known result from continuous time random walk theory [53,123]. The H-function simplifies if the exponent α is a rational number. For instance, for $\alpha = 1/2$, it can be rewritten in terms of the Meijer G-function

$$P(x,t) = \frac{1}{\sqrt{8\pi^3 K_{1/2} t^{1/2}}} G_{0,3}^{3,0}\left[\left(\frac{x^2}{16 K_{1/2} t^{1/2}}\right)^2 \,\middle|\, \begin{matrix} - \\ 0, \frac{1}{4}, \frac{1}{2}\end{matrix}\right], \tag{75}$$

that is found in symbolic mathematics packages such as Mathematica. Using the representation in Fig. 12, we show the subdiffusive propagator with its pronounced cusps at the site of the initial condition, in comparison to the smooth Gaussian propagator of Brownian motion.

Boundary Value and First Passage Time Problems

In the presence of reflecting, absorbing, or mixed, boundary conditions, the narrow jump length distribution of subdiffusive processes makes it possible to use analogous techniques to calculate the propagator as known from normal diffusion. Thus, separation into eigenmodes or the method of images can be applied, the only difference entering through time-dependence of the modes. A very convenient way is to employ the Brownian result for a given geometry, subordinating that result to obtain the behavior for a subdiffusing particle, according to Eq. (68).

In the presence of an absorbing boundary an important quantity characterizing the dynamics of the system is the survival probability

$$\mathcal{S}_\alpha(t) = \int_{\mathbb{D}} \mathcal{P}(x,t)\mathrm{d}x , \tag{76}$$

where $\mathbb{D}$ denotes the interval over which $\mathcal{P}(x,t)$ is defined. The initial value of the survival probability is $\mathcal{S}_\alpha(0) = 1$,

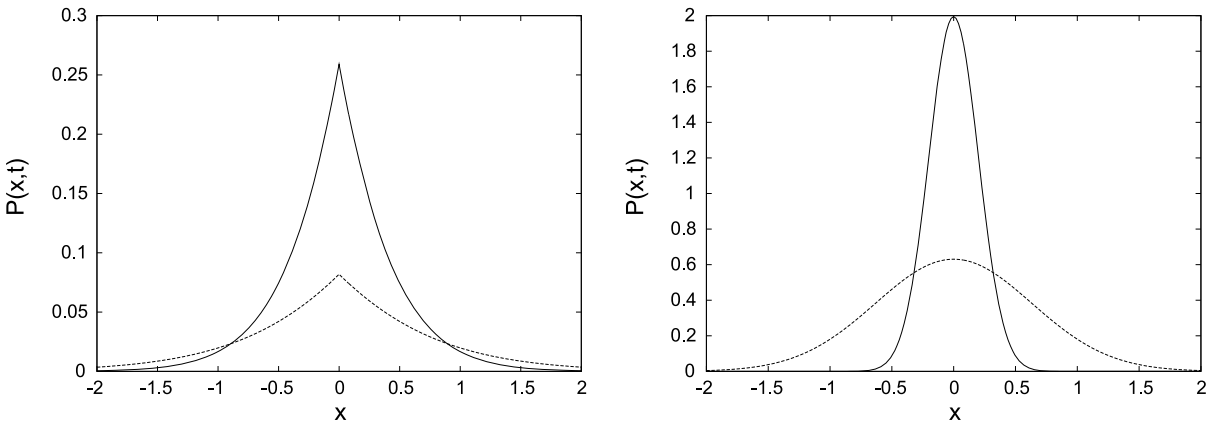

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 12
Propagator $P(x, t)$ of the fractional diffusion equation with $\alpha = 1/2$ (*left*) and the normal diffusion equation (*right*) for consecutive times

and it decays to zero for long times. From $\mathcal{S}_\alpha(t)$, we can define the first passage time density

$$p_\alpha(t) = -\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{S}_\alpha(t)\,. \tag{77}$$

The mean first passage time is

$$\mathcal{T} = \int_0^\infty p_\alpha(t) t \mathrm{d}t\,. \tag{78}$$

Note that from Eqs. (77) and (68), we obtain the relation

$$p_\alpha(u) = p_1\left(u^\alpha\right) \tag{79}$$

between the subdiffusive and Brownian results in Laplace space. Thus, these are connected by a relation similar to the subordination (68), with kernel $\exp\left(-su^\alpha\right)$ instead of $\mathcal{E}_\alpha(s, u)$ [79,80].

For the three most prominent cases of first passage time problems, we obtain the following subdiffusive generalizations:

(i) For subdiffusion in the semi-infinite domain with an absorbing wall at the origin and initial condition $P(x, 0) = \delta(x - x_0)$, it was found that [75]

$$p(t) \sim \frac{x_0}{|\Gamma(-\alpha/2)|K_\alpha^{1/2}} t^{-1-\alpha/2}\,, \tag{80}$$

that is, the decay becomes a flatter power-law than in the Markovian case where $p(t) \sim t^{-3/2}$.

(ii) Subdiffusion in the semi-infinite domain in the presence of an external bias towards the absorbing boundary falls off faster, but still in power-law manner [6,79,102]:

$$p(t) \sim t^{-1-\alpha}\,. \tag{81}$$

In strong contrast to the biased Brownian case, we now end up with a process whose characteristic time scale diverges. This is exactly the mirror of the multiple trapping model, that is, the classical motion events become repeatedly interrupted such that immobilization time dominates the process. In contrast, for $\alpha = 1$ the result

$$p_1(t) = \frac{x_0}{\sqrt{4\pi K_\alpha t^3}} \exp\left(-\frac{(x_0 - Vt)^2}{4Kt}\right) \tag{82}$$

is valid, producing the classical form $\mathcal{T} = x_0/V$ for the mean first passage time.

(iii) Subdiffusion in a finite box [75]:

$$p(t) \sim t^{-1-\alpha}\,, \tag{83}$$

that is, this process leads to the same scaling behavior for longer times as found for the biased semi-infinite case (ii).

The latter two results should be compared to the classical Scher–Montroll finding for the first passage time density of biased motion in a finite system of size L with absorbing boundary conditions. In that case, the first passage time density exhibits two power-laws

$$p(t) \sim \begin{cases} t^{\alpha-1}\,, & t < \tau \\ t^{-1-\alpha}\,, & t > \tau \end{cases} \tag{84}$$

the sum of whose exponents equals -2 [92,101,102]. Here, τ is a system size dependent time scale [92].

Fractional Ornstein–Uhlenbeck Process

The Ornstein–Uhlenbeck process corresponds to the motion in an harmonic potential $V(x) = \frac{1}{2}m\omega^2 x^2$ giving rise

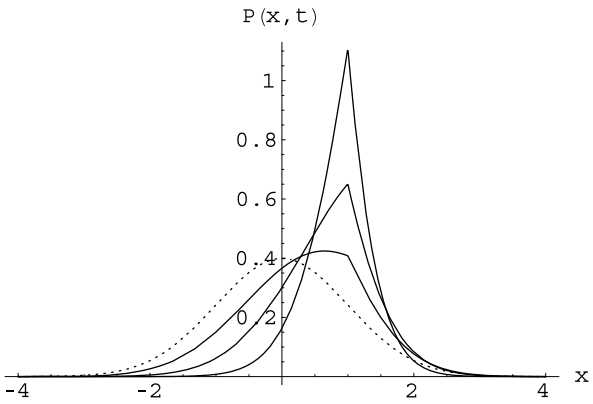

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 13
Time evolution of the PDF of the fractional Ornstein–Uhlenbeck process ($\alpha = 1/2$). The initial condition was chosen as $\delta(x-1)$. Note the strongly persistent cusp at the location of the initial peak. Dimensionless times: 0.02, 0.2, 2. The *dashed line* corresponds to the Boltzmann equilibrium

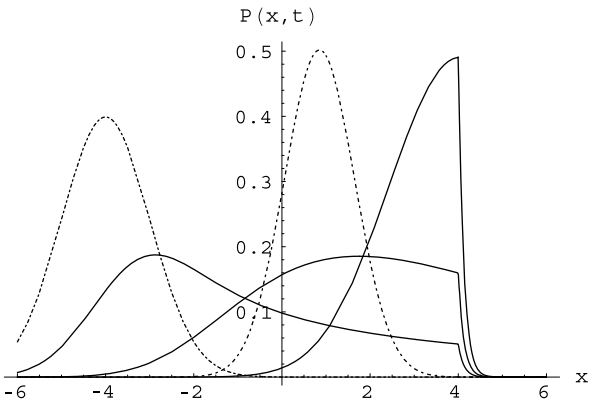

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 14
Time evolution of the PDF of the fractional Ornstein–Uhlenbeck process with superposed constant force of dimensionless strength $V = -4$ ($\alpha = 1/2$). The initial condition was chosen as $\delta(x-4)$. Dimensionless times: 0.02, 0.2, 2. The *dashed line* corresponds to the Brownian solution at times 0.5 and 50 (in essence, the stationary state). Again, note the cusps due to the initial condition, causing a strongly asymmetric shape of the PDF in contrast to the Gaussian nature of the Brownian counterpart

to the restoring force field $F(x) = -m\omega^2 x$, that is, to the dynamical equation

$$\frac{\partial}{\partial t}P(x,t) = {}_0D_t^{1-\alpha}\left(\frac{\partial}{\partial x}\frac{\omega^2 x}{\eta_\alpha} + K_\alpha\frac{\partial^2}{\partial x^2}\right)P(x,t)\,. \quad (85)$$

From separation of variables and the definition of the Hermite polynomials [1], one finds the series solution for the fractional Fokker–Planck equation with the Ornstein–Uhlenbeck potential [78,84],

$$P(x,t) = \sqrt{\frac{m\omega^2}{2\pi k_B T}}\sum_{n=0}^{\infty}\frac{1}{2^n n!}E_\alpha\left(-\frac{n\omega^2 t^\alpha}{\eta_\alpha}\right) H_n\left(\frac{\sqrt{m}\omega x_0}{\sqrt{2k_B T}}\right)\times H_n\left(\frac{\sqrt{m}\omega x}{\sqrt{2k_B T}}\right)\exp\left(-\frac{m\omega^2 x^2}{2k_B T}\right) \quad (86)$$

plotted in Fig. 13. Individual spatial eigenmodes follow the ordinary Hermite polynomials of increasing order, while their temporal relaxation is of Mittag–Leffler form, with decreasing internal time scale $(\eta_\alpha/[n\omega^2])^{1/\alpha}$. Numerically, the solution (86) is somewhat cumbersome to treat. In order to plot the PDF $P(x,t)$ in Fig. 13, it is preferable to use the closed form solution (we use dimensionless variables)

$$P(x,t) = \frac{1}{\sqrt{2\pi\left(1-\mathrm{e}^{-2t}\right)}}\exp\left(-\frac{\left(x-x_0\mathrm{e}^{-t}\right)^2}{2\left(1-\mathrm{e}^{-2t}\right)}\right) \quad (87)$$

of the Brownian case, and the transformation (68) to construct the fractional analogue.

Figure 13 shows the distinct cusps at the position of the initial condition at $x_0 = 1$. The relaxation to the final Gaussian Gibbs–Boltzmann PDF can be seen from the sequence of three consecutive times. Only at stationarity does the cusp give way to the smooth Gaussian shape of the equilibrium PDF. By adding an additional linear drift V to the harmonic restoring force, the drift term in the FFPE (67) changes to $-\partial(x-V)P(x,t)/\partial x$, and the exponential in expression (87) takes on the form $\exp(-[x-V-(x_0-V)\mathrm{e}^{-t}]/[2(1-\mathrm{e}^{-2t})])$. As displayed in Fig. 14, the strong persistence of the initial condition causes a highly asymmetric shape in the PDF, whereas the Brownian solution shown in dashed lines retains its symmetric Gaussian profile.

Let us finally address the moments of the fractional Ornstein–Uhlenbeck process, Eq. (86). These can be readily obtained either from the Brownian result with the integral transformation (68), or from integration $\int \mathrm{d}x x^n\cdot$ of the FFPE (67). For the first and second moments one obtains:

$$\langle x(t)\rangle = x_0 E_\alpha\left(-\frac{\omega^2 t^\alpha}{\eta_\alpha}\right) \quad (88)$$

and

$$\langle x(t)^2\rangle = x_{\mathrm{th}}^2 + \left(x_0^2 - x_{\mathrm{th}}^2\right)E_\alpha\left(-\frac{2\omega^2 t^\alpha}{\eta_\alpha}\right), \quad (89)$$

respectively. The first moment starts off at the initial position, x_0, and then falls off in a Mittag–Leffler pattern,

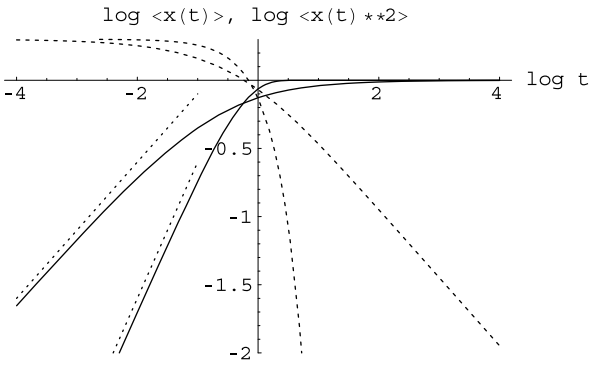

Levy Statistics and Anomalous Transport: Levy Flights and Subdiffusion, Figure 15

First ($x_0 = 2$, *dashed line*) and second ($x_0 = 0$, *full line*) moment of the fractional Ornstein–Uhlenbeck process ($\alpha = 1/2$), in comparison to the Brownian case. $\log_{10}$—$\log_{10}$ scale. The *dotted straight lines* show the initial (sub)diffusive behavior with slopes 1/2 and 1, in the special case $x_0 = 0$ chosen for the second moment

reaching the terminal inverse power-law $\sim t^{-\alpha}$. The second moment turns from the initial value x_0^2 to the thermal value $x_{\rm th}^2 = k_{\rm B}T/(m\omega^2)$. In the special case $x_0 = 0$, the second moment measures initial force-free diffusion due to the initial exploration of the flat apex of the potential. We graph the two moments in Fig. 15 in comparison to their Brownian counterparts.

Future Directions

Anomalous diffusion is becoming widely recognized in a variety of fields. Apart from the anomalous spreading of tracers and the consequences for the propagator, additional questions such as ageing and weak ergodicity breaking become important for the understanding of experiments and their modeling on complex systems. At the same time, the theory of anomalous processes is expanding. For instance, regarding questions on the weak ergodicity breaking, the calculation of multipoint moments or the correct introduction of cutoffs are currently being worked on, to complete the world *Beyond Brownian Motion* [55].

Acknowledgments

RM acknowledges partial financial support through the Natural Sciences and Engineering Research Council (NSERC) of Canada and the Canada Research Chairs program of the Government of Canada. AVC acknowledges partial support from the Deutsche Forschungsgemeinschaft (DFG).

Bibliography

1. Abramowitz M, Stegun I (1972) Handbook of Mathematical Functions. Dover, New York
2. Amblard F, Maggs AC, Yurke B, Pargellis AN, Leibler S (1996) Subdiffusion and anomalous local viscoelasticity in actin networks. Physical Review Letters 77:4470–4473. Erratum: Physical Review Letters 81:1136–1136. Reply: Physical Review Letters 81:1135–1135
3. Andronow AA, Chaikin CE, Lefschetz S (1949) Theory of oscillations. University Press, Princeton
4. Atkinson RPD, Rhodes CJ, Macdonald DW, Anderson RM (2002) Scale-free dynamics in the movement patterns of jackals. OIKOS 98:134–140
5. Baeumer B, Benson DA, Meerschaert MM (2005) Advection and dispersion in time and space. Phys A-Stat Mech Appl 350:245–262
6. Barkai E (2001) Fractional Fokker–Planck equation, solution, and application. Phys Rev E 63:046118
7. Barkai E (2002) CTRW pathways to the fractional diffusion equation. Chem Phys 284:13–27
8. Barkai E (2006) Aging in subdiffusion generated by a deterministic dynamical system. Phys Rev Lett 96:104101
9. Barkai E, Silbey RJ (2000) Fractional Kramers equation. J Phys Chem B 104:3866–3874
10. Barkai E, Metzler R, Klafter J (2000) From continuous time random walks to the fractional Fokker–Planck equation. Phys Rev E 61:132–138
11. Bel G, Barkai E (2006) Weak ergodicity breaking in the continuous-time random walk. Phys Rev Lett 94:240602
12. Bogoliubov NN, Mitropolsky YA (1961) Asymptotic methods in the theory of non-linear oscillations. Hindustan Publ. Corp., Delhi (distr. by Gordon & Breach, New York)
13. Bouchaud JP (1992) Weak ergodicity breaking and aging in disordered-systems. J de Phys 2:1705–1713
14. Bouchaud JP, Georges A (1990) Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. Phys Rep 195:127–293
15. Brockmann D, Geisel T (2003) Particle dispersion on rapidly folding random heteropolymers. Phys Rev Lett 91:048303
16. Brown R (1828) A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. Philos Mag 4:161–173
17. Carus TL (1975) De rerum natura (50 BCE): On the nature of things. Harvard University Press, Cambridge
18. Chandrasekhar S (1943) Stochastic problems in physics and astronomy. Revs Mod Phys 15:1–89
19. Chechkin AV, Gonchar VY, Szydlowski M (2002) Fractional kinetics for relaxation and superdiffusion in a magnetic field. Phys of Plasmas 9:78–88
20. Chechkin AV, Metzler R, Gonchar VY, Klafter J, Tanatarov LV (2003) J Phys A-Math Gen 36:L537–L544
21. Chechkin AV, Klafter J, Gonchar VY, Metzler R, Tanatarov LV (2003) Phys Rev E 67:010102(R)
22. Chechkin AV, Gonchar VY, Klafter J, Metzler R, Tanatarov LV (2004) J Stat Phys 115:1505–1535
23. Chechkin AV, Gonchar VY, Klafter J, Metzler R (2005) Europhys Lett 72:348–354

24. Chechkin AV, Gonchar VY, Klafter J, Metzler R (2005) Phys Rev E 72:010101
25. Chechkin AV, Gonchar VY, Klafter J, Metzler R (2006) Fundamentals of Lévy flight processes. Adv Chem Phys 133B:439–496
26. Chechkin AV, Sliusarenko AYU, Metzler R, Klafter J (2007) Barrier crossing driven by Lévy noise: Universality and the role of noise intensity. Phys Rev E 75:041101
27. Compte A, Metzler R (1997) The generalized Cattaneo equation for the description of anomalous transport processes. J Phys A-Math Gen 30:7277–7289
28. Davis H (1962) Introduction to nonlinear differential and integral equations. Dover Publications, New York
29. Ditlevsen PD (1999) Anomalous jumping in a double-well potential. Phys Rev E 60:172–179
30. Ditlevsen PD (1999) Observation of alpha-stable noise induced millennial climate changes from an ice-core record. Geophys Res Lett 26:1441–1444
31. Eliazar I, Klafter J (2003) Lévy-driven Langevin systems: Targeted stochasticity. J Stat Phys 111:739–768
32. Eliazar I, Klafter J (2004) On the first passage of one-sided Levy motions. Phys A-Stat Mech Appl 336:219–244
33. Eliazar I, Koren T, Klafter J (2007) Searching circular DNA strands. J Phys-Condens Matter 19:065140
34. Fa KS, Lenzi EK (2003) Power law diffusion coefficient and anomalous diffusion: Analysis of solutions and first passage time. Phys Rev E 67:061105
35. Feller W (1971) An Introduction to Probability Theory and its Appl, vol 2. Wiley, New York
36. Fogedby HC (1994) Lévy flights in random-environments. Phys Rev Lett 73:2517–2520
37. Fogedby HC (1998) Lévy flights in quenched random force fields. Phys Rev 58:1690–1712
38. Frisch U, Frisch H (1995) Universality of escape from a half space for symmetrical random walks. In: Shlesinger MF, Zaslavsky GM, Frisch U (eds) Lévy flights and related topics in physics. Lecture Notes in Physics, vol 450:262–268. Springer, Berlin
39. Gitterman M (2000) Mean first passage time for anomalous diffusion. Phys Rev E 62:6065–6070
40. Gnedenko BV, Kolmogorov AN (1954) Limit Distributions for Sums of Random Variables. Addison-Wesley, Reading
41. Golding I, Cox EC (2006) Physical nature of bacterial cytoplasm. Phys Rev Lett 96:098102
42. Gorenflo R, Mainardi F, Moretti D, Pagnini G, Paradisi P (2002) Fractional diffusion: probability distributions and random walk models. Phys A-Stat Mech Appl 305:106–112
43. Gu Q, Schiff EA, Grebner S, Wang F, Schwarz R (1996) Non-Gaussian transport measurements and the Einstein relation in amorphous silicon. Phys Rev Lett 76:3196–3199
44. Hänggi P, Talkner P, Borkovec M (1990) Reaction-rate theory – 50 years after Kramers. Revs Mod Phys 62:251–341
45. He Y, Burov S, Metzler R, Barkai E (2008) Random Time-Scale Invariant Diffusion and Transport Coefficients. Phys Rev Lett 101:058101
46. Hughes BD (1995) Random Walks and Random Environments, vol 1. Oxford University Press, Oxford
47. Hughes BD (2002) Anomalous diffusion, stable processes, and generalized functions. Phys Rev E 65:035105
48. Imkeller P, Pavlyukevich I (2006) Lévy flights: transitions and meta-stability. J Phys A-Math Gen 39: L237–L246
49. Ingenhousz J (1785) Nouvelles expériences et observations sur divers objets de physique. T. Barrois le jeune, Paris
50. Jespersen S, Metzler R, Fogedby HC (1999) Lévy flights in external force fields: Langevin and fractional Fokker–Planck equations and their solutions. Phys Rev E 59:2736–2745
51. Kappler E (1931) Versuche zur Messung der Avogadro–Loschmidtschen Zahl aus der Brownschen Bewegung einer Drehwaage. Ann Phys 11:233–256
52. Katori H, Schlipf S, Walther H (1997) Anomalous dynamics of a single ion in an optical lattice. Phys Rev Lett 79:2221–2224
53. Klafter J, Zumofen G (1994) Probability distributions for continuous-time random-walks with long tails. J Phys Chem 98:7366–7370
54. Klafter J, Blumen A, Shlesinger MF (1987) Stochastic pathway to anomalous diffusion. Phys Rev A 35:3081–3085
55. Klafter J, Shlesinger MF, Zumofen G (1996) Beyond Brownian motion. Phys Today 49:33–39
56. Klimontovich YUL (1992) Turbulent motion and the structure of chaos: a new approach to the statistical theory of open systems. Kluwer, Dordrecht
57. Koren T, Lomholt MA, Chechkin AV, Klafter J, Metzler R (2007) Leapover lengths and first passage time statistics for Lévy flights. Phys Rev Lett 99:160602
58. Koren T, Chechkin AV, Klafter J (2007) On the first passage time and leapover properties of Lévy motion. Phys A 379: 10–22
59. Koren T, Lomholt MA, Chechkin AV, Klafter J, Metzler R Unpublished
60. Langevin P (1908) The theory of Brownian movement. C R Hebd Seances Acad Sci 146:530–533
61. Lenzi EK, Mendes RS, Fa KS, Malacarne LC, da Silva LR (2003) Anomalous diffusion: Fractional Fokker–Planck equation and its solutions. J Math Phys 44:2179–2185
62. Levandowsky M, White BS, Schuster FL (1997) Random movements of soil amebas. Acta Protozool 36:237–248
63. Lévy P (1954) Théorie de l'addition des variables aléatoires. Gauthier-Villars, Paris
64. Lomholt MA, Ambjornsson T, Metzler R (2005) Optimal target search on a fast-folding polymer chain with volume exchange. Phys Rev Lett 95:260603
65. Lomholt MA, Zaid IM, Metzler R (2007) Subdiffusion and weak ergodicity breaking in the presence of a reactive boundary. Phys Rev Lett 98:200603
66. Lubelski A, Sokolov IM, Klafter J (2008) Non ergodicity Mimics Inhomoneneity in Single Particle Tracking. Phys Rev Lett 100:250602
67. Luchko Y, Gorenflo R (1998) Scale-invariant solutions of a partial differential equation of fractional order. Fractional Calc Appl Anal 1:63–78
68. Lutz E (2004) Power-law tail distributions and nonergodicity. Phys Rev Lett 93:190602
69. Mainardi F, Luchko Y, Pagnini G (2001) The fundamental solution of the space-time fractional diffusion equation. Fractional Calc Appl Anal 4:153–192
70. Meerschaert MM, Benson DA, Scheffler HP, Becker-Kern P (2002) Governing equations and solutions of anomalous random walk limits. Phys Rev E 66:060102
71. Metzler R (2000) Generalized Chapman–Kolmogorov equation: A unifying approach to the description of anomalous transport in external fields. Phys Rev E 62:6233–6245

72. Metzler R (2001) Non-homogeneous random walks, generalized master equations, fractional Fokker–Planck equations, and the generalized Kramers–Moyal expansion. Eur Phys J B 19:249–258
73. Metzler R, Compte A (2000) Generalized diffusion-advection schemes and dispersive sedimentation: A fractional approach. J Phys Chem B 104:3858–3865
74. Metzler R, Klafter J (2000) Accelerating Brownian motion: A fractional dynamics approach to fast diffusion. Europhys Lett 51:492–498
75. Metzler R, Klafter J (2000) Boundary value problems for fractional diffusion equations. Phys A-Stat Mech Appl 278:107–125
76. Metzler R, Klafter J (2000) From a generalized Chapman–Kolmogorov equation to the fractional Klein–Kramers equation. J Phys Chem B 104:3851–3857
77. Metzler R, Klafter J (2000) Subdiffusive transport close to thermal equilibrium: From the Langevin equation to fractional diffusion. Phys Rev E 61:6308–6311
78. Metzler R, Klafter J (2000) The random walk's guide to anomalous diffusion: a fractional dynamics approach. Phys Rep 339:1–77
79. Metzler R, Klafter J (2003) When translocation dynamics becomes anomalous. Biophys J 85:2776–2779
80. Metzler R, Klafter J (2004) The restaurant at the end of the random walk: recent developments in the description of anomalous transport by fractional dynamics. J Phys A-Math Gen 37: R161–R208
81. Metzler R, Nonnenmacher TF (1998) Fractional diffusion, waiting-time distributions, and Cattaneo-type equations. Phys Rev E 57:6409–6414
82. Metzler R, Nonnenmacher TF (2002) Space- and time-fractional diffusion and wave equations, fractional Fokker–Planck equations, and physical motivation. Chem Phys 284: 67–90
83. Metzler R, Sokolov IM (2002) Superdiffusive Klein–Kramers equation: Normal and anomalous time evolution and Levy walk moments. Europhys Lett 58:482–488
84. Metzler R, Barkai E, Klafter J (1999) Anomalous diffusion and relaxation close to thermal equilibrium: A fractional Fokker–Planck equation approach. Phys Rev Lett 82:3563–3567
85. Metzler R, Barkai E, Klafter J (1999) Deriving fractional Fokker–Planck equations from a generalized master equation. Europhys Lett 46:431–436
86. Monthus C, Bouchaud JP (1996) Models of traps and glass phenomenology. J Phys A-Math Gen 29:3847–3869
87. Montroll EW, Weiss GH (1965) Random Walks on Lattices. II. J Math Phys 6:167–181
88. Montroll EW, West BJ (1979) In: Montroll EW, Lebowitz JL (eds) Fluctuation phenomena. North-Holland, Amsterdam, pp 61–176
89. Oldham KB, Spanier J (1974) The fractional calculus. Academic Press, New York
90. Perrin J (1908) Molecular agitation and the Brownian movement. C R Hebd Seances Acad Sci 146:967–970
91. Perrin J (1909) Brownian motion and molecular reality. Ann Chem Phys 18:5–114
92. Pfister G, Scher H (1978) Dispersive (non-Gaussian) transient transport in disordered solids. Adv in Phys 27:747–798
93. Podlubny I (1998) Fractional differential equations. Academic Press, San Diego
94. Priyatinska A, Saichev AI, Woyczynski WA (2005) Models of anomalous diffusion: the subdiffusive case. Phys A 349:375–420
95. Ramos-Fernandez G, Mateos JL, Miramontes O, Cocho G, Larralde H, Ayala-Orozco B (2004) Lévy walk patterns in the foraging movements of spider monkeys (Ateles geoffroyi). Behav Ecol Sociobiol 55:223–230
96. Redner S (2001) A guide to first-passage processes. Cambridge University Press, Cambridge
97. Risken H (1989) The Fokker-Planck equation. Springer, Berlin
98. Saichev AI, Zaslavsky GM (1997) Fractional kinetic equations: solutions and applications. Chaos 7:753–764
99. Samko SG, Kilbas AA, Marichev OI (1993) Fractional Integrals and Derivatives, Theory and Applications. Gordon and Breach, New York
100. Samorodnitsky G, Taqqu MS (1994) Stable non-gaussian random processes: Stochastic models with infinite variance. Chapman and Hall, New York
101. Scher H, Montroll EW (1975) Anomalous transit-time dispersion in amorphous solids. Phys Rev B 12:2455–2477
102. Scher H, Margolin G, Metzler R, Klafter J, Berkowitz B (2002) The dynamical foundation of fractal stream chemistry: the origin of extremely long retention times. Geophys Res Lett 29:1061
103. Schneider WR, Wyss W (1989) Fractional diffusion and wave equations. J Math Phys 30:134–144
104. Sokolov IM, Metzler R (2003) Towards deterministic equations for Levy walks: The fractional material derivative. Phys Rev E 67:010101
105. Sokolov IM, Metzler R (2004) Non-uniqueness of the first passage time density of Levy random processes. J Phys A-Math Gen 37: L609–L615
106. Sokolov IM, Mai J, Blumen A (1997) Paradoxical diffusion in chemical space for nearest-neighbor walks over polymer chains. Phys Rev Lett 79:857–860
107. Sokolov IM, Klafter J, Blumen A (2002) Fractional kinetics. Phys Today 55:48–54
108. Sparre Andersen E (1953) On the fluctuations of sums of random variables. Math Scand 1:263–285
109. Sparre Andersen E (1954) On the fluctuations of sums of random variables II. Math Scand 2:194–223
110. Tolić-Norrelytke I, Munteanu EL, Thon G, Berg-Sørensen K (2004) Anamalous Diffusion in Living Yeast Cells. Phys Rev Lett 93:078102
111. Uchaikin VV (2002) Multidimensional symmetric anomalous diffusion. Chem Phys 284:507–520
112. Uchaikin VV, Zolotarev VM (1999) Chance and stability. Stable distributions and their applications. VSP, Utrecht
113. van Kampen NG (1981) Stochastic Processes in Physics and Chemistry. North-Holland, Amsterdam
114. Visser AW, Thygesen UH (2003) Random motility of plankton: diffusive and aggregative contributions. J Plankton Res 25:1157–1168
115. Viswanathan GM, Afanasyev V, Buldyrev SV, Murphy EJ, Prince PA, Stanley HE (1996) Lévy flight search patterns of wandering albatrosses. Nature 381:413–415
116. Viswanathan GM, Buldyrev SV, Havlin S, da Luz MGE, Raposo EP, Stanley HE (1999) Optimizing the success of random searches. Nat 401:911–914
117. Weiss GH (1994) Aspects and Applications of the Random Walk. North-Holland, Amsterdam

118. West BJ, Nonnenmacher TF (2001) An ant in a gurge. Phys Lett A 278:255–259
119. West BJ, Seshadri V (1982) Linear systems with Lévy fluctuations. Phys A 113:203–216
120. West BJ, Grigolini P, Metzler R, Nonnenmacher TF (1997) Fractional diffusion and Levy stable processes. Phys Rev E 55:99–106
121. Wong IY, Gardel ML, Reichman DR, Weeks ER, Valentine MT, Bausch AR, Weitz DA (2004) Anomalous diffusion probes microstructure dynamics of entangled F-actin networks. Phys Rev Lett 92:178101
122. Zaslavsky GM (2002) Chaos, fractional kinetics and anomalous transport. Phys Rep 371:461–580
123. Zumofen G, Klafter J (1993) Scale-invariant motion in intermittent chaotic systems. Phys Rev E 47:851–863
124. Zumofen G, Klafter J (1994) Spectral random walk of a single molecule. Chem Phys Lett 219:303–309
125. Zumofen G, Klafter J (1995) Absorbing boundary in one-dimensional anomalous transport. Phys Rev E 51:2805–2814
126. Zumofen G, Klafter J (1995) Laminar-localized-phase coexistence in dynamical systems. Phys Rev E 51:1818–1821

Linear and Non-linear Fokker–Planck Equations

Till D. Frank
University of Connecticut, Storrs, USA

Article Outline

Glossary

Linear Here, linear with respect to a probability density.

Nonlinear Here, nonlinear with respect to a probability density.

Markov process Process for which it is sufficient to have information about the presence in order to make best predictions about the future. Additional information about the past will not improve the predictions.

Martingale process $\hat{Z}$ Process for which the future mean value $\langle Z(t+\Delta t)\rangle$ of a set of realizations $Z^{(i)}$ that is passing at presence t through a certain common state z is the state z: $\langle Z(t+\Delta t)\rangle_{Z(t)=z} = z$. Additional information about states z' visited at times t' prior to t is irrelevant.

Definition of the Subject

Let $\hat{X}$ denote a stochastic process defined on the space Ω and the time interval $[t_0, \infty]$, where t_0 denotes the initial time of the process. We assume that the process $\hat{X}$ can be described in terms of a random variable $X \in \Omega$. More precisely, let $X(t)$ denote the time-dependent evolution of the random variable X for $t \geq t_0$. Then, we assume that the process $\hat{X}$ can be described in terms of the infinitely large set of realizations $X^{(i)}(t)$ of $X(t)$ with $i = 1, 2, \ldots$. The realizations $i = 1, 2, \ldots$ constitute a statistical ensemble. At every time t the probability density P of $\hat{X}$ can be computed from the realizations $X^{(i)}(t)$, that is, from the ensemble by means of

$$P(x,t) = \langle \delta(x - X(t)) \rangle \ , \tag{1}$$

where $\langle\cdot\rangle$ denotes ensemble averaging and $\delta(\cdot)$ is the delta function. We assume that at time t_0 the process is distributed like u. That is, the function $u(x)$ describes the initial probability density of $\hat{X}$ and we have $P(x, t_0) = u(x)$. In general, the evolution of P depends on how the process is distributed at initial time t_0. In order to emphasize this point we will use in what follows the notation $P(x, t; u)$. That is, we interpret Eq. (1) as a conditional probability density with the constraint given by the initial distribution u:

$$P(x,t;u) = \langle \delta(x - X(t)) \rangle_{\langle \delta(x - X(t_0)) \rangle = u(x)} \ . \tag{2}$$

We may also say that we study a family of stochastic processes [80]. Each family member has a label or name which is given by u. For example, consider three experiments in which the evolution of dust particles in the air is observed for Gaussian, Lévy, and Cauchy initial distributions, respectively. It is known that dust particles perform a so-called Brownian random walk. So we would distinguish the three members $\hat{X}_1, \hat{X}_2, \hat{X}_3$ of our family of Brownian walk processes by the names of their initial distributions: Gauss, Lévy, and Cauchy.

Let us consider a stochastic process $\hat{X}$ whose evolution of its probability density P is defined by a partial differential equation of the form

$$\frac{\partial}{\partial t} P(x,t;u) = \left[-\frac{\partial}{\partial x} D_1(x,t) + \frac{\partial^2}{\partial x^2} D_2(x,t) \right] P(x,t;u) \, , \tag{3}$$

where D_1 and D_2 are functions of state x and time t. The functions D_1 and D_2 are referred to as drift- and diffusion coefficients and constitute the Fokker–Planck operator

$$L^0(x,t) = -\frac{\partial}{\partial x} D_1(x,t) + \frac{\partial^2}{\partial x^2} D_2(x,t) \,. \tag{4}$$

The evolution equation (3) is linear with respect to P. In this sense Eq. (3) is a linear partial differential equation. Irrespective of this feature the coefficients D_1 and D_2 may depend in a highly nonlinear fashion on the state x. For example, we may have $D_1 = -x + x^3$.

For appropriately chosen coefficients D_1 and D_2 Eq. (3) describes the probability density P of a Markov process. In this case Eq. (3) is referred to as a Fokker–Planck equation. More precisely, if $\hat{X}$ is a Markov diffusion process whose probability density P is defined by Eq. (3) then Eq. (3) is called a Fokker–Planck equation. Note that roughly speaking a Markov *diffusion* process is a Markov process characterized by a partial differential operator that can be truncated after the second-order partial derivative (see Sect. "Kramers–Moyal Expansion"). In order to distinguish between linear and nonlinear Fokker–Planck equations, we will use the phrase 'linear Fokker–Planck equation' instead of 'Fokker–Planck equation'.

Let us generalize Eq. (3) by assuming that the drift- and diffusion coefficients depend on the probability density P. In this case, Eq. (3) becomes

$$\frac{\partial}{\partial t} P(x,t;u) = \left[-\frac{\partial}{\partial x} D_1(x,t,P(x,t;u)) + \frac{\partial^2}{\partial x^2} D_2(x,t,P(x,t;u)) \right] P(x,t;u) \,. \tag{5}$$

Likewise, the operator (4) is generalized to

$$L(x,t,P(x,t;u)) = -\frac{\partial}{\partial x} D_1(x,t,P(x,t;u)) + \frac{\partial^2}{\partial x^2} D_2(x,t,P(x,t;u)) \,. \tag{6}$$

Equation (5) is nonlinear with respect to $P(x,t;u)$. Since the structure of the differential operator in the bracket of Eq. (5) is equivalent to the structure of the differential operator (4), evolution equations of the form (5) are frequently called nonlinear Fokker–Planck equations. In this context it is important to realize that the phrase 'nonlinear Fokker–Planck equation' does not necessarily imply that we are dealing with a Markov process. The phrase 'nonlinear Fokker–Planck equation' simply means that we are dealing with a nonlinear partial differential equation involving a partial differential operator that exhibits the structure of a Fokker–Planck operator.

Linear and Non-linear Fokker–Planck Equations, Table 1
Definition of linear and nonlinear Fokker–Planck equations based on structure, existence of solutions, and Markov property

Fokker–Planck equations	Linear	Nonlinear
Structure	Eq. (3) –(4)	Eq. (5) –(6)
Solutions	exist	do not necessarily exist
Corresponding processes are Markov processes?	yes	maybe

Note again that if an evolution equation of the form (3) is referred to as a Fokker–Planck equation then it is tacitly assumed that there exists a stochastic process defined by that equation and that this process is a Markov process. Table 1 summarizes how to define linear and nonlinear Fokker–Planck equations by means of structure, existence of solutions, and Markov property.

Linear Fokker–Planck equations are an indispensable tool to describe stochastic processes in a variety of disciplines, see Fig. 1. The theoretical concept of Markov diffusion processes related to linear Fokker–Planck equations is well-established. Researchers, applied scientist, technicians, research and development engineers in general and financial engineers in particular are usually aware that the particular linear Fokker–Planck model they are using belongs to class of Markov models. That is, the world of linear Fokker–Planck equations is closed and connected.

Nonlinear Fokker–Planck equations are used in a variety of fields that are as diverse as the application fields of linear Fokker–Planck equations. Unfortunately, so far, there is no well-established theory connecting all kinds of nonlinear Fokker–Planck equations. There is not even an academic consent about how to define them at all. This is why in Table 1 we used a very general and less constraining definition for nonlinear Fokker–Planck equations. Concepts of nonlinear Fokker–Planck equations are often developed for particular purposes and are not put into other contexts. That is, theoretical results and other achievements are often tailored to serve special needs and are not discussed in a larger framework. Even worse, so far, a well-established link between linear and nonlinear Fokker–Planck equations that applies to the variety of nonlinearities found in the literature does not exist. In sum, the world of nonlinear Fokker–Planck equations is disconnected. Different types of nonlinear Fokker–Planck equations and different application fields of nonlinear Fokker–Planck equations are often not related to each other and nonlinear Fokker–Planck equations are only loosely connected with their linear 'relatives', see Fig. 1. For example, it will most likely take years before a group of liquid crystal scientist will apply a new approach to solve numerically

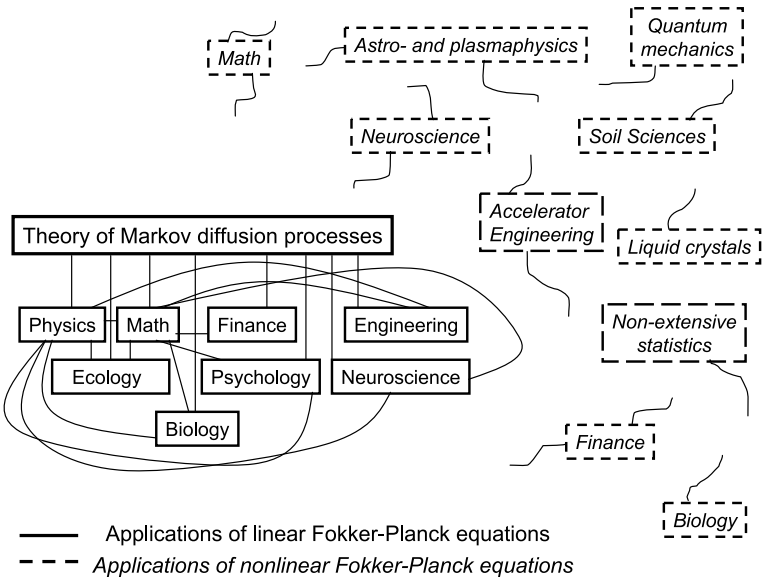

Linear and Non-linear Fokker–Planck Equations, Figure 1
Connected and disconnected applications of linear and nonlinear Fokker–Planck equations

a nonlinear Fokker–Planck equations that was discovered by a group of accelerator scientists. In addition, most likely this new solution method will not be discussed in the context of linear Fokker–Planck equations at all.

Therefore, there is a need for developing an unifying approach to nonlinear Fokker–Planck equations that involves the concept of linear Fokker–Planck equations and applies to all types of nonlinearities and in doing so applies to all kinds of scientific disciplines. Some first efforts in this regard have been made previously [3,31,32,69,70,71,80,85,87,88,127,128,199,200]. In the following sections we will review these efforts, present them in a consistent way, and in doing so make a further effort into this direction.

Introduction

Linear and nonlinear Fokker–Planck equations are widely used to describe stochastic phenomenon, see Fig. 1.

Linear Fokker–Planck equations [98,108,184] have been introduced by Fokker [66] and Planck [176]. In physics, linear Fokker–Planck equations have been used for example to describe Brownian motion, that is, the diffusion of dust particles in air or fluid layers [180]. Linear Fokker–Planck equations have been applied in engineering sciences for example to describe fluctuations in electronic circuits [98]. Linear Fokker–Planck equations have been frequently used in chemistry to model stochastic aspects of chemical reactions [107,228]. In finance, one of the most important applications of the Fokker–Planck theory is option pricing by means of the so-called Black–Scholes model [172]. Linear Fokker–Planck equations of biology systems [101] have been concerned for example with so-called Brownian motors [106,181]. Population diffusion [167] and group behavior [60,157,193] in ecological systems and stochastic neuronal processing [116] are further examples of application fields of linear Fokker–Planck equations. In psychology linear Fokker–Planck models have been proposed for decision making [19,179] and group behavior [194].

Many applications of nonlinear Fokker–Planck equation are related to several benchmark models: the Desai–Zwanzig model [52], the liquid crystal model, [55,115], the Kuramoto–Shinomoto–Sakaguchi model [3,143,207], the Vlasov model, and the nonlinear diffusion equation [12,173]. Let us highlight some of these benchmark models.

Desai–Zwanzig Model

The Desai–Zwanzig model

$$\frac{\mathrm{d}}{\mathrm{d}t}P(x,t;u) = -\frac{\partial}{\partial x}\Bigg[h(x) - \kappa\left(x - \int xP(x,t;u)\,\mathrm{d}x\right) + Q\frac{\partial}{\partial x}\Bigg]P(x,t;u) \quad (7)$$

for $\kappa, Q > 0$ has been proposed by Desai and Zwanzig [52] and Kometani and Shimizu [139] to study collective phenomena in self-organizing systems.

- A Lyapunov functional approach to the Desai–Zwanzig model has been introduced by Shiino [195,196] and since then has found several generalizations [31,32,47,68,80,88,89,127,128,135,192,197,198,199,200]. With such a Lyapunov functional at hand, the stability of stationary probability densities, collective phenomena and bifurcations can be studied by means of Lyapunov's direct method.
- The original Desai–Zwanzig model and Various modifications of it have been discussed [45,146,148].
- The additive noise term in Eq. (7) has been replaced by a multiplicative noise term [117] in order to study the interplay between the nonlinearity and the multiplicative noise [17,159,239].
- Fluctuation-dissipation theorems for stochastic processes described by the Desai–Zwanzig model have been derived [59,76,158].
- The Desai–Zwanzig model has frequently been used as a mean field approximation of spatially distributed systems with diffusive coupling. By means of such a mean field approximation analytical result have been derived and compared with numerical simulations [96,97,225,226,227].

Liquid Crystal Model

The nonlinear Fokker–Planck equation proposed by Hess [115] and Doi and Edwards [55] reads

$$\frac{\partial}{\partial t}P(\mathbf{x},t;u) = D_r\mathbf{L}\cdot\left\{\mathbf{L} + \frac{1}{kT}\left[\mathbf{L}\,e(\mathbf{x},P)\right]\right\} P(\mathbf{x},t;u) \quad (8)$$

with $D_r, k, T > 0$ and $\mathbf{L} = \mathbf{x} \times \partial/\partial\mathbf{x}$. The function $e(\mathbf{x}, P)$ describes the self-consistent potential of the Maier–Saupe mean field force. For processes $\hat{X}$ that exhibit cylindrical symmetry $e(\mathbf{x}, P)$ reads

$$e(\theta,P) = -U_0 kT\frac{3\cos^2\theta - 1}{2}\left\langle\frac{3\cos^2\theta - 1}{2}\right\rangle, \quad (9)$$

where θ is related to the unit vector $\mathbf{x}$ by $\mathbf{x} = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$. Equation (8) and generalization of it have been extensively studied in the literature [64,65,118,120,121,145] (see also [163] in general and Sect. 6.3.2 in [163] in particular). We will return to this model in Sect. "Liquid Crystal Model".

Winfree and Kuramoto Model

Winfree's seminal studies on synchronization among animal populations [234,235] supported the interest in the nonlinear Fokker–Planck equation

$$\frac{\partial}{\partial t}P(x,t;u) = \left\{-\frac{\partial}{\partial x}\left[h(x) - \kappa \cdot\int\sin(x-y)P(y,t;u)\,\mathrm{d}y\right] + Q\frac{\partial^2}{\partial x^2}\right\}P(x,t;u), \quad (10)$$

that has been proposed by Kuramoto and co-workers [143]. In Eq. (10) $h(x)$ is a 2π-periodic function and κ, $Q > 0$.

- While the Kuramoto–Shinomoto–Sakaguchi model involves an interaction term $\int a(x,y)P(y,t)\,\mathrm{d}y$, the model originally proposed by Winfree exhibits a coupling term of the form $\int a(y)P(y,t)\,\mathrm{d}y$. Models of this latter kind have also been address in [8,147,178].
- The Kuramoto–Shinomoto–Sakaguchi model describes an ensemble of phase oscillators. The eigenfrequencies of the phase oscillators do not occur in Eq. (10) because Eq. (10) describes an ensemble of phase oscillators exhibiting the same eigenfrequency ω. In this case the common eigenfrequency ω can then be eliminated by means of a variable transformation into a rotating frame [80]. However, in general, we may think of ensembles of coupled phase oscillators with different eigenfrequencies. In this context, the question arises as to what extent oscillators with different eigenfrequencies synchronize their behavior [2,7,21,38,141,174,188,209].
- Coupled phase oscillator models of the form (10) have been used to describe associative memories [237,238].
- Just as for the Desai–Zwanzig model, the interplay between multiplicative noise and the nonlinearity of the Kuramoto–Shinomoto–Sakaguchi model has been investigated in several studies [136,170,182,183].
- The *sine*-coupling term in Eq. (10) has been replaced by higher-order coupling functions $\sin(2z)$, $\sin(3z)$, ... [6,40,41,111,143]. In this context, Daido proposed the so-called order function [40,41] that generalizes the notion of cluster phases and cluster amplitudes [143]. This order function has also been related to experimental data [240].
- The Kuramoto–Shinomoto–Sakaguchi has found clinical applications in the context of Parkinsonian disease [216,217,218] (see also [215]).

Vlasov–Fokker–Planck Model

Vlasov–Fokker–Planck models frequently describe particle systems with electromagnetic interactions between charged particles. A typical examples of a Vlasov–Fokker–

Planck equation is shown here [13,137]:

$$\frac{\partial}{\partial t}P(\mathbf{v},t;u) = -\sum_{i=1}^{3}\frac{\partial}{\partial v_i}D_i(\mathbf{v},P)P + \sum_{i,k=1}^{3}\frac{\partial^2}{\partial v_i \partial v_k}D_{ik}(\mathbf{v},P)P\,. \quad (11)$$

Equation (11) involving the drift- and diffusion-coefficients

$$D_i(\mathbf{v},P) = a\frac{\partial}{\partial v_i}\int_{\Omega}\frac{P(\mathbf{v}',t;u)}{|\mathbf{v}-\mathbf{v}'|}\,\mathrm{d}^3v'\,,$$

$$D_{ik}(\mathbf{v},P) = b\frac{\partial^2}{\partial v_i v_k}\int_{\Omega}\left|\mathbf{v}-\mathbf{v}'\right|P(\mathbf{v}',t;u)\,\mathrm{d}^3v'\,. \quad (12)$$

Time-dependent solutions [5,149,160,187,213] for Eq. (12) and generalization of Eq. (12) that account for additional drift forces [26,62] have been studied. In particular, numerical methods using short-time propagators (see Sect. "Short-Time Propagator") have been developed for Vlasov–Fokker–Planck equations of the form (12) [56,57,205]. Such nonlinear Vlasov–Fokker–Planck equations play important roles in plasma physics [13,137,160] and astrophysics [16,144]. In general, astrophysical problems often require a stochastic description in terms of nonlinear Fokker–Planck equations [33,200,206]. Finally note that Vlasov–Fokker–Planck models have been used in accelerator physics and accelerator engineering to examine instabilities in particle beams [72,82,113,114,203,211,230].

Nonlinear Diffusion Equation, Nonextensive Thermostatistics, and Semi–classical Descriptions of Quantum Systems

The nonlinear diffusion equation [12,173] reads

$$\frac{\partial}{\partial t}P(x,t;u) = -\frac{\partial}{\partial x}h(x)P(x,t;u) + \frac{\partial^2}{\partial x^2}D(P(x,t;u))\,, \quad (13)$$

where $D(P)$ is a diffusion coefficient that depends on the probability density $P(x,t)$ of $X(t)$. In the original version of the nonlinear diffusion equation, the drift coefficient $h(x)$ vanishes and the diffusion coefficient is proportional to a power of P. In general, there might be a more complicated dependence of D on P [37,43].

- Since fluid flow through porous materials is a important application of the nonlinear diffusion equation, nonlinear diffusion plays a crucial role in soil sciences [14]. In biology, nonlinear diffusion equations of the form (13) seem to capture particular aspects of population diffusion [104,167].
- The nonlinear diffusion equation (13) provides a link to stochastic processes subjected to Tsallis's nonextensive thermostatistics [1,219,220,221]. For $D(P) \propto P^q$ Eq. (13) becomes

$$\frac{\partial}{\partial t}P(x,t;u) = -\frac{\partial}{\partial x}h(x)P(x,t;u) + Q\frac{\partial^2}{\partial x^2}P^q(x,t;u)\,. \quad (14)$$

  Plastino and Plastino showed that stationary distributions of Eq. (14) correspond to canonical distribution that can be derived in a nonextensive framework [177]. Equation (14) has turned out to be a testbed for various analytical and numerical studies [22,31,32,35,58,200,222]. Alternative nonlinear Fokker–Planck equation related to the Tsallis statistics have been derived from master equations in [39,161]. In addition, Eq. (14) has more recently discussed in finance in the context of a generalized Black–Scholes model for option pricing [23,24,25,229] and fat tail distributions [36,156].
- The nonlinear diffusion equation (13) is also related to semi-classical descriptions of quantum mechanical systems. For an appropriate choice of D nonlinear Fokker–Planck equations for Fermi–Bose and Einstein–Dirac systems have been derived from Eq. (13) [80,85]. Alternative forms of nonlinear Fokker–Planck equations have been derived from quantum mechanical Boltzmann equations [127,129,130], on the basis of energy balance equations [223,224], and by means of other techniques [75,79,125,206]. We will return to semi-classical quantum mechanical descriptions in Sect. "Semi-classical Description of Quantum Systems".

In addition nonlinear Fokker–Planck equations have been turned out to be useful models to describe stochastic aspects of Josephson arrays [105,233], Landau damping [210], arrays of semi-conductor lasers [142], charge density waves [20], and neurons coupled by Hodgkin–Huxley equations [109,110].

Stochastic systems composed of different kinds of interacting subsystems or species have been modeled in terms of multivariate nonlinear Fokker–Planck equations [44,95,119]. For example, the collective behavior of coupled relaxation oscillators has been studied [236]. Networks of neural oscillators as defined by the Wilson–Cowan model, [191], the two-dimensional Morris–Lecar system [109], the FitzHugh–Nagumo equations [112,126,169], and the Hidmarsh–Rose equations [186] have been studied.

The dynamics of mean field coupled phase oscillators under the impact of inertia effects [4] and related models

have attracted considerable attention. Bridge vibrations induced by pedestrian walking have discussed in this context recently [61,208]. Models for circadian rythms have been examined [42].

Solutions of the Kadar–Parisi–Zhang equation have been examined by means of nonlinear Fokker–Planck equations [99,151]. In doing so, the growth of surfaces and roughening phenomena have been studied.

Wetting processes [51], interacting Brownian motors [15,189], and spatially distributed phase oscillators [133,134] have been analyzed by means of the nonlinear Fokker–Planck perspective.

In the mathematical literature, a seminal study on nonlinear Fokker–Planck equations of the Burgers equation type was due to McKean Jr. [153]. In particular, the convergence of stochastic processes described by multivariate linear Fokker–Planck equations to processes described by nonlinear Fokker–Planck equations [29,45,46,53,54,67,94,103,124,153,154,155,162,168,175,185] and martingales of stochastic processes defined by nonlinear Fokker–Planck equations have been addressed [54,67,94,102,103,124,154,155,168]. Moreover, the propagation of molecular chaos has been studied [20,154,155]. The convergence of transient solutions of nonlinear Fokker–Planck equations to stationary ones has been examined by means of functionals that are similar to the Lyapunov functionals introduced by Shiino (see above) [9,10,11,27,28]. In addition, from a purely mathematical perspective nonlinear Fokker–Planck equations should be consider as nonlinear parabolic partial differential equations that have been discussed in several textbooks [91].

In what follows we will show that there is a common theoretical framework that unifies most of the aforementioned studies on nonlinear Fokker–Planck equations and includes the theory of linear Fokker–Planck equations as a special case. This common theoretical framework is rooted in the notion of Markov processes and martingales.

Time–Dependent Solutions and First Order Statistics

Linear Case

Equation (3) defines the evolution of P given an initial distribution $u(x)$. The norm of the probability density P equals unity for all times provided that the norm of $u(x)$ equals unity. That is, if $\int_\Omega u(x)\mathrm{d}x = 1$ holds we have $\int_\Omega P(x,t;u)\mathrm{d}x = 1$ for $t \geq t_0$. We can see this by integrating Eq. (3) with respect to x. For appropriate boundary conditions it can be shown by partial integration that the right hand side vanishes which implies that $\mathrm{d}[\int_\Omega P(x,t;u)\,\mathrm{d}x]/\mathrm{d}t = 0$ holds. The formal solution of Eq. (3) reads

$$P(x,t;u) = \exp\left\{\int_{t_0}^{t} \mathrm{d}z\, L^0(x,z)\right\} u(x)\,, \tag{15}$$

where L^0 is defined in Eq. (4). Equation (15) can be used to solve Eq. (3) numerically (see Vol. 1, Sect. 6.5 in [108]). Let t_n denote a discrete time point $t_n = t_0 + n\Delta t$ with $n = 0, 1, 2, \ldots$, where Δt is the interval of a single time step and should be small. Let us define $P_n(x;u) = P(x,t_n;u)$. Then, we have

$$P_{n+1}(x;u) = \left\{1 + \Delta t L^0(x,t_n)\right\} P_n(x;u) \tag{16}$$

with $P_0 = u(x)$. If $\hat{X}$ corresponds to an autonomous processes, then the coefficient D_1 and D_2 do not depend on t. In this case P_n can be expressed in terms of u as

$$P_n(x;u) = \left[1 + \Delta t L^0(x)\right]^n u(x)\,. \tag{17}$$

Numerical solutions converge to exact solutions in the limit $\Delta t \to 0$.

Nonlinear Case

For appropriately chosen drift- and diffusion coefficients Eq. (5) exhibits time-dependent solutions P. By analogy with the linear case, these solutions are normalized to unity provided that appropriate boundary conditions hold and that the initial probability density is normalized to unity. Solutions of Eq. (5) are formally defined by

$$P(x,t;u) = \exp\left\{\int_{t_0}^{t} \mathrm{d}z\, L(x,z,P(x,z;u))\right\} u(x)\,. \tag{18}$$

The time-dependent solutions P can be computed numerically by analogy to the linear case discussed above. That is, the probability densities $P_n(x;u) = P(x,t_n;u)$ on the discrete time grid $t_0, t_0 + \Delta t, t_0 + 2\Delta t, \cdots$ can be computed from

$$P_{n+1}(x;u) = \{1 + \Delta t L(x,t_n,P_n(x;u))\}\, P_n(x;u) \tag{19}$$

with $P_0 = u(x)$ and $n = 0, 1, 2, \ldots$. If drift- and diffusion coefficients do not explicitly depend on t, we find that the operator L still depends on t because it depends on the time-dependent solution P that in turn depends on t. Consequently, it is not trivial to generalize Eq. (17) to the nonlinear case. If the drift- and diffusion coefficients do not depend explicitly on time t and the process converges to a stationary one, then the nonlinear Fokker–Planck operator L does not depend on time. This implies that the stationary probability density P_{st} satisfies

$$P_{\mathrm{st}} = \{1 + \Delta t L(x,P_{\mathrm{st}})\}\, P_{\mathrm{st}} + O(\Delta t^2)\,. \tag{20}$$

Finally, note that we do not necessarily need to define the

formal solution with respect to the initial probability density u as in Eq. (18). We can solve the nonlinear Fokker–Planck equation on the time interval $[t_0, t]$ by splitting the solution in two intervals $[t_0, t']$ and $[t', t]$. Then, we obtain

$$P(x, t; u) = \exp\left\{\int_{t'}^{t} \mathrm{d}z\, L(x, z, P(x, z; u))\right\} P(x, t'; u)\,. \quad (21)$$

Equation (21) can be solved iteratively by means of Eq. (19) yielding a mapping $T_{\Delta t}$: $P(x, t; u) = T_{t-t'}[P(x, t'; u)]$ with $\Delta t = t - t'$.

Markov Property, 2nd Order and Higher Order Statistics

Conditional Probability Densities

Let $p(x)$ define the probability density of the time-dependent random variable X at time t given that X assumed at earlier times $t', t'', t''', \ldots$ with $t \geq t' > t'' > t''' > \ldots$ particular values $x', x'', x''', \ldots$. Then p is defined by

$$p = \langle \delta(x - X(t)) \rangle_{X(t')=x', X(t'')=x'', X(t''')=x''', \ldots}\,. \quad (22)$$

In order to point out the information that we need to compute p, we write

$$p = p(x, t|x', t'; x'', t''; x''', t''', \ldots)\,. \quad (23)$$

A conditional probability density is a relation that gives us estimates about future events and tells us what we need to know in order to be able to calculate these estimates. In our example given by Eq. (23) we see that we need the information of random values X at times $t' > t'' > t''' > \cdots$ in order to make a prediction about the statistics or probability density of X at time t. Alternatively, we may say that the conditional probability density depends on a list of variables. In the context of Markov processes this alternative viewpoint however gives raise to a problem that will be discussed below.

If $\hat{X}$ is a Markov process then the information about the stochastic process available at one particular time t' is sufficient to make predictions about the future $t \geq t'$. Adding more information about how the process evolved in the past before t' does not improve these predictions. That is, the information of events at time t' are sufficient to make statistical estimates about events at time $t \leq t'$. An alternative definition of a Markov process is that a Markov process exhibits a conditional probability density $p(x, t|\cdot)$ that depends only on one time point prior to t. That is, according to the first definition we look from time t' into the future, whereas according to the second definition we look in the opposite direction: we look from time t into the past.

For example, in order to describe the probability density $p(T)$ of the temperature T in Boston (USA) on December 1st, 2007, given that on November 1st, 2007, the temperature was 2 degrees Celsius and on October 1st, 2007, the temperature was 3 degrees Celsius, we would define the conditional probability density $p(T,\ \text{Dec 1st 2007} \mid T = 2,\ \text{Nov 1st 2007};\ T = 3,\ \text{Oct 1st 2007})$. If the temperature T as a function of time t is a Markov process, it is sufficient to know the temperature at November 1st in order to compute the probability density $p(T)$ at December 1st. For example, we would obtain the same function $p(T)$ for the conditions (i) and (ii) with (i) $T = 2°C$ on Nov 1st and $T = 3°C$ on Oct 1st and (ii) $T = 2°C$ on Nov 1st and $T = 5°C$ on Oct 1st. That is, we would have

$$\begin{aligned} &p(T,\ \text{Dec 1st 2007} \mid T = 2,\ \text{Nov 1st 2007};\\ &\qquad\qquad T = 3,\ \text{Oct 1st 2007})\\ &= p(T,\ \text{Dec 1st 2007} \mid T = 2,\ \text{Nov 1st 2007};\\ &\qquad\qquad T = 5,\ \text{Oct 1st 2007})\\ &= p(T,\ \text{Dec 1st 2007} \mid T = 2,\ \text{Nov 1st 2007})\,. \end{aligned} \quad (24)$$

The information about the October temperature is irrelevant. In this sense the conditional probability density would depend on the November temperature but would not depend on the October temperature.

A problem that arises in the context of the definition of Markov process is as follows. Suppose that there is a purely deterministic dynamical aspect involved in a stochastic process. In our example about Boston temperatures we may think of the annual period changes of the temperature that are related to the annual changes in distance and declination angle between earth and sun. Let us assume that distance and declination angle change periodically in a purely deterministic fashion such that the distance and declination angle at November 1st can be computed from the distance and declination angle at January 1st by a simple one-to-one mapping. Then the question arises: does the temperature in Boston on December 1st depend on the distance and declination angle of November 1st as suggested by $p(T,\ \text{Dec 1st 2007} \mid T = 25,\ \text{Nov 1st 2007})$ or does it depend on the distance and declination angle of January 1st. In the former case we have a Markov conditional probability density. In the latter case we would need to write p like

$$p(T,\ \text{Dec 1st 2007} \mid T = 25,\ \text{Nov 1st 2007};\ \text{distance and angle,}\ \text{Jan 1st 2007})$$

indicated that we are dealing with a non-Markovian process. The situation becomes even worse if we take into consideration that the earth-sun distance and the declination

angle at January 1st 2008 can be computed from the information known at November 1st using our simple one-to-one mapping. Therefore, we may say that the temperature estimate for December 1st, 2007, depends on a future event, namely, the earth-sun distance and declination angle given at January 1st, 2008. The conditional probability density would assume the form

$$p(T,\ \text{Dec 1st 2007} \mid T = 25,\ \text{Nov 1st 2007};\ \text{distance and angle,}\ \text{Jan 1st 2008})\,,$$

which would suggest again that we are dealing with a non-Markovian and – to a certain extent – non-causal process. We can solve this problem by realizing that purely deterministic relationships in time that represent external driving forces are irrelevant for the distinction between Markov and non-Markov processes. We can completely determine such external driving force by a parameter set $\{t_0, A_1, A_2, \dots\}$ that describe the initial state of the driving forces. Although this initial state is related to the initial time t_0 of the Markov process the conditional probability density does not actually depend (i. e. it does not explicitly depend) on t_0. Likewise the conditional probability density does not actually depend on the parameters $\{A_1, A_2, \dots\}$. The information that we have at time t' includes the information about the driving force at time t' and therefore the information about the driving force at all times $t \in [t_0, \infty)$. Consequently, the information at time t' is sufficient to predict how the driving force will evolve in the future at times $t \le t'$. There is no need to assess information about events prior to t' or information about events that will happen in the future at times larger than t' in order to determine the evolution of the deterministic driving force.

Let us summarize. A stochastic process $\hat{X}$ is called a Markov process if information about the process at time t' is sufficient to make predictions about future events. This implies that the conditional probability density p defined in Eq. (23) can be simplified like

$$p(x, t \mid x', t'; x'', t''; x''', t''', \dots) = p(x, t \mid x', t')\,. \quad (25)$$

Note that we may say that p depends only on the state x' related to the time t' in the sense that the information at time t' is sufficient to predict how X will be distributed at time $t \le t'$. We may say it is sufficient to *best* predict future events where *best* refers to the fact that adding additional information about the past does not improve our predictions.

Let us illustrate this issue by another example. Let $p(x, t \mid X = \theta)$ denote the probability density of X at time t given that X equals the function θ in the interval $[t_0, t']$ with $t' \le t$. If X describes a Markov process we have

$$\begin{aligned} p(x, t \mid X = \theta) &= p(x, t \mid \theta(t'), t') \\ &= p(x, t \mid x', t') \end{aligned} \quad (26)$$

with $x' = \theta(t')$.

Linear Fokker–Planck Equations

As mentioned in Sect. "Definition of the Subject", linear Fokker–Planck equations describe Markov processes [98, 184]. Markov processes related to linear Fokker–Planck equations of the form (4) have conditional probability densities defined by

$$\frac{\partial}{\partial t} p(x, t|x', t') = L^0(x, t)\, p(x, t|x', t') \quad (27)$$

with $\lim_{t \to t'} p(x, t|x', t') = \delta(x - x')$. The conditional probability density p is also called the fundamental solution or Green's function of the Fokker–Planck equation (3). In general a stochastic process $\hat{X}$ is completely defined in terms of the joint probability density

$$\begin{aligned} P(x_n, t_n; x_{n-1}, t_{n-1}; \dots; x_0, t_0) = \langle &\delta(x_n - X(t_n)) \\ &\cdot \delta(x_{n-1} - X(t_{n-1})) \cdots \delta(x_0 - X(t_0)) \rangle\,, \end{aligned} \quad (28)$$

where n can assume arbitrarily large integer numbers. In particular, if $\hat{X}$ is a Markov process then this joint probability density can be computed from p and u like

$$\begin{aligned} P(\cdot) &= p(x_n, t_n|x_{n-1}, t_{n-1}) \\ &\cdot p(x_{n-1}, t_{n-1}|x_{n-2}, t_{n-2}) \cdots p(x_1, t_1|x_0, t_0) u(x_0)\,. \end{aligned} \quad (29)$$

Consequently, the linear Fokker–Planck equation (3) defines completely a Markov process via the associated evolution equation (27) and the initial distribution u.

In particular, the time-dependent probability densities $P(x, t; u)$ and $P(x, t'; u)$ with $t \ge t'$ are related to each other by means of a *linear* functional

$$P(x, t; u) = \int_\Omega p(x, t|x', t') P(x', t'; u)\, \mathrm{d}x'\,. \quad (30)$$

That is, the Green's function p induces a functional that is linear with respect to $P(x', t'; u)$.

Langevin Equations of Linear Fokker–Planck Equations

The stochastic trajectories $X(t)$ of the Markov process $\hat{X}$ defined by Eq. (3) can be computed from the Ito–Langevin equation [34,98,184]

$$\frac{\mathrm{d}}{\mathrm{d}t} X(t) = D_1(X(t), t) + \sqrt{D_2(X(t), t)}\, \Gamma(t)\,, \quad (31)$$

where $\Gamma(t)$ denotes a Langevin force normalized to the

delta function like $\langle \Gamma(t)\Gamma(t')\rangle = 2\delta(t-t')$. From the Langevin equation (31) it follows again that we are dealing with a Markov process. Information about one reference time t' is sufficient to compute the future behavior of the trajectory $X(t)$ with $t \geq t'$. On a discrete time grid the stochastic trajectories or realizations of $\hat{X}$ can be computed iteratively like [138,184]

$$X_{n+1} = X_n + \Delta t\, D_1(X_n, t_n) + \sqrt{\Delta t\, D_2(X_n, t_n)}\,\epsilon_n \quad (32)$$

with $X(t_n) = X_n$, $t_n = t_0 + n\Delta t$, and $n = 0, 1, 2, \ldots$ Here, ϵ_n are independent Gaussian distributed random numbers with vanishing mean and variance 2. That is, we have $\langle \epsilon_n \rangle = 0$ and $\langle \epsilon_n \epsilon_m \rangle = 2\delta_{nm}$, where δ_{nm} is the Kronecker symbol. Finally, the probability density W of ϵ_n at every step n is given by

$$W(\epsilon_n) = \frac{1}{\sqrt{4\pi}} \exp\left\{-\frac{\epsilon_n^2}{4}\right\} . \quad (33)$$

Note that we do not necessarily need to start the iteration scheme at t_0. The scheme (32) can be started at any time step n. Moreover, in order to compute the subsequent time steps it is sufficient to have information about the random variable X at time t_n. Consequently, the sequence $X_n, X_{n+1}, X_{n+2}, \ldots$ computed from Eq. (32) describes a trajectory of a Markov process.

Strongly Nonlinear Fokker–Planck Equations

In Sect. "Definition of the Subject" we pointed out that there is some kind of asymmetry in the definition of linear and nonlinear Fokker–Planck equations. While a linear Fokker–Planck equation defines a stochastic process, a nonlinear Fokker–Planck equation defines at best the evolution of a probability density $P(x,t)$. That is, if solutions of Eq. (5) exist for $u \in U$, then Eq. (5) defines the evolution of first-order statistical properties of a stochastic process $\hat{X}$ such as the time-dependent probability density, the mean and the variance of the process $\hat{X}$. In any case, Eq. (5) does not define second- and higher-order statistical quantities such as correlation functions and conditional probability densities [77]. In particular, the time-dependent solutions P of Eq. (5) in general can not be used to construct Green's functions of Markov processes because they do not necessarily correspond to Green's functions of Markov processes [73]. Note that this is not a peculiarity of stochastic processes defined by nonlinear Fokker–Planck equations. In fact, time-dependent solutions P of linear Fokker–Planck equations Eq. (3) involving explicitly time-dependent coefficients D_1 and D_2 do not necessarily correspond to Green's functions. Mathematically speaking, let $P(x,t;u=\delta(x-x_0))$ denote the probability density of a process $\hat{X}$ defined by a non-autonomous linear Fokker–Planck equation or by a nonlinear Fokker–Planck equation and let $p(x,t|x',t')$ denote the conditional probability density of that process $\hat{X}$, then we have [73]

$$P\big(x,t;u=\delta\big(x-x'\big)\big)_{t_0=t'} \quad \text{is not necessarily equivalent to } p\big(x,t \mid x',t'\big) , \quad (34)$$

where $P(x,t;u=\delta(x-x'))_{t_0=t'}$ means that we take the time-dependent solution $P(x,t;u=\delta(x-x_0))$ and replace in this expression x_0 by x' and t_0 by t'.

Let us return to the issue how to define a stochastic process $\hat{X}$ on the basis of a nonlinear Fokker–Planck equation (5). In order to do so, we need to define appropriate constraints such that out of all possible stochastic processes that exhibit a time-dependent probability density P defined by Eq. (5) one particular process is selected. In what follows, we will discuss one particular set of constraints [77]. As we will see the stochastic processes thus defined exhibit the Markov property.

Let U denote a set of initial probability densities u. That is, U is a set of functions or a space of functions. Let $P(x,t;u)$ denote the solution of the nonlinear Fokker–Planck equation

$$\frac{\partial}{\partial t} P(x,t;u) = L(x,t,P(x,t;u))P(x,t;u) \quad (35)$$

with

$$L(x,t,P(x,t;u)) = -\frac{\partial}{\partial x} D_1(x,t,P(x,t;u)) + \frac{\partial^2}{\partial x^2} D_2(x,t,P(x,t;u)) \quad (36)$$

for an initial distribution $u \in U$. Let us introduce the associated drift- and diffusion coefficients $\tilde{D}_1$ and $\tilde{D}_2$ by

$$\begin{aligned} \tilde{D}_1(x,t;u) &= D_1(x,t,P(x,t;u)) , \\ \tilde{D}_2(x,t;u) &= D_2(x,t,P(x,t;u)) . \end{aligned} \quad (37)$$

That is, for any $u \in U$ Eq. (35) is solved analytically or by numerical iteration (19). The solution is substituted into the drift- and diffusion coefficients D_1 and D_2. The coefficients thus obtained are the functions $\tilde{D}_1$ and $\tilde{D}_2$ associated to D_1 and D_2. Let us assume that for all $u \in U$ the evolution equation

$$\frac{\partial}{\partial t} p(x,t|x',t') = \left[-\frac{\partial}{\partial x} D_1(x,t;u) + \frac{\partial^2}{\partial x^2} D_2(x,t;u) \right] \cdot p(x,t|x',t') \quad (38)$$

has a fundamental solution or Green's function. Then, this solution p and its corresponding initial distribution u define a Markov process. In [77] nonlinear Fokker–

Planck equations that induce evolution equations (38) with fundamental solutions were called strongly nonlinear Fokker–Planck equations. Note that nonlinear Fokker–Planck equations (5) do not necessarily exhibit the property of being strongly nonlinear. Note also that in some applications it might be worth to define carefully the set U of initial probability densities u such that a nonlinear Fokker–Planck equation under consideration becomes strongly nonlinear.

As indicated above the time-dependent probability densities P of a nonlinear Fokker–Planck equation depend on the initial distribution u. Likewise, the associated coefficients $\tilde{D}_1$ and $\tilde{D}_2$ depend on u. As a result the conditional probability densities $p(x, t|x', t')$ depend on u as well. For this reason, the notation $p(x, t|x', t'; u)$ has been suggested. Unfortunately, this notation is likely to cause confusion because one might think that p depends not only on the time t' but also on the initial time t_0 which seems to be incompatible with the notion of a Markov conditional probability density [83,152]. In fact, this confusion results from the second alternative way to define Markov processes that has been discussed above. The evolution of the function $P(x, t; u)$ is a purely deterministic one. That is, $P(x, t; u)$ represent a deterministic driving force for the purpose of computing the conditional probability density. The distribution u is just a parameter which determines the initial value of this driving force. In this context note again that the conditional probability density of a Markov process in general depends on parameters and in particular can depend on the initial time t_0 and other parameters $\{A_1, A_2, \dots\}$ that define the initial state of a driving force. Consequently, the notation $p(x, t|x', t'; u)$ does not imply a contradiction with the notion of a Markov process. For example, the non-autonomous Langevin equation

$$\frac{\partial}{\partial t}X(t) = -\gamma X(t) - A\cos(\omega(t - t_0)) + \sqrt{Q}\Gamma(t) \quad (39)$$

with $\gamma, A, \omega, Q > 0$ defines a Markov process that is driven by a harmonic force $-A\cos(\omega(t - t_0))$. That is, the harmonic force has amplitude A at the beginning of the process. The conditional probability density of that process depends on the parameter γ, ω, Q but also on the parameters t_0 and A which correspond to the initial amplitude and time. We have (see Sect. 3.7.3 in [80])

$$p = p(x, t \mid x', t', \gamma, \omega, Q, A, t_0)\,. \quad (40)$$

Nevertheless, in what follows we will develop a slightly different notation for conditional probability densities p of strongly nonlinear Fokker–Planck equations that is in line with the first definition of Markov processes discussed above and will be helpful to elucidate that the functions p reflect indeed Markov processes.

Let us exploit first the fact that if Eq. (35) is a strongly nonlinear Fokker–Planck equations then time-dependent solutions $P(x, t; u)$ of Eq. (35) exist for $u \in U$ and are related to their initial probability densities u by a one-to-one mapping T_t. That is, for every t we have $P(x, t; u) = T_t[u(x)]$. For an explicit construction of the map T_t see for example Eq. (18). Likewise, we have $P(x, t'; u) = T_{t'}[u(x)]$. Using the inverse of T we can map u to P like $u(x) = T_{t'}^{-1}[P(x, t; u)]$. Substituting these expressions into $p(x, t \mid x', t'; u)$, we obtain $p(x, t \mid x', t'; T_{t'}^{-1}[P(x, t'; u)]$. This result demonstrates that the information about the stochastic process $\hat{X}$ at time t' is sufficient to predict the future at $t > t'$. We can regard the conditional probability density p as a function that does not depend explicitly on u but it depends explicitly on the state of the driving force P at time t'. In line with this remark, we introduce the conditional probability densities $p(x, t \mid x', t', P(x', t'; u))$.

Let us dwell on the interpretation of the conditional probability density $p(x, t \mid x', t', P(x', t'; u))$. To this end, we need to discuss briefly the notion of a particular conditional averaging that is important in this context and will become important later on as well. Let us assume that we make observations of realizations of a stochastic process $\hat{X}$ for which the following two conditions hold: (i) $X(t') = x'$ and (ii) the ensemble of all realization is distributed like P at time t'. Next, we average across all observations that we make under these conditions. In doing so, we average under the constraints

$$X(t') = x' \text{ and } \left\langle \delta(x' - X(t')) \right\rangle = P(x', t'; u)\,. \quad (41)$$

In order to indicate that such a structured constraint should hold, we will use the notation

$$(\cdot)_{X(t')=x';\,\langle\delta(x'-X(t'))\rangle=P(x',t';u)}\,. \quad (42)$$

The conditional probability density $p(x, t \mid x', t', P(x', t'; u))$ is a special case in which the delta function is averaged under the constraint (41). We have

$$p(x, t \mid x', t', P(x', t'; u)) = \langle \delta(x - X(t)) \rangle_{X(t')=x';\,\langle\delta(x'-X(t'))\rangle=P(x',t';u)}\,. \quad (43)$$

Summarizing the results we have derived so far, we see that strongly nonlinear Fokker–Planck equations define Markov processes whose

- time-dependent probability densities $P(x, t; u)$ are defined by Eq. (35)

- and conditional probability densities $p(x,t|x',t',P(x',t';u))$ are defined by

$$\frac{\partial}{\partial t}p(x,t\mid x',t',P') = L(x,t,P)\,p(x,t\mid x',t',P') \quad (44)$$

with L given by Eq. (36), $P = P(x,t;u)$, and $P' = P(x',t';u)$. Note that by multiplying Eq. (44) with $P(x',t';u)$ and integrating with respect to x' we get Eq. (35) which in turn defines the evolution of $P(x,t;u)$. Consequently, Eq. (44) defines both the evolution of $P(x,t;u)$ and $p(x,t\mid x',t',P')$. Note also that the solution of Eq. (44) formally reads

$$p(x,t\mid x',t',P(x',t';u)) = \exp\left\{\int_{t'}^{t} \mathrm{d}z\, L(x,z,P(x,z;u))\right\}\delta(x-x') \quad (45)$$

and depends on the evolution of $P(x,z;u)$ for $z\in[t',t]$. In fact, as indicated above p depends only on $P(x',t',u)$. To see this recall that the formal solution (21) can be obtained by means of the iterative method (19) such that we can write $P(x,z;u) = T_{z-t'}[P(x,t';u)]$. Substituting this solution into Eq. (45) we get

$$p(x,t\mid x',t',P(x',t';u)) = \exp\left\{\int_{t'}^{t} \mathrm{d}z\, L\big(x,z,T_{z-t'}[P(x,t';u)]\big)\right\}\delta(x-x')\,. \quad (46)$$

In addition, we find that the solution (45) does not explicitly depend on u.

We arrive at the following conclusion: conditional probability densities $p(x,t|\cdot)$ of Markov processes described by strongly nonlinear Fokker–Planck equations depend only on the value of individual realizations at one prior time $t'\le t$ and on the probability density P defined by all realizations at the very same prior time t'.

Equation (45) can be simplified for stationary Markov processes with operators L that do not depend explicitly on time t. Then the conditional probability density in the stationary case can be computed from

$$p(x,t\mid x',t',P_{\mathrm{st}}(x')) = \exp\big\{(t-t')\,L(x,P_{\mathrm{st}}(x))\big\}\,\delta(x-x')\,. \quad (47)$$

where $P_{\mathrm{st}}(x)$ denotes a stationary probability density out of a set of stationary probability densities defined by $LP_{\mathrm{st}} = 0$. Note that in this context P_{st} plays the role of an initial distribution u.

Just as in the linear case, the conditional probability density p in combination with the initial distribution u completely defines the stochastic process $\hat{X}$. In particular, the joint probability density $P(x_n,t_n;x_{n-1},t_{n-1};\ldots;x_0,t_0)$ can be computed from p and u like

$$P(\cdot) = p(x_n,t_n\mid x_{n-1},t_{n-1},P_{n-1})\cdot p(x_{n-1},t_{n-1}\mid x_{n-2},t_{n-2},P_{n-2})\cdots \cdots p(x_1,t_1\mid x_0,t_0,u)u(x_0)\,, \quad (48)$$

with $P_n = P(x_n,t_n;u)$, $P_{n-1} = P(x_{n-1},t_{n-1};u)$, and so on.

In particular, the time-dependent probability densities $P(x,t;u)$ and $P(x,t';u)$ with $t\ge t'$ are related to each other by means of a *nonlinear* functional

$$P(x,t;u) = \int_{\Omega} p(x,t\mid x',t',P(x',t';u))\,P(x',t';u)\,\mathrm{d}x'\,, \quad (49)$$

where p is defined by Eq. (46). That is, the Green's function p induces a functional that is nonlinear with respect to $P(x',t';u)$.

Langevin Equations of Strongly Nonlinear Fokker–Planck Equations

The stochastic trajectories $X(t)$ of the Markov process $\hat{X}$ defined by Eq. (44) can be computed from two-layered Langevin equations (see Sect. 3.4 in [80]) or alternatively from the self-consistent Ito–Langevin equation

$$\frac{\mathrm{d}}{\mathrm{d}t}X(t) = D_1(X(t),t,P(X(t),t;u)) + \sqrt{D_2(X(t),t,P(X(t),t;u))}\,\Gamma(t)\,, \quad (50)$$

where $\Gamma(t)$ denotes the Langevin force introduced earlier. Note that the expression $P(X(t),t;u)$ means that the function $P(x,t;u)$ is evaluated at the state x that is given by the random variable X at time t. That is, we may write $P(X(t),t;u) = P(x,t;u)|_{x=X(t)}$. From the Langevin equation (50) we can read off that we are dealing with a Markov process. Information about one reference time t' in terms of the state $X(t')$ of a realization and the distribution of the ensemble as given by the probability density $P(x,t';u)$ is sufficient to compute the future behavior of the trajectory $X(t)$ with $t\ge t'$.

The Langevin equation (50) may be implemented on a computer using the iterative scheme

$$X_{n+1} = X_n + \Delta t\, D_1(X_n,t_n,P(X_n,t_n;u)) + \sqrt{\Delta t\, D_2(X_n,t_n,P(X_n,t_n;u))}\,\epsilon_n \quad (51)$$

with $X(t_n) = X_n$, $t_n = t_0 + n\Delta t$, $n = 0,1,2,\ldots$, and ϵ_n given as statistically independent Gaussian distributed random numbers with vanishing mean and variance 2 (see

above). The expression $P(X_n, t; u)$ can be computed from the realizations generated by the iteration scheme (51). Let $X_n^{(i)}$ denote the ith realization at time step n. Then, the stochastic trajectories $X(t)$ can numerically be computed by simulating an ensemble of realizations $i = 1, \ldots, N$ like

$$X_{n+1}^{(i)} = X_n^{(i)} + \Delta t\, D_1(X_n^{(i)}, t_n, P_n(X_n^{(i)})) + \sqrt{\Delta t\, D_2(X_n^{(i)}, t_n, P_n(X_n^{(i)}))}\, \epsilon_n^{(i)}\,, \quad (52)$$

where $\epsilon_n^{(i)}$ are statistically independent Gaussian random numbers with respect to both indices n and i and P_n is computed from the set $\{X_n^{(1)}, \ldots, X_n^{(N)}\}$ of realizations using standard kernel estimators. For example, we may use

$$P_n(x) = \frac{1}{Ns\sqrt{2\pi}} \sum_{i=1}^{N} \exp\left\{-\frac{(x - X_n^{(i)})^2}{2s^2}\right\}\,, \quad (53)$$

with $s = N^{-1/5}\sigma_e(t_n)$ where $\sigma_e(t_n)$ is the standard deviation of the empirical ensemble $\{X_n^{(1)}, \ldots, X_n^{(N)}\}$ [80,84,204]. Just as in the case of Langevin equations of linear Fokker–Planck equations, the scheme (52) can be started at any time step n provided that we have information about P_n and X_n. In particular, if we start at a step $n > 0$ we see that the information about the initial distribution is irrelevant. Consequently, the sequence $X_n, X_{n+1}, X_{n+2}, \ldots$ computed from the time-discrete Langevin equation (52) related to the nonlinear Fokker–Planck equation (44) describes a trajectory of a Markov process.

Finally, note that self-consistent Langevin equations can be evaluated analytically in order to determine second-order statistical properties of a stochastic process defined by a strongly nonlinear Fokker–Planck equation [22,135].

Short-Time Propagator

The Green's function for short time intervals is frequently called the short-time propagator and can be derived from the time-discrete Ito-Langevin (51). Equation (51) relates the random variable ϵ_n that is distributed like $W(\epsilon_n)$, see Eq. (33), to the random variable X_{n+1}. In general, if X_{n+1} is a function of ϵ_n then the probability density $W'(x_{n+1})$ of X_{n+1} is given by

$$W'(x_{n+1}) = W(\epsilon_n)\frac{\mathrm{d}\epsilon_n}{\mathrm{d}x_{n+1}}\,. \quad (54)$$

In particular, if X_{n+1} is computed from ϵ_n for a particular value x_n and probability density P, then we obtain the short-time conditional probability density

$$p_s(x_{n+1} \mid x_n, P(x_n, t_n; u)) = W(\epsilon_n)\,\frac{\mathrm{d}\epsilon_n}{\mathrm{d}x_{n+1}}\,. \quad (55)$$

Equation (51) can be transformed into

$$\epsilon_n = \frac{X_{n+1} - X_n + \Delta t\, D_1(X_n, t_n, P(X_n, t_n; u))}{\sqrt{\Delta t\, D_2(X_n, t_n, P(X_n, t_n; u))}}\,. \quad (56)$$

Substituting Eq. (56) into Eq. (55) we obtain

$$p_s(x_{n+1} \mid x_n, P(x_n, t_n; u)) = \frac{\exp\left\{-\dfrac{[x_{n+1} - x_n + \Delta t\, D_1(x_n, t_n, P(x_n, t_n; u))]^2}{4\Delta t\, D_2(x_n, t_n, P(x_n, t_n; u))}\right\}}{\sqrt{4\pi\Delta t\, D_2(x_n, t_n, P(x_n, t_n; u))}}\,. \quad (57)$$

Using the time-continuous framework and the replacements $n \to t'$, $x_n \to x'$, $n+1 \to t = t' + \Delta t$, $x_{n+1} \to x$ and likewise $P(x_n, t_n; u) \to P(x', t'; u) = P'$, we obtain the short-time propagator (see Sect. 2.8.1 in [80])

$$p_s(x, t \mid x', t', P') = \frac{\exp\left\{-\dfrac{[x - x' + \Delta t\, D_1(x', t', P')]^2}{4\Delta t\, D_2(x', t', P')}\right\}}{\sqrt{4\pi\Delta t\, D_2(x', t', P')}}\,. \quad (58)$$

The short-time propagator has originally been proposed by Wehner and Wolfer [232] and can be used to solve nonlinear Fokker–Planck equations numerically [56,57,205]. To this end, the short-time propagator is substituted into Eq. (48) and subsequently Eq. (48) is integrated over all variables $x_{n-1}, \ldots, x_0$. Thus we obtain $P(x, t_n; u)$ for $t_n = n\Delta t$. In the context of linear Fokker–Planck equation the construction of solutions by means of short-time propagators is referred to as path integral approach [98,108,184]. We will return to a similar path integral approach in Sect. "Semi-classical Description of Quantum Systems".

Expectation values of functions f can be computed from (58) like

$$\langle f(X(t))\rangle_{X(t')=x';\,\langle\delta(x'-X(t'))\rangle=P(x',t';u)} = \int_\omega f(x) p_s(x, t \mid x', t', P')\,\mathrm{d}x\,, \quad (59)$$

which holds for small intervals $\Delta t = t - t'$. The short-time propagator illustrates again the Markov property of solutions of the strongly nonlinear Fokker–Planck equation (44). The information about x' and P' at time t' is sufficient to make predications in terms of expectation values that the stochastic process will assume at time $t = t' + \Delta t$.

Chapman–Kolmogorov Equation, Kramers–Moyal Expansion and Drift-Diffusion Estimates

Linear Fokker–Planck equation can be derived using the Kramers–Moyal expansion of the Chapman–Kolmogorov equation [98,184]. The definition of the expansion coeffi-

cients in turn can be used to estimate the Kramers–Moyal coefficients in general and the drift- and diffusion-coefficients of linear Fokker–Planck equations in particular from experimental data [92,93]. We will show in this section that if a stochastic process defined by a nonlinear Fokker–Planck equation can be embedded into a Markov process using the concept of strongly nonlinear Fokker–Planck equations, then we can proceed as in the linear case. Taking a slightly different perspective, we may say that there are Markov processes that involve conditional probability densities of the form $p(x, t \mid x', t', P(x', t'; u))$ and can be characterized in terms of generalized Kramers–Moyal expansion coefficients.

Chapman–Kolmogorov Equation

Let $\hat{X}$ denote a stochastic Markov process with conditional probability density $p(x, t \mid x', t', P(x', t'; u))$. Then as discussed in the previous section the joint probability $P(x, t; x', t'; x'', t''; u)$ can be expressed by

$$P(x, t; x', t'; x'', t''; u) = p(x, t \mid x', t', P(x', t'; u)) \cdot p(x', t' \mid x'', t'', P(x'', t''; u))\, P(x'', t''; u)\,. \quad (60)$$

Integrating with respect to x' and dividing by $P(x'', t''; u)$ yields the generalized Chapman–Kolmogorov equation

$$p(x, t|x'', t'', P(x'', t''; u)) = \int_{\Omega} p(x, t|x', t', P(x', t'; u)) \cdot p(x', t' \mid x'', t'', P(x'', t''; u))\, \mathrm{d}x'\,. \quad (61)$$

Note that in what follows we will use the notation

$$\begin{aligned} P &= P(x, t; u)\,, \\ P' &= P(x', t'; u)\,, \\ P'' &= P(x'', t''; u)\,. \end{aligned} \quad (62)$$

If we need to express probability densities P different from those listed in Eq. (62), we will write down if necessary their arguments explicitly. For example, we will write $P(x, t'; u)$ to express the probability density $\langle \delta(x - X(t')) \rangle$ for a stochastic process $\hat{X}$ with initial distribution u.

Using the notation of Eq. (62), we can write the joint probability density (60) like

$$P(x, t; x', t'; x'', t''; u) = p(x, t \mid x', t', P')\, p(x', t' \mid x'', t'', P'')\, P'' \quad (63)$$

and the generalized Chapman–Kolmogorov equation (61) becomes

$$p(x, t \mid x'', t'', P'') = \int_{\Omega} p(x, t \mid x', t', P')\, p(x', t' \mid x'', t'', P'')\, \mathrm{d}x'\,. \quad (64)$$

Kramers–Moyal Expansion

In this section the Kramers–Moyal expansion for linear Fokker–Planck equations as discussed in [184] will be generalized to the nonlinear case. Consider the conditional probability density $p(x, t|x', t', P')$ for $t = t' + \Delta t$. Then, we have

$$p(x, t' + \Delta t \mid x', t', P') = \int_{\Omega} \delta(y - x)\, p(y, t' + \Delta t \mid x', t', P')\, \mathrm{d}y\,. \quad (65)$$

The variables x and x' denote arbitrary states in Ω. However, let us consider next states y that are close to x' such that $\epsilon = y - x'$ is small. Using $y - x = \epsilon + x' - x$, we obtain

$$p(x, t' + \Delta t \mid x', t', P') = \int_{\Omega} \delta(x' - x + \epsilon) \cdot p(x' + \epsilon, t' + \Delta t \mid x', t', P')\, \mathrm{d}\epsilon\,. \quad (66)$$

Use

$$\delta(x' - x + \epsilon) = \delta(x' - x) + \sum_{1}^{\infty} \frac{\epsilon^n}{n!} \left(\frac{\partial}{\partial x'} \right)^n \delta(x' - x). \quad (67)$$

Then, Eq. (66) becomes

$$p(x, t' + \Delta t|x', t', P') = \delta(x' - x) + \sum_{1}^{\infty} \int_{\Omega} \mathrm{d}\epsilon \frac{\epsilon^n}{n!} \cdot p(x' + \epsilon, t' + \Delta t|x', t', P') \left(\frac{\partial}{\partial x'} \right)^n \delta(x' - x)\,. \quad (68)$$

Multiplying Eq. (68) with $p(x', t'|x'', t'', P'')$ and integrating with respect to x' yields on the left hand side

$$\begin{aligned} \mathrm{LHS} &= \int_{\Omega} p(x, t' + \Delta t|x', t', P')\, p(x', t'|x'', t'', P'')\, \mathrm{d}x' \\ &= p(x, t' + \Delta t|x'', t'', P'') \end{aligned} \quad (69)$$

and on the right hand side

$$\begin{aligned} &\mathrm{RHS} \\ &= p(x, t'|x'', t'', P'') \\ &+ \sum_{1}^{\infty} \int_{\Omega} \mathrm{d}x' \int_{\Omega} \mathrm{d}\epsilon \frac{\epsilon^n}{n!} p(x' + \epsilon, t' + \Delta t|x', t', P') \\ &\cdot p(x', t'|x'', t'', P'') \frac{\partial^n \delta(x' - x)}{\partial x'^n} \\ &= p(x, t'|x'', t'', P'') \\ &+ \sum_{1}^{\infty} \int_{\Omega} \mathrm{d}x' \delta(x' - x) \frac{\partial^n}{\partial x'^n} (-1)^n \\ &\cdot \int_{\Omega} \mathrm{d}\epsilon \frac{\epsilon^n}{n!} p(x' + \epsilon, t' + \Delta t|x', t', P') p(x', t'|x'', t'', P'') \end{aligned}$$

$$= p(x,t'|x'',t'',P'') + \sum_1^\infty \left(-\frac{\partial}{\partial x}\right)^n \int_\Omega d\epsilon \frac{\epsilon^n}{n!} p(x+\epsilon, t'+\Delta t\,|\,x,t',P(x,t';u)) \cdot p(x,t'|x'',t'',P'') \,. \quad (70)$$

Note that we used the Chapman–Kolmogorov equation (61) in order to evaluate the left hand side (69) and we used Eq. (61) as well as partial integration in order to evaluate the right hand side (70). Let us define the moments $M_n(x,t,\Delta t,P(x,t;u))$ by

$$M_n(x,t,\Delta t,P) = \int_\Omega d\epsilon \frac{\epsilon^n}{n!} p(x+\epsilon, t+\Delta t|x,t,P) \quad (71)$$

or using $\epsilon + x = z$ by

$$M_n(x,t,\Delta t,P) = \int_\Omega dz \frac{(z-x)^n}{n!} p(z,t+\Delta t|x,t,P) \,. \quad (72)$$

Combining the left and right hand sides given by Eqs. (69) and (70), respectively, we obtain

$$p(x,t'+\Delta t\,|\,x'',t'',P'') = p(x,t'\,|\,x'',t'',P'') + \sum_1^\infty \left(-\frac{\partial}{\partial x}\right)^n \cdot M_n(x,t',\Delta t,P(x,t';u))\, p(x,t'|x'',t'',P'') \,. \quad (73)$$

To improve readability let us replace t' by t and subsequently t'' by t'. Thus, we obtain

$$p(x,t+\Delta t\,|\,x',t',P') = p(x,t\,|\,x',t',P') + \sum_1^\infty \left(-\frac{\partial}{\partial x}\right)^n \cdot M_n(x,t,\Delta t,P)\, p(x,t|x',t',P') \,. \quad (74)$$

This is the time-discrete version of the Kramers–Moyal expansion of the generalized Chapman–Kolmogorov equation (64). Note that M_n depends on $P(x,t;u)$, whereas p depends on $P(x',t';u)$. Next, we define the Kramers–Moyal coefficients

$$D_n(x,t,P) = \lim_{\Delta t\to 0} \frac{M_n}{\Delta t} = \lim_{\Delta t\to 0} \frac{1}{\Delta t} \int_\Omega dx \frac{(z-x)^n}{n!} \cdot p(z,t+\Delta t\,|\,x,t,P) \,. \quad (75)$$

Dividing Eq. (74) by Δt and taking the limiting case $\Delta t \to 0$, Eq. (74) becomes the time-continuous generalized Kramers–Moyal expansion

$$\frac{\partial}{\partial t} p(x,t\,|\,x',t',P') = \sum_1^\infty \left(-\frac{\partial}{\partial x}\right)^n D_n(x,t,P) p(x,t\,|\,x',t',P') \,. \quad (76)$$

Note that by generalizing the Kramers–Moyal expansion to the nonlinear case we found immediately that the coefficients D_n depend on $P(x,t;u)$, whereas the conditional probability density p depends on $P(x',t';u)$. Note also that in the special case $D_n = 0$ for $n \geq 3$ the Kramers–Moyal expansion (76) yields the nonlinear Fokker–Planck equation (44). Note also that since we have $M_n(\Delta t = 0) = 0$ for all n Kramers–Moyal coefficients can also be defined by

$$D_n(x,t,P) = \left.\frac{\partial M_n}{\partial \Delta t}\right|_{\Delta t=0} = \int_\Omega dx \frac{(z-x)^n}{n!} \left.\frac{\partial}{\partial u} p(z,u|x,t,P)\right|_{u=t} \,. \quad (77)$$

Drift-Diffusion Estimates

The definition of the Kramers–Moyal coefficients can be exploited to extract the drift- and diffusion coefficients of nonlinear Fokker–Planck equations from time series data. Accordingly, the drift coefficient D_1 and the diffusion coefficient D_2 are defined by

$$D_1(x,t,P) = \lim_{\Delta t\to 0} \frac{1}{\Delta t} \int_\Omega dx (z-x) p(z,t+\Delta t|x,t,P) \,,$$
$$D_2(x,t,P) = \lim_{\Delta t\to 0} \frac{1}{2\Delta t} \int_\Omega dx \frac{(z-x)^2}{2} p(z,t+\Delta t|x,t,P) \,. \quad (78)$$

The limiting case Δt may be approximated by the smallest time step that is available in the data set:

$$D_1(x,t,P) \approx \frac{1}{\Delta t} \int_\Omega dx (z-x)\, p(z,t+\Delta t\,|\,x,t,P) \,,$$
$$D_2(x,t,P) \approx \frac{1}{2\Delta t} \int_\Omega dx \frac{(z-x)^2}{2} p(z,t+\Delta t|x,t,P) \,, \quad (79)$$

where Δt denotes now the sampling interval between two data points. Note that on the basis of the alternative definition (77) higher-order approximations can also be defined [171]. The conditional averages can be approximated by empirical conditional averages computed from a finite set of realizations $X^{(1)}(t), X^{(2)}(t), \ldots, X^{(N)}(t)$. Thus, we obtain

$$D_1(x,t,P) \approx \frac{1}{\Delta t} \frac{1}{\sum_{i\in I(t,x)} 1} \cdot \sum_{i\in I(t,x)} [X^{(i)}(t+\Delta t) - X^{(i)}(t)] \,,$$
$$D_2(x,t,P) \approx \frac{1}{2\Delta t} \frac{1}{\sum_{i\in I(t,x)} 1} \cdot \sum_{i\in I(t,x)} [X^{(i)}(t+\Delta t) - X^{(i)}(t)]^2 \,, \quad (80)$$

where $I(t,x)$ is the set of indices i for which $X^{(i)}(t) \approx x$. In the case of Markov processes described by linear Fokker–Planck equations the argument P in the coefficients can be dropped and the above drift-diffusion estimates reduce to the estimates proposed in [92,93] that have recently found many applications [18,122,212,231]. For Markov processes described by nonlinear Fokker–Planck equations we need to compute the conditional averages for different probability densities P. To this end, we may vary the initial distribution u of a stochastic process. For a stochastic process with a particular distribution of X at time t we will obtain the coefficients D_1 and D_2 only for that particular distribution. Using the kernel estimate method mentioned above, we obtain

$$D_1\left(x,t,P \approx \frac{1}{Ns\sqrt{2\pi}}\sum_{i=1}^{N}\exp\left\{-\frac{(x-X^{(i)}(t))^2}{2s^2}\right\}\right)$$
$$\approx \frac{1}{\Delta t}\frac{1}{\sum_{i\in I(t,x)}1}\sum_{i\in I(t,x)}[X^{(i)}(t+\Delta t)-X^{(i)}(t)] \quad (81)$$

and

$$D_2\left(x,t,P \approx \frac{1}{Ns\sqrt{2\pi}}\sum_{i=1}^{N}\exp\left\{-\frac{(x-X^{(i)}(t))^2}{2s^2}\right\}\right)$$
$$\approx \frac{1}{2\Delta t}\frac{1}{\sum_{i\in I(t,x)}1}\sum_{i\in I(t,x)}[X^{(i)}(t+\Delta t)-X^{(i)}(t)]^2 \quad (82)$$

with $s = N^{-1/5}\sigma_e(t_n)$, where $\sigma_e(t_n)$ is the standard deviation of the empirical ensemble $\{X^{(1)}(t),\ldots,X^{(N)}(t)\}$. Note that in general the Kramers–Moyal coefficients of Markov processes induced by conditional probability densities of the form $p(x,t|x',t',P')$ can be estimated using

$$D_n\left(x,t,P \approx \frac{1}{Ns\sqrt{2\pi}}\sum_{i=1}^{N}\exp\left\{-\frac{(x-X^{(i)}(t))^2}{2s^2}\right\}\right)$$
$$\approx \frac{1}{n!\Delta t}\frac{1}{\sum_{i\in I(t,x)}1}\sum_{i\in I(t,x)}[X^{(i)}(t+\Delta t)-X^{(i)}(t)]^n\,. \quad (83)$$

Alternatively, parametric estimate methods may be used. For example, we may be interested in estimating the exponent q of a Markov process defined by the Plastino–Plastino model (see Sect. "Nonextensive Systems" below)

$$\frac{\partial}{\partial t}p(x,t|x',t',P')$$
$$=\left[\frac{\partial}{\partial x}\gamma x + Q\frac{\partial^2}{\partial x^2}P(x,t;u)^{q-1}\right]p(x,t|x',t',P') \quad (84)$$

with $\gamma, Q, q > 0$. Then, the diffusion coefficient $D_2(P) = QP^{q-1}$ involves the parameter Q and q. Using Eq. (82) and taking the logarithm, we get

$$\ln Q + (q-1)\ln\left\{\frac{1}{Ns\sqrt{2\pi}}\sum_{i=1}^{N}\exp\left\{-\frac{(x-X^{(i)}(t))^2}{2s^2}\right\}\right\}$$
$$\approx \ln\left\{\frac{1}{2\Delta t}\frac{1}{\sum_{i\in I(t,x)}1}\sum_{i\in I(t,x)}[X^{(i)}(t+\Delta t)-X^{(i)}(t)]^2\right\}. \quad (85)$$

For example, at a particular time t Eq. (85) can be evaluated for different states x_i. In that case Eq. (85) assumes the form $\ln Q + (q-1)A_1(x_i) = A_2(x_i)$. Then, the expressions $\ln Q$ and $q-1$ (and in doing so the parameters Q and q) can be estimated from a linear regression [90].

Martingales

Let $Z(t)$ denote a functional of a stochastic process $\hat{X}$ defined for $t \geq t_0$. In this chapter we will put $t_0 = 0$. Then, Z is a martingale of $\hat{X}$ if

$$\langle Z(t)\rangle_{X=\theta} = Z(t') \quad (86)$$

holds for $t \geq t'$, where θ is a realization of the random variable X on the interval $[0,t']$ (see Sect. 1.3 in [131]). That is, the constraint $X = \theta$ means $X(s) = \theta(s)$ holds for $s \in [0,t']$. Roughly speaking, a martingale is a random variable for which the best predictor of its future mean value is the present value. With regard to Eq. (86) the prediction of the future mean value is $\langle Z(t)\rangle$, whereas the present value of Z is $Z(t')$. Alternatively, we may say that the information at one time t' about the value of the martingale Z is sufficient to predict the mean value of the martingale Z for future times $t \geq t'$. Note that this alternative point of view is closely related with the first definition of Markov processes discussed in the previous section.

For linear Fokker–Planck equations there is a close link between martingales and Markov property. Accordingly, a stochastic process is a Markov process defined by a linear Fokker–Planck equation with drift and diffusion coefficients D_1 and D_2 if and only if a particular random variable Z that involves the Fokker–Planck operator is a martingale (see Sect. 15.1 in [132]). In the mathematical literature this link has also been studied in the context of nonlinear Fokker–Planck equations [54,67,94,102,103,124,154,155,168].

Our aim in this section is to make the martingale approach more accessible to scientists working in physics, applied mathematicians, and related disciplines. To this end, we will in what follows illustrate this link between martingales and Markov processes defined by strongly nonlinear Fokker–Planck equation by means of standard techniques frequently used in physics.

Theorem 1 *Let $\hat{X}$ be a stochastic process with initial probability density $u(x)$ and conditional probability density*

$$p(x,t|x',t';P') = \langle \delta(x - X(t)) \rangle_{X(t')=x';\langle \delta(x'-X(t'))\rangle = P(x',t';u)} \,. \quad (87)$$

Then, $\hat{X}$ is a Markov process defined by the nonlinear Fokker–Planck equation

$$\frac{\partial}{\partial t} p(x,t|x',t';P') = L(x,t,P)p(x,t|x',t';P') \quad (88)$$

with

$$L(x,t,P) = -\frac{\partial}{\partial x} D_1(x,t,P) + \frac{\partial^2}{\partial x^2} D_2(x,t,P) \quad (89)$$

if and only if $Z(t)$ defined by

$$Z(t) = f(X(t)) - \int_0^t L_B f[X(z),z,P]\, dz \quad (90)$$

with

$$L_B(x,t,P) = D_1(x,t,P)\frac{\partial}{\partial x} + D_2(x,t,P)\frac{\partial^2}{\partial x^2} \quad (91)$$

is a martingale of X for smooth functions f. In the context of linear Fokker–Planck equations the operator L_B is the Fokker–Planck backwards operator [98,184]. Note that in our context we refer to f as a smooth function if it has continuous second-order derivatives. Note also that above and in what follows we will frequently use the notation (54). Note finally that in the above theorem the notion $L_B f[X(z),z,P]$ should be interpreted like

$$L_B f[X(z),z,P] = L_B(X(z),z,P) f(X(z)) = \{L_B(x,t,P)f(x)\}_{x=X(z),t=z} \,. \quad (92)$$

That is, first we carry out the differentiations defined by the operator L_B. Subsequently, we replace in the result the state variable x by the value of the random variable X at time z. Moreover we replace t by z. Let us prove the theorem in two parts.

From Strongly Nonlinear Fokker–Planck Equations to Martingales

Let us prove in this section that a Markov process defined by a strongly nonlinear Fokker–Planck equation exhibits the martingale Z. To this end we first compute the conditional mean of the random variable Z defined in Eq. (90). Thus, we obtain

$$\begin{aligned}\langle Z(t)\rangle_{X=\theta} &= \langle f(X(t))\rangle_{X=\theta} - \int_0^t dz\, \langle L_B f[X(z),z,P]\rangle_{X=\theta} \\ &= \langle f(X(t))\rangle_{X=\theta} - \int_{t'}^t dz\, \langle L_B f[X(z),z,P]\rangle_{X=\theta} \\ &\quad - \int_0^{t'} ds L_B f[\theta(s),s] \,. \end{aligned} \quad (93)$$

The Markov property implies that the constraints can be relaxed. That is, for every functional $g(t)$ of $X(t)$ with $t \geq t'$ we have $\langle g(t)\rangle_{X=\theta} = \langle g(t)\rangle_{X(t')=x';P'}$, where x' is given by $x' = \theta(t')$. We have indicated here that the average may depend on how the process is distributed at time t'. Consequently, Eq. (93) becomes

$$\begin{aligned}\langle Z(t)\rangle_{X=\theta} &= \langle f(X(t))\rangle_{X(t')=x';\langle \delta(x'-X(t'))\rangle = P'} \\ &\quad - \int_{t'}^t dz\, \langle L_B f[X(z),z,P]\rangle_{X(t')=x',\langle \delta(x'-X(t'))\rangle = P'} \\ &\quad - \int_0^{t'} ds L_B f[\theta(s),s,P] \,. \end{aligned} \quad (94)$$

Multiplying the Fokker–Planck equation (88) with $f(x)$ and integrating with respect to x, we obtain

$$\frac{\partial}{\partial t}\int_\Omega f(x)p(x,t|x',t';P')\,dx = \int_\Omega f(x)L(x,t,P)p(x,t|x',t';P')\,dx \,. \quad (95)$$

By means of partial integration we find that $\int_\Omega f(x)\, L(x,t,P)\, p(x,t\,|\,x',t';P')\,dx = \int_\Omega p(x,t\,|\,x',t';P')\, L_B(x,t,P)\, f(x)\,dx$. As a result, Eq. (95) can be transformed into

$$\begin{aligned}\frac{\partial}{\partial t}\int_\Omega f(x)p(x,t|x',t',P')\,dx &= \int_\Omega p(x,t|x',t',P')L_B(x,t,P)f(x)\,dx \\ &= \langle L_B(x,t,P)f(x)\rangle_{X(t')=x';P'} \,. \end{aligned} \quad (96)$$

Using Eq. (96), we obtain

$$\langle L_B f[X(z),z]\rangle_{X(t')=x';\langle \delta(x-X(t'))\rangle = P'} = \frac{\partial}{\partial z}\int_\Omega f(x)p(x,z|x',t';P')\,dx \,. \quad (97)$$

Consequently, the following integral transformation holds:

$$\begin{aligned} I &= \int_{t'}^t dz\, \langle L_B f[X(z),z,P]\rangle_{X(t')=x';P'} \\ &= \int_{t'}^t dz \frac{\partial}{\partial z}\int_\Omega dx\, f(x)p(x,z|x',t',P') \\ &= \langle f(X(t)\rangle_{X(t')=x';P'} - f(x') \,. \end{aligned} \quad (98)$$

Substituting Eq. (98) into Eq. (94), we get

$$\langle Z(t)\rangle_{X=\theta} = f(x') - \int_0^{t'} \mathrm{d}s L_B f[\theta(s), s, P] \,. \tag{99}$$

By definition the function $Z(t')$ for $X(s) = \theta(s)$ given in $s \in [0, t']$ reads

$$Z(t') = f(x') - \int_0^{t'} \mathrm{d}s\, L_B f[\theta(s), s, P] \,. \tag{100}$$

Consequently, we have our final result

$$\langle Z(t)\rangle_{X=\theta} = Z(t') \tag{101}$$

and the proof is completed.

From Martingales to Strongly Nonlinear Fokker–Planck Equations

Let us prove next that the martingale (90) defines a Markov process of a strongly nonlinear Fokker–Planck equation. Evaluating Eq. (90) by analogy to Eq. (93) gives us

$$\langle Z(t)\rangle_{X=\theta} = \langle f(X(t))\rangle_{X=\theta} - \int_{t'}^{t} \mathrm{d}z\, \langle L_B f[X(z), z, P]\rangle_{X=\theta} - \int_0^{t'} \mathrm{d}s\, L_B f[\theta(s), s, P] \,. \tag{102}$$

Substituting this result into Eq. (86) and substituting Eq. (100) into Eq. (86), we see that Eq. (86) becomes

$$\langle f(X(t))\rangle_{X=\theta} = f(x') + \int_{t'}^{t} \mathrm{d}z\, \langle L_B f[X(z), z]\rangle_{X=\theta} \,. \tag{103}$$

Equation (103) can equivalently be written as

$$\int_\Omega f(x) p(x, t|X = \theta)\, \mathrm{d}x = f(x') + \int_{t'}^{t} \mathrm{d}z \int_\Omega \mathrm{d}x\, p(x, z|X = \theta) L_B f[x, z, P] \tag{104}$$

with $p(x, z|X = \theta) = \langle \delta(x - X(z))\rangle_{X=\theta}$. Using partial integration, we can show that the operator L_B and the differential operator L are related to each other like

$$\int_\Omega \mathrm{d}x\, p(x, z|X = \theta) L_B f[x, z, P] = \int_\Omega \mathrm{d}x\, f(x) L(x, z, P) p(x, z|X = \theta) \,. \tag{105}$$

Substituting this result into Eq. (104) yields

$$0 = \int_\Omega \mathrm{d}x\, f(x) \Big\{ p(x, t|X = \theta) - \delta(x - x') - \int_{t'}^{t} \mathrm{d}z\, L(x, z, P) p(x, z|X = \theta) \Big\} \,. \tag{106}$$

This holds for arbitrary smooth functions f. Since f is arbitrary, the expression in the brackets $\{\cdot\}$ of Eq. (106) must vanish and we obtain

$$p(x, t|X = \theta) = \delta(x - x') + \int_{t'}^{t} \mathrm{d}z\, L(x, z, P) p(x, z|X = \theta) \,. \tag{107}$$

Differentiating Eq. (107) with respect to t gives us

$$\frac{\partial}{\partial t} p(x, t|X = \theta) = L(x, t, P) p(x, t|X = \theta) \,. \tag{108}$$

Multiplying with the probability density $P(X = \theta)$ and performing a functional integration with respect to the path θ, we obtain

$$\frac{\partial}{\partial t} P(x, t) = L(x, t, P) P(x, t) \,. \tag{109}$$

The formal solutions of Eqs. (108) and (109) read

$$p(x, t|X = \theta) = \exp\left\{ \int_{t'}^{t} \mathrm{d}z\, L(x, z, P) \right\} \delta(x - x') \tag{110}$$

and

$$P(x, t) = \exp\left\{ \int_{t'}^{t} \mathrm{d}z\, L(x, z, P) \right\} P(x, t') \,. \tag{111}$$

We see that a solution of Eq. (108) under the initial condition $p(x, t|X = \theta) = \delta(x - x')$ for $t \to t'$ with $x' = \theta(t')$ only depends on $\theta(t')$ but does not depend on $\theta(s)$ for $s < t'$. Consequently, X is a Markov process. However, L depends on P. From Eq. (110) it is clear that the conditional probability density p depends on the time-dependent probability density P for $z \in [t', t]$. Since $P(x, t; u)$ for $t \geq t'$ can be computed from $P(x, t'; u)$ as shown in Eq. (111) we conclude that p depends only on $P(x, t'; u)$ and does not depend on the evolution of P on the whole interval $[t', t]$. Therefore, we have $p(x, t|X = \theta) = p(x, t|x', t', P')$. Substituting this result into Eq. (108), we see that Eq. (108) becomes a strongly nonlinear Fokker–Planck equation

$$\frac{\partial}{\partial t} p(x, t|x', t'; P') = L(x, t, P) p(x, t|x', t'; P') \,. \tag{112}$$

Examples

Shimizu–Yamada Model

The Shimizu–Yamada model [201,202] corresponds to the Desai–Zwanzig model (7) for a linear single-particle force $h(x) = -\gamma x$. The evolution of the conditional probability

density p is defined by

$$\frac{\partial}{\partial t}p(x,t\mid x',t',P') = \left[\frac{\partial}{\partial x}\gamma x+\kappa\left(x-\int_\Omega xP(x,t;u)\,\mathrm{d}x\right)+Q\frac{\partial^2}{\partial x^2}\right]\cdot p(x,t\mid x',t',P') \quad (113)$$

with $\Omega=\mathbb{R}$ and $\gamma,\kappa,Q>0$. Multiplying Eq. (113) with $P(x',t';u)$ and integrating with respect to x' yields the evolution equation for $P(x,t;u)$:

$$\frac{\partial}{\partial t}P(x,t;u) = \left[\frac{\partial}{\partial x}\gamma x+\kappa\left(x-\int_\Omega xP(x,t;u)\,\mathrm{d}x\right)+Q\frac{\partial^2}{\partial x^2}\right]\cdot P(x,t;u)\,. \quad (114)$$

See also [77] and Sect. 3.10 in [80]. From Eq. (114) it follows that the mean value $m(t)=\int_\Omega xP(x,t;u)\,\mathrm{d}x$ decays exponentially like

$$m(t)=m(t_0)\exp\{-\gamma(t-t_0)\} \quad (115)$$

with $m(t_0)=\int_\Omega xu(x)\,\mathrm{d}x$. Substituting Eq. (115) into Eqs. (113) and (114), we realize that a solution $P(x,t;u)$ and a Green's function p exist for any initial probability density $u(x)$. Therefore, the Shimizu–Yamada model is a strongly nonlinear Fokker–Planck equation and describes a Markov process.

It can be shown that the conditional probability density $p(x,t\mid x',t';u)$ reads (see [77] and Sect. 3.10 in [80])

$$p(x,t\mid x',t';u) = \frac{\exp\left\{-\dfrac{[x-g(t,t',t_0,u)-x'm(t,t')]^2}{2K(t,t')}\right\}}{\sqrt{2\pi K(t,t')}} \quad (116)$$

with

$$m(t,t')=\exp\left\{-(\gamma+\kappa)(t-t')\right\}\,, \quad (117)$$

$$K(t,t')=\frac{Q}{\gamma+\kappa}\left[1+\exp\left\{-2(\gamma+\kappa)(t-t')\right\}\right]\,, \quad (118)$$

and

$$g=\left[\exp\{-\gamma(t-t_0)\}-\exp\{-(\gamma+\kappa)t+\gamma t_0+\kappa t'\}\right]\cdot\int_\Omega xu(x)\,\mathrm{d}x\,. \quad (119)$$

The mean value $m(t)$ acts as a self-organized driving force or time-dependent order parameter of the stochastic process. Since there is a one-to-one mapping of $m(t)$ to $m(t')$ with $t'<t$, we can eliminate the parameter u in $p(x,t|x',t';u)$ as argued in Sect. "Strongly Nonlinear Fokker–Planck Equations". Substituting Eq. (115) into Eq. (119), we obtain

$$g(t,t',P(x,t';u))=\exp\{-\gamma(t-t')\}\cdot\left[1-\exp\{-\kappa)(t-t')\}\right]\int_\Omega xP(x,t';u)\,\mathrm{d}x \quad (120)$$

or

$$g(t,t',\langle X(t')\rangle)\exp\{-\gamma(t-t')\}\cdot\left[1-\exp\{-\kappa)(t-t')\}\right]\langle X(t')\rangle\,. \quad (121)$$

Consequently, the conditional probability density $p(x,t\mid x',t',P')$ reads

$$p(x,t\mid x',t';P') = \frac{\exp\left\{-\dfrac{[x-g(t,t',\langle X(t')\rangle)-x'm(t,t')]^2}{2K(t,t')}\right\}}{\sqrt{2\pi K(t,t')}}\,. \quad (122)$$

Dynamic Takatsuji Model

The dynamic Takatsuji model for the conditional probability density p is defined by

$$\frac{\partial}{\partial t}p(x,t\mid x',t';P') = \left[\frac{\partial}{\partial x}(\gamma+c)x-\sqrt{c}\tanh\left(\sqrt{c}\int_\Omega xP(x,t;u)\,\mathrm{d}x\right)+Q\frac{\partial^2}{\partial x^2}\right]\cdot p(x,t\mid x',t';P') \quad (123)$$

with $x\in\Omega=\mathbb{R}$ and $c,Q>0,\gamma\in\mathbb{R}$. Likewise, the probability density $P(x,t;u)$ satisfies

$$\frac{\partial}{\partial t}P(x,t;u) = \left[\frac{\partial}{\partial x}(\gamma+c)x-\sqrt{c}\tanh\left(\sqrt{c}\int_\Omega xP(x,t;u)\,\mathrm{d}x\right)+Q\frac{\partial^2}{\partial x^2}\right]\cdot P(x,t;u)\,. \quad (124)$$

For details see [78,214]. From Eq. (124) it follows that the first moment $M_1(t)=\langle X\rangle$ can be computed from

$$\frac{\mathrm{d}}{\mathrm{d}t}M_1(t)=-(\gamma+c)M_1+\sqrt{c}\tanh[\sqrt{c}M_1(t)]\,. \quad (125)$$

For arbitrary initial distributions u solutions of $M_1(t)$ exist and are smooth functions of t. Substituting these solutions into Eqs. (123) and (124), we see that solutions of Eqs. (123) and (124) in terms of Green's functions p and probability densities P exist as well. Consequently, the dynamic Takatsuji model belongs to the class of strongly nonlinear Fokker–Planck equations and describes a Markov process.

Since $p(x,t|x',t',P')$ depends on P', the expected mean value of $X(t)$ of realizations that assume the value x' at time t' depends on the distribution of the ensemble at time t'. Let us illustrate this issue. The conditional mean value under consideration reads

$$\langle X(t)\rangle_{X(t')=x',P'} = \int x p(x,t \mid x',t',P')\,\mathrm{d}x\,. \quad (126)$$

Multiplying Eq. (123) with x and integrating with respect to x, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}\langle X(t)\rangle_{X(t')=x',P'} = -(\gamma+c)\langle X(t)\rangle_{X(t')=x',P'} + \sqrt{c}\tanh[\sqrt{c}M_1(t)]\,. \quad (127)$$

The solution reads

$$\langle X(t)\rangle_{X(t')=x',P'} = x'\exp\{-(\gamma+c)(t-t')\} + \sqrt{c}\int_{t'}^{t}\tanh[\sqrt{c}M_1(z)]\,\mathrm{d}z\,, \quad (128)$$

where $M_1(z)$ is the solution of Eq. (125) for initial value $M_1(t') = \int_\Omega xP(x,t';u)\,\mathrm{d}x$. Let I denote the integral $I = \sqrt{c}\int_{t'}^{t}\tanh[\sqrt{c}M_1(z)]\,\mathrm{d}z$. Then, I depends on $M_1(t')$, c, γ, t' and t: $I = I(t,t',M_1(t'),c,\gamma)$. Consequently, Eq. (128) can be cast into the form

$$\langle X(t)\rangle_{X(t')=x',P'} = x'\exp\{-(\gamma+c)(t-t')\} + I\left(t,t',\int x'P(x',t';u)\,\mathrm{d}x',c,\gamma\right) \quad (129)$$

or

$$\langle X(t)\rangle_{X(t')=x',P'} = x'\exp\{-(\gamma+c)(t-t')\} + I(t,t',\langle X(t')\rangle,c,\gamma)\,. \quad (130)$$

Equation (130) illustrates that in order to predict future conditional mean values of a Takatsuji process $\hat{X}$ at times t it is sufficient to have at one time $t' \le t$ information about the state value x' of a realization of $\hat{X}$ and the mean value $\langle X(t')\rangle$ of all realizations of $\hat{X}$.

Note that the trajectories $Z(t)$ of martingale processes $\hat{Z}$ induced by the Takatsuji process $\hat{X}$ are given by

$$Z(t) = f(X) - \int_0^t \mathrm{d}s\left\{\left[-(\gamma+c)X(s) + \sqrt{c}\tanh(\sqrt{c}\langle X(s)\rangle\right]\frac{\partial f}{\partial X(s)} + Q\frac{\partial^2 f}{\partial X^2(s)}\right\} \quad (131)$$

for arbitrary smooth functions f. We can exploit these martingale processes in order to compute conditional expectations. For example, for $f(y) = y$ from the martingale property (86) it follows that

$$\langle X(t)\rangle_{X(t')=x',P'} = x' - \int_{t'}^{t}\mathrm{d}s\,(\gamma+c)\,\langle X(s)\rangle - \sqrt{c}\tanh[\sqrt{c}\,\langle X(s)\rangle]\,. \quad (132)$$

Differentiating this relation with respect to t we obtain Eq. (127) again and so we can compute the conditional expectation (130).

Liquid Crystal Model

Liquid crystals exhibit nematic-isotrope phase transitions [30,48,50]. At high temperatures the liquid crystal macromolecules exhibit an orientational disorder. The liquid crystal is said to be in the isotropic phase. Below a critical temperature the macromolecules show some degree of orientational order. The degree of orientational order is often measured by the Maier–Saupe order parameter S [150].

A nonlinear Fokker–Planck equation that describes the stochastic behavior of the liquid crystal in the isotropic and nematic phases and can to a certain extent also describe the phase transition was proposed by Doi and Edwards [55] and Hess [115] and is shown above in Eq. (8). Equation (8) describes the random walk of a orientation of liquid crystal molecules, where the orientation is given by a vector $\mathbf{x}$ that points to the surface of a unit sphere. For liquid crystals with an axial symmetry the liquid crystal model can be simplified. The simplified model describes the random walk of the molecule alignment with the symmetry axis. The random variable is defined on $X \in \Omega = [0,1]$. For sake of simplicity we will extend the range of definition to the interval $\Omega = [-1,1]$ and require that distributions are symmetric. For $X = 0$ the molecule has an orientation perpendicular to the symmetry axis. If $X = 1$ or $X = -1$ the molecule points exactly in the direction of the symmetry axis. In this symmetric case, the probability density P of X satisfies [64]

$$\frac{\partial}{\partial t}P(x,t;u) = \frac{\partial}{\partial x}(1-x^2)\cdot\left[-\frac{9}{2}\kappa x\left(\int x^2P(x,t;u)\,\mathrm{d}x - \frac{1}{3}\right) + D_r\frac{\partial}{\partial x}\right]P(x,t;u) \quad (133)$$

with $\kappa, D_r > 0$. Equation (133) as well as the original equation (8) are regarded as description for an ensemble of macromolecules that perform rotational Brownian motion [55]. Since Brownian motion is a Markov process it is reasonable to construct on the basis of Eq. (133) a model for a many-body system that exhibits a Markov process. In line with our discussion in Sect. "Markov Property, 2nd

Order and Higher Order Statistics" we assume that the conditional probability density p satisfies [81]

$$\frac{\partial}{\partial t}p(x,t|x',t',P') = \frac{\partial}{\partial x}(1-x^2) \cdot\left[-\frac{9}{2}\kappa x\left(\int x^2 P(x,t;u)\,\mathrm{d}x - \frac{1}{3}\right) + D_r\frac{\partial}{\partial x}\right] \cdot p(x,t|x',t',P')\,. \quad (134)$$

Note that the expression in the bracket $(\cdot)$ is related to the Maier–Saupe order parameter which reads in the symmetric case

$$S(t) = \frac{1}{2}\left(3\int x^2 P(x,t;u)\,\mathrm{d}x - 1\right)\,. \quad (135)$$

Due to the boundary conditions $X \in [-1,1]$ the order parameter S and consequently the bracket $(\cdot)$ is bounded. This implies that solutions P and p of Eqs. (133) and (134) exist and that the liquid crystal model (133)–(134) describes a Markov process. The self-consistent Ito–Langevin equation of this Markov process reads [81]

$$\frac{\mathrm{d}}{\mathrm{d}t}X(t) = \frac{9\kappa}{2}\left(1-X(t)^2\right)X(t)\left(\left\langle X(t)^2\right\rangle - \frac{1}{3}\right) - 2D_r X(t) + \sqrt{D_r(1-X(t)^2)}\,\Gamma(t)\,. \quad (136)$$

Trajectories $Z(t)$ of martingale processes $\hat{Z}$ of the liquid crystal model are defined by [83]

$$Z(t) = f(X(t)) - \int_0^t \mathrm{d}s\Big[\,\cdot\,\Big]f(X(s)) \quad (137)$$

with

$$\Big[\,\cdot\,\Big] = \left[\left(\frac{9\kappa}{2}(1-X(s)^2)X(s)\left(\left\langle[X(s)]^2\right\rangle - \frac{1}{3}\right) - 2D_r X(s)\right)\frac{\partial}{\partial X(s)} + D_r\frac{\partial^2}{\partial X^2(s)}\right]. \quad (138)$$

In the stationary case the short-time autocorrelation function $C(\Delta t) = \langle X(t)X(t+\Delta t)\rangle_{\mathrm{st}}$ reads [81]

$$C(\Delta t) = \frac{2S+1}{3} - \frac{2D_r(1-S)}{3}\Delta t + O(\Delta t^2)\,, \quad (139)$$

where S denotes the order parameter (see above) in the stationary case. That is, we have $S = (3\langle X^2\rangle_{\mathrm{st}} - 1)/2$. Consequently, C depends on S. This has important implications for the hysteresis loop of the nematic-isotropic phase transition. Let us assume that if we decrease the temperature of a liquid crystal we find the transition from the isotrope to the nematic phase with $S = 0 \to S > 0$ at the critical temperature $T_{c,\mathrm{low}}$. In contrast, if we increase the temperature of a liquid crystal we find the transition from the nematic to the isotropic phase with $S > 0 \to S = 0$ at the the slightly higher critical temperature $T_{c,\mathrm{high}}$. Then, in the temperature interval $[T_{c,\mathrm{low}}, T_{c,\mathrm{high}}]$ the liquid crystal exhibit two autocorrelation functions

$$C_{\mathrm{isotrope}}(\Delta t) = \frac{1}{3} - \frac{2D_r}{3}\Delta t\,, \quad (140)$$

$$C_{\mathrm{nem}}(\Delta t) = \frac{2S(T)+1}{3} - \frac{2D_r\,(1-S(T))}{3}\Delta t \quad (141)$$

which hold up to terms of order Δt^2. Equations (140)–(141) illustrate that we are dealing with a system that exhibits two kinds of Markov processes that we may label or name 'isotropic' and 'nematic'. The modeling approach by means of strongly nonlinear Fokker–Planck equations indicates that these Markov processes are just different members of a family of Markov processes that naturally emerge in the self-organized liquid crystal. That is, the two Markov processes are not related to two different systems but they are just two different faces of the very same self-organizing many-body system.

Let us compute the conditional mean value of molecules that are perpendicular to the symmetry axis. To this end, we consider the random walk of the orientation angle ϕ defined by $X(t) = \sin\phi(t)$. Using the Stratonovich–Langevin equation of Eq. (134) (see [81]), we obtain a self-consistent Langevin equation for ϕ:

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi = \frac{9}{4}\kappa\sin(2\phi(t))\left(\left\langle\sin^2(\phi(t))\right\rangle - \frac{1}{3}\right) - D_r\tan\phi(t) + D_r\Gamma(t)\,. \quad (142)$$

For short time intervals $\Delta t = t - t'$ and appropriate small noise amplitudes D_r we assume that $\phi(t) \approx 0$ if $\phi(t') \approx 0$. Linearizing Eq. (142) at $\phi = 0$ yields

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi(t) = \left\{\frac{9}{4}\kappa\left(\left\langle\sin^2(\phi(t))\right\rangle - \frac{1}{3}\right) - D_r\right\}\phi(t) + D_r\Gamma(t)\,. \quad (143)$$

The conditional expectation value $\int \phi\, p(\phi,t|\phi',t',P)\mathrm{d}\phi$ for short time intervals Δt can then be computed from Eq. (143) simply by averaging both side of Eq. (143) under the constraint $\phi(t') = \phi'$ and $\phi(t')$ distributed like P'. Thus, we obtain

$$\langle\phi(t)\rangle_{\phi(t')=\phi',P'} = \phi'\left[1 + \Delta t\left\{\frac{9}{4}\kappa\left(\left\langle\sin^2(\phi(t'))\right\rangle - \frac{1}{3}\right) - D_r\right\}\right] \quad (144)$$

or

$$\langle \phi(t)\rangle_{\phi(t')=\phi',P'} = \phi' \left[1 + \Delta t \left\{\frac{3}{2}\kappa\, S(t') - D_r\right\}\right] . \tag{145}$$

These estimates hold for small intervals Δt, sufficiently small noise amplitudes D_r, and orientation angles $\phi' \approx 0$. Again, in line with our general discussion in the preceding sections we see that the conditional expectation $\langle \phi(t)\rangle_{\phi(t')=\phi',P'}$ can be computed provided that for $t' < t$ the distribution of $\phi(t')$ or at least the order parameter $S(t')$ is known and the angle ϕ' is selected.

Semi-classical Description of Quantum Systems

A stochastic treatment of semi-classical quantum systems by means of nonlinear Fokker–Planck equations that can be cast into the form of Eqs. (5) and (6) has been proposed and analyzed in several studies [28,31,85,125,127,129,130]. Accordingly, a Fermi or Bose particle with mass 1 that moves in a one-dimensional space with velocity v exhibits in the stationary case a Fermi–Dirac or Bose–Einstein distribution of the kinetical energy $E_{\text{kin}} = v^2/2$. The free diffusion of the particle can be described by the nonlinear Fokker–Planck equations [85]

$$\frac{\partial}{\partial t}P(v,t;u) = \frac{\partial}{\partial v}\gamma v\left[1 \mp P(v,t,u)\right]P(v,t;u) + Q\frac{\partial^2}{\partial v^2}P(v,t;u) , \tag{146}$$

where the upper sign holds for Fermi particles, the lower for Bose particles. The parameters γ and Q represent damping and fluctuation strength and are related to the temperature T by the fluctuation dissipation theorem $Q/\gamma = 1/(k_B T)$, where k_B is the Boltzmann constant. The stationary probability density $P_{\text{st}}(v)$ of Eq. (146) reads

$$P_{\text{st}}(v) = \frac{1}{\exp\{(E_{\text{kin}} - \mu)/(k_B T)\} \pm 1} , \tag{147}$$

where μ is a normalization constant that can be interpreted as chemical potential. The transient solution $P(v, t; u)$ can be obtained by solving the integral equation [83, 155]

$$P(v,t;u) = \int_\Omega \mathrm{d}v_0\, G_B(t,t_0,v,v_0)u(v_0,t_0) + \int_{t_0}^{t}\mathrm{d}s \cdot \int_\Omega \mathrm{d}v'\, G_B(t,s,v,v')\gamma\frac{\partial}{\partial v'}\left[1 \mp P(v',s;u)\right]P(v',s;u) \tag{148}$$

with

$$G_B(t,t',v,v') = \frac{1}{\sqrt{2\pi Q(t-t')}}\exp\left\{-\frac{(v-v')^2}{2Q(t-t')}\right\} , \tag{149}$$

where G_B is the Gaussian propagator of Brownian motion. In the limit $t \to \infty$ the transient solution $P(v, t; u)$ approaches $P_{\text{st}}(v)$ [88,127]. Equation (148) is a useful description for numerical approaches. Using partial integration Eq. (148) can be written in the form

$$P(v,t;u) = \int_\Omega \mathrm{d}v_0\, G_B(t,t_0,v,v_0)u(v_0,t_0) + \int_{t_0}^{t}\mathrm{d}s \int_\Omega \mathrm{d}v'\, G_B(t,s,v,v')\gamma\frac{v-v'}{Q(t-s)} \cdot \left[1 \mp P(v',s;u)\right]\, P(v',s;u) . \tag{150}$$

This integral relation can be solved iteratively. In contrast to the iterative procedure discussed in Sect. "Time–Dependent Solutions and First Order Statistics", there is no need to compute derivatives. That is, we are dealing with some kind of path integral approach here that is similar to the numerical path integral approach involving short-time propagators, see Sect. "Short-Time Propagator".

In order to describe quantum particles that exhibits a Markov process we may exploit the approach outlined in Sect. "Markov Property, 2nd Order and Higher Order Statistics". Accordingly, the Markov conditional probability density of the quantum particle satisfies

$$\frac{\partial}{\partial t}p(v,t \mid v',t',P') = \left\{\frac{\partial}{\partial v}\gamma v\left[1 \mp P(v,t;u)\right] + Q\frac{\partial^2}{\partial v^2}\right\} \cdot p(v,t \mid v',t',P') \tag{151}$$

and the self-consistent Langevin equation reads

$$\frac{\mathrm{d}}{\mathrm{d}t}v(t) = -\gamma v(t)(1 \mp P(v(t),t;u)) + \sqrt{Q}\,\Gamma(t) . \tag{152}$$

From a martingale perspective we see that stochastic trajectories $v(t)$ induce for arbitrary smooth functions f the martingale $\hat{Z}$ with trajectories

$$Z(t) = f(v(t)) - \int_0^t \mathrm{d}s \cdot \left[-\gamma v(s)\left[1 \mp P(v(s),s;u)\right]\frac{\partial}{\partial v} + Q\frac{\partial^2}{\partial v^2}\right] f(v)\Big|_{v=v(s)} . \tag{153}$$

The Markov short-time propagator for $t = t' + \Delta t$ reads

$$p(v,t|v',t',P') = \sqrt{\frac{1}{2\pi Q\Delta t}}\exp\left\{-\frac{\left[v - v' + \Delta t\gamma v'[1 \mp P']\right]^2}{2Q\Delta t}\right\} , \tag{154}$$

and can be computed from the information about the distribution P' of $v(t')$ and the state v that was observed for particular realizations of the process $\hat{v}$. In the stationary case p reads for small time intervals $\Delta t = t - t'$:

$$p(v, t \mid v', t'; P_{\mathrm{st}}(v')) = \sqrt{\frac{1}{2\pi Q \Delta t}} \exp\left\{ -\frac{\left[v - v' + \Delta t \gamma v' (1 \mp P_{\mathrm{st}}(v'))\right]^2}{2Q\Delta t} \right\} \tag{155}$$

with P_{st} defined by Eq. (147).

Nonextensive Systems

Nonextensive thermostatistical systems have been related to the Tsallis entropy [1,219]

$$S_q = \frac{1}{q-1} \int_\Omega \left[P(v)^q - P(v)\right] \mathrm{d}v \,, \tag{156}$$

where q measures the degree of nonextensivity. Diffusion processes in nonextensive thermostatistical systems can be regarded as generalized Ornstein–Uhlenbeck processes that satisfy the nonlinear Fokker–Planck equation [177] (see also [22,31,32,35,58,85,86,200,222])

$$\frac{\partial}{\partial t} P(v, t; u) = \frac{\partial}{\partial v} \gamma v P(x, t; u) + Q \frac{\partial^2}{\partial v^2} P(v, t; u)^q \,, \tag{157}$$

where v is the velocity of a particle with mass 1 that moves in one spatial dimension. In the asymptotic domain $P(v, t; u)$ approaches a stationary Tsallis distribution

$$P_{\mathrm{st}}(v) = \frac{D_{\mathrm{st}}}{\left[1 + \gamma(1-q)v^2/[2qQ\, D_{\mathrm{st}}^{q-1}]\right]^{1/(1-q)}} \tag{158}$$

for $q \in (1/3, 1)$ with $D_{\mathrm{st}} = [\gamma/(2qQz_q^2)]^{1/(1+q)}$ and $z_q = \sqrt{\pi/(1-q)}\Gamma[(1+q)/[2(1-q)]]/\Gamma[1/(1-q)]$ [86]. The process described by Eq. (157) is said to evolve in a nonextensive thermodynamic framework because its stationary probability density (158) maximizing the entropy measure (156). As discussed in Sect. "Markov Property, 2nd Order and Higher Order Statistics" the Markov conditional probability density p satisfies

$$\frac{\partial}{\partial t} p(v, t \mid v', t', P') = \frac{\partial}{\partial x} \gamma v \, p(v, t \mid v', t', P') + Q \frac{\partial^2}{\partial v^2} P(v, t; u)^{q-1} \, p(v, t \mid v', t', P') \,. \tag{159}$$

The self-consistent Langevin equation of the Markov diffusion process reads (see also [22])

$$\frac{\mathrm{d}}{\mathrm{d}t} v(t) = -\gamma v(t) + \sqrt{QP(v(t), t; u)^{q-1}}\, \Gamma(t) \,. \tag{160}$$

Any stochastic path $v(t)$ computed from Eq. (160) yields for arbitrary smooth functions f the martingale

$$Z(t) = f(v(t)) - \int_0^t \mathrm{d}s \left[-\gamma v(s) \frac{\partial}{\partial v(s)} + QP(v(s), s; u)^{q-1} \frac{\partial^2}{\partial v(s)^2} \right] f(v(s)) \,. \tag{161}$$

The autocorrelation $C = \langle v(t) v(t') \rangle$ in the transient domain for $u(v_0) = \delta(v - v_0)$ reads [74]

$$C(t, t', v_0, t_0) = M_2(t', t_0, v_0) \exp\{-\gamma(t - t')\} \tag{162}$$

with

$$\begin{aligned} M_2(t', t_0, v_0) &= K(t', t_0) + M_1^2(t', t_0, v_0) \,, \\ K(t', t_0) &= \frac{1}{3q-1} \left[\frac{2qQ[z_q]^{(1-q)}}{\gamma} \cdot (1 - \exp\{-(1+q)\gamma(t' - t_0)\}) \right]^{2/(1+q)} , \\ M_1(t', t_0, v_0) &= v_0 \exp\{-\gamma(t' - t_0)\} \,. \end{aligned} \tag{163}$$

The autocorrelation function C depends on t_0. This is not in contradiction with the Markov property of the underlying process as discussed in Sect. "Markov Property, 2nd Order and Higher Order Statistics". In particular, we may eliminating the initial condition. Then, Eq. (162) simply reads

$$C(t, t') = \left\langle v^2(t') \right\rangle \exp\{-\gamma(t - t')\} \,. \tag{164}$$

and holds for arbitrary initial probability densities u.

Linear Nonequilibrium Thermodynamics

Linear and nonlinear Fokker–Planck equations alike can be approached from the principles of linear nonequilibrium thermodynamics [49,100,140]. For stochastic processes to which linear nonequilibrium thermodynamics applies the probability density $P(x, t; u)$ of a process evolves such that the free energy functional $F[P]$ decreases as a function of time t. More precisely, following a study by Compte and Jou [35] it has been proposed that P satisfies the nonlinear Fokker–Planck equations of the form [32,70,80,190]

$$\frac{\partial}{\partial t} P = \frac{\partial}{\partial x} P \tilde{M} \frac{\partial}{\partial x} \frac{\delta F}{\delta P} \,, \tag{165}$$

where $\tilde{M}$ is an appropriately defined mobility coefficient and $\delta F/\delta P$ denotes the variational derivative of F. Note that this thermodynamic approach is closely related to the GENERIC approach developed in [63,123,164,165,166].

For example, the Desai–Zwanzig model (8) can be expressed in terms of Eq. (165) for $\tilde{M} = 1$ and [80,196]

$$F = \langle V \rangle + U_{\mathrm{MF}} - QS_{\mathrm{BGS}} \,. \tag{166}$$

Here, V is the potential of the force h (i. e. we have $V(x) = -\int h(x)\,\mathrm{d}x$), U_{MF} is the mean field energy given by $U_{\mathrm{MF}} = -\kappa\sigma^2/2$ (where σ^2 is the variance of the process) and S_{BGS} is the Boltzmann–Gibbs-Shannon entropy

$$S_{\mathrm{BGS}} = -\int_{\Omega} P(x,t;u)\ln P(x,t;u)\,\mathrm{d}x \,, \tag{167}$$

where we have put the Boltzmann constant equal to unity. The liquid crystal model (8) can be written as Eq. (165) with [81]

$$F = -\frac{\kappa}{2}S^2 - D_r S_{\mathrm{BGS}} \,, \tag{168}$$

where S is the Maier–Saupe order parameter (135) and $\tilde{M} = 1 - x^2$ and $-\kappa S^2/2$ is the Maier–Saupe mean field energy. The Kuramoto–Shinomoto–Sakaguchi model (10) can equivalently be expressed in term of Eq. (165) and $\tilde{M} = 1$ and see Sect. 5.4 in [80]

$$F = \langle V \rangle - \frac{\kappa}{2}r^2 - QS_{\mathrm{BGS}} \,, \tag{169}$$

where r is the cluster phase defined by $r = |\langle \exp\{-iX(t)\}\rangle|$ and $-\kappa r^2/2$ is a measure for the mean field energy among the phase oscillators described by the model. The Takatsuji model (124) involves a constant mobility coefficient $\tilde{M} = 1$ and the free energy functional [80]

$$F = \frac{\gamma + c}{2}\left\langle X^2 \right\rangle - \ln\cosh\left(\sqrt{c}\,\langle X\rangle\right) - QS_{\mathrm{BGS}} \,. \tag{170}$$

The Plastino–Plastino model (14) related to the nonextensive Tsallis entropy (156) is given by Eq. (165) and $\tilde{M} = 1$ with [80]

$$F = \langle V \rangle - QS_q \,, \tag{171}$$

where V is the potential of the gradient force h. For an appropriate choice of $\tilde{M}$ the quantum mechanical nonlinear Fokker–Planck equations (146) can be cast into the form Eq. (165) with

$$F = \langle V \rangle - QS_{\mathrm{FD,BE}} \,, \tag{172}$$

where $S_{\mathrm{FD,BE}}$ is the quantum mechanical entropy of the Fermi–Dirac or Bose–Einstein statistics and V is the potential of the function $h(x)$ again. For details see [80,85].

From the perspective of linear nonequilibrium thermodynamics linear and nonlinear Fokker–Planck equations can be distinguished by means of the thermodynamic flux [35,70,80,83]

$$J = -\tilde{M}P\frac{\partial}{\partial x}\frac{\delta F}{\delta P} \,. \tag{173}$$

Note that in this approach the thermodynamic flux is equivalent to the probability current [80]. As can be seen from Eq. (172) the flux is associated to the free energy F on the one hand and the evolution equation (165) on the other hand by

$$\frac{\partial}{\partial t}P = -\frac{\partial}{\partial x}J \,. \tag{174}$$

If J is linear with respect to P, then the corresponding Fokker–Planck equation is linear with respect to P as well. If J is nonlinear with respect to P, then we are dealing with a nonlinear Fokker–Planck equation. The question whether J is linear or not is answered by nature herself [83]. For the Brownian particle motion we have $\tilde{M} = \gamma$ and

$$F = \frac{1}{2}\left\langle v^2 \right\rangle - QS_{\mathrm{BGS}} \,, \tag{175}$$

which yields

$$J = -\gamma v P - Q\frac{\partial}{\partial x}P \,. \tag{176}$$

J is linear and the corresponding Fokker–Planck equation is linear as well. For self-organizing systems J is nonlinear because F involves a mean field energy term that is nonlinear with respect to P see, for example, Eqs. (166), (168), (169), and (170). Likewise, for quantum and nonextensive systems we find that J is nonlinear because F involves quantum and nonextensive entropies such as $S_{\mathrm{FD,BE}}$ and S_q, see, for example, Eqs. (171) and (172).

Summary and Future Directions

From the previous discussion in Sect. "Linear Nonequilibrium Thermodynamics", it is clear that modeling approaches based on nonlinear Fokker–Planck equations are rooted in the theory of collective phenomena and self-organization, on the one hand, and in the theory of quantum mechanical and nonextensive systems, on the other. In contrast, linear Fokker–Planck equations are tailored to address the stochastic properties of systems composed of non-interacting subsystems when equating the material subsystem ensemble with the ensemble of statistical realizations. Note that – of course – linear Fokker–Planck equations can also be applied to discuss stochastic properties of self-organizing systems. However, in such cases either the stochastic behavior of order parameters by means of low dimensional linear Fokker–Planck equations is

discussed [108] or linear high-dimensional or even functional Fokker–Planck equations are involved [98].

We showed in Sect. "Markov Property, 2nd Order and Higher Order Statistics" that both linear and nonlinear Fokker–Planck equations exhibit Green's functions and Langevin equations. The fact that a nonlinear evolution equation can give rise to a Green's function may be counter-intuitive because Green's functions are associated with linearity. In fact, the evolution equation of the Green's function p is linear with respect to p. The nonlinearity is in the evolution equation for the time-dependent probability density P but not in the evolution equation of the Green's function p. In this context, we would like to re-iterate what we pointed out in Sect. "Strongly Nonlinear Fokker–Planck Equations": time-dependent solutions P do not necessarily correspond to Green's function p [73].

In the mathematical literature the theory of Markov processes that involve conditional probability densities of the form $p(x, t|x', t', P')$ has been discussed for decades (see references in Sect. "Definition of the Subject"). In line with these studies we suggest to refer to Markov processes with conditional probability densities of the form $p(x, t|x', t', P')$ as *nonlinear* Markov processes or *nonlinear* families of Markov processes [77]. Likewise, we suggest to refer to Markov processes whose conditional probability densities do not depend on P' as *linear* Markov processes or *linear* families of Markov processes. Using this terminology we would say that strongly nonlinear Fokker–Planck equations describe nonlinear Markov diffusion processes and vice versa nonlinear Markov diffusion processes can be expressed in terms of strongly nonlinear Fokker–Planck equations.

In physics and related disciplines the relevance of nonlinear Markov processes has to be explored in the future. That is, except for mathematics the theory of Markov process constructed from conditional probability measures of the form $p(x, t|x', t', P')$ is still in its infancy. The Chapman–Kolmogorov equation and the Kramers–Moyal expansion presented in Sect. "Markov Property, 2nd Order and Higher Order Statistics" provide promising departure points for future studies in this regard.

In the present study we pointed out that there are a few 'umbrella' concepts that apply to linear and nonlinear Fokker–Planck equations alike: the concepts of Markov diffusion processes, martingales, and linear nonequilibrium thermodynamics. Therefore, future studies that follow the ideas of the present study could change the state of the art illustrated in Fig. 1 into a scenario as shown in Fig. 2. In doing so, a closely connected world of linear and nonlinear Fokker–Planck equations that is governed by a small set of powerful principles could emerge.

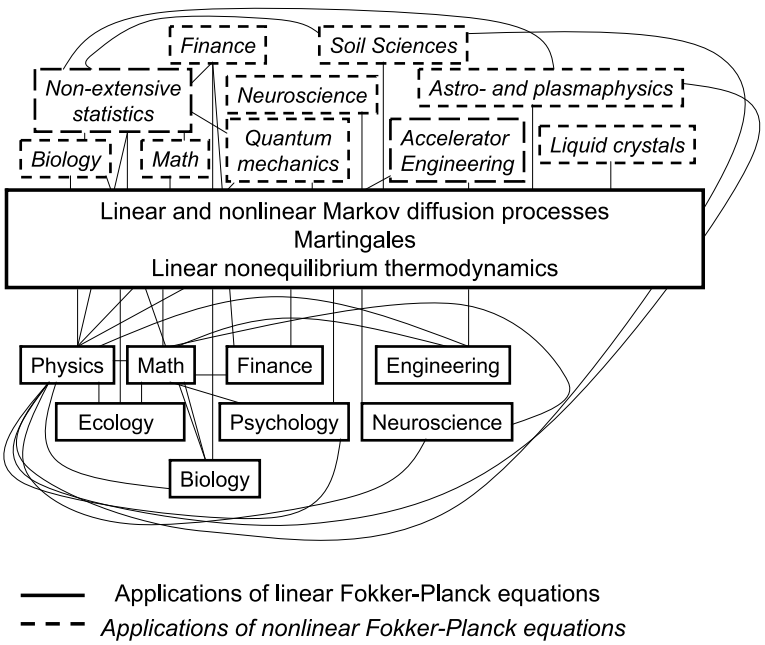

Linear and Non-linear Fokker–Planck Equations, Figure 2
Connected applications of linear and nonlinear Fokker–Planck equations

Bibliography

1. Abe S, Okamoto Y (2001) Nonextensive statistical mechanics and its applications. Springer, Berlin
2. Acebron JA, Bonilla LL, De Leo S, Spigler R (1998) Phys Rev E 57:5287
3. Acebron JA, Bonilla LL, Vicente CJP, Ritort F, Spigler R (2005) Rev Mod Phys 77:137
4. Acebron JA, Spigler R (1998) Phys Rev Lett 81:2229
5. Allen EJ, Victory HD Jr (1994) Physica A 209:318
6. Aonishi T, Okada M (2002) Phys Rev Lett 88:024102
7. Arenas A, Perez Vincente CJ (1994) Phys Rev E 50:949
8. Ariaratnam JT, Strogatz SH (2001) Phys Rev Lett 86:4278
9. Arnold A, Bonilla LL, Markowich P (1996) Transp Theor Stat Phys 25:733
10. Arnold A, Markowich P, Toscani G (2000) Transp Theor Stat Phys 29:571
11. Arnold A, Markowich P, Toscani G, Unterreiter A (2001) Commun Part Diff Eq 26:43
12. Aronson DG (1986) In: Dobb A, Eckmann B (eds) Nonlinear diffusion problems – Lecture notes in mathematics, vol 1224. Springer, Berlin, pp 1–46
13. Balescu R (1975) Equilibrium and nonequilibrium statistical mechanics. Wiley, New York
14. Barenblatt GI, Entov VM, Ryzhik VM (1990) Theory of fluid flows through natural rocks. Kluwer Academic Publisher, Dordrecht
15. Becker V, Engel A (2007) Phys Rev E 75:031118
16. Binney J, Tremaine S (1987) Galatic dynamics. Princeton University Press, Princeton, New Jersey
17. Birner T, Lippert K, Müller R, Kühnel A, Behn U (2002) Phys Rev E 65:046110
18. Bödeker HU, Röttger MC, Liehr A, Frank TD, Friedrich R, Purwins HG (2003) Phys Rev E 67:056220
19. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) Psychol Rev 113:700
20. Bonilla LL (1987) J Stat Phys 46:659
21. Bonilla LL, Vicente CJP, Rubi JM (1993) Stat J Phys 70:921
22. Borland L (1998) Phys Rev E 57:6634
23. Borland L (2002) Phys Rev Lett 89:098701
24. Borland L (2008) Anomalous fluctuations in complex systems: plasmas, fluids and financial markets. In: Riccardi C, Roman HE (eds) Special Review Book for Research Singpost, Transworld Research Network, KeralaMadison, 2008, in press
25. Borland L, Bouchaud JP (2004) Quantitative finance 4:499
26. Bychenkov VY, Rozmus W, Tikhonchuk VT (1995) Phys Rev Lett 75:4405
27. Carillo JA, Jüngel A, Markowich PA, Toscani G, Unterreiter A (2001) Monatshefte für Mathematik 113:1
28. Carrillo JA, Rosado J, Salvarani F (2008) Appl Math Letters 21:148
29. Cepa E, Lepingle D (1997) Probab Theory Relat Fields 107:429
30. Chandrasekhar S (1977) Liquid crystals. Cambridge University Press, Cambridge
31. Chavanis PH (2003) Phys Rev E 68:036108
32. Chavanis PH (2004) Physica A 340:57
33. Chavanis PH, Rosier C, Sire C (2002) Phys Rev E 66:036105
34. Coffey WT, Kalmykov YP, Waldron JT (2004) The Langevin equation. World Scientific, New Jersey
35. Compte A, Jou D (1996) Phys J Math A Gen 29:4321
36. Cortines AAG, Riera R (2007) Physica A 377:181
37. Crank J (1975) The mathematics of diffusion. Clarendon Press, Oxford
38. Crawford JD (1995) Phys Rev Lett 74:4341
39. Curado EMF, Nobre FD (2003) Phys Rev E 67:021107
40. Daido H (1996) Physica D 91:24
41. Daido H (1996) Phys Rev Lett 77:1406
42. Daido H (2001) Phys Rev Lett 87:048101
43. Daly E, Porporato A (2004) Phys Rev E 70:056303
44. Dano S, Hynne F, De Monte S, d'Ovidio F, Sorensen PG, Westerhoff H (2001) Faraday Discuss 120:261
45. Dawson DA (1983) J Stat Phys 31:29
46. Dawson DA (1993) Lecture notes in mathematics 1541. In: Dawson DA, Maisonneuve B, Spencer J (eds) Springer, Berlin, pp 6–16
47. Dawson DA, Gärtner J (1989) Large deviations, free energy functional and quasi-potential for a mean field model of interacting diffusions (American Mathematical Society, Providence)
48. de Gennes P, Prost J (1993) The physics of liquid crystals. Clarendon Press, Oxford
49. de Groot SR, Mazur P (1962) Non-equilibrium thermodynamics. North-Holland, Amsterdam
50. de Jeu WH (1980) Physical properties of liquid crystalline materials. Gordon and Beach, New York
51. de los Santos F, Telo da Gamma MM, Munoz MA (2003) Phys Rev E 67:021607
52. Desai RC, Zwanzig R (1978) Stat J Phys 19:1
53. Ding X (1994) Adv Appl Prob 26:1022
54. Djehiche B, Kaj I (1995) Ann Probab 23:1414
55. Doi M, Edwards SF (1988) The theory of polymer dynamics. Clarendon Press, Oxford
56. Donoso JM, Salgado JJ (2006) J Phys A 39:12587
57. Donoso JM, Salgado JJ, Soler M (2005) J Phys A 38:9145
58. Drazer G, Wio HS, Tsallis C (2000) Phys Rev E 61:1417
59. Drozdov AN, Morillo M (1996) Phys Rev E 54:931
60. Ebeling W, Sokolov IM (2004) Statistical thermodynamics and stochastic theory of nonequilibrium systems. World Scientific, Singapore
61. Eckhardt B, Ott E, Strogatz SH, Abrams DM, McRobie A (2007) Phys Rev E 75:021110
62. Epperlein EM, Rickard GJ, Bell AR (1988) Phys Rev Lett 61:2453
63. Espanol P, Serrano M, Öttinger HC (1999) Phys Rev Lett 83:4542
64. Felderhof BU (2003) Physica A 323:88
65. Fialkowiski M, Hess S (2000) Physica A 282:65
66. Fokker AD (1914) Annalen der Physik 43:810
67. Fontbona J (2003) J Funct Analysis 200:198
68. Frank TD (2001) Phys Lett A 280:91
69. Frank TD (2001) Phys Lett A 290:93
70. Frank TD (2002) Phys Lett A 305:150
71. Frank TD (2002) Physica A 310:397
72. Frank TD (2003) Phys Lett A 319:173
73. Frank TD (2003) Physica A 320:204
74. Frank TD (2004) Eur Phys J B 37:139
75. Frank TD (2004) J Phys A 37:3561
76. Frank TD (2004) Phys Lett A 329:475
77. Frank TD (2004) Physica A 331:391
78. Frank TD (2004) Physica D 195:229
79. Frank TD (2005) Math Comput Mod 42:1057
80. Frank TD (2005) Nonlinear Fokker–Planck equations: Fundamentals and applications. Springer, Berlin

81. Frank TD (2005) Phys Rev E 72:041703
82. Frank TD (2006) Phys Rev ST-AB 9:084401
83. Frank TD (2007) Physica A 382:453
84. Frank TD (2008) Physica A 387:773
85. Frank TD, Daffertshofer A (1999) Physica A 272:497
86. Frank TD, Daffertshofer A (2000) Physica A 285:351
87. Frank TD, Daffertshofer A (2001) Physica A 292:392
88. Frank TD, Daffertshofer A (2001) Physica A 295:455
89. Frank TD, Daffertshofer A, Peper CE, Beek PJ, Haken H (2001) Physica D 150:219
90. Frank TD, Friedrich R (2005) Physica A 347:65
91. Friedman A (1969) Partial differential equations. Holt, Rinehart and Winston, Inc., New York
92. Friedrich R, Peinke J (1997) Phys Rev Lett 78:863
93. Friedrich R, Siegert S, Peinke J, Lück S, Seifert M, Lindemann M, Raethjen J, Deuschl G, Pfister G (2000) Phys Lett A 271:217
94. Gärtner J (1988) Math Nachr 137:197
95. Gang H, Haken H, Fagen X (1996) Phys Rev Lett 77:1925
96. Garcia-Ojalvo J, Parrondo JMR, Sancho JM, van den Broeck C (1996) Phys Rev E 54:6918
97. Garcia-Ojalvo J, Sancho JM (1999) Noise in spatially extended systems. Springer, New York
98. Gardiner CW (1997) Handbook of stochastic methods, 2nd edn. Springer, Berlin
99. Giada L, Marsili M (2000) Phys Rev E 62:6015
100. Glansdorff P, Prigogine I (1971) Thermodynamic theory of structure, stability, and fluctuations. Wiley, New York
101. Goel NS, Dyn NR (1974) Stochastic models in biology. Academic Press, New York
102. Graham C (1990) Appl Math Optim 22:75
103. Greven A (2005) Interacting stochastic systems. In: Deuschel J, Greven A (eds) Springer, Berlin, pp 209–246
104. Gurtin ME, MacCamy RC (1977) Mathematical Biosciences 33:35
105. Hadley P, Beasley MR, Wiesenfeld K (1988) Phys Rev B 38:8712
106. Hänggi P, Marchesoni F, Nori F (2005) Annalen der Physik 14:51
107. Hänggi P, Talkner P, Borkovec M (1990) Rev Mod Phys 62:251
108. Haken H (2004) Synergetics: Introduction and advanced topics. Springer, Berlin
109. Han SK, Kurrer C, Kuramoto Y (1995) Phys Rev Lett 75:3190
110. Hansel D, Mato G, Meunier C (1993) Europhysics Letters 23: 367
111. Hansel D, Mato G, Meunier C (1993) Phys Rev E 48:3470
112. Hasegawa H (2003) Phys Rev E 67:041903
113. Heifets S (2001) Phys Rev ST-AB 4:044401
114. Heifets S (2003) Phys Rev ST-AB 6:080701
115. Hess S (1976) Naturforschung Z A 31:1034
116. Holden AV (1976) Models of the stochastic activity of neurons. In: Lecture notes in biomathematics, vol 12. Springer, Berlin
117. Horsthemke W, Lefever R (1984) Noise-induced transitions. Springer, Berlin
118. Hütter M, Karlin IV, Öttinger HC (2003) Phys Rev E 68:016115
119. Ichiki A, Ito H, Shiino M (2007) Physica E 40:402
120. Ilg P, Karlin IV, Öttinger HC (1999) Phys Rev E 60:5783
121. Ilg P, Kröger M, Hess S (2005) Phys Rev E 71:051201
122. Jafari GR, Fazeli SM, Ghasemi F, Allaei SMV, Tabar MRR, Zad AI, Kavei G (2002) Phys Rev Lett 91:226101
123. Jelic A, Hütter M, Öttinger HC (2006) Phys Rev E 74:041126
124. Jourdain B (2000) Methodology and computing in applied probability 2:69
125. Kadanoff LP (2000) Statistical physics: statics, dynamics and renormalization. World Scientific, Singapore
126. Kanamaru T, Horita T, Okabe Y (2001) Phys Rev E 64:031908
127. Kaniadakis G (2001) Phys Lett A 288:283
128. Kaniadakis G (2001) Physica A 296:405
129. Kaniadakis G, Quarati P (1993) Phys Rev E 48:4263
130. Kaniadakis G, Quarati P (1994) Phys Rev E 49:5103
131. Karlin S, Taylor HM (1975) A first course in stochastic processes. Academic Press, New York
132. Karlin S, Taylor HM (1981) A second course in stochastic processes. Academic Press, New York
133. Kawamura Y (2007) Phys Rev E 76:047201
134. Kawamura Y, Nakao H, Kuramoto Y (2007) Phys Rev E 75:036209
135. Kharchenko DO, Kharchenko VO (2005) Physica A 354:262
136. Kim S, Park SH, Ryu CS (1997) Phys Rev Lett 78:1616
137. Klimontovich YL (1986) Statistical physics. Harwood Academic Publ, New York
138. Kloeden PE, Platen E (1992) The numerical solution of stochastic differential equations. Springer, Berlin
139. Kometani K, Shimizu H (1975) J Stat Phys 13:473
140. Kondepudi D, Prigogine I (1998) Modern thermodynamics. Wiley, New York
141. Kostur M, Luczka J, Schimansky-Geier L (2002) Phys Rev E 65: 051115
142. Kozyreff G, Vladimirov AG, Mandel P (2000) Phys Rev Lett 85:3809
143. Kuramoto Y (1984) Chemical oscillations, waves, and turbulence. Springer, Berlin
144. Lancellotti C, Kiessling M (2001) Astrophys J 549:L93
145. Larson RG, Öttinger HC (1991) Macromolecules 24:6270
146. Li JH, Huang ZQ, Xing DY (1998) Phys Rev E 58:2838
147. Li JH, Hänggi P (2001) Phys Rev E 64:011106
148. Lo CF (2005) Phys Lett A 336:141
149. MacDonald WM, Rosenbluth MN, Chuck W (1957) Phys Rev 107:350
150. Maier W, Saupe A (1958) Z Naturforschung A 13:564
151. Marsili M, Bray AJ (1996) Phys Rev Lett 76:2750
152. McCauley JL, Gunaratne GH, Bassler KE (2006) Physica A 382: 445
153. McKean HP Jr (1969) Propagation of Chaos for a class of nonlinear parabolic equations. In: Aziz AK (ed) Lectures in differential equations, vol II. Van Nostrand Reinhold Company, New York, pp 177–193
154. Meleard S (1996) Asymptotic behavior of some interacting particle systems: McKean-Vlasov and Boltzmann models. In: Graham C, Kurtz TG, Meleard S, Potter PE, Pulvirenti M, Talay D (eds) Probabilistic models for nonlinear partial differential equations. Springer, Berlin, pp 42–95
155. Meleard S, Coppoletta SR (1987) Stoch Process Appl 26:317
156. Michael F, Johnson MD (2003) Physica A 320:525
157. Mikhailov AS, Zanette DH (1999) Phys Rev E 60:4571
158. Morillo M, Gomez-Ordonez J, Casado JM (1995) Phys Rev E 52: 316
159. Müller R, Lippert K, Kühnel A, Behn U (1997) Phys Rev E 56: 2658
160. Nicholson DR (1983) Introduction to plasma theory. Wiley, New York
161. Nobre FD, Curado EMF, Rowlands G (2004) Physica A 334: 109
162. Oelschläger K (1989) Probab Th Rel Fields 82:565

163. Öttinger HC (1996) Stochastic processes in polymeric fluids. Springer, Berlin
164. Öttinger HC (2005) Beyond equilibrium thermodynamics. Wiley, New Jersey
165. Öttinger HC (2007) Phys Rev Lett 99:130602
166. Öttinger HC, Grmela M (1997) Phys Rev E 56:6633
167. Okubo A, Levin SA (2001) Diffusion and ecological problems: modern perspectives. Springer, Berlin
168. Overbeck L (1996) Ann Probab 24(2):743
169. Park K, Lai YC, Liu Z, Nachman A (2004) Phys Lett A 326:391
170. Park SH, Kim S (1996) Phys Rev E 53:3425
171. Patanarapeelert K, Frank TD, Friedrich R, Beek PJ, Tang IM (2006) Phys Lett A 360:190
172. Paul W, Baschnagel J (1999) Stochastic processes: from physics to finance. Springer, Berlin
173. Peletier LA (1981) In: Amann H, Bazley N, Kirchgässner K (eds)(1981) Applications of nonlinear analysis in the physical science. Pitman Advanced Publishing Program, Boston, pp 229–241
174. Pikovsky A, Rosenblum M, Kurths J (2001) Synchronization: a universal concept in nonlinear sciences. Cambridge University Press, Cambridge
175. Pilipenko AY (2005) Urkainian Math J 57:1507
176. Planck M (1917) Sitzungsber. Preuss Akad Wissens, Berlin, p 324
177. Plastino AR, Plastino A (1995) Physica A 222:347
178. Quinn DA, Rand RH, Strogatz SH (2007) Phys Rev E 75:036218
179. Ratcliff R, Gomez P, McKoon G (2004) Psychol Rev 111:159
180. Reif F (1965) Fundamentals of statistical and thermal physics. McGraw-Hill Book Company, New York
181. Reimann P (2002) Phys Rep 361:57
182. Reimann P, Kawai R, Van den Broeck C, Hänggi P (1999) Europhysics Letters 45:545
183. Reimann P, Van den Broeck C, Kawai R (1999) Phys Rev E 60:6402
184. Risken H (1989) The Fokker–Planck equation – Methods of solution and applications. Springer, Berlin
185. Rogers LCG, Shi Z (1993) Probab Theroy Relat Fields 95:555
186. Rosenblum MG, Pikovsky AS (2004) Phys Rev Lett 92:114102
187. Rosenbluth M, MacDonald WM, Judd DL (1957) Phys Rev 107:1
188. Sakaguchi H (1988) Prog Theor Phys 79:39
189. Savel'ev S, Marchesoni F, Nori F (2003) Phys Rev Lett 91:010601
190. Scarfone MA, Wada T (2007) Physica A 384:305
191. Schuster HG, Wagner P (1990) Biol Cybern 64:77
192. Schwämmle V, Nobre FD, Curado E (2007) Phys Rev E 76:041123
193. Schweitzer F (2003) Brownian agents and active particles. Springer, Berlin
194. Schwämmle V, Gonzahles MC, Moreira AA, Andrade JS, Hermann HJ (2007) Phys Rev E 75:066108
195. Shiino M (1985) Phys Lett A 112:302
196. Shiino M (1987) Phys Rev A 36:2393
197. Shiino M (2001) J Math Phys 42:2540
198. Shiino M (2002) J Math Phys 43:2654
199. Shiino M (2002) J Phys Soc 40:1037
200. Shiino M (2003) Phys Rev E 67:056118
201. Shimizu H (1974) Prog Theor Phys 52:329
202. Shimizu H, Yamada T (1972) Prog Theor Phys 47:350
203. Shobuda Y, Hirata K (2001) Phys Rev E 64:067501
204. Silverman BW (1986) Density estimation for statistics and data analysis. Chapman and Hall, London
205. Soler M, Martinez FC, Donoso JM (1992) Stat J Phys 69:813
206. Sopik J, Sire C, Chavanis PH (2006) Phys Rev E 74:011112
207. Strogatz SH (2000) Physica D 143:1
208. Strogatz SH, Abrams D, McRobie A, Eckhardt B, Ott E (2005) Nature 438:43
209. Strogatz SH, Mirollo RE (1991) Stat J Phys 63:613
210. Strogatz SH, Morillo RE, Matthews PC (1992) Phys Rev Lett 68:2730
211. Stupakov GV, Breizman BN, Pekker MS (1997) Phys Rev E 55:5976
212. Sura P, Barsugli J (2002) Phys Lett A 305:304
213. Takai M, Akiyama H, Takeda S (1981) J Phys Soc Japan 50:1716
214. Takatsuji M (1975) Biol Cybern 17:207
215. Tass PA (1999) Phase resetting in medicine and biology – Stochastic modelling and data analysis. Springer, Berlin
216. Tass PA (2001) Europhysics Letters 55:171
217. Tass PA (2003) Biol Cybern 89:81
218. Tass PA (2006) Biol Cybern 94:58
219. Tsallis C (1988) Stat J Phys 52:479
220. Tsallis C (1997) Phys World 10(7):42
221. Tsallis C (2004) Physica D 193:3
222. Tsallis C, Bukman DJ (1996) Phys Rev E 54:R2197
223. Tsekov R (1995) J Phys A Math Gen 28:L557
224. Tsekov R (2001) Int J Mol Sci 2:66
225. van den Broeck C, Parrondo JMR, Armero J, Hernandez-Machado A (1994) Phys Rev E 49:2639
226. van den Broeck C, Parrondo JMR, Toral R (1994) Phys Rev Lett 73:3395
227. van den Broeck C, Parrondo JMR, Toral R, Kawai R (1997) Phys Rev E 55:4084
228. van Kampen NG (1981) Stochastic processes in physics and chemistry. North-Holland, Amsterdam
229. Vellekoop M, Nieuwenhuis H (2007) Quantitative finance 7:563
230. Venturini M, Warnock R (2002) Phys Rev Lett 89:224802
231. Waechter M, Riess F, Schimmel T, Wendt U, Peinke J (2004) Eur Phys J B 41:259
232. Wehner MF, Wolfer WG (1987) Phys Rev A 35:1795
233. Wiesenfeld K, Colet P, Strogatz SH (1996) Phys Rev Lett 76:404
234. Winfree AT (1967) Theor J Biol 16:15
235. Winfree AT (2001) The geometry of biological time, 2nd edn. Springer, Berlin
236. Yamaguchi Y, Shimizu H (1984) Physica D 11:213
237. Yamana M, Shiino M, Yoshioka M (1999) Phys J A Math Gen 32:3525
238. Yoshioka M, Shiino M (2000) Phys Rev E 61:4732
239. Zaikin AA, Garcia-Ojalvo J, Schimansky-Geier L, Kurths J (2002) Phys Rev Lett 88:010601
240. Zhai Y, Kiss IZ, Daido H, Hudson JL (2005) Physica D 205:57

Link Analysis and Web Search

Johannes Bjelland,
Geoffrey Canright,
Kenth Engø-Monsen
Telenor R&I, Fornebu, Norway

Article Outline

Glossary

Graph A set of vertices connected by directed or undirected edges.

Web graph The World Wide Web represented as a graph. The webpages are vertices and hyperlinks are edges. The Web graph is called a *directed graph* since the hyperlinks are directed links from one page to another.

Adjacency matrix A matrix A which represents the structure of a graph. For the Web graph, with directed links, A is typically defined such that $A_{ij} = A_{i\to j} = 1$ if page i points to page j, and 0 otherwise.

Link analysis In Link Analysis the relations between nodes in a graph are studied. In large complex networks link analysis is a tool to analyze a node's relations to all other nodes in the network. Based on the link structure of the graph, several node-properties can be found, for instance: Is the node well connected to the rest of the network? Can node A be reached from node B? How many edges separate A from B? Is a node central in its network? Is it contained in a sub-network? With link analysis these and other node properties are found. In short, link analysis helps us to extract a node's role in a graph.

Web-page ranking Link analysis is used to measure the quality of a webpage. A link is interpreted as a recommendation. Link analysis scores are used by search engines when hits from queries (Web pages) are presented to the user in the form of a ranked list.

PageRank The most well-known Web-page ranking algorithm, developed at Stanford University by Larry Page and Sergey Brin. Page and Brin are the founders of Google, which employs the PageRank-algorithm.

Personalization of ranking-scores When the web-page link analysis score is calculated, some pages may be boosted according to criteria other than just the link-structure of the Web graph. For example, pages relevant to a certain topic (which is preferred by the user) might be boosted to receive a better ranking.

Link-spamming A method where a spammer tries to boost a webpage's rank in a search engine hit list. This is done by constructing several pages (a 'link spam farm') and manipulating the link-structure in such a way that a targeted page will get a boost in its link analysis score, and hence a higher ranking in a search engine hit list.

Crawling A *crawler* (or *spider*) is a program that maps the contents of the World Wide Web by starting at some initial webpage(s) and jump from page to page via hyperlinks in a systematic manner. All pages reachable from the startpage(s) by clicking hyperlinks are then visible to the crawler. From each page visited, hyperlinks and useful information is extracted. Search engines use webcrawlers to build their searchable databases (*index*). The part of the WWW which has been visited by a crawler and stored in a search-engine's database is called the *indexable Web* or the *visible Web*.

Node indegree (outdegree) The number of links pointing to (from) a node. The idea is readily generalized to a set of nodes such as an SCC (see below).

Strongly connected component (SCC) A maximal set of nodes in a directed graph, such that there is a directed path from any node in the SCC to any other.

Source (sink) SCC An SCC with zero indegree (outdegree).

Definition of the Subject

The World Wide Web consists of a vast amount of information in the form of web-pages. Almost all subjects are covered somewhere in the billions of pages reachable from your computer. To access this enormous library from our computers is a wonderful thing – but the tremendous size of the web is also a problem. With an ever-growing number of pages covering all sorts of subjects, it becomes very hard to find the one, relevant piece of information you need. Web link analysis is a tool which assists us with this task. It helps us to rank web documents, and so to find the most important documents among the huge amount of irrelevant information available online. The input to the web link analysis is a network of hyperlinked documents, where the nodes are documents, and the directed edges are hyperlinks from one page to another. The output is a vector containing a numerical score which represents the importance of each page in the Web graph. A typical use of this score is to assist ranking of a search engine query result. The most well-known example is the search engine

Google, which uses a web-link analysis algorithm developed by Brin and Page [68]. In this article we will present an overview of the principal existing methods used to calculate importance scores for Web pages, using the Web graph as input. We will discuss the theory behind the different approaches, compare them to each other, and discuss some selected practical and technical problems with Web link analysis.

Introduction

Since the early days of the World Wide Web (WWW) there has been an explosion in the number of hyperlinked documents available online. Lawrence and Giles [56] estimated the publicly indexable web to contain about 800 million pages in 1999. In 2005 the number was estimated to be over 11.5 billion pages [39]. The task of sorting and organizing all this information is an impossible task for humans – machines must somehow be taught to do this job. Web link analysis is just one of the ways that machines help us to find the needle in the enormous haystack that is the World Wide Web.

Consider a reader who has just gotten a machine to help her to find a document of interest. More precisely: the reader *hopes* that the machine has delivered one or more documents of interest. This is an extremely common situation in the Information Age; and it presents interesting challenges – some of which we will address in this article. The reader normally has two types of criteria in mind for a 'good' document. First, it should be *relevant* to the user's current focus or need – for example the purchase of new cars in Malaysia. Second, the document should be of high *quality*. Clearly, each of these criteria is important for the reader. Furthermore, each is difficult for a machine to evaluate; but the second one – quality – is perhaps the most difficult. After all, human readers can debate endlessly what constitutes or exemplifies high quality, and what does not; how then can a machine say something useful about this question?

Web link analysis (or simply 'link analysis') – which we often abbreviate simply as LA – is a tool which helps us to calculate the quality of web pages. Consider two web pages A and B. The basic idea behind LA is that if page A contains a hyperlink to page B, this is interpreted as a recommendation of page B by page A. This interpretation is plausible since many (but not all) hyperlinks are laid down by human writers – typically the author of the pointing page. By using a hyperlink, the author is sending the readers away to another page. If one assumes the author has read this page and found the contents interesting enough to hyperlink it, this would increase our subjective idea of the quality of the targeted page. Two of the most well-known LA algorithms that use this approach are PageRank [68] and HITS [52]. These and other algorithms will be discussed in Sect. "Approaches to Web Link Analysis".

We note briefly here that there are close connections between Web link analysis and the older field termed *bibliometrics* – in particular, a subfield called *citation analysis*. The latter involves the analysis of the "web of citations" for scholarly publications, with the aim of estimating the importance or impact of each article. Citations (like hyperlinks) are one-way links, so that the citation graph is a directed graph. However – unlike the case for hyperlinks – citations which point "forward in time" are very rare, so that the citation graph is more treelike than the Web graph. Garfield [37] defines the 'impact factor' for a publication in a way that basically counts the inlinks (citations) for a given period. Hence this approach is close to one in Web link analysis which we will call *'link popularity'* below. For a good discussion of more sophisticated, whole-graph approaches to citation analysis, see [52].

The WWW contains billions of pages connected with hyperlinks. A page's LA score can be regarded as a form of *collective recommendation* from all the other pages. High LA score means that the page has a high degree of collective recommendation. It is then perhaps more precise to replace the term 'quality' with 'importance' [19] or 'authority' [52] on the WWW.

A common way to retrieve information from the large and complex WWW is to submit a query to a search engine. Lawrence and Giles [56] found that search-related pages ranked among the top 10 most visited sites on the web. Also, from studying the most frequently used search words (queries), one can see that many users use search engines to replace bookmarks – typing in the name or (part or even all of) the url of a well known page. Thus, search engines are a normal part of life for many of us. A search engine contains a searchable database *(index)* of extracted information from pages found on the World Wide Web. The largest search engines have indexed billions of web-pages [39]. If a query is not too obscure, the user will be presented with a large list of hits more or less relevant to the query. For instance, a query 'car' returns around 3 billion hits from the Yahoo search engine at the time of writing this document [88]. The huge majority of the hits are useless for the person submitting the search. Some hits may even be spam – pages constructed with the purpose to get as many hits as possible. (Search engine spamming is discussed in Sect. "Sensitivity of Link-Analysis Scores".)

For a human, it is impossible to browse billions of

search engine hits to find the most relevant and high-quality pages. The importance score from link analysis is used in the process of ranking the search results presented to the user, according to their quality or importance on the web. This importance score is typically combined with a *text relevance score* which evaluates each hit with respect to its relevance to the query. Hits are then ranked according to the combined score, so that pages that are both important and relevant appear first on the hit list. An example of a successful implementation of LA in a search engine is the earlier mentioned PageRank used by Google. The PageRank of a web page is Google's way of measuring how important this page is, as compared to all other pages in the enormous Web.

In this article we present an overview of web search and link analysis. In link analysis it is useful to define the World Wide Web (WWW) as a network. The web-documents are nodes and the hyperlinks are edges. This network is then represented by a directed graph. Properties of general directed graphs, and also of the Web graph, are discussed in Sect. "Structure of Directed Graphs". A subsection of Sect. "Structure of Directed Graphs" is devoted to the 'Sink Problem' – a numerical problem that may arise when calculating scores for groups of pages which basically have many inlinks but no outlinks. In Sect. "Approaches to Web Link Analysis", existing approaches to web link analysis are discussed. Among these are the already mentioned and well-known methods PageRank and HITS, plus two other methods 'Link popularity' and 'SALSA' [58], and also a recent algorithm called 'T-Rank' [15,16]. In Sect. "The 'Sink Problem'", the 'sink problem' is revisited and discussed in relation to the methods described in Sect. "Approaches to Web Link Analysis". Another topic that has gotten considerable study is the sensitivity of the link analysis scores. How do network topology changes, sink remedies and spam affect the scores? This is discussed in Sect. "Sensitivity of Link-Analysis Scores". Then, in Sect. "Personalization of Link Analysis Scores" we discuss approaches to personalizing link analysis scores, and in Sect. "Peer-to-Peer Search and Link Analysis" we look at distributed approaches to link analysis, including the application to peer-to-peer search. Finally, we sum up with a look at some future directions of search and link analysis in Sect. "Future Directions".

Structure of Directed Graphs

In this section we briefly review the structure of directed graphs – which is different in important ways from that of undirected graphs. We also briefly present the structure of the Web graph: a directed, dynamic, largely self-organized graph, composed of Web documents connected via hyperlinks. The Web graph is the raw material which is used by Web link analysis algorithms to generate LA scores.

General

As earlier mentioned, the World Wide Web is a network of pages connected by hyperlinks. In mathematics and computer science such a network is called a graph. The webpages are nodes and the hyperlinks are edges. Since the hyperlinks have a direction, the webgraph is a directed graph. In this subsection a few useful basic graph-theoretic terms will be explained. These and several other properties of graphs are described in textbooks on graph theory and in Web graph articles, for instance [13,21], and [34].

We start with the simpler case – an undirected (or symmetric) graph. Undirected links are termed 'symmetric' because an undirected link $\mathbf{A - B}$ implies a path from A to B *and* from B to A, and hence is equivalent to $\mathbf{A \leftrightarrow B}$. The *degree* of a node in an undirected graph is the number of edges it has. We typically study *connected* graphs – i. e., graphs which do not decompose into several pieces, without removing one or more links. Connected symmetric graphs are very simple: for *any* two nodes i and j in the graph, there is a path from i to j. The *distance* from i to j is the number of edges we have to follow to reach j starting at i.

Directed graphs are built up of one-way links: $\mathbf{A \rightarrow B}$. This statement does not exclude the possibility of reciprocal one-way links, i. e., that both $\mathbf{A \rightarrow B}$ and $\mathbf{A \leftarrow B}$ are present; however, neither is this guaranteed. The above definition of a connected graph may also be applied to directed graphs. This is equivalent to the more conventional definition – that a directed graph is connected if, when all links are made symmetric, there is a path from i to j for any two nodes i and j in the graph.

A node in a directed graph has both edges pointing out from it and edges pointing in. The number of edges pointing respectively out and in is called the node's *outdegree* and *indegree*. The directionality of the links gives directed graphs features not seen in symmetric (undirected) graphs. For example, in general, when we do *not* ignore the direction of the links, there is *not* necessarily a (directed) path from i to j, for any two nodes i and j in a connected directed graph. This is clear from the 'simplest possible' directed graph, the two-node graph $\mathbf{A \rightarrow B}$: there is a path from A to B, but not from B to A.

This single (negative) statement distinguishes directed graphs from undirected graphs. Now we build from this property, in a positive way. One can find a set C of nodes in a directed graph, such that, for any two nodes i and j

in the set C, there is a directed path from i to j. This set is called *strongly connected*; and if we make this set as large as possible – without killing its property of being strongly connected – then we get a set (subgraph) which is called a *strongly connected component* or SCC. SCCs are very useful in thinking about directed graphs. Any directed graph may be uniquely decomposed into its strongly connected components – that is, every node belongs uniquely to one and only one SCC. Furthermore, given two SCCs C1 and C2, any links connecting the two are either (i) all in the same direction, e. g. **C1→C2**, or (ii) nonexistent. A two-way connection such as **C1↔C2** is impossible, since then the two SCCs collapse to one. For the same reason, there are no cyclic connections of the form **C1→C2→C3→**…**→C1**, since the entire cycle comprises in fact a single SCC.

Thus, movement along directed paths in a directed graph takes two forms: 'circulation' from any node to any other *within* an SCC; and one-way or 'downhill' flow *between* SCCs. This overall one-way flow in directed graphs, which is found at the coarse (inter-SCC) scale (see Fig. 1), is quite distinct from the purely circulatory nature of connected undirected graphs; and it has interesting consequences. For example: unless the directed graph is composed of a single SCC, there must always be one or more SCCs which have no inlinks from other SCCs; these are called *source* SCCs. Similarly, there must be one or more 'ends' to the downhill flow – SCCs without outlinks, termed *sink* SCCs.

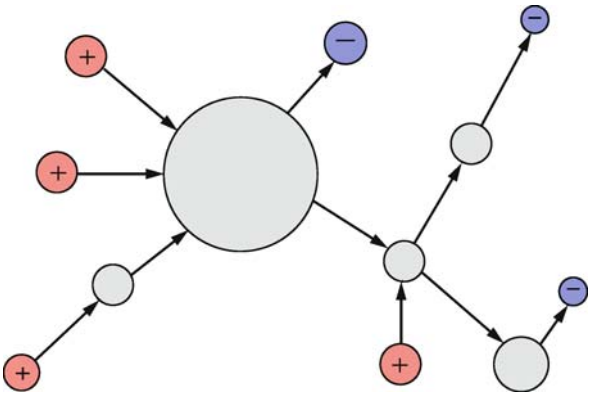

Link Analysis and Web Search, Figure 1
Generic structure of a directed graph. Each strongly connected component or SCC (see text) is indicated with a *circle* here. The (one-way) direction of links between SCCs is indicated with *arrows*. We see that the *arrows* take one from 'source' SCCs (marked with '+'), via (typically) intermediate SCCs, to 'sink' SCCs (marked with '−')

If we now recall the real, concrete Web graph, our intuition might lead us to say that 'most' Web pages participate in some kind of large-scale community of pages, with circulation from any page to any other. We will see that this intuitive picture is not particularly misleading. Yet, at the same time, we reiterate that because the Web graph is a directed graph, which is almost certainly not a single SCC – it *must* have one or more sets of pages for which there are no inlinks to the set. These are the source SCCs of the Web graph. Similarly, the Web graph is guaranteed to have one or more sink SCCs – sets of pages such that a reader who arrives to the set via a hyperlink will not find any path leading out again. Furthermore, we cannot expect just one single, almost-all-encompassing intermediate SCC lying between these sources and sinks. The Web graph is built up by billions of uncoordinated individual acts; and its structure (as we will see in the next subsection) reflects its anarchic and self-organized nature.

The Web Graph

The Web graph is a directed graph. The edges are hyperlinks which point from page A to page B. Using a normal web-browser, there is no way to jump from page B to A if the hyperlink points from A to B. The Web graph is also extremely huge, and is continuously growing and changing. It expands by the addition of nodes (webpages) and of new edges (hyperlinks). Anyone can add pages and links to the WWW, so the structure of the Web graph becomes very complex. The Web graph has grown into a massive world-wide 'organism' [32]. An interesting feature of the Web graph in relation to its complex structure is the indegree-distribution. As the age of a page increases, the page might become more and more popular, such that other pages start linking to it. Studies have shown that the indegree-distribution of the Web graph follows a *power-law distribution* [9] and [52]. A discrete random variable X follows a power-law distribution if the probability of X taking value x is proportional to $x^{-\gamma}$, with $\gamma > 0$ being the power-law exponent. Barabasi and Albert [9] found that this feature was a consequence of new nodes constantly being added, and that new vertices are preferentially attached to already well-connected nodes.

Broder et al. [21] studied a set of 200 million webpages with 1.5 billion links and confirmed the power-law scaling of indegrees, with $\gamma = 2.1$. In addition they found that the Web graph has a structure schematically described by Fig. 2. This figure is now generally accepted as describing the structure of the World Wide Web.

The structure in Fig. 2 has been called a 'bow tie' – for

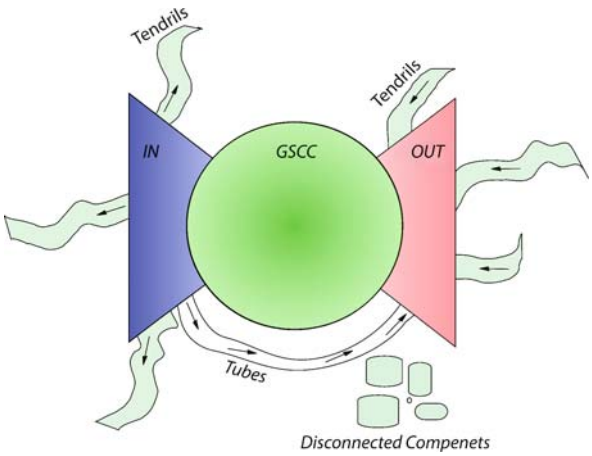

Link Analysis and Web Search, Figure 2
The bow-tie structure of the WWW as found by Broder et al., reproduced after [21]

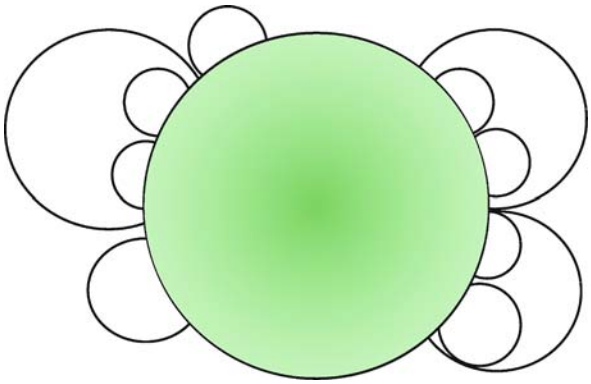

Link Analysis and Web Search, Figure 3
Daisy. The structure of the WWW as found by Donato et al. The IN and OUT set are highly fragmented and lie close (in hops) to the highly connected CORE. Figure reproduced after [32]

obvious reasons. The 'knot' of the bow-tie, the CORE, is a giant SCC (GSCC), which lies *centrally* (in an important sense) in the Web graph. The IN component represents another huge set of nodes. A Web surfer who starts in the IN component may reach the giant SCC from a node in IN; but no node in IN is reachable from the GSCC. IN should be thought of as many SCCs, many of which (but not all) are source SCCs. Similarly, nodes in OUT may be reached from the GSCC – but there is no path from any node in OUT into the GSCC. Also, we see that some paths which leave IN do not reach the GSCC; instead they terminate in 'out-Tendrils' to which IN is linked. OUT has its corresponding 'in-Tendrils' (which must include some source SCCs).

Still we have not exhausted the possible categories of Web page positions in the Web graph. We see from Fig. 2 that there are paths from IN to OUT which never reach the GSCC. Nodes in such paths lie in the 'Tubes'. Finally, there are pages which are not even connected to the large, connected structure which is composed of IN, GSCC, OUT, Tendrils, and Tubes. These pages are noted as 'Disconnected components' in Fig. 2.

The corresponding sizes of the components in the connected, 200M subset of the WWW were found to be: 44M in the IN-set, 56M in the GSCC, 44M in the Tendrils and 44M in the OUT-set. In the same way as the distributions of indegrees, the sizes of SCCs inside the main components were also found to follow a power-law distribution.

In Donato et al. [32] the inner structure of the "bow-tie" was studied, and a refined picture of the Web graph was suggested. They found that the main components (IN, OUT, GSCC=CORE, TENDRILS) have differences on a finer level. IN and OUT were found to be highly fragmented, while the CORE was well interconnected. No large SCC was found in the IN and OUT components. Most of the nodes contained in these sets belong to very small SCCs, and lie within a short distance from the GSCC. The resulting picture that emerges from [32] is shown in Fig. 3. We see that the IN and OUT sets are fragmented into small and poorly interconnected SCCs, hanging from the large and well-connected GSCC.

The CORE of the Web graph is highly interconnected and has "Small World" properties [3,6]. Graphs found in several man-made systems have these properties. For example the power grid of the western US, the collaboration graph of film actors and the neural network of the worm Caenorhabditis elegans, all share the property of being "small-world" networks [3]. A 'small world' graph is highly clustered (so that the idea of 'neighborhoods' makes sense), while at the same time the minimum distance between any two nodes in the network is short [85]. Albert et al. estimate an average of 19 clicks between any two pages of the WWW (at that time estimated to contain $8 * 10^8$ pages), based on a crawl of the nd.edu site [6]. Adamic finds an average distance of 4.2 between any sites in a GSCC containing about 65 000 sites [3]. Using a 200M nodes crawl of the WWW, Broder et al. finds an average distance of roughly 16 [21]. This means that even if the Web graph is large and complex, the diameter is small, and a huge amount of information is in principle reachable within a few clicks.

To implement a successful web link analysis, the link structure of the web has to be mapped.

Crawling is used to build large indexes of the WWW. These indexes are used by search engines, when a query is submitted, to find the hits which contain the search terms.

In addition, the web link structure is obtained by crawling, and stored. A crawler starts at some page (or set of pages), collects the information and outlinks on this page, and continues to the next by following one of the collected outlinks. Given the directed "bow-tie" structure of the Web graph, its size and its complexity, crawling is not a trivial task. A crawl starting on a page in the OUT set will never find pages in the GSCC or IN, since the crawler only sees the outlinks on any given page. More generally, a crawl starting in any SCC will only be able to reach "downstream" SCC's (see discussion on SCC's and Fig. 1).

Knowledge of the webgraph's structure is used to improve crawling strategies and to create effective crawlers [63]. Results from [78] imply that the structure of the obtained graph might depend on the crawler used.

In summary: Figs. 2 and 3 reveal the Web graph as a somewhat wild and woolly community of Web pages, which is built up by the combined, but largely uncoordinated, actions of humans and machines. It is the task of Web link analysis to extract meaningful and useful information from this structure – one which, furthermore, is constantly growing and changing. This is a formidable task; and yet, remarkably enough, it is not impossible.

The 'Sink Problem'

If we view hyperlinks as recommendations, then there are clearly pages which receive little or no collective recommendation from the Web graph. For example, pages lying in source SCCs may recommend one another; but no page from 'outside' this set recommends any of these. Similarly, IN is in some sense more weakly recommended than OUT. The point is that the asymmetric nature of hyperlinks seems to favor 'selfish' pages (and sets of pages) which manage to accumulate recommendations, without giving out recommendations. Put differently: the arrows themselves in Fig. 1 show that there is a real tendency for 'recommending weight' to accumulate on the 'downstream' side. In some cases (as we will see below) this tendency takes the extreme form that 'upstream' SCCs may receive *no* weight (and hence a score of zero), even though they are far from being isolated, and may be, by most criteria, important.

We will call this generic tendency the 'sink problem'. We use the word 'problem' because some approaches to link analysis give poor or useless results unless some corrective measure is applied – a 'sink remedy'. We will look at the sink problem in more detail in Sect. "Sensitivity of Link-Analysis Scores", after we have described several distinct link analysis algorithms in the next Section.

Approaches to Web Link Analysis

In the previous section we have sketched the structure of the Web graph. We see that this structure is heavily influenced by the one-way nature of the hyperlinks. In this section, we examine many of the studied approaches to link analysis. The problem is simply stated: given a page p in the Web graph, how can we generate a meaningful 'importance score' for p, which is based on the *position* of p in the entire graph?

Link Popularity

Here the idea is very simple: a Web page is a good page if it has many pages pointing to it – and hence recommending it. Therefore, one can simply count the inlinks to a given Web page to get a quantitative measure of its quality or importance. The resulting score k^{in} is commonly termed the 'link popularity' of the page – since it is entirely determined by the number of 'votes' given to a page.

This idea is intuitively reasonable. Also, it is very easy to calculate: once the full Web graph is obtained, the 'invisible' inlinks become fully visible, and the counting of these is a computational task which can be done very easily. The obtaining of k_i^{in} for every page i is equivalent to multiplying a matrix times a vector once. To see this, let A be the *adjacency matrix* of the Web graph, such that $A_{ij} = A_{i\to j} = 1$ if i points to j, and 0 otherwise. Then

$$k_i^{\text{in}} = \sum_j A_{j\to i} = (A^{\text{T}}\mathbf{1})_i \,. \tag{1}$$

Here A^{T} is the transpose of A, and $\mathbf{1}$ is a vector of ones.

The idea of link popularity is simple, and it has a certain utility. However, link popularity is particularly vulnerable to the problem of *link spam*, by which 'phony' pages and links are used to boost the importance scores of 'real' pages. Clearly, it is easy and cheap to set up a large number of fake pages which all point to a given target page T – thus giving T a high (and meaningless) link popularity score. What is clearly needed is some LA approach which takes into consideration the *quality* of the 'voting' pages responsible for the high popularity score. The approaches described in the following subsections attempt to do this, in distinct ways.

We conclude this subsection by mentioning *anchor text scoring*. This is a hybrid of link popularity and text relevance: a page P gets a high anchor text score, relative to the search term 'bla', if there are many other pages with inlinks to P, and those other pages associate the anchor text 'bla' with the links to P. It is intuitively reasonable that a short text, which is used by a page P_{out} to describe a page P_{in}, may be a good succinct description of the

linked-to page P_{in} (but not always – see the discussion on 'Google bombing' below).

HITS

The fundamental idea that may be expected to correct for the effects of link spam – largely, if not completely – is to insist that the *quality* of the pointing pages should be taken into account in assigning the importance of the pointed-to page. This means, among other things, that worthless dummy pointing pages should contribute nothing to the score of the pointed-to page.

This fundamental idea forms the foundation for all the other methods for LA that we will discuss. The idea clearly can be expressed in many ways; but it is important to note that, since the quality of the pointing pages is used to find the quality of the pointed-to pages, one finds oneself in fact in possession of a circular definition. We may express this circularity most succinctly as follows:

Good pages are pointed to by good pages. (2)

The HITS approach of Jon Kleinberg [52] is an implementation of the idea embodied in (2), but in a form in which every document gets *two* types of importance scores – one for the extent to which it is a good '*hub*', and one expressing the extent to which it is a good '*authority*'. These two roles may be expressed (again circularly) in terms of each other, as follows:

Good hubs point to good authorities. (3)

Good authorities are pointed to by good hubs. (4)

Note that (3) and (4) are redundant – that is, (4) is simply a restatement of (3). Together, these statements amount to a circular definition of two kinds of quality: 'authoritativeness' and 'hubness'. Kleinberg's idea was that – just as with people – some documents may be viewed as good sources to consult, because of the special knowledge that they have about some theme. These documents are then the authorities. Similarly, other documents may be the best places to go if one wants to find out where are the best authorities. These 'other' documents are hubs – they are like the friend that does not know so much about repairing cars, but does know where to find a good mechanic.

Kleinberg proposed that we view all Web pages as having – to some degree – both of these functions. We can quantify their two roles – hub and authority – by translating (3) and (4) into equations. We let $h(i)$ be the hub score of node i, and $a(i)$ its authority score. Then we relate the two as follows:

$$h(i) = (\text{const}) \times \sum_j A_{i\to j} a(j) = (\text{const}) \times (Aa)_i \quad (5)$$

$$a(i) = (\text{const}) \times \sum_j A_{j\to i} h(j) = (\text{const}) \times (A^{\mathrm{T}} h)_i \,. \quad (6)$$

Here we have used the symbol a to represent the vector of authority scores, while h is the vector of hub scores. We insert an undetermined constant (const) in each equation because, without it, a consistent solution is in general impossible – as we shall soon see. Equations (5) and (6) are simply the expression of (3) and (4), respectively, in mathematical form. For example, in (5), we estimate the hub score of page i by simply using the authority scores of all pages j that i points to – consistent with the verbal expression (3).

Equations (5) and (6) may be rewritten as

$$h = \left(\frac{1}{c_h}\right) Aa = \left(\frac{1}{|Aa|}\right) Aa \quad (7)$$

$$a = \left(\frac{1}{c_a}\right) A^{\mathrm{T}} h = \left(\frac{1}{|A^{\mathrm{T}} h|}\right) A^{\mathrm{T}} h \,. \quad (8)$$

Here (following Kleinberg) we choose the constant, in each case, to give a vector of unit length as a result. Kleinberg in fact presents the process of assigning authority and hub scores as an iterative application of (7) and (8), and shows that this iterative process converges. Calling the converged score vectors a^* and h^*, we find from (7) and (8) that these must be solutions to

$$\begin{aligned} h^* &= \left(\frac{1}{c_h^*}\right) Aa^* = \left(\frac{1}{c_h^*}\right) A \left(\frac{1}{c_a^*}\right) A^{\mathrm{T}} h^* \\ &= \left(\frac{1}{c_h^* c_a^*}\right) AA^{\mathrm{T}} h^* \end{aligned} \quad (9)$$

$$\begin{aligned} a^* &= \left(\frac{1}{c_a^*}\right) A^{\mathrm{T}} h^* = \left(\frac{1}{c_a^*}\right) A^{\mathrm{T}} \left(\frac{1}{c_h^*}\right) Aa^* \\ &= \left(\frac{1}{c_a^* c_h^*}\right) A^{\mathrm{T}} Aa^* \,. \end{aligned} \quad (10)$$

Thus we see that the hub and authority score vectors h^* and a^* must be *eigenvectors* of, respectively, the matrices

$$\boldsymbol{H} = AA^{\mathrm{T}} \quad (11)$$

and

$$\boldsymbol{A} = A^{\mathrm{T}} A \,. \quad (12)$$

We note that both the Hub matrix $\boldsymbol{H}$ and the Authority matrix $\boldsymbol{A}$ are symmetric matrices. Also, it is clear from

our discussion of Eqs. (5) and (6) that right multiplication by the matrix A sends weights "backwards" – i. e., against the arrows, hence towards the pointing node – while right multiplication with A^{T} sends the scores forwards – with the arrows, towards the pointed-to node.

We can see this by examining a typical matrix element of (for example) $\boldsymbol{H}$:

$$\boldsymbol{H}_{ij} = \left(AA^{\mathrm{T}}\right)_{ij} = \sum_k A_{ik}A^{\mathrm{T}}_{kj} = \sum_k A_{i\to k}A_{k\leftarrow j}\,. \quad (13)$$

Thus, when $\boldsymbol{H}$ acts (by right multiplication) on a vector h of candidate hub scores, it will send scores from a node j, first forward to j's outlinks k (which are thus candidate authority nodes), then backwards to these nodes' inlinks – for example, i – which is thus rated for its role as a hub. In short, $\boldsymbol{H}$ allows nodes to send one another (via two hops) their hub scores; and a node j with a high hub score (which thus "thinks" it points to good authorities) will send (via action of $\boldsymbol{H}$) larger hub weight to other nodes i which point to the same nodes (good authority nodes) as j does. This sending is iterated until convergence – at which point all nodes agree on their respective hub scores. The same logic applies to authority scores; one simply exchanges, in the above discussion, authority ↔ hub, forward ↔ backward, pointing ↔ pointed-to, etc.

As noted after Eqs. (9) and (10), the converged scores h^* and a^* are eigenvectors of (respectively) $\boldsymbol{H}$ and $\boldsymbol{A}$. However, we know that these matrices, for a Web graph with N nodes, will be $N \times N$ matrices, which may have as many as N distinct eigenvectors [38]. Thus we must ask: which eigenvector do we want?

The answer is simple. As noted by Kleinberg, the iterative process of sending scores will converge (almost always) to a unique vector, called the *principal eigenvector* of the matrix in question ($\boldsymbol{H}$ or $\boldsymbol{A}$). The principal eigenvector is that corresponding to the largest eigenvalue. For nonnegative matrices such as $\boldsymbol{H}$ or $\boldsymbol{A}$ (and also for A), the largest eigenvalue is positive, and the principal eigenvector is nonnegative (positive or zero) at every node Theorem 2.7 in [83]. In fact, as long as the graph represented by $\boldsymbol{H}$ or $\boldsymbol{A}$ is connected, the principal eigenvector will be strictly positive at each node – because these matrices are symmetric. Hence we can expect a positive, nonzero hub and authority score for each node as a result of this process.

Kleinberg designed the HITS method to work on subgraphs of the WWW. That is, he assumed that one is given (from some search engine) a hit list that is large, because the topic of the search is not very narrow. The hit list likely includes many good authorities, but also many uninteresting pages. To find the good authorities (and hubs) for the given topic, he proposed to take the top-ranked t pages from the hit list (thus forming the *root set* of pages), and then to augment this root set by including (i) all pages pointed to by the root set, (ii) some or all pages pointing to the root set, and (iii) all links internal to the resulting, augmented, *start set*. Kleinberg argued that the resulting start set should contain most of the good authorities for the topic in question. Because of this property of using hyperlinks to 'zero in' on a topic, the method acquired the name Hypertext-Induced Topic Selection or HITS.

We note here that Kleinberg's choice of using a topic-focused subgraph – built up from a hit list – is not logically necessary for the application of the mathematics of the HITS method, as defined by Eqs. (9) and (10). We can think of this choice (subgraph taken from a hit list) as a *'real-time'* approach, since it requires that LA be performed *after* the generation of a hit list (thus, in response to a query, and so in real time). Alternatively, one can apply the HITS method (as is done with PageRank, see below) to the *entire* Web graph. When the whole graph is used we call the approach an *'offline'* method, i. e., one which is carried out in background mode, and which gives query-independent results. Thus, the reader should keep in mind that any method for LA (mapping a directed graph to a score) may be applied either as a real-time or as an offline approach.

The key ideas of HITS are by now (we hope) clear:

(i) Define the quality of Web pages in a circular way – that is, in terms of the quality of other Web pages.
(ii) Use the link structure of the Web graph to express these definitions as equations.
(iii) The resulting equations call for the principal eigenvector of a matrix; this eigenvector gives the desired quality scores.

These three basic steps underlie all of the methods we will discuss, beyond the most basic method of link popularity.

PageRank

Now we discuss a method which was originally devised to be used on the entire Web graph. The PageRank method [68] gives a single score for each Web page – its PageRank. This score is best regarded as a kind of authority score, in that it is generated (like link popularity) from inlinks.

We begin with the basic concept as expressed in (2). However, we will only consider one type of "goodness" or quality score; hence we need not expand (2) to (3)+(4).

Instead we express (2) mathematically as:

$$p_i = \sum_{j \to i} \left(\frac{1}{k_j^{\text{out}}} \right) p_j = \sum_j \left(\frac{1}{k_j^{\text{out}}} \right) A_{ij}^{\mathrm{T}} p_j$$
$$= \sum_j \left(A^{\mathrm{T}} \right)_{ij}^{(\text{col})} p_j \,. \tag{14}$$

The first equality in (14) tells us that – as prescribed by (2) – node i gets its PageRank score p_i from all nodes j pointing to i. However, it also says that nodes j do not send their full quality score (weight) to i; instead, sending nodes j must divide their weight by the number of outlinks k_j^{out} which they have. Thus, the PageRank approach, as expressed in (14), is rather "democratic": every page can only "vote" once (using its current score). That is, a good page cannot give its full weight to each of its outlinks, because the *total* weight it sends must be conserved – hence, divided up among its outlinks. This is expressed in another way by saying that the forward-sending matrix A^{T} is *normalized*. More precisely, $(A^{\mathrm{T}})^{(\text{col})}$ is a *column-normalized* version of A^{T}, since the entries in each (nonzero) column are adjusted so as to sum to one.

Note that there is no constant prefactor in (14). The reason is that we know the constant before we do any calculation; and it is one. We can (as before) reformulate (14) into an eigenvector equation:

$$p = \left(A^{\mathrm{T}} \right)^{(\text{col})} p \,; \tag{15}$$

but in this case we know that the (principal) eigenvalue is 1 – because of the normalization property of $(A^{\mathrm{T}})^{(\text{norm})}$ [13].

On the face of it the PageRank approach (14)–(15) appears to be simpler than the HITS approach: just one equation, giving one set of scores p. This simplicity has a price however.

We recall that the HITS approach gives a positive score (actually, two positive scores) to every page, as long as the effective graph ($\boldsymbol{H}$ or $\boldsymbol{A}$) is connected. This nice positive definite property is guaranteed for the principle eigenvector of any strongly connected graph, and hence comes from the fact that the effective links (in $\boldsymbol{H}$ or $\boldsymbol{A}$) are symmetric. PageRank uses the simpler operator $(A^{\mathrm{T}})^{(\text{col})}$,which is *not* in general symmetric. This means that, even given a connected Web graph, one can expect in general exactly zero weight for a large fraction of all the pages. This is an undesirable outcome: it means that, for the set of nodes with exactly zero weight, (i) one is forced to regard them as having zero quality, and (ii) LA provides one with no basis whatsoever for comparing them.

One can in fact show [16,17] that all pages in the Web graph – except those lying in sink SCCs – will get zero weight in the principal eigenvector of $(A^{\mathrm{T}})^{(\text{col})}$. This fact agrees with our intuition: all SCCs lose weight to downstream SCCs, except sinks (which have no downstream SCCs). Thus we have come up against the 'sink problem' mentioned earlier – and in a rather severe form that clearly justifies its name.

Brin and Page [19] offer their own solution to this sink problem. The normalization of their forward operator allows them to think of the score p_i as a *probability* – more precisely, the probability that a random walker, following outlinks from one page to another, will (in the limit of many hops) visit page i. They then allow this "random Web surfer" to hop (again at random) from any page to any other, with probability ε/N. To avoid violating weight (i. e., probability) conservation, they correspondingly adjust the probability of following an outlink from j to be $(1-\varepsilon)(1/k_j^{\text{out}})$. In terms of matrices, this amounts to adding an all-to-all 'random surfer matrix' (or 'teleportation matrix') $\hat{R}$ to the given matrix $(A^{\mathrm{T}})^{(\text{col})}$, and weighting the sum to retain probability conservation. The PageRank eigenvalue equation then becomes

$$p = \left((1-\varepsilon) \left(A^{\mathrm{T}} \right)^{(\text{col})} + \varepsilon \hat{R} \right) p \,. \tag{16}$$

The all-to-all matrix $\hat{R}$ (with $\hat{R}_{ij} = 1/N$) makes the effective network which is represented by the sum $((1-\varepsilon)(A^{\mathrm{T}})^{(\text{col})} + \varepsilon \hat{R})$ a single SCC. Hence there are no sinks in this effective network, and thus no sink problem: all entries in p are positive definite.

We note that the term $\hat{R}p$ in (16) can be simplified: $\hat{R}p = (1/N)\left[\sum_i p_i \right] \mathbf{1} \equiv v$ (where $\mathbf{1}$ is a vector of ones). That is, it gives a vector v (the 'teleportation vector') which is known, because $\sum_i p_i = 1$ is known. Because the teleportation matrix is completely unbiased, the teleportation vector is also unbiased, or 'flat'. We will look at applications for a biased teleportation vector in Sect. "Personalization of Link Analysis Scores".

We note that the factor $(1-\varepsilon)$ in (16) is often called the *'damping factor'*. This reflects the fact that a common approach to computing PageRank scores is the *power method* [48] – in which one finds the principal eigenvector of a matrix by repeating the operation (matrix) × (vector) until convergence. One finds that convergence is faster for smaller values of the damping factor [55]; this is related to the fact that the eigenvalue gap (discussed below) increases with ε, and hence with decreasing damping factor. Brin and Page [19] report obtaining good convergence after 50–100 iterations of the power method, with $\varepsilon = 0.85$. Langville and Meyer [55] discuss the practical consequences (in terms of score accuracy) of these choices.

Kleinberg [52] also used the power method to find hub and authority scores. Here – because the principal eigenvalue is larger than 1 – it is necessary to periodically rescale the length of the converging vector, as one iterates (matrix) × (vector); otherwise the vector's length will grow arbitrarily large.

SALSA

The 'SALSA' algorithm (Stochastic Approach for Link-Structure Analysis) was investigated by Lempel and Moran [58]. In the SALSA approach, the normalized Hub operator is defined as

$$\widetilde{\boldsymbol{H}} = A^{(\mathrm{row})}(A^{\mathrm{T}})^{(\mathrm{col})} \tag{17}$$

and the normalized Authority operator is correspondingly

$$\widetilde{\boldsymbol{A}} = (A^{\mathrm{T}})^{(\mathrm{col})}A^{(\mathrm{row})} \,. \tag{18}$$

Here $A^{(\mathrm{row})}$ is the row-normalized (i. e., by indegree) adjacency matrix, and $(A^{\mathrm{T}})^{(\mathrm{col})}$ is the column-normalized transpose which we have seen before in the PageRank algorithm. The logic is thus (we hope) clear: SALSA seeks to combine the two-level approach of HITS with the normalization of PageRank. For example, (18) says that authority weights are sent via two hops: first backwards to hub candidates, using $A^{(\mathrm{row})}$ – with the row (indegree) normalization ensuring that each candidate authority can only use its full authority weight once in 'voting' for candidate hubs. Then $(A^{\mathrm{T}})^{(\mathrm{col})}$ allows these candidate hubs to 'vote' once for new candidate authorities, by dividing the weight they have gotten among their outlinks.

The SALSA approach clearly offers a new set of basic choices for link analysis. However – remarkably – the eigenvectors of the SALSA operators $\widetilde{\boldsymbol{H}}$ and $\widetilde{\boldsymbol{A}}$ can be proven [58] to give the same scores as simple link popularity. That is, to within vector normalization, the score of a node in the dominant eigenvector of $\widetilde{\boldsymbol{A}}$ is its indegree, and the score of a node in the dominant eigenvector of $\widetilde{\boldsymbol{H}}$ is its outdegree.

This result is very similar to a familiar result for undirected graphs. That is, it is known [62] that, when A is symmetric (as is the case for undirected graphs), the dominant eigenvector of the column-normalized $A^{(\mathrm{col})}$ is (except again for vector normalization) simply the node degree. Since column normalization gives weight conservation, one can also interpret the (vector normalized) scores as probabilities. Hence, this known result for undirected graphs says that, in the limit of a very long random walk, the probability of visiting each node is proportional to the node degree. Similarly, the SALSA results say that a random walk which always follows a forward hop with a backward hop (and vice versa) will (in the limit of very long time) reach a node after a forward hop with probability proportional to that node's indegree, and will reach a node after a backward hop with a probability proportional to that node's outdegree.

This result is even more remarkable in light of the fact that the naïve generalization of these ideas to the one-level, normalized case (PageRank) does *not* work. That is, it is not the case that, in the dominant eigenvector of $(A^{\mathrm{T}})^{(\mathrm{col})}$, each node's score is proportional to that node's indegree.

We can offer an intuitive understanding of all these results, based on a simple idea. Suppose we have a special directed graph satisfying the following constraint:

$$k_i^{\mathrm{in}} = k_i^{\mathrm{out}} \,, \quad \forall i \,. \tag{19}$$

That is, every node in the graph has (number of inlinks) = (number of outlinks). Now we show that, for this graph, the vector $\vec{k} = \vec{k}^{\mathrm{in}} = \vec{k}^{\mathrm{out}}$ is the dominant eigenvector of $(A^{\mathrm{T}})^{(\mathrm{col})}$. Suppose that each node i has weight k_i at some step of the weight passing process. We then apply $(A^{\mathrm{T}})^{(\mathrm{col})}$ to send weights again; this makes each node send out exactly one unit of weight over each outlink. At the receiving end, each node j receives exactly $k_j^{\mathrm{in}} = k_j$ units of weight. Thus the weight distribution is unchanged by the weight passing – it is an eigenvector. That is, the node degree gives the eigenvector (hence LA score), for this special graph, for the PageRank operator without teleportation. We then note that every undirected graph always satisfies Eq. (19) – hence we see that the node degree is also the principal eigenvector for a normalized symmetric graph. Finally we consider the SALSA approach. Suppose every node i has weight k_i^{in} after a Forward iteration. Then the next hop (Backward) will send exactly one unit of weight (Backward) over each of i's inlinks – for every node i – so that each receiving node j receives exactly k_j^{out} units. It is easy to see that the next Forward operation takes us back to our starting point. Hence again we get the (in- and out-) degree as the eigenvector of (respectively) $\widetilde{\boldsymbol{A}}$ and $\widetilde{\boldsymbol{H}}$.

T-Rank

The three approaches presented above (HITS, PageRank, and SALSA) are by far the most well studied approaches to Web link analysis. Now we attempt to place all three (and a new one) in a simple, common picture.

HITS and SALSA were first proposed as a *postprocessing* (real-time) form of LA: one first generates a hit list using conventional text relevance analysis, and then generates and analyzes a topic-focused subgraph from this hit list. PageRank, in contract, is a preprocessing, 'offline' form of LA, in which the entire Web graph is analyzed,

and the resulting LA (PageRank) scores are *search-independent* – they are 'global'.

Hence we have a fundamental choice to make in choosing a LA approach: global and offline, or postprocessing on a topic-focused subgraph. This choice clearly has important implications, both in terms of the results obtained from LA, and in terms of practical considerations. In this review we will be essentially 'agnostic' with regard to this choice. That is, we focus primarily on the mathematical structure of the LA algorithm, and on the consequences of that structure, without addressing in detail the question of whether the algorithm is applied to a subgraph or to the entire Web graph.

The mathematical structure of HITS is simply stated. The approach is 'two-level' – in that each type of score (authority and hub) is passed from node to node in two hops – and it does not use matrix normalization. SALSA is also two-level, but uses matrix normalization, so that the weight passing is weight conserving. Finally, PageRank is a one-level approach – authority scores are passed directly from a node to its neighbors, in a single hop – and it uses matrix normalization.

Thus, the fundamental mathematical choices distinguishing these three approaches boil down to two binary choices: one- or two-level, and normalized or not.

We summarize this very simple analysis in Table 1. An alert reader will note that two binary choices generate four possible approaches; and we have included, in Table 1, a very recently studied approach (*'T-Rank'*) in the upper right-hand box [15,16]. T-Rank is a non-normalized, one-level approach. In other words, the vector of T-Rank scores t is taken to be the principal eigenvector of A^{T}:

$$A^{\mathrm{T}} t = \Lambda t \,, \tag{20}$$

where Λ is the principal eigenvalue of A^{T}. T-Rank has also been studied with the teleportation sink remedy, in which case we have

$$\Lambda' t = \left(A^{\mathrm{T}} + \varepsilon \hat{R}\right) t \,. \tag{21}$$

Here there is no need for a factor $(1 - \varepsilon)$, since there is no attempt to normalize the composite matrix $(A^{\mathrm{T}} + \varepsilon \hat{R})$.

Link Analysis and Web Search, Table 1
A simple tabulation of some fundamental choices in link analysis, including a recently studied approach (T-Rank) which seeks to fill the 'hole' in the upper right

	Normalized	Non-normalized
1-level Forward	PageRank	T-Rank
2-level (compound) H and A	SALSA	HITS

The T-Rank approach has received very little study, especially in comparison to the 'big three'. However we see that it is, in an a priori sense, equally worthy of study: a thorough understanding of link analysis should include an understanding of all the options in Table 1. Hence we point out here the existence of T-Rank – both logically and in [16]. The latter paper includes a preliminary comparison of all four entries in Table 1; a principal (tentative) conclusion from this comparison is that changing columns has a more profound effect on the results than does changing rows. In other words, the choice of whether or not to normalize the weight passing has a stronger effect on the nature of the eigenvector than does the choice of one- versus two-level weight passing. These effects were quantified in [16] in terms of the strength of the 'tightly-knit community' (TKC) effect [58] – in which a relatively small elite of mutually reinforcing pages dominates the scores – and in terms of the tendencies of the scores to correlate strongly with node degree. More specifically, the right column of Table 1 (non-normalized approaches) tended to give a stronger TKC effect, while the normalized approaches (left column) tended to show a stronger correlation of LA scores with node degree (which was of course complete for SALSA).

Another benefit of seeing LA in terms of Table 1 is that one can begin to imagine *intermediate* approaches. For example, the choice to normalize or not need not be strictly binary – one can instead imagine, in the top row of Table 1, a continuous tuning of the degree of normalization, from full normalization (PageRank) to no normalization (T-Rank). A simple (non-unique) interpolation formula which accomplishes this is given in [16]. Here it is also shown that the strength of the TKC effect may be continuously varied, by varying the strength of the normalization.

Of course, the picture given in Table 1 is highly simplified, and the reality is not so tidy. We mention for example the study of Ding et al. [31], who also placed the big three approaches in a common framework. This framework however was distinct from that of Table 1, and T-Rank was not mentioned. The tendency of LA scores to correlate with node degree was also noted here.

Another feature which is missing from the Table is the choice of Forward or Backward propagation. This choice is irrelevant for the two-level approaches (which use both types of propagation, but in alternation), but it is significant for the one-level approaches. Ding et al. [31] mention using one-level backward propagation in order to generate a kind of hub score without resorting to a two-level approach. Mathematically, this involves simply using the adjacency matrix A – normalized, or not – rather than its

transpose. This one-level hub-score approach has received little study. Mathematically, it is the same as the authority approach, since a matrix and its transpose are distinguished only in the details. In practice, of course, with the real Web graph generating the matrix, the choice of forward or backward propagation is significant.

The 'Sink Problem'

We have argued above that, in some cases, the sink problem can have a strong and undesirable effect on LA scores. We revisit this question in this section, examining each of the approaches discussed in Sect. "Approaches to Web Link Analysis".

HITS and SALSA

We can discuss HITS and SALSA simultaneously here, because each of these approaches solves the sink problem in the same fashion – i. e., by building an effective, two-hop graph which is symmetric. Symmetric graphs have no sinks or sources, hence no sink problem.

This solution is not necessarily problem-free however, because the resulting effective graph may be disconnected [66]. This can happen if, for example, two subgraphs are not reachable, one from the other, via a forward hop followed by a backward hop (i. e., by application of the Hub operator). The possibility of the effective graph being disconnected is discussed in [58]; here there is offered a simple solution, with weights distributed among the disconnected components according to the number of nodes in each component.

PageRank

The PageRank approach as expressed in Eq. (16) has a sink remedy built in; and typically, by the term 'PageRank' one means this 'remedied' combination of normalized forward weight passing – using $(A^{\mathrm{T}})^{(\mathrm{col})}$ – and teleportation – using $\hat{R}$. In the absence of $\hat{R}$ (i. e., with $\varepsilon = 0$), the PageRank approach has a severe sink problem, as noted in [8,14,17], and [16]. This may be understood rather simply: the PageRank eigenvector (with $\varepsilon = 0$) represents the asymptotic (in time) probability distribution for a random walker which moves over the directed links without help of the random-hop option. Clearly, such walks become 'stuck' in the sinks in the limit of large time. Hence, in the absence of any sink remedy, the eigenvector will have zero weight everywhere on the graph except in sink SCCs.

For this reason, PageRank is only practical for nonzero ε. Much effort has been expended [55] towards the goal of finding the 'best' value for this parameter. Clearly, a compromise must be taken between two extremes, neither of which is attractive: for $\varepsilon = 0$, all weight is in sinks, while for $\varepsilon = 1$, the eigenvector is uniform and carries no information – all nodes are equally important. It is perhaps remarkable that eigenvectors for values of ε lying between these two meaningless extremes can in fact give highly useful and meaningful information about the status of Web pages in the Web graph.

The sink problem impacts PageRank in another way. We consider a 'dangling page' D which has only inlinks and no outlinks. This page D is thus a single-node sink SCC. Furthermore, as long as there is at least one such dangling node in the entire Web graph, the graph's adjacency matrix is not normalizable: page D is represented by a column of zeroes in A^{T}, and this column cannot be made to sum to one in $(A^{\mathrm{T}})^{(\mathrm{col})}$. Such dangling nodes represent (in the language of Bianchini et al. [14]) a source of 'energy loss': weight which is sent to them simply vanishes in the next iteration. Dangling nodes do occur in the Web graph, so that the problem is not academic. One idea [68] is to remove them; this idea has the corresponding problem that removing some dangling nodes may in turn create new ones. A common solution is to implement a 'back button': one allows for the possibility (with suitable normalization) of a hop to *any* other node from a dangling page D – which thus ceases to be a sink. With the back-button solution in place, there are no zero columns in the modified A^{T}, and so normalization is possible.

T-Rank

As noted above, the non-normalized, one-level approach has received little study. Some limited results may be found in [16]. Here it is shown that the sink problem is much less severe when the normalization constraint is dropped, in that the eigenvector – *without* sink remedy – has positive weight in the GSCC, and in the OUT set. This is in fact guaranteed to be the case, given the (almost certainly true) assumption that the GSCC has the largest eigenvalue of all SCCs. Also, numerical results are presented which suggest that adding a random teleportation matrix with weight ε has little effect on the eigenvector. In short: there is no evidence that $\varepsilon = 0$ is a singular point for T-Rank, as it is for PageRank.

It is also shown in [16] that two novel types of sink remedy are possible for the non-normalized approach (and that neither is possible with the normalized approach). One of them, termed 'source pumping', has the unique property that it is the *only* way to ensure a nonzero eigenvector over the entire graph, without making the graph strongly connected. We note that sink remedies are

of interest for T-Rank because, without some such remedy, the IN set, TENDRILS, and TUBES have exactly zero weight.

Sensitivity of Link-Analysis Scores

There are at least three real-world problems which complicate the application of link analysis to the Web. First, there is the fact that one is always working with incomplete information – it is impossible to crawl the entire Web. Secondly – in the case of one-level approaches – one normally applies a sink remedy such as the teleportation operator. This changes the (incompletely) known Web graph into another. Third, there is the very real likelihood that many of the links and even pages in the Web graph are placed with intent to influence the link analysis score – that is, there is "link spam".

It is thus important to examine how various link analysis approaches are affected by these three kinds of perturbations. We take them up in this Section, in order, starting with the incomplete information.

Sensitivity to Small Topology Changes

The Web graph which is input for LA differs from the "real" Web graph in that it lacks many existing pages and links. Furthermore, this input graph has links and pages which no longer exist, due to imperfect updating. Hence we must ask, how important is this source of error? More specifically, how much does the change from the "real" graph W to the measured graph W' affect the scores and rankings obtained from LA?

Early work in this direction was done by Ng et al. [66,67]. Here HITS and PageRank were examined for stability against small topology changes. In [66] it was shown that the change in PageRank scores is bounded above by $2\Sigma/\varepsilon$, where Σ is the sum of PageRank scores of the pages suffering the modification, and ε is the teleportation weight. Thus, for "reasonable" ε and for perturbations avoiding the most important pages, the bound is encouraging. In the same paper, the stability of HITS was related to its eigengap – the difference between the principal eigenvalue and the second largest. Experimental results implied that HITS was more sensitive than PageRank (with $\varepsilon = 0.2$). Thus we see a stabilizing effect from the teleportation operator. Motivated by this, [67] proposed a "randomized HITS" approach (HITS plus random teleportation), and gave evidence that it is more stable than the original version.

In [57], and in [14], the bound of Ng et al. is tightened to $2(1-\varepsilon)\Sigma/\varepsilon$. Also, [57] studied the stability of SALSA, finding it stable as long as the effective graph is connected, and proposed a stabilized form of SALSA, using (again) random teleportation.

Borodin et al. [18] studied similarities between, and stability of, HITS, variations on HITS, and node degree (INDEGREE), in a very general framework, and found that almost all algorithms (including HITS) are unstable by their definition and assumptions. Donato et al. [32] examined again the stability of HITS, for a class of graphs termed product graphs, and found that (under some assumptions) HITS scores were both stable, and very close to node degree (authority $\approx$ indegree). The product graph model allows for the possibility of realistic structures such as a power law.

Lempel and Moran [58] point out that stability of ranking is not equivalent to stability of scoring – since, in many cases, small changes in scores can lead to large changes in rank. They found further that SALSA is rank stable on the class of graphs for which it is scoring stable, while PageRank and HITS are not.

We emphasize here that we have not sought to give an exhaustive guide to the literature in this section. There exist many definitions of stability, and many possible perturbations. However, some simple facts are clear. The eigengap for PageRank can be tuned by tuning ε, and is in fact equal to ε for uniform teleportation [44]. Thus a stable behavior can be practically ensured for PageRank. Basic HITS in contrast has no such tuning parameter and so cannot guarantee a large enough gap – unless of course teleportation is also added here. SALSA is stable in the sense that it is local – so that local topology changes cannot affect the entire eigenvector. Finally, T-Rank has not been studied in regard to this question.

Sensitivity to Sink Remedies

Now we consider the sensitivity of PageRank to variation of the teleportation probability ε. Here we find a clear and simple picture: PageRank scores become unstable as one approaches the 'un-remedied' limit $\varepsilon \to 0$. This conclusion is already implicit in the results of Ng et al., cited above. In addition, Langville and Meyer [55] pointed out that the eigengap for PageRank vanishes as $\varepsilon \to 0$. A bound for derivatives of PageRank scores with respect to ε was presented; this bound blows up as $\varepsilon \to 0$. Boldi et al. [17] found closed forms for all orders of derivatives for PageRank scores, and graphically demonstrated the instability of the scores as $\varepsilon \to 0$. Similar results were found by Avrachenkov et al. [8].

All of these results are consistent with the fact (discussed above) that the PageRank distribution is pathological at $\varepsilon = 0$ – along with the well known fact that

it gives sensible results for ε not too small. The distribution of scores is also pathological (uniform) as $\varepsilon \to 1$. As expressed in [17]: "… PageRank oscillates between a meaningless uniform distribution [$\varepsilon \to 1$] and a meaningless distribution concentrated mostly in irrelevant nodes [$\varepsilon \to 0$]." This makes it all the more remarkable that PageRank gives, seemingly, stable and useful results when ε is held away from these limits.

Sensitivity to Spam

Because getting noticed has a monetary value, and because link analysis (applied to Web search) helps determine which pages are noticed, it is also true that understanding (and being able to manipulate) link analysis is worth money. It is also a good research problem. For these two reasons, the related arts of committing and of fighting *link spam* have received a great deal of attention.

The problem was discussed already in Sect 6.1 in [68] – in which use of a nonuniform teleportation vector was proposed. In another early work, Davison [30] discusses using machine learning to classify links and to recognize "nepotistic" (spam) links.

A fairly standard model for link spam was given in Gyöngyi and Garcia-Molina [40]. One has a set of *target pages* for which it is desired to boost the scores, and a set of *boosting pages* whose main or only function is to help the target pages. Also, the resulting *spam farm* may in addition exploit the benefit gained from laying in *hijacked links*. These are links from "real" (external to the spam farm) Web pages, but they are pages of a type in which the spammer is able to place outlinks pointing to the spam farm. (Examples of hijackable pages are unmoderated discussion fora and directories.) This work found (assuming PageRank for LA) optimal structures for a spam farm with a single target page and k boosting pages, and also for an alliance of two such spam farms. Bianchini et al. [14] presented similar results, but using a more general model for a spam farm. Du et al. [35] present a modification of the simple optimal rules of Gyöngyi and Garcia-Molina [40], obtained by relaxing an assumption of the earlier authors.

Gyöngyi and Garcia-Molina [41] presented a useful taxonomy of types of Web spam. (Note the name of the conference in this reference, underscoring the coming to maturity of the field!) They also presented some experimental data indicating that (i) about 7–18% of Web pages, in the years 2002–3, were spam pages, and (ii) this fraction seemed to be growing over time. In addition, there is a limited discussion of spamming HITS, in which it is clear that boosting a hub score is easier than boosting an authority score (as it is easier to choose one's outlinks than one's inlinks). An earlier paper [42] proposes the TrustRank algorithm for combating spam. Here, reputable "seed" pages are used to bias the teleportation vector in a PageRank approach, thus allowing the trustedness of the seeds to propagate throughout the distribution. (A version of TrustRank may apparently be tested live at http://trustrank.org/.) A similar method is described in Benczúr et al. [11]. Here, suspicious pages are recognized by their non-typical distribution of inlinks, and are then penalized by (again) biasing the teleportation vector. The difference here is that one punishes "bad" pages, rather than rewarding "good" ones. Taking this idea even further, one can compute a "BadRank", based on the principle that pages pointing to "bad" pages cannot be very "good". BadRank is thus a reverse propagation of PageRank, with seeds in the teleportation vector being known bad pages. BadRank has appeared on the Web in the form of rumors (http://pr.efactory.de/e-pr0.shtml), and has been studied by Liang et al. [61].

Zhang et al. [89] studied spam in the form of *collusion*: coordinated link structure manipulation by a group of targets (similar to the alliances mentioned above, but more general). For some simple models of collusion, it was found that the boost to be expected (in PageRank) is proportional to $(1/\varepsilon)$ – consistent with the early result of Ng et al. [66,67]. This suggested in turn a novel spam recognition mechanism: colluding nodes should be recognizable by an abnormally high sensitivity of their scores. The proposed remedy is again to bias the probabilities in the teleportation vector.

Another variant of spam is termed 'link bombing' or 'Google bombing' [2]. Here there is normally a coordinated use of anchor text, in order to boost the ranking of a target with respect to the anchor text in the boosting links. A famous example is the boosting of a web-biography of the US President with respect to the text "miserable failure". Although the purpose here is negative, clearly the same mechanism can be used to give boosts that the target would desire. Adali et al. show that uncoordinated individual attacks give optimal bombing.

The spam farm described above (with boosting pages and target pages) falls under the more general category of 'sybil strategies' [58]. In a 'sybil attack' on a *reputation system*, an individual entity masquerades as multiple simultaneous identities, in order to enhance its own reputation. Cheng and Friedman [26] show that no symmetric reputation function can be sybilproof. Here 'symmetric' means invariant under a relabeling of the agents (nodes, Web pages). This result is intuitively reasonable, since 'sybilproof' is a strong criterion (unlimited duplication is allowed), and since symmetric reputation systems (such as

PageRank) are essentially anonymous – i. e., dependent solely on the link structure. Cheng and Friedman [27] examine the manipulability of PageRank by sybil attacks, and find bounds on the attainable boost which are linear in the number of sybils – and hence consistent with the results of Garcia-Molina [40] and Bianchini et al. [14].

A recent paper [10] examines the possible gains in spam fighting by combining a variety of techniques. Included among these is a new one, Truncated PageRank, in which the normal damping factor $(1 - \varepsilon)$ is set to zero for neighbors within T hops. This approach is motivated by the fact that most spam links are close to the target page of the spam farm; also, it is consistent with the finding of Zhang et al. [89] that spam pages are more sensitive to damping factor. The authors report a spam detection rate of 80%, using the best found combination of (many) methods. This rate is improved further in the work of Castillo et al. [23] – to over 88% detection, at a cost of about 6% false positives. Castillo et al. use a machine-learning approach which combines both link attributes (in fact, all of the attributes used by Becchetti et al. [10]) and attributes based on examining the *content* of pages. Regarding the link attributes, Castillo et al. report evidence supporting the idea that spam and non-spam pages are not well interlinked: they find that spam pages are mostly linked to by other spam pages, while non-spam pages are mostly linked to by other non-spam pages. This statistical tendency was used to give a significant improvement in spam detection.

Clearly, Web link spam has received a great deal of attention. Much of this is no doubt based on the perception that Google's PageRank is and remains an important factor in its (highly successful and influential) approach to ranking. There is no sign of a "silver bullet" against link spam; and it seems likely that the fight between link spammers and anti-link-spam warriors will continue for as long as Web link analysis is useful.

Personalization of Link Analysis Scores

Web search engines, and in particular the approaches to Web link analysis discussed so far, have a "one size fits all" approach to ranking search results. Obviously, it is of interest to be able to skew the rankings, based on the preferences of the individual user. In this section we discuss some approaches to personalizing Web link analysis. We note that the teleportation vector is also commonly known as the "personalization vector", because it presents an obvious and simple mechanism (already pointed out in [68]) for skewing scores obtained from the PageRank algorithm. Also, we note that, since the skewing is normally expressed in terms of the content of the preferred pages, personalization is often equivalent to topic-sensitive scoring.

Rafiei and Mendelzon [74] propose methods – one a modification of PageRank, and one a modification of HITS – for assessing each page's reputation with respect to a topic. In each case, a biased random walk is implemented; in the PageRank case, the modification is small, while for HITS, one modifies the distribution of Forward hops to bias for the topic of interest. Chang et al. [25] also modify HITS, but in a different way. Here one or more preferred pages are given as a target, and the weights of the adjacency matrix are modified by adding the gradient towards the target.

A related (link reweighting) approach was studied in Richardson and Domingos [75]; here both the teleportation links and the hyperlinks of a PageRank approach are reweighted according to their relevance to a given target query (topic). Hence one obtains Query-dependent PageRank. It is proposed to precompute such scores for all one-word queries in the lexicon; and some discussion is given of the resulting scaling problems. Haveliwala [43] proposes instead to precompute such scores for a chosen set of "basis topics", and then to resolve any query in terms of these basis topics (16 in number in the implementation). Jeh and Widom [46] offer a more sophisticated approach in the same spirit, in which a set of partial vectors are precomputed, while the basis vectors are efficiently reconstructed as needed, such that the effective number of basis vectors can be large (tens of thousands). Haveliwala et al. [45] present a nice, concise comparison of the latter two approaches [43,46], along with the 'BlockRank' approach of [49].

The work by Fogaras et al. [36] offers a solution which allows for 'full personalization', i. e., personalization for *any* possible interest profile. This seemingly intractable goal is achieved by providing only an approximate solution, but for a complete set of basis vectors. The latter are generated by performing N random walks starting from each page in the Web graph, and the approximation is due to using only a finite N. The resulting approach requires storage which is linear in the size of the Web graph. Fogaras et al. found good precision for their method on a graph with 80 million nodes, using $N = 1000$.

We see in the above approaches a variety of ways to compute personalized (i. e., topic-sensitive) link analysis scores. There is no lack of ideas, and some good approaches to the scaling problem have been offered. We are not however aware of any commercial activity along these lines.

Another approach to personalizing Web search is to *decentralize* it. This allows for the possibility of local views

of the Web, biased by local interests, and represented by profiles which can act as guides for queries. We discuss this approach – with the scope limited to link analysis – in the next section.

Peer-to-Peer Search and Link Analysis

There are at least two motivations for considering distributed link analysis. One is the need to ameliorate the severe scaling problems – particularly those associated with crawling – that accompany a centralized approach to the vast Web. This could be done for example by spreading the load among Web servers. Another motivation is the desire to build a *peer-to-peer (P2P) search engine*, in which a distributed set of users share the work of crawling, indexing, query routing, and (not least) link analysis. These two scenarios are not necessarily exclusive. Furthermore, in the case of link analysis – and at the level of detail we present in this section – there is little difference between these scenarios; each boils down to seeking to compute global LA scores in a distributed fashion. A short and readable introduction to the world of P2P search may be found in Nejdl [64]; see also Part VI in Steinmetz and Wehrle [80].

Early work by Sankaralingam et al. [76,77] emphasized the need to solve the eigenvector equations (here PageRank) in an asynchronous fashion. This means in turn that convergence of iterative methods, such as the power method, is no longer guaranteed. In Sankaralingam et al. [77] it is pointed out that the PageRank system with its eigengap bounded from below satisfies the necessary and sufficient conditions for asynchronous convergence. The method employed was asynchronous iterated passing of updated weights (scores), with local estimates of convergence. We note that, in this work, it is assumed that all local nodes can find the hosts for every outlink that they store, by using hashed URLs as page IDs, and a structured distributed hash table assigning IDs to nodes. Shi et al. [79] present a similar approach, in which weights are handled differently according to whether they come from "outside" or "inside" the host. They show theoretically that the distributed system converges monotonically, and show experimentally that it converges to the same result as the centralized one. Chirita et al. [28] present a related approach, with the new feature that the distributed PageRank calculation is personalized.

Aberer and Wu [1,86] present a distributed approach, in which *rankings*, rather than scores, are aggregated from the local level to the global level. The local rankings are obtained using PageRank on the local graph. In Wu and Aberer [86], one first computes SiteRank (using the PageRank algorithm on the 'site graph'), and then (based partially on this) computes the local 'DocRank' for each document. The final global ranking for a document is obtained using the ranking aggregation method. In [87], it is shown that the distributed two-level approach converges to the same result as the equivalent centralized version. Wang and DeWitt [84] present a closely related two-level approach. Also, Zhu et al. [90] use a two-level approach, but based on iterative aggregation-disaggregation; they are able to prove convergence for this approach. We note that none of these works address the problem of asynchronous computation.

The 'JXP' method is presented in [69,70,71,72], and [73]. Here the peers do not initially know of any links or pages other than those they themselves host. They then learn about links to their own pages by random meetings with other peers. An iterated PageRank calculation is run, in which no other topology information is retained by any peer; instead, the rest of the Web graph is modeled using a 'world node' W. Monotonic convergence to the global PageRank vector is proven in [70], with the assumption that all peers know the total number of pages. [71] addresses the problem of cheating peers, while [72] show that informed (rather than random) strategies for meetings can enhance the rate of convergence.

Finally, we mention three, rather distinct, recent works. Jelasity et al. [47] present a general analysis of asynchronous distributed eigenvector calculations, including graphs whose spectral radius is larger than 1. They present methods for controlling the length of the iterated vector, and for estimating the principal eigenvalue ρ (when it is not known to be 1). While they do not prove convergence for the $\rho > 1$ case, they present numerical experiments showing good convergence, even in the face of severe noise (message delay and message loss).

Costache et al. [29] present a distributed ranking system for a collaborative work environment, in which users have privacy concerns which lead them to provide only incomplete information to other users. Furthermore, the link analysis (ObjectRank) deviates from PageRank in that links are weighted, based on metadata about the linked nodes. Several approaches are presented for handling partial information exchange. Interestingly, approaches using a "world node" are judged to give the best results in terms of respecting privacy without significant loss of accuracy.

Kempe and McSherry [50] present a decentralized approach for finding the k principal eigenvectors of a weighted symmetric matrix. The approach uses a decentralized form of orthogonal iteration in order to both normalize and orthogonalize the k "working vectors"; and a proof of convergence is given. Regarding application to Web link analysis, we note that PageRank and T-Rank rely

on asymmetric matrices, and also on some form of teleportation, so that this approach may not be used for these types of LA without further work. The HITS matrices are symmetric and weighted, so that the method may be applied directly to these matrices. One would however need a distributed protocol for computing the elements of these matrices – a not insuperable problem.

Thus we see that distributed LA – with main focus on PageRank – has been well studied, and a number of promising approaches have been identified. Perhaps the most interesting questions for this line of work are not rooted in the link analysis itself, but are rather about the future viability of decentralized search scenarios.

Future Directions

In this final section we offer some brief and speculative discussion about future directions. Many of the topics have appeared above; a few have not.

Personalization

We have seen that there exist nice research results pointing the way to personalizing Web search via personalizing link analysis scores. Google offers personalized search (see http://googleblog.blogspot.com/2007/02/personally-speaking.html), but it is not clear whether this personalization uses personalized LA or not. We note further that, whenever a machine is asked to read the mind of a user, there is some danger that the result is as much irritating as satisfying. Hence we offer this question: is there a future for personalized Web search? If so, will it involve personalized link analysis?

P2P

The arguments for peer-to-peer search are similar to those for other peer-to-peer systems: tapping the resources of billions of idling devices (technical), breaking free of centralized controls (political), and empowering the individual (humanistic). One vision is that each peer can represent a living, curious person, who has a personalized index based on her own interests, and a network of related peers based on related interests. In short, there are many motivations offered for the idea of P2P search. These visions have as their counterpart the seemingly innocuous and satisfying situation that users get what they want, instantly, and for free, with current centralized systems such as Google. Thus we find the same question as for personalization: is there a future for these visionary ideas? And, again, we expect that the answer will depend as much on human factors as on technical aspects.

T-Rank

The LA literature is dominated by the "big three". PageRank (as employed by Google) appears to have great commercial success. The HITS approach was developed in the context of the IBM CLEVER project (see http://alme1.almaden.ibm.com/cs/k53/clever.html and Chakrabarty et al. [24]), and stimulated a great deal of further work. Methods which build on the HITS approach are apparently in use [54] by the search engine Teoma (now Ask.com – see http://about.ask.com/en/docs/about/ask_technology.shtml). SALSA (as we have seen) gives results equivalent to scoring by in- or outdegree; the former (link popularity) is used by a number of search engines.

We have seen that placing the big three in a simple table suggests a fourth alternative, termed T-Rank. T-Rank appears to fit the same kind of technical niche as does PageRank – that is, it seems best suited for offline, global importance scoring. T-Rank has only recently appeared, and it is too early to say whether it represents a viable alternative to PageRank. Once the one-level non-normalized approach is considered, however, several new possibilities arise: interpolation – allowing for continuous tuning between full normalization and no normalization, with the resulting tuning of other properties – and novel sink remedies.

Mobile Search

The number of mobile devices is exploding, and has already surpassed the number of PCs. Mobile users also want information; but their priorities tend to be different from those of stationary users (more active, less patient), and the mobile devices themselves present new technical limitations and possibilities (small screen, lower bandwidth, GPS, etc.). How will Web search change as it goes mobile? And most importantly for our own theme, what will happen to link analysis? Will the "mobile Web" be simply a transcoding of today's Web, with hyperlinks intact? Perhaps the kinds of documents mobile users want to see will be so different from today's Web as to amount to a new Web – and one without hyperlinks.

Link Analysis Without Hyperlinks

All of our discussion thus far is based on the assumption that one has a Web graph which is built up from hyperlinks. There are many information domains which have few or no hyperlinks. Also, there is no guarantee that the present, unmanaged, system of collectively generated hyperlinks will persist in a form that makes link analysis useful. We need only consider the explosive growth of

user-generated content in the form of pictures and videos: whether they are placed in the 'visible Web' or not, they are unlikely to ever have the kind of rich hyperlink structure that has made LA work so well on the Web. Thus it is of interest to consider the idea of link analysis in the absence of hyperlinks. The beauty of link analysis is that it places a document in *context* – something that (we believe) is easier for a machine to do than it is to evaluate a document's *content*. Therefore, one should look at methods for placing a document in the context of the entire set of documents (or a chosen subset), by building links between the given documents and analyzing the resulting graph. An obvious choice is to use some kind of distance or similarity measure between pairs of documents; such measures give weighted, symmetric links – in contract to the unweighted, directed hyperlinks of the Web graph. A more elaborate approach, involving merging a Web graph (with hyperlinks) with a structured ontology, may be found in Stoyanovich et al. [81]. Here, the ontology and the Web graph together form a larger directed graph, to which a two-level authority propagation approach is applied; but the ideas here also have application in scenarios where the hyperlinks are few or absent.

Link Analysis Without the Web

Web link analysis has motivated and generated a great deal of understanding of, and insight into, the structure and properties of directed graphs. We believe that there exist possibilities for exporting of the ideas and methods presented here to fields which do not involve Web pages at all, but do involve directed graphs in a significant way. One possibility is the spreading of computer and mobile-phone viruses using address books: the graph on which the spreading takes place is directed, due to the one-way nature of such address lists [51,65]. One example of such cross-fertilization may be found in [22].

Another area in which asymmetric links play a role is in the related fields of social network analysis and innovation diffusion. It is common in these fields to treat links as symmetric, but in many cases this is far from realistic. Once asymmetry is taken into account, then most of the (eigenvector-based) notions of centrality need to be generalized. In such cases, HITS, PageRank, and T-Rank – both in the Forward or 'authority' sense, and in the Backward or 'hub' sense – might offer useful tools for quantifying and understanding the roles of the nodes in these networks.

Bibliography

Primary Literature

1. Aberer K, Wu J (2003) A framework for decentralized ranking in web information retrieval. In: Web technologies and applications. Lecture Notes in computer science, vol 2642. Springer, Berlin, p 594
2. Adali S, Liu S, Magdon-Ismail M (2005) Optimal link bombs are uncoordinated. In: Proceedings of the first International workshop on adversarial information retrieval on the web (AIRWeb) Chiba, May 2005, pp 58–69; available online at http://airweb.cse.lehigh.edu/2005/proceedings.pdf
3. Adamic LA (1999) The small world web. In: Lecture notes in computer science, vol 1696. Proceedings of the third European conference on research and advanced technology for digital libraries. Springer, London, pp 443–452
4. Aktas MS, Nacar MA, Menczer F (2006) Using hyperlink features to personalize web search. In: Advances in web mining and web usage analysis. Lecture notes in computer science, vol 3932. Springer, Berlin
5. Aktas M, Nacar M, Menczer F (2004) Personalizing PageRank based on domain profiles. In: Proceedings of the KDD workshop on web mining and web usage analysis (WebKDD) Seattle, August 22–25 2004. Springer, Berlin
6. Albert R, Jeong H, Barabasi AL (1999) Diameter of World-Wide Web. Nature 410:130–131
7. Altman A, Tennenholtz M (2007) An Axiomatic Approach to Personalized Ranking Systems. In: Veloso MM (ed) Proceedings of the 20th International joint conference on artificial intelligence (IJCAI) Hyderabad, January 6–12 2007. AAAI Press, pp 1187–1192
8. Avrachenkov K, Litvak N, Pham KS (2007) Distribution of PageRank mass among principle components of the web. In: Bonato A, Chung FRK (eds) Proceedings of the 5th International Workshop on algorithms and models for the web-graph (WAW) San Diego, December 11–12 2007. Lecture notes in computer science, vol 4863. Springer
9. Barabasi AL, Albert R (1999) Emergence of Scaling in Random Networks. Science 286:509–512
10. Becchetti L, Castillo C, Donato D, Leonardi S, Baeza-Yates R (2006) Link-based characterization and detection of web spam. In: Proceedings of the second International workshop on adversarial information retrieval on the web (AIRWeb) Seattle, August 10 2006, available online at http://airweb.cse.lehigh.edu/2006/proceedings.pdf
11. Benczúr A, Csalogány K, Sarlós T, Uher M (2005) Spamrank: fully automatic link spam detection. In: Proceedings of the first International workshop on adversarial information retrieval on the web (AIRWeb) Chiba, May 2005, pp 58–69; available online at http://airweb.cse.lehigh.edu/2005/proceedings.pdf, held in conjunction with WWW2005
12. Berberich K, Bedathur SJ, Weikum G, Vazirgiannis M (2007) Comparing apples and oranges: Normalized pagerank for evolving graphs. In: Williamson CL, Zurko ME, Patel-Schneider PF, Shenoy PJ (eds) Proceedings of the 16th International conference on World Wide Web (WWW) Banff, May 8–12 2007. ACM, pp 1145–1146
13. Berman A, Plemmons RJ (1979) Nonnegative Matrices in the Mathematical Sciences. Academic Press, New York
14. Bianchini M, Gori M, Scarselli F (2005) Inside PageRank. ACM Trans Inter Tech 5(1):92–128
15. Bjelland J, Canright GS, Engø-Monsen K (2008) Web Link Analysis: estimating a document's importance from its context. Telektronikk 1:95–113; also in http://delis.upb.de/paper/DELIS-TR-0631.pdf

16. Bjelland J, Burgess M, Canright G, Engø-Monsen K (2008) Eigenvectors of Directed Graphs and Importance Scores: Dominance, T-Rank, and Sink Remedies. DELIS technical report TR-0629, http://delis.upb.de/paper/DELIS-TR-0629.pdf
17. Boldi P, Santini M, Vigna S (2005) PageRank as a function of the damping factor. In: Proc. of the Fourteenth International World Wide Web Conference, Chiba. ACM Press, New York
18. Borodin A, Rosenthal JS, Roberts GO, Tsaparas P (2005) Link Analysis Ranking: Algorithms, Theory and Experiments. ACM Trans Internet Technol (TOIT) 5(1):131–152
19. Brin S, Page L (1998) The anatomy of a large-scale hypertextual Web search engine. In: Proceedings of the seventh international conference on World Wide Web. Comput Netw ISDN Syst 30(1–7):107–117
20. Brinkmeier M (2006) PageRank revisited. ACM TOIT 6(3): 257–279
21. Broder A, Kumar R, Maghoul F, Raghavan P, Rajagopalan S, Stata S, Tomkins A, Wiener J (2000) Graph Structure in the web. Comput Netw 33:309–320
22. Canright G, Engø-Monsen K (2007) A pumping model for the spreading of computer viruses. In: Bandara AK, Burgess M (eds) Inter-domain management. Proceedings of the first International conference on autonomous infrastructure, management and security (AIMS) Oslo, June 21–22 2007. Lecture notes in computer science, vol 4543. Springer, pp 133–144
23. Castillo C, Donato D, Gionis A, Murdock V, Silvestri F (2007) Know your neighbors: Web spam detection using the web topology. In: Proceedings of the 30th annual International ACM SIGIR conference on research and development in information retrieval (SIGIR) Amsterdam, July 23–27 2007. ACM, New York, pp 423–430
24. Chakrabarti S, Dom B, Gibson D, Kleinberg J, Kumar SR, Raghavan P, Rajagopalan S, Tomkins A (1999) Hypersearching the web. Sci Am 280(6)
25. Chang H, Cohn D, McCallum A (2000) Creating customized authority lists. In: Langley P (ed) Proceedings of the 17th International conference on machine learning (ICM) Stanford University, Stanford, June 29 – July 2 2000. Morgan Kaufmann, pp 127–134
26. Cheng A, Friedman E (2005) Sybilproof reputation mechanisms. In: Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of Peer-To-Peer Systems. Philadelphia, 22 August. P2PECON '05. ACM, New York, pp 128–132
27. Cheng A, Friedman EJ (2006) Manipulability of PageRank under Sybil strategies. In: Proceedings of the first workshop on the economics of networked systems (NetEcon) Ann Arbor, June 11 2006, available at http://www.cs.duke.edu/nicl/netecon06/papers/proceedings.pdf
28. Chirita P, Nejdl W, Scurtu O (2004) Knowing where to search: Personalized search strategies for peers in P2P networks. SIGIR Workshop on Peer-to-Peer Information Retrieval, Sheffield
29. Costache S, Nejdl W, Paiu R (2007) Personalizing PageRank-Based Ranking over Distributed Collections. Advanced Information Systems Engineering. In: Proceedings of the 19th International Conference, CAiSE 2007, Trondheim, 11–15 June. Springer, Berlin, pp 111–126
30. Davison BD (2000) Recognizing nepotistic links on the Web. In: Artificial Intelligence for Web Search. AAAI Press, Menlo Park, pp 23–28
31. Ding C, He X, Husbands P, Zha H, Simon HD (2002) PageRank, HITS and a unified framework for link analysis. In: Proceedings of the 25th Annual international ACM SIGIR Conference on Research and Development in information Retrieval, Tampere, 11–15 August. SIGIR '02. ACM, New York, pp 353–354
32. Donato D, Leonardi S, Millozzi S, Tsaparas P (2005) Mining the inner structure of the Web graph. In: Proc of the 8th WebDB. J Phys A Math Theor 41:224017
33. Donato D, Leonardi S, Tsaparas P (2005) Stability and Similarity of Link Analysis Ranking Algorithms. ICALP, Lisbon
34. Donato D, Laura L, Leonardi S, Millozzi S (2007) The Web as a graph: How far we are. ACM Trans Internet Technol 7:1
35. Du Y, Shi Y, Zhao X (2007) Using spam farm to boost PageRank. In: Proceedings of the 3rd international Workshop on Adversarial information Retrieval on the Web, Banff, 08 May. AIRWeb '07 vol 215. ACM, New York, pp 29–36
36. Fogaras D, Rácz B, Csalogány K, Sarlós T (2005) Towards Scaling Fully Personalized PageRank: Algorithms, Lower Bounds, and Experiments. Internet Math 2(3):333–358
37. Garfield E (1972) Citation analysis as a tool in journal evaluation. Science 178:471–479
38. Golub GH, Van Loan CF (1996) Matrix Computations, 3rd edn. Johns Hopkins University Press, Baltimore
39. Gulli A, Signorini A (2005) The indexable web is more than 11.5 billion pages. In: WWW 2005 Conference Proceedings, Chiba. ACM Press, New York, pp 902–903
40. Gyöngyi Z, Garcia-Molina H (2005) Link spam alliances. In: Proceedings of the 31st international Conference on Very Large Data Bases, Trondheim, 30 August–02 September. Very Large Data Bases. VLDB Endowment, pp 517–528
41. Gyöngyi Z, Garcia-Molina H (2005) Web spam taxonomy. In: First International Workshop on Adversarial Information Retrieval on the Web, AIRWeb, Chiba
42. Gyöngyi Z, Garcia-Molina H, Pedersen J (2004) Combating web spam with trustrank. In: Proceedings of the Thirtieth International Conference on Very Large Data Bases, VLDB, Toronto, pp 576–587. Morgan Kaufmann Publishers, San Francisco
43. Haveliwala TH (2002) Topic-sensitive PageRank. Proc. of WWW Conf, Honolulu. ACM Press, New York, pp 517–526
44. Haveliwala TH, Kamvar SD (2003) The second eigenvalue of the Google matrix. Stanford University Technical Report 2003-20
45. Haveliwala T, Kamvar S, Jeh G (2003) An analytical comparison of approaches to personalizing PageRank. Technical report 2003-35, Stanford University. http://citeseer.ist.psu.edu/haveliwala03analytical.html
46. Jeh G, Widom J (2003) Scaling personalized web search. In: WWW '03: Proceedings of the 12th international conference on World Wide Web, New York. ACM Press, New York, pp 271–279
47. Jelasity M, Canright G, Engø-Monsen K (2007) Asynchronous distributed power iteration with gossip-based normalization. In: Kermarrec A-M, Bougé L, Priol T (eds) Euro-Par 2007. Lecture Notes in Computer Science, vol 4641. Springer, Berlin, pp 514–525
48. Kamvar SD, Haveliwala TH, Manning CD, Golub GH (2003) Extrapolation methods for accelerating PageRank computations. In: Proc 12th Intl Conf World Wide Web. ACM Press, New York, pp 261–270
49. Kamvar SD, Haveliwala TH, Manning CD, Golub GH (2003) Exploiting the block structure of the web for computing pagerank. Technical Report 2003–17, Stanford University
50. Kempe D, McSherry F (2008) A decentralized algorithm for spectral analysis. J Comput Syst Sci 74(1):70–83

51. Kephart JO, White SR (1991) Directed-graph epidemiological models of computer viruses. In: Proceedings of the IEEE symposium on security and privacy, Oakland, May 1991. IEEE Computer Society Press, Piscataway, pp 343–359
52. Kleinberg JM (1999) Authoritative sources in a hyperlinked environment. J ACM 46(5):604–632
53. Kleinberg JM, Kumar R, Raghavan P, Rajagopalan S, Tomkins AS (1999) The web as a graph: measurements, models and methods. In: Proceedings of the 19th ACM SIGMOD-SIGACT-SIGART symposium on principles of database systems (PODS) Dallas, May 15–18 2000. ACM, New York, pp 1–10
54. Langville AN, Meyer CD (2004) The Use of Linear Algebra by Web Search Engines. IMAGE Newsl 33:2–6
55. Langville AN, Meyer CD (2005) Deeper Inside PageRank. Internet Math 1(3):335–380
56. Lawrence S, Giles CL (1999) Accessibility of information on the web. Nature 400:107–109
57. Lee HC, Borodin A (2003) Perturbation of the hyper-linked environment. In: Warnow T, Zhu B (eds) Proceedings of the 9th annual International conference on computing and combinatorics (COCOON) Big Sky, July 25–28 2003. Lecture notes in computer science, vol 2697. Springer, pp 272–283
58. Lempel R, Moran S (2000) The stochastic approach for link-structure analysis (SALSA) and the TKC effect. Comput Netw 33(1–6):387–401
59. Lempel R, Moran S (2005) Rank Stability and Rank Similarity of Link-Based Web Ranking Algorithms in Authority Connected Graphs. Inf Retr 8:245–264
60. Levine BN, Shields C, Margolin NB (2006) A Survey of Solutions to the Sybil Attack. Tech report 2006–052, University of Massachusetts Amherst, Amherst
61. Liang C, Ru L, Zhu X, R-SpamRank: A Spam Detection Algorithm Based on Link Analysis. J Comput Inf Syst, to appear
62. Motwani R, Raghavan P (1996) Randomized Algorithms. ACM Press, New York
63. Najork M, Wiener JL (2001) Breadth-first search crawling yields high-quality pages. In: Proc 10th Intl WWW Conf. Elsevier Science, Hong-Kong, pp 114–118
64. Nejdl W (2004) How to Build Google2Google – An (Incomplete) Recipe. The Semantic Web – ISWC 2004, 1–5. Lecture Notes in Computer Science, vol 3298. Springer, Berlin
65. Newman MEJ, Forrest S, Balthrop J (2002) Email networks and the spread of computer viruses. Phys Rev E 66:35–101
66. Ng A, Zheng A, Jordan M (2001) Link analysis, eigenvectors and stability. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence. Seattle, pp 903–910
67. Ng AY, Zheng AX, Jordan MI (2001) Stable algorithms for link analysis. In: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval. ACM Press, New York, pp 258–266
68. Page L, Brin S, Motwani R, Winograd T (1998) The PageRank citation ranking: Bringing order to the web. Technical Report, Stanford University, Stanford
69. Parreira JX, Weikum G (2005) JXP: Global authority scores in a P2P network. In: Doan AH, Neven F, McCann R, Bex GJ (eds) Proceedings of the 8th International workshop on the web and databases (WebDB) Baltimore, collocated with ACM SIGMOD/PODS 2005, June 16–17 2005. ACM Press, New York, pp 31–36
70. Parreira JX, Donato D, Michel S, Weikum G (2006) Efficient and decentralized PageRank approximation in a peer-to-peer web search network. In: Dayal U, Whang K, Lomet D, Alonso G, Lohman G, Kersten M, Cha SK, Kim Y (eds) Proceedings of the 32nd international Conference on Very Large Data Bases, Seoul, 12–15 September. Very Large Data Bases. VLDB Endow, pp 415–426
71. Parreira JX, Donato D, Castillo C, Weikum G (2007) Computing trusted authority scores in peer-to-peer web search networks. In: Proceedings of the 3rd international Workshop on Adversarial information Retrieval on the Web, Banff, 8 May. AIRWeb '07, vol 215. ACM, New York, pp 73–80
72. Parreira JX, Michel S, Weikum G (2007) p2pDating: Real life inspired semantic overlay networks for Web search. Inf Process Manag 43(3):643–664
73. Parreira XP, Castillo C, Donato D, Michel S, Weikum G (2008) The Juxtaposed approximate PageRank method for robust PageRank approximation in a peer-to-peer web search network. VLDB J 17(2). Springer, New York
74. Rafiei D, Mendelzon AO (2000) What is this page known for? Computing Web page reputations. In: Proceedings of the 9th International world wide web conference on computer networks. Int J Comput Telecommun Netw. North-Holland, Amsterdam, pp 823–835
75. Richardson M, Domingos P (2002) The intelligent surfer: Probabilistic combination of link and content information in pagerank. Adv Neural Inf Process Syst 14:1441–1448
76. Sankaralingam K, Sethumadhavan S, Browne JC (2003) Distributed Pagerank for P2P Systems. In: Proceedings of the 12th IEEE International symposium on high performance distributed computing (HPDC) Seattle, June 22–24 2003. IEEE Computer Society, Washington, pp 58–68
77. Sankaralingam K, Yalamanchi M, Sethumadhavan S, Browne JC (2003) Pagerank Computation and Keyword Search on Distributed Systems and P2P Networks. J Grid Comput 1(3): 291–307
78. Serrano MA, Maguitman A, Boguna M, Fortunato S, Vespignani A (2007) Decoding the structure of the WWW: A comparative analysis of web crawls. ACM Trans Web 1(2):10
79. Shi S, Yu J, Yang GW, Wang DX (2003) Distributed Page Ranking in Structured P2P Networks. In: Proceedings of the 32nd International conference on parallel processing (ICPP) Kaohsiung, October 6–9 2003. IEEE Computer Society, Washington, pp 179–186
80. Steinmetz R, Wehrle K (2005) Peer-to-peer Systems and Applications. Springer, Heidelberg
81. Stoyanovich J, Bedathur SJ, Berberich K, Weikum G (2007) EntityAuthority: Semantically enriched graph-based authority propagation. In: Proceedings of the 10th International workshop on the web and databases (WebDB) Beijing, June 15 2007. ACM, New York
82. Sydow M (2004) Random surfer with back step. In: Proceedings of the 13th International world wide web conference on alternate track papers and posters (WWW Alt) New York, May 19–21 2004. ACM Press, New York, pp 352–353
83. Varga RS (1962) Matrix iterative Analysis. Prentice-Hall, Upper Saddle River, USA
84. Wang Y, DeWitt D (2004) Computing pagerank in a distributed internet search system. In: Nascimento MA, Özsu MT, Kossmann D, Miller RJ, Blakeley JA, Schiefer KB (eds) Proceedings of the 30th International conference on very large data bases (VLDB) Toronto, August 31 – September 3 2004. VLDB Endowment, pp 420–431

85. Watts D, Strogatz S (1998) Collective dynamics of 'small world' networks. Nature 393:440–442
86. Wu J, Aberer K (2004) Using SiteRank for decentralized computation of Web document ranking. Adaptive Hypermedia and adaptive Web-based Systems. LNCS, vol 3137. Springer, Berlin, pp 265–274
87. Wu J, Aberer K (2005) Using a Layered Markov Model for Distributed Web Ranking Computation. In: Proceedings of the 25th IEEE International Conference on Distributed Computing Systems. ICDCS 10(10):533–542
88. Yahoo.com (2008) Search query 'car' at http://www.yahoo.com. Accessed 13 Feb 2008
89. Zhang H, Goel A, Govindan R, Mason K, van Roy B (2004) Making eigenvector-based reputation systems robust to collusion. In: Proceedings of the third Workshop on Web Graphs (WAW). Lecture Notes in Computer Science, vol 3243. Springer, Rome, pp 92–104
90. Zhu Y, Ye S, Li X (2005) Distributed pagerank computation based on iterative aggregation-disaggregation methods. In: Proceedings of the 14th ACM International conference on information and knowledge management (CIKM) Bremen, October 31 – November 5 2005. ACM Press, New York, pp 578–585

Books and Reviews

Bianchini M, Gori M, Scarselli F (2005) Inside PageRank. ACM Trans Inter Tech 5(1):92–128
Chakrabarti S (2002) Mining the Web: Analysis of Hypertext and Semi Structured Data. Morgan Kaufmann, San Francisco
Langville AN, Meyer CD (2005) Deeper Inside PageRank. Internet Math 1(3):335–380
Langville AN, Meyer CD (2006) Google's PageRank and Beyond: The Science of Search Engine Rankings. Princeton University Press, Princeton
Thelwall M (2004) Link Analysis: An Information Science Approach. Academic Press, Amsterdam

Logic and Geometry of Agents in Agent-Based Modeling

Samson Abramsky
Oxford University Computing Laboratory, Oxford, UK

Article Outline

Glossary

Agent A computational or biological entity which can perform actions which affect its environment, and observe actions performed by the environment or other agents.

Agent-based modeling Modeling a complex system in terms of multiple agents interacting with each other.

Interaction A pattern of actions in a multi-agent system. Each action is performed by some agent, and may be observed by others.

Information flow The process whereby information held by one agent or part of a system is transferred to another, possibly in some modified form. Generally mediated by interaction: information flow is caused by agents performing actions, and other agents observing those actions.

Compositionality Describing a complex system or object in a structured fashion, as built up by applying certain operations hierarchically, starting from a stock of basic types of system. Compositional definitions of functions of systems are those which respect this hierarchical structure.

Linear logic A sub-structural logic in which the operations of copying and deleting premises are not allowed in general.

Combinators Basic operations in a function algebra. Combinatory logic is a variable-free formulation of functional computation.

Definition of the Subject

Agent-based Modeling has become of increasing importance in Computer Science, and also in mathematical modeling and simulation. The idea is that the behavior of a complex system can be described as arising from the interaction of multiple agents, with each other and with the environment, using simple local rules. It is widely recognized that building a sound and widely applicable theory for such systems will require an inter-disciplinary approach, and the development of new mathematical and computational concepts.

In this article, agents and interaction will be studied from the perspective of Logic and Computer Science. It will be shown how ideas about Logical Dynamics, Games and Geometry of Interaction, which have been developed over the past two decades, lead towards a structural theory of agents and interaction. This provides a basis for powerful logical methods such as compositionality, types and high-level calculi, which have proved so fruitful in Computer Science, to be applied in this domain.

This approach should be contrasted with the more familiar approaches to agent-based modeling using primitive agents such as cellular automata. The main focus of such approaches is systems modeling, where the agents are essentially discrete counterparts to traditional PDE or ODE models of dynamical systems. Ultimately, these ap-

proaches should be seen as complementary, and the goal is to combine them. In the present article, the emphasis is on presenting some basic ideas of the novel approach to complex systems using logical and compositional methods.

Introduction

In this article, a logical and geometric perspective on the information flow arising from agent interaction will be described. This is a novel and distinctive approach within the emerging field of agent-based modeling. While it is still in its formative stages, there are already some key structural insights which it offers which will be fundamental to any deep and comprehensive theory of agents.

Firstly, a fundamental methodological point will be discussed, which has played a crucial rôle in Computer Science for several decades, but has yet to achieve the recognition in general scientific modeling which it deserves: the importance of **compositionality**.

Compositionality

A methodological principle from Computer Science (and Logic) of **major** potential importance for mathematical modeling throughout the sciences.

- **Traditional approach**: Whole-system (monolithic) analysis of given systems. A key rôle is played by structuring templates, e. g. "Find the Hamiltonian".
- **Compositional approach**: Start with a fixed set of basic (simple) building blocks, and **constructions** for building new (in general more complex) systems out of given sub-systems, and build up the required complex system with these.

More formally, compositionality can be expressed **algebraically**:

$$S = \omega(S_1, \ldots, S_n) .$$

The system S is described as being built up from sub-systems $S_1, \ldots, S_n$ by the operation ω. There is also a **logical** perspective:

$$\frac{S_1 \models \phi_1, \ldots, S_n \models \phi_n}{\omega(S_1, \ldots, S_n) \models \phi}$$

(Read $S \models \phi$ as "system S satisfies the property ϕ"). Here **properties** ϕ of the compound system S can be inferred by verifying properties $\phi_1, \ldots, \phi_n$ for the simpler sub-systems $S_1, \ldots, S_n$.

Some Key Points

- The compositional view of complex systems sees them as built up, not just by one-level composition of basic agents: $S = \|_{i \in I}\, a_i$, which is the usual scenario in current agent-based modeling and simulation, but **hierarchically**:

  $$S = \omega_1(\omega_2(a_1, a_2, a_3), \omega_1(a_4, a_5)) .$$

  It tracks the properties of the sub-compound systems all the way up (or down) the tree of syntax.
- The **available repertoire** of system constructors is also an important aspect of the modeling here, leading to questions of **expressiveness** and **functional completeness**.
- This paradigm has played a major role in Computer Science over the past four decades, and is already starting to be applied in
  - quantum computing (quantum programming languages)
  - biological modeling (process calculi)
  - business modeling (idem)

  and will surely be applied more widely and deeply, in economics as well as physical and biological sciences.
- This kind of modeling carries in its train a range of powerful **analytical techniques**: types, semantics, verification, model-checking, etc.

Computation as Interaction

The second key point to be emphasized is the emergence of **interaction** as a key notion in computation, and increasingly in a wide range of scientific fields, whether in physics (quantum information), biology (at a range of scales and levels, from molecular interactions to evolutionary theory), economics (game theory), linguistics (dialogical analysis), etc. Moreover, these developments are feeding back into logic and philosophy.

Note that interaction happens **between agents**—a theory of interaction amounts to an **agent dynamics**.

Changing Views of Computation The scene can be set by recalling how perspectives on computation have changed since the first computers appeared. The early practice of computing can be pictured as in Fig. 1. This is the era of stand-alone machines and programs: computers are served by an elite priesthood, and have only a narrow input-output interface with the rest of the world.

First-Generation Models of Computation These models live on the existing intellectual inheritance from discrete mathematics and logic. **Time** and **processes** lurk in the background, but are largely suppressed.

Given this limited vision of computing, there is a very natural abstraction of computation, in which programs are

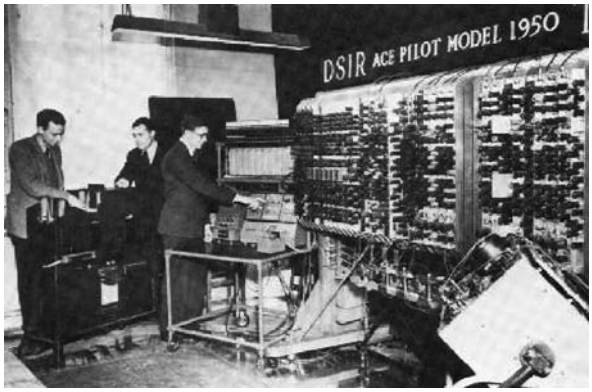

Logic and Geometry of Agents in Agent-Based Modeling, Figure 1
Computing "in the isolation ward"

seen as computing **functions** or **relations** from inputs to outputs.

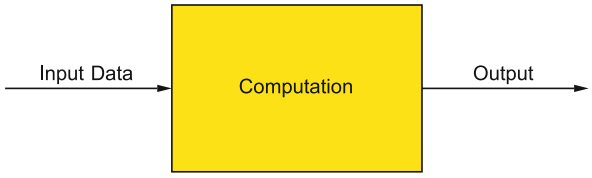

Computation in the Age of the Internet Instead of isolated systems, with rudimentary interactions with their environment, the standard unit of description or design becomes a **process** or **agent**, the essence of whose behavior is **how it interacts** with its environment.

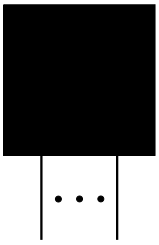

The interaction between the System and the Environment can be conceptualized as a **two-person game**. A program specifying how the System should behave in the face of all possible actions by the Environment is then a **strategy** for the player corresponding to the System.

Interaction Complex behavior arises as the global effect of a **system** of **interacting agents** (or processes).

The key building block is the agent. The key operation is **interaction**—plugging agents together so that they interact with each other.

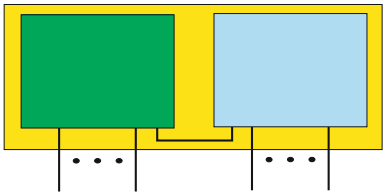

Who is the System? Who is the Environment? This **symmetry** between System and Environment carries a first clue that there is some structure here; it will lead to a key **duality**, and a deep connection to logic.

Towards a Compositional Approach to Complex Systems This conceptual model works at all "scales" :

- Macro-scale: processes in operating systems, software agents on the Internet, transactions.
- Micro-scale: how programs are implemented (subroutine call-return protocols, register transfer) all the way down into hardware.

It is applicable both to **design** (synthesis) and to **description** (analysis); to **artificial** and to **natural** information-processing systems.

There are of course large issues lurking here, e. g. in the realm of "Complex Systems": **emergent behavior** and even **intelligence**. Is is helpful, or even feasible, to understand this complexity **compositionally**? New conceptual tools, new theories, are needed to help us analyze and synthesize these systems, to help us to **understand** and to **build**.

Towards a Logic and Geometry of Interaction

Towards a "Logic of Interaction"

Specifying and reasoning about the behavior of computer programs takes us into the realm of logic. For the first-generation models, logic could be taken "as it was"—static and timeless. For the second-generation models, getting an adequate account—a genuine "logic of interaction"—may require a fundamental reconceptualization of logic itself. This radical revision of the view of logic is happening anyway—prompted partly by the applications, and partly by ideas arising within logic.

The Static Conception of Logic The usual "static" notion of tautology is as "a statement which is vacuously true because it is compatible with all states of affairs"

$$A \vee \neg A\,.$$

"It is raining **or** it is not raining"—truth-functional semantics. This is illustrated (subversively) in Fig. 2. But what could a **dynamic notion of tautology** look like?

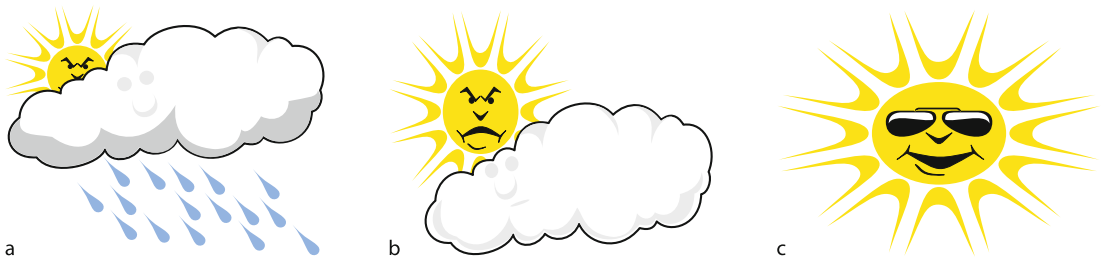

Logic and Geometry of Agents in Agent-Based Modeling, Figure 2
Tertium non datur?

Logic and Geometry of Agents in Agent-Based Modeling, Figure 3
How to beat a Grandmaster

The Copy-Cat Strategy Consider the following little fable, illustrated by Fig. 3:

> How to beat an International Chess Grandmaster by the power of pure logic

The idea is to rely on logic, rather than on any talent at Chess. We arrange to play two games of Chess with the grandmaster, say Gary Kasparov, once as White and once as Black. Moreover, we so arrange matters that we start with the game in which we play as Black. Kasparov makes his opening move; we respond by playing the **same** move in the **other** game—this makes sense, since we are playing as White there. Now Kasparov responds (as Black) to our move in that game; and we copy that response back in the first game. We simply proceed in this fashion, copying the moves that our opponent makes in one board to the other board. The net effect is that **we play the same game twice—once as White, and once as Black**. (We have essentially made Kasparov play against himself). Thus, whoever wins that game, we can claim a win in one of our games against Kasparov! (Even if the game results in a stalemate, we have done as well as Kasparov over the two games—surely still a good result!) Of course, this idea has nothing particularly to do with Chess. It can be applied to any two-person game of a very general form. The use of Chess-boards to illustrate the discussion will continue, but this underlying generality should be kept in mind.

What are the salient features which can be extracted from this example?

A dynamic tautology There is a sense (which will shortly be made more precise) in which the copy-cat strategy can be seen as a **dynamic version** of the tautology $A \vee \neg A$. Note, indeed, that an essential condition for being able to play the copy-cat is that the rôles of the two players are inter-changed on one board as compared to the other. Note also the disjunctive quality of the argument that we must win in one or other of the two games. But the copy-cat strategy is a **dynamic**

process: a two-way channel which maintains the correlation between the plays in the two games.

Conservation of information flow The copy-cat strategy does not **create** any information; it reacts to the environment in such a way that information is conserved. It ensures that exactly the same information flows out to the environment as flows in from it. Thus one gets a sense of logic appearing in the form of **conservation laws for information dynamics**.

The power of copying Another theme which appears here, and of which more will be seen later, concerns the surprising power of simple processes of copying information from one place to another. Indeed, as shall eventually be seen, such processes are **computationally universal**.

The geometry of information flow From a dynamical point of view, the copy-cat strategy realizes a channel between the two game boards, by performing the **actions** of copying moves. But there is also some implicit **geometry** here. Indeed, the very idea of two boards laid out side by side appeals to some basic underlying spatial structure. In these terms, the copy-cat channel can also be understood geometrically, as creating a graphical link between these two spatial locations. These two points of view are complementary, and link the logical perspective to powerful ideas arising in modern geometry and mathematical physics.

Further evidence that the copy-cat strategy embodies more substantial ideas than might at first be apparent, can be obtained by varying the scenario. Consider now the case where we play against Kasparov on **three boards**; one as Black, two as White.

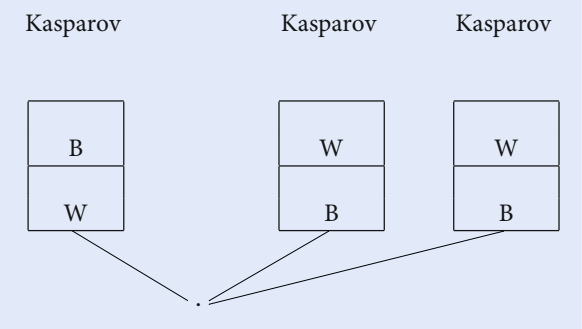

Does the Copy-Cat strategy still work here? In fact, it can easily be seen that it does **not**. Suppose Kasparov makes an opening move m_1 in the left-hand board where he plays as White; we copy it to the board where we play as White; he responds with m_2; and we copy m_2 back to the board where Kasparov opened. So far, all has proceeded as in our original scenario. But now Kasparov has the option of playing a **different** opening move, m_3 say, in the rightmost board. We have no idea how to respond to this move; nor can we copy it anywhere, since the board where we play as White is already "in use". This shows that these simple ideas already lead us naturally to the setting of a **resource-sensitive** logic, in which in particular the Contraction Rule, which can be expressed as $A \to A \wedge A$ (or equivalently as $\neg A \vee (A \wedge A)$) cannot be assumed to be valid.

What about the other obvious variation, where we play on two boards as White, and one as Black?

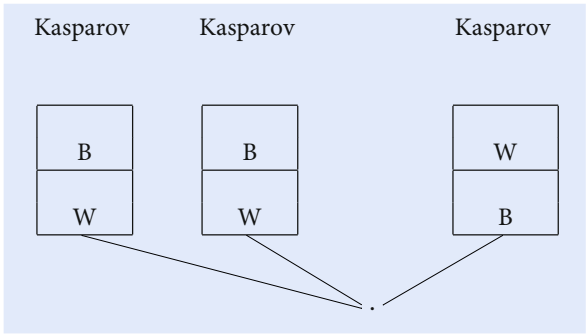

It seems that the copy-cat strategy **does** still work here, since we can simply ignore one of the boards where we play as White. However, a geometrical property of the original copy-cat strategy has been lost, namely a **connectedness** property, that information flows to every part of the system. This at least calls the corresponding logical principle of Weakening, which can be expressed as $A \wedge A \to A$, (or equivalently as $\neg A \vee \neg A \vee A$) into question.

These remarks indicate that we are close to the realm of Linear Logic and its variants; and, mathematically, to the world of monoidal (rather than cartesian) categories.

Game Semantics These ideas find formal expression in **Game Semantics**. Games play the role of:

- Interface types for computation modules
- Propositions with dynamic content.

In particular, 2-person games capture the duality of:

- Player vs. Opponent
- System vs. Environment.

Agents are Strategies In this setting, agents or processes can be modeled as **strategies** for playing the game. These strategies **interact** by playing against each other. A notion of correctness is obtained which is **logical** in character in terms of the idea of **winning** strategy—one which is guaranteed to reach a successful outcome however the environment behaves. This in a sense replaces (or better, **refines**) the logical notion of "truth": winning strategies are the dynamic version of tautologies (more accurately, of their **proofs**).

Building Complex Systems by Combining Games It will now be seen how games can be combined to produce more complex behaviors while retaining control over the interface. This provides a basis for the **compositional** understanding of systems of interacting agents—understanding the behavior of a complex system in terms of the behavior of its parts. This is crucial for both analysis and synthesis, i. e. for both description and design. These operations for building games can be seen as (dynamic forms of) "type constructors" or "logical connectives". (The underlying logic here will in fact be Linear Logic).

Duality—"Linear Negation" $A^{\perp}$—interchange rôles of Player and Opponent (reflecting the symmetry of interaction).

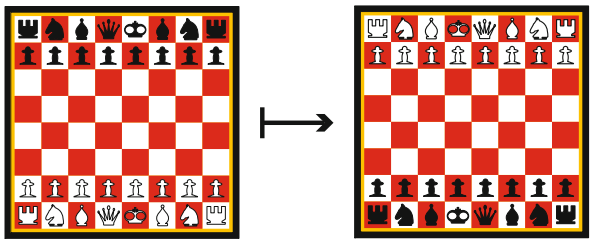

Note that, with this interpretation, negation is involutive:

$$A^{\perp\perp} = A\,.$$

Tensor—"Linear Conjunction"

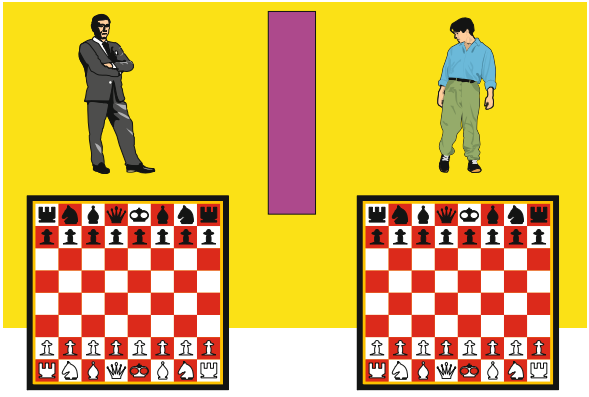

$A \otimes B$

The idea here is to combine the two game boards into one system, **without any information flow between the two sub-systems**. (This is the significance of the "wall" separating the two players, who will be referred to as Gary (Kasparov) and Nigel (Short)). This connective has a conjunctive quality, since we must independently be able to play (and to win) in each conjunct. Note however, that there is no constraint on information flow for the environment, as it plays against this compound system.

Par—"Linear Disjunction"

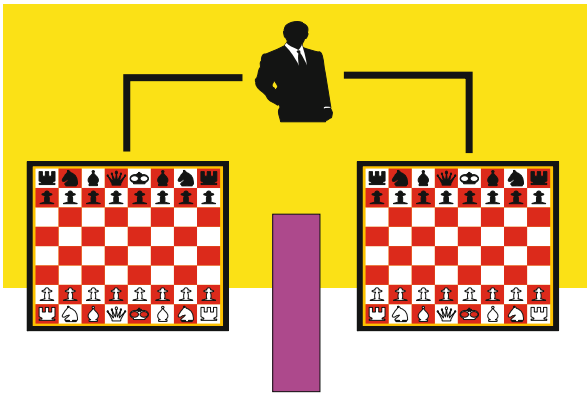

$A \parr B$

In this case, there are two boards, but one player (who shall be referred to as the Copy-Cat), indicating that we **do** allow information flow for this player between the two game boards. This for example allows information revealed in one game board by the Opponent to be used against him on the other game board—as exemplified by the copy-cat strategy. However, note that the wall appears on the environment's side now, indicating that the environment is constrained to play separately on the two boards, with no communication between them.

Thus there is a De Morgan duality between these two connectives, mediated by the Linear negation:

$$(A \otimes B)^{\perp} = A^{\perp} \parr B^{\perp}\,, \qquad (A \parr B)^{\perp} = A^{\perp} \otimes B^{\perp}\,.$$

The idea is that on one side of the mirror of duality (Player/System for the Tensor, Opponent/Environment for the Par), there is the constraint of no information flow, while on the other side, there is information flow.

The Copy-Cat strategy can now be reconstrued in logical terms:

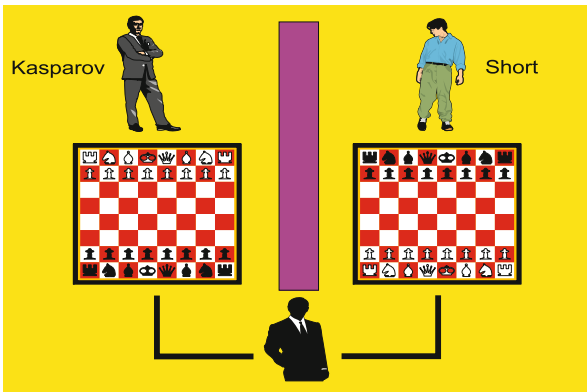

It can be seen that it is indeed a winning strategy for $A^{\perp} \parr A$. Moreover, $A \multimap B$ ("Linear implication") can be defined by

$$A \multimap B \equiv A^{\perp} \parr B\,,$$

(*cf.* $A \supset B \equiv \neg A \vee B$.) Then the copy-cat strategy becomes the canonical proof of the most basic tautology of all: $A \multimap A$.

The information flow possibilities of Par receive a more familiar logical interpretation in terms of the Linear implication; namely, that information about the antecedent can be used in proving the consequent (and conversely with respect to their negations, if proof by contraposition is considered).

Thus an entire "linearized" logical structure opens up before us, with a natural interpretation in terms of the dynamics of information flow.

Interaction A key step in the development now arises: the modeling of **interaction** itself. Constructors create "potentials" for interaction; the operation of plugging modules together so that they can communicate with each other **releases** this potential into **actual computation**. Consider the diagram shown in Fig. 4. Here two separate sub-systems are shown, each with a compound structure, expressed by the **logical types of their interfaces**. What these types tell us is that these systems are **composable**; in particular, the **output type** of the first system, namely B, matches the input type of the second system. Note that this "logical plug-compatibility" makes essential use of the duality, just as the copy-cat strategy did. What makes Gary (the player for the first system) a fit partner for interaction with Nigel (the player for the second system), is that they have **complementary views** of their locus of interaction, namely B. Gary will play in this type "positively", as Player (he sees it as B), while Nigel will play "negatively", as Opponent (he sees it as $B^{\perp}$). Thus each will become part of the environment of the other—part of the potential environment of each will be realized by the other, and hence part of the **potential** behavior of each will become **actual** interaction.

This leads to a dynamical interpretation of the fundamental operation of **composition**, in mathematical terms:

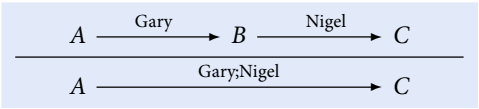

or of the **Cut rule**, in logical terms:

$$\text{Cut:}\quad \frac{\vdash \Gamma, A \quad \vdash A^{\perp}, \Delta}{\Gamma, \Delta}$$

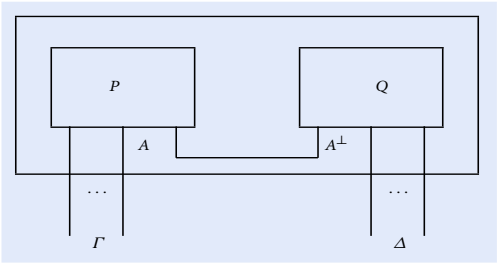

Composition as Interaction

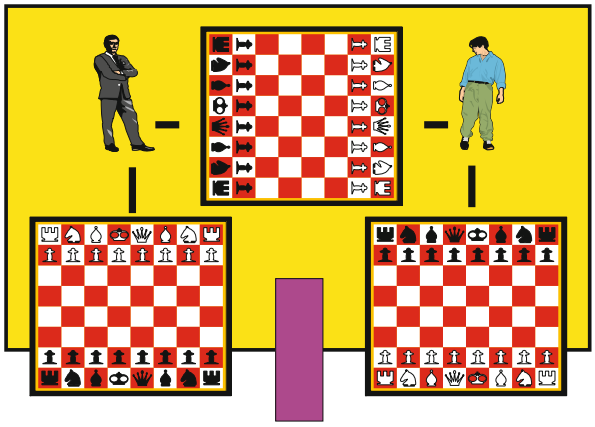

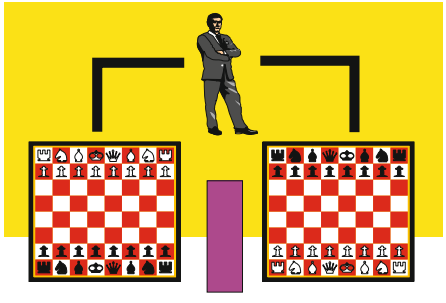

$A^{\perp} \parr B \equiv A \multimap B\,,$

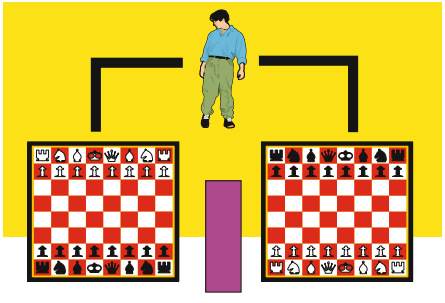

$B^{\perp} \parr C \equiv B \multimap C\,.$

Logic and Geometry of Agents in Agent-Based Modeling, Figure 4
Two composable systems

The picture here shows the new system formed by plugging together the two sub-systems. The "external interface" to the environment now shows just the left hand board A as input, and the right hand board C as output. The Cut formula B is hidden from the environment, and becomes the locus of interaction inside the black box of the system. Suppose that the Environment makes some move m in C. This is visible only to Nigel, who as a strategy for $B \multimap C$ has a response. Suppose this response m_1 is in B. This is a move by Nigel as Player in $B^{\perp}$, hence appears to Gary as a move by Opponent in B. Gary as a strategy for $A \multimap B$ has a response m_2 to this move. If this response is again in B, Nigel sees it as a response by the environment to his move, and will have a response again; and so on. Thus there is a sequence of moves $m_1, \ldots, m_k$ in B, ping-ponging back and forth between Nigel and Gary. If, eventually, Nigel responds to Gary's last move by playing in C, or Gary responds to Nigel's last move by playing in A, then this provides the response of the **composed strategy** Gary; Nigel to the original move m. Indeed, all that is visible to the Environment is that it played m, and eventually some response appeared, in A or C.

Moreover, if both Nigel and Gary are winning strategies, then so is the composed strategy; and the composed strategy will not get stuck forever in the internal ping-pong in B. To see this, suppose for a contradiction that it did in fact get stuck in B. Then there would be an infinite play in B following the winning strategy Gary for **Player** in B, and the **same** infinite play following the winning strategy Nigel for Player in $B^{\perp}$, hence for **Opponent** in B. Hence the same play would count as a win for both Player and Opponent. This yields the desired contradiction.

Discussion Game Semantics in the sense discussed in this section has had an extensive development over the past decade and a half, with a wealth of applications to the semantics of programming languages, type theories and logics [9,11,12,13,14,16,31]. More recently, there has been an algorithmic turn, and some striking applications to verification and program analysis [2,17,24,38].

From the point of view of the general analysis of Information, there are the following promising lines of development:

- Game semantics provides a promising arena for exploring the combination of quantitative and qualitative theories of information. In particular, it provides a setting for quantifying information flow between agents. Important quantitative questions can be asked about **rate of information flow** through a strategy (representing a program, or a proof); how can a system gain **maximum** information from its environment while providing **minimal** information in return; robustness in the presence of **noise**, etc.
- As in the discussion of the copy-cat strategy, there is an intuition of logical principles arising as **conservation laws for information flow**. (And indeed, in the case of Multiplicative Linear Logic, the proofs correspond exactly to "generalized copy-cat strategies"). Can this intuition be developed into a fully-fledged theory? Can logical principles be **characterized** as those expressing the conservation principles of this information flow dynamics?
- There is also the hope that the more structured setting of game semantics will usefully constrain the exuberant variety of possibilities offered by process algebra, and allow a sharper exploration of the logical space of possibilities for information dynamics. This has already been borne out in part, by the success of game semantics in exploring the space of programming language semantics. It has been possible to give crisp characterizations of the "shapes" of computations carried out within certain **programming disciplines**: including purely functional programming [16,31], stateful programming [11,12], general references [15], programming with non-local jumps and exceptions [35,36], non-determinism [29], probability [21], concurrency [25,26], names [18], polymorphism [10,30] and more. See [13] for an overview (now rather out of date). There has also been a parallel line of development of giving **full completeness** results for a range of logics and type theories, characterizing the "space of proofs" for a logic in terms of informatic or geometric constraints which pick out those processes which are proofs for that logic [9,14,19,20,22,37]. This allows a new look at such issues as the boundaries between classical and constructive logic, or the fine structure of polymorphism and second-order quantification.
- This also gives some grounds for optimism that what **computational processes** are can be captured—in a "machine-independent", and moreover "geometrical", non-inductive way—, **without** referring back to Turing machines or any other explicit machine model.
- In the same spirit as for computability, can **polynomial-time computation** and other complexity classes be characterized in such terms?

Emergent Logic: The Geometry of Information Flow

The geometric structure inherent in interaction and information flow, and its connections to the dynamic view of interaction, will now be considered. Recent work has shown how common structure arises in a variety of contexts: in **logic** (geometric representations of proofs);

geometry (diagram algebras, especially the Temperley–Lieb algebra, with connections to knot theory and geometric topology); **computation** (especially functional computation); and **quantum mechanics** (quantum information protocols, exploiting the information flow inherent in quantum entanglement).

A rapid tour will be made through these disparate areas, focusing on how common structure arises. The exposition will rely extensively on **graphical calculi** which provide an intuitive and visually appealing window onto the various formalisms to be encountered. These calculi have a substantial mathematical content, founded on the diagrammatic approach to tensor categories; further details can be found in the references.

Logic Firstly Multiplicative Linear Logic [27], the logic of the linear connectives $\otimes$, $\parr$, $(\_)^\perp$ which have already been encountered, will be considered, as a basic paradigmatic example. A key insight [27] is that the essential information in a proof in this system is given by a pairwise matching of the occurrences of positive and negative literals in the sequent—a **proof structure**. E.g., the two possible proof structures for the sequent $\alpha^\perp \parr \alpha^\perp, \alpha \otimes \alpha$ are:

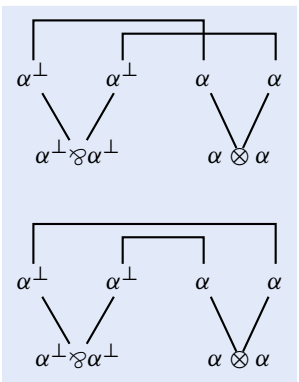

A **geometric** view of proof structures, as certain graphs obtained from the forest of formula trees in the sequent by drawing arcs—"axiom links"—between the paired literal occurrences. Alternatively, a more **dynamic** view can be taken, and proof structures can be represented as **involutions** on the set of literal occurrences in the sequent. These must be **literal preserving**, i. e. an occurrence of a literal l must be mapped to an occurrence of its dual $l^\perp$. Such functions represent a flow of information around the sequent. They can be viewed as **copy-cat strategies**.

Every sequent proof determines a proof structure. The fundamental question is: **which proof structures arise from sequent proofs**? A first answer is **geometric** (or topological) in character. For each proof structure, a **switching graph** can be obtained by deleting, for each occurrence of a subformula $A \parr B$, exactly one of the arcs $A \,—\, A \parr B \,—\, B$ connecting it to its immediate subformulas. If all such switching graphs are trees, the proof structure is said to be a **proof net**.

A second form of characterization is **interactive**. An orthogonality relation can be defined, between permutations on the set of occurrences of literals in the sequent Γ: $f \perp g$ if fg is cyclic. The idea is that f is a candidate proof net, while g is an attempted counter-proof—a passage through the literals induced by a choice of switching graph. Note that the alternating composition of f and g expresses **interaction** between f and g, thought of as **strategy** and **counter-strategy**. It generates a path along which information flows around the system.

A semantics of MLL proofs can be given by specifying, for each formula A, a set S of permutations on the set of literal occurrences $|A|$, such that $S = S^{\perp\perp}$, where $S^\perp = \{g \mid \forall f \in S.\ f \perp g\}$. For a literal, the unique permutation (the identity) is specified.

$$S(A \otimes B) = \{f + g \mid f \in S(A) \wedge g \in S(B)\}^{\perp\perp}$$
$$S(A \parr B) = S(A^\perp \otimes B^\perp)^\perp .$$

Here $f + g$ is disjoint union of permutations, expressing the absence of information flow (or **information independence**).

Theorem 1 (Sequentializability [27] and Full Completeness [9]) *Let f be a literal-respecting involution on $|\Gamma|$. The following are equivalent: (i) f is the permutation assigned to a sequent proof of Γ; (ii) f is a proof net; (iii) $f \in S(\Gamma)$.*

This shows that **the geometric and interactive characterizations of the space of proofs coincide**.

Further Developments From a computational perspective, the multiplicative connectives embody **concurrency** and **causal independence**. The scope of the enterprise is greatly expanded when the other **levels of connective** in Linear Logic are incorporated:

Additives: The additive conjunction and disjunction allow **causality**, **conflict** and **case analysis/conditionals** to be expressed. The interaction between the additive and multiplicative levels is rather subtle. The theory sketched above for the multiplicatives has in large part been lifted to Multiplicative Additive Linear Logic [14], but a number of key questions and issues needed for a deeper analysis remain to be investigated.

Exponentials: The multiplicative fragment only allows linear time computation to be expressed (under the Curry–Howard paradigm). For a full analysis of com-

putationally expressive systems, it is necessary to allow for **copying** and **deleting**, as regulated by the **exponential connectives** of Linear Logic. Existing results have extended the multiplicative theory to various systems of typed λ-calculus (corresponding to various forms of Intuitionistic Multiplicative Exponential Linear Logic), and have begun to investigate systems which, by constraining the exponential types in certain natural ways, capture significant **complexity classes**, especially **PTIME**.

Diagram Algebras It will now be indicated how this apparently very specialized corner of Proof Theory in fact connects directly to a broad topic arising in Representation Theory and Knot Theory, with connections to Mathematical Physics. On the one hand, some structure will be **lost**, by obliterating the distinction between $\otimes$ and $\parr$; this corresponds to moving from $*$**-autonomous** to **compact closed** categories. This means that the formula tree structure can be dispensed with altogether; it is simply a matter of connecting up literal occurrences, which shall be drawn as "joining up the dots". Motivation: compact closed categories show up in many contexts of interest!

On the other hand, rather than one-sided sequents, general arrows or two-sided sequents will be represented diagrammatically. This means arrows can be represented as

$$A_1 \otimes \cdots \otimes A_n \longrightarrow B_1 \otimes \cdots \otimes B_m \ ,$$

where each A_i and B_j is a literal, by involutions on $\{1, \ldots, n\} + \{1, \ldots, m\}$, which are literal-preserving in the extended sense that **opposite** literals are connected in the domain or in the codomain, while occurrences of the **same** literal in the domain and the codomain are connected . An advantage of this representation is that composition is expressed very transparently, by "stacking" arrows.

Example The composition of

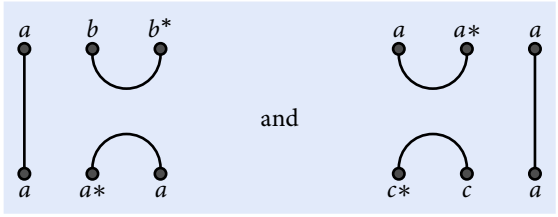

is given by

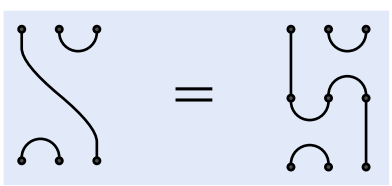

Temperley–Lieb Algebra The **Temperley–Lieb algebra** played a central role in the **Jones polynomial invariant of knots** [32] and ensuing developments. It was originally presented, rather forbiddingly, in terms of abstract generators and relations [32]. It was recast in beautifully elementary and conceptual terms by Louis Kauffman as a **planar diagram algebra** [33].

Generators:

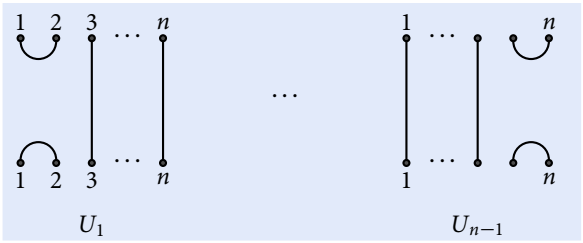

Relations:

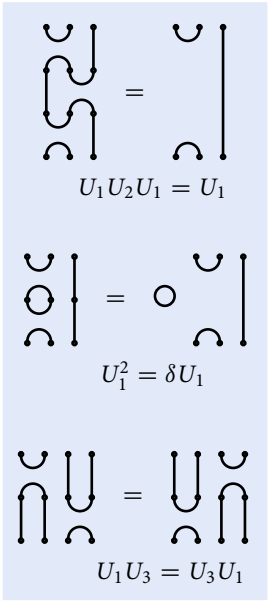

The general form of an element of the algebra (actually of the basic multiplicative monoid: the algebra is then constructed freely over this as the "monoid algebra") is obtained by "joining up the dots" in a **planar** fashion. Multiplication xy is defined by identifying the bottom row of x with the top row of y. In general loops may be formed—these are "scalars", which can float freely across these figures, represented symbolically by δ above.

How does this connect to knots? A key conceptual insight is due to Kauffman, who saw how to recast the Jones polynomial in elementary combinatorial form in terms of his **bracket polynomial**. The basic idea of the bracket polynomial is expressed by the following equation:

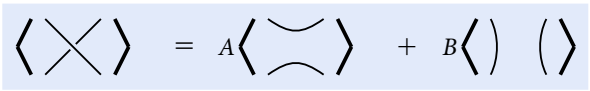

Each over-crossing in a knot or link is evaluated to a weighted sum of the two possible planar smoothings. With suitable choices for the coefficients A and B (as Laurent polynomials), this is invariant under the second and third Reidemeister moves. With an ingenious choice of normalizing factor, it becomes invariant under the first Reidemeister move—and yields the Jones polynomial! What this means algebraically is that the braid group has a representation in the Temperley–Lieb algebra—the above bracket evaluation shows how the basic generators of the braid group are mapped into the Temperley–Lieb algebra. Every knot arises as the closure of a braid; the invariant arises by mapping the open braid into the Temperley–Lieb algebra, and taking the trace there.

Moreover, it turns out that this connection can itself carry interesting information between the Computer Science ideas and the geometry and algebra. Indeed, using Computer Science methods it is possible to give the first **direct presentation** (no quotients) of the Temperley–Lieb algebra, using logical methods. In fact, the elements of the Temperley–Lieb algebra are completely determined by the relations they induce on the "dots"; and **planarity** can be characterized using only the ordering relations on the two rows of dots. Moreover, the multiplication of the algebra can be described as a form of Cut Elimination, using the methods developed in the "Geometry of Interaction" [1, 8,28]. This exactly corresponds to characterizing the **geometric composition of diagrams** as above in terms of an **information flow dynamics**, just as in the case of Multiplicative Linear Logic. This also shows that **planarity is an invariant of Cut Elimination**, and raises interesting questions about computational expressiveness under topological constraints such as planarity, and the computational significance of braiding.

Applicative Computation This ideas in turn apply directly to **applicative computation**, offering the same combination of geometric/diagrammatic and information-flow tools of analysis. This is illustrated firstly with a **planar combinator**

$$\mathbf{B} \equiv \lambda x.\lambda y.\lambda z.\, x(yz)\colon (B \to C) \to (A \to B) \to (A \to C)$$

characterized by the equation

$$\mathbf{B}abc = a(bc)\,.$$

The interpretation of the open term

$$x\colon B \to C\,,\quad y\colon A \to B\,,\quad z\colon A \vdash x(yz)\colon C$$

is as follows:

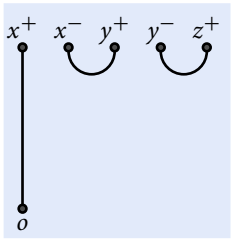

Here x^+ is the output of x, and x^- the input, and similarly for y. The output of the whole expression is o. When the variables are abstracted, the following caps-only diagram is obtained:

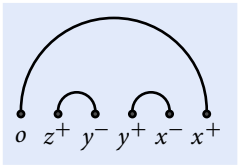

Now consider an application $\mathbf{B}abc$:

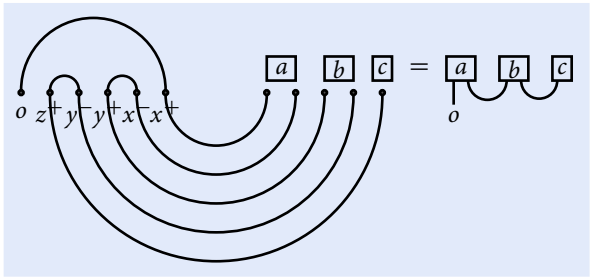

If the constraint on planarity in the diagrams is relaxed, a similar representation can be given of the **commutation combinator**:

$$\mathbf{C} \equiv \lambda x.\,\lambda y.\,\lambda z.\, xzy\colon (A \to B \to C) \to B \to A \to C$$

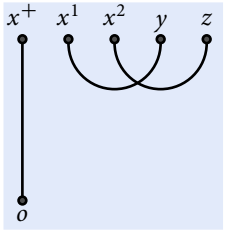

More generally, the **Brauer algebra** (1931) arises if the planarity condition on the TL algebra is removed. This plays an important role in the representation theory of the Orthogonal group ('Schur–Weyl duality'), which is now part of a whole genre of 'diagram algebras' in Representation Theory.

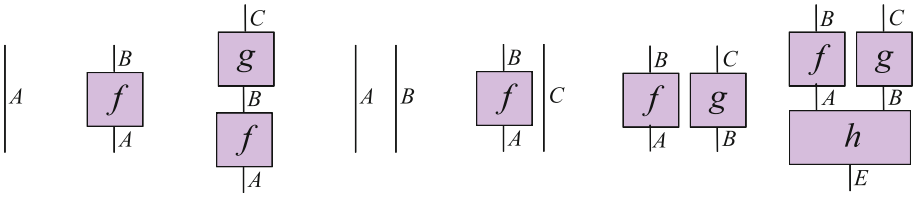

Logic and Geometry of Agents in Agent-Based Modeling, Figure 5
Graphical calculus for monoidal categories

With **BCI** combinators one can interpret **Linear λ-calculus**. The Kelly–Laplaza construction of the free compact closed category [34] can be retrieved by a straightforward generalization of these ideas.

Quantum Computation The **same** graphical calculus and underlying algebraic structure which have been tracked through logic, geometry and computation has been applied to quantum information and computation, yielding an incisive analysis of **quantum information flow**, and powerful and illuminating methods for reasoning about quantum informatic processes and protocols [6]. This "strongly compact closed graphical calculus" can be seen as a very substantial 2-dimensional extension of Dirac's *bra-ket* notation [23]. In the graphical calculus, physical processes are depicted by boxes, and the inputs and outputs of these boxes are labeled by *types* which indicate the kind of system on which these boxes act, e. g. one qubit, several qubits, classical data, etc.: see Fig 5. Algebraically, these correspond to $1_A : A \to A$, $f : A \to B$, $g \circ f$, $1_A \otimes 1_B$, $f \otimes 1_C$, $f \otimes g$, $(f \otimes g) \circ h$, respectively. (The convention in these diagrams is that the 'upward' vertical direction represents progress of time.)

Kets, Bras and Scalars A special role is played by boxes with either no input or no output, corresponding to **states** and **costates** respectively (cf. Dirac's kets and bras [23]), which are depicted by triangles. **Scalars** then arise naturally by composing these elements (cf. inner-product or Dirac's bra-ket):

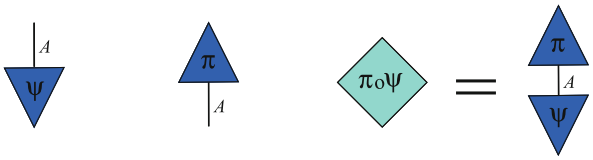

Bell States and Costates The cups and caps which have already appeared in their various guises as axiom and cut links, or in abstraction and application, now take on the rôle of **Bell states and costates**, the fundamental building blocks of quantum entanglement. (Mathematically, they arise as the transpose and co-transpose of the identity, which exist in any finite-dimensional Hilbert space by "map-state duality"). They are represented with triangles encasing the cups and caps to emphasize their **operational character** in the physical interpretation: they represent **preparation and test** of the Bell state.

The formation of **names** and **conames** of arrows (i. e. map-state and map-costate duality) is conveniently depicted thus:

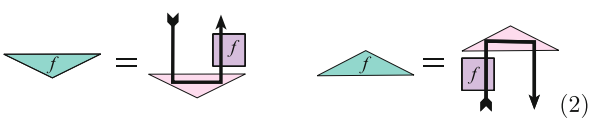

(2)

The key lemma in exposing the quantum information flow in (bipartite) entangled quantum systems [6]:

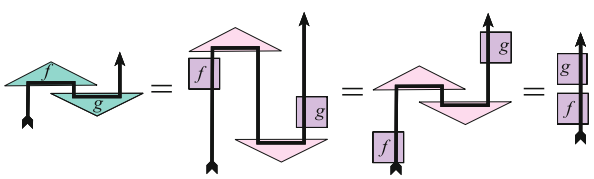

Note in particular the interesting phenomenon of "apparent reversal of the causal order" . While on the left, physically, the state labeled g is firstly prepared, and then the costate labeled f is applied, the global effect is *as if* f itself is applied first, and only then g.

Derivation of Quantum Teleportation This is the most basic application of compositionality in action. The basic quantum mechanical potential for teleportation can be read off immediately from the geometry of Bell states and

costates:

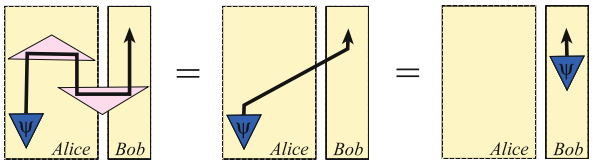

This is not quite the whole story, because of the non-deterministic nature of measurements. But it suffices to introduce a unitary correction. Using the lemma, the full description of teleportation becomes:

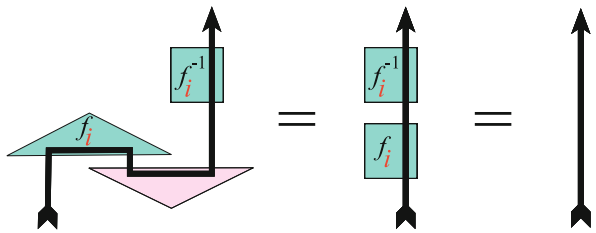

Further Directions

This article has described what is still very much an emerging area of research, rather than surveying an established field. Based on the progress which has already been made, a number of promising directions for future work are apparent, which may lead to important contributions to the general study of agent-based systems:

- The way in which, as illustrated in the previous subsection, the **same** structures arise in a wide diversity of situations, strongly suggests that a Logic of Interaction can be developed, which combines genuine depth with wide applicability.
- Further study of the links between logic and geometry, and of the significance of geometric constraints on informational contexts, looks promising.
- The program of full completeness, and giving equivalent geometric and interactive characterizations of constrained type systems which capture important complexity classes, may open up new insights into currently intractable problems.
- The connections between informatic and physical structures and constraints which have been found in the work on quantum information are promising both as new directions in the foundations of quantum mechanics and the quantum/classical boundary, and in elucidating fundamental structures of informatics.

Bibliography

Primary Literature

1. Abramsky S (1996) Retracing some paths in process algebra. In: Proc of CONCUR 96. Lectures notes in computer science, vol 1119. Springer, Berlin, pp 1–17
2. Abramsky S (2002) Algorithmic game semantics: a tutorial introduction. In: Proof and system-reliability. Kluwer, Dordrecht
3. Abramsky S (2004) High-level methods for quantum computation and information. In: Proc of the 19th annual IEEE symposium on logic in computer science. IEEE Computer Science Press, Los Alamitos pp 410–414
4. Abramsky S (2005) Abstract scalars, loops, and free traced and strongly compact closed categories. In: Proc of CALCO 2005. Springer lecture notes in computer science, vol 3629. Springer, Berlin, pp 1–31
5. Abramsky S (2007) Temperley–Lieb algebras and geometry of interaction. In: Chen G, Kauffman L, Lamonaco S (eds) Mathematics of quantum computing and technology. Chapman and Hall/CRC, Boca Raton, pp 413–458
6. Abramsky S, Coecke B (2004) A categorical semantics of quantum protocols. In: Proc of the 19th Annual IEEE symposium on logic in computer science: LICS 2004. IEEE Computer Society, Los Alamitos, pp 415–425
7. Abramsky S, Coecke B (2005) Abstract physical traces. Theory Appl Categ 14:111–124
8. Abramsky S, Jagadeesan R (1994) New foundations for the geometry of interaction. Inform Comput 111:53–119
9. Abramsky S, Jagadeesan R (1994) Games and full completeness for multiplicative linear logic. J Symb Log 59:543–574
10. Abramsky S, Jagadeesan R (2005) A game semantics for generic polymorphism. Ann Pure Appl Log 133:3–37
11. Abramsky S, McCusker G (1997) Linearity, sharing and state. In: O'Hearn P, Tennent RD (eds) Algol-like languages. Birkhauser, Basel, pp 317–348
12. Abramsky S, McCusker G (1999) Full abstraction for idealized Algol with passive expressions. Theor Comput Sci 227: 3–42
13. Abramsky S, McCusker G (1999) Game semantics. In: Computational logic: Proc of the 1997 Marktoberdorf Summer School. Springer, Berlin, pp 1–56
14. Abramsky S, Mellies P-A (1999) Concurrent games and full completeness. In: Proc of the 14th international symposium on logic in computer science. Computer Society Press of the IEEE, Los Alamitos, pp 431–442
15. Abramsky S, Honda K, McCusker G (1998) A fully abstract game semantics for general references. In: Proc of the 13th international symposium on logic in computer science. Computer Society Press of the IEEE, Los Alamitos, pp 334–344
16. Abramsky S, Jagadeesan R, Malacaria P (2000) Full abstraction for PCF. Inform Comput 163:409–470
17. Abramsky S, Ghica DR, Murawski AS, Ong C-HL (2004) Applying game semantics to compositional software modeling and verification. In: Proc TACAS'04. LNCS, vol 2988. pp 421–435
18. Abramsky S, Ghica DR, Murawski AS, Stark IDB, Ong C-HL (2004) Nominal games and full abstraction for the nu-calculus. In: Proc LICS'04. IEEE Computer Society Press, Los Alamitos, pp 150–159
19. Blute R, Scott PJ (1998) The Shuffle Hopf Algebra and noncommutative full completeness. J Symb Log 63(4):1413–1436
20. Blute R, Hamano M, Scott PJ (2005) Softness of hypercoherences and MALL full completeness. Ann Pure Appl Log 131(1–3):1–63

21. Danos V, Harmer R (2002) Probabilistic game semantics. ACM Trans Comput Log 3(3):359–382
22. Devarajan H, Hughes D, Plotkin G, Pratt V (1999) Full completeness of the multiplicative linear logic of Chu spaces. In: Proc of the 14th Annual IEEE symposium on logic in computer science. pp 234–242
23. Dirac PAM (1947) Principles of quantum mechanics. Oxford University Press, Oxford
24. Ghica DR, McCusker G (2000) Reasoning about idealized algol using regular languages. In: Proc ICALP'00. LNCS, vol 1853. pp 103–116
25. Ghica DR, Murawski AS (2004) Angelic semantics of fine-grained concurrency. In: Proc FOSSACS'04. LNCS, vol 2987. pp 211–225
26. Ghica DR, Murawski AS (2006) Compositional model extraction for higher-order concurrent programs. In: Proc TACAS'06. LNCS
27. Girard J-Y (1987) Linear logic. Theor Comput Sci 50(1):1–102
28. Girard J-Y (1989) Geometry of interaction I: interpretation of system F. In: Ferro R et al (eds) Logic Colloquium '88. Elsevier, Amsterdam, pp 221–260
29. Harmer R, McCusker G (1999) A fully abstract game semantics for finite nondeterminism. In: Proc 14th Annual IEEE symposium on logic in computer science. IEEE Computer Society Press, Los Alamitos
30. Hughes D (2000) Hypergame semantics: full completeness for system F. D Phil Mathematical Sciences. Oxford University, Oxford
31. Hyland JME, Ong C-HL (2000) On full abstraction for PCF: i. models, observables and the full abstraction problem, ii. dialogue games and innocent strategies, iii. a fully abstract and universal game model. Informat Comput 163:285–408
32. Jones VFR (1985) A polynomial invariant for links via von Neumann algebras. Bull Amer Math Soc 129:103–112
33. Kauffman LH (1990) An invariant of regular isotopy. Trans Amer Math Soc 318(2):417–471
34. Kelly GM, Laplaza ML (1980) Coherence for compact closed categories. J Pure Appl Algebra 19:193–213
35. Laird J (1997) Full abstraction for functional languages with control. In: Proc of the 12th Annual symposium on logic in computer science, LICS '97. Extended abstract
36. Laird J (2001) A fully abstract games semantics of local exceptions. In: Proc of the 16th Annual symposium on logic in computer science, LICS '01. Extended abstract
37. Loader R (1994) Models of Lambda calculi and linear logic. D Phil thesis. Oxford University, Oxford
38. Murawski AS, Ong C-HL, Walukiewicz I (2005) Idealized Algol with ground recursion and DPDA equivalence. In: Proc ICALP'05. LNCS, vol 3580. pp 917–929

Books and Reviews

Girard J-Y, Lafont Y, Taylor P (1989) Proof and types. Cambridge Tracts in Theoretical Computer Science

Hindley JR, Seldin JP (1986) Introduction to combinators and the λ-calculus. Cambridge University Press, Cambridge

Kauffman LH (1994) Knots in physics. World Scientific Press, Singapore

Troelstra AS (1992) Lectures on linear logic. Center for the Study of Language and Information Lecture Notes No. 29

Lyapunov–Schmidt Method for Dynamical Systems

André Vanderbauwhede
Department of Pure Mathematics and Computer Algebra, Ghent University, Gent, Belgium

Article Outline

Glossary

Bifurcation A change in the local or global qualitative behavior of a system under the change of one or more parameters; the critical parameter values at which the change takes place are called bifurcation values.

Saddle-node bifurcation A bifurcation where two equilibria, usually one stable and the other unstable, merge and then disappear. A similar bifurcation can also happen with periodic orbits.

Pitchfork bifurcation A bifurcation where an equilibrium loses stability and at the same time two new stable equilibria emerge; this type of bifurcation frequently appears in systems with symmetry.

Period-doubling bifurcation A bifurcation where a fixed point of a mapping (discrete system) loses stability and a stable period-two point emerges. In a continuous system this corresponds to a periodic orbits losing stability while a new periodic orbit, with approximately twice the period of the original one, bifurcates.

Equivariant system A system with symmetries; these symmetries form a group of (usually linear) transformations commuting with the vectorfield or the mapping generating the system.

Reversible system A symmetric system where some of the symmetries involve a reversal of the time.

Definition of the Subject

The mathematical description of the state and the evolution of systems used to model all kind of applications requires in many cases a large or even an infinite number of

variables. In contrast, a number of basic phenomena observed in such systems (in particular bifurcation phenomena) are known to depend only on a very small number of variables. A simple example is the Hopf bifurcation, where the system changes from an equilibrium state to a periodic regime when the parameters cross a certain critical boundary, a transition which in a typical system essentially only involves two of the state variables. There are several ways in which the original system can be reduced to a much smaller system which captures the phenomenon one wants to analyze. There is a dynamical approach, usually in the form of some kind of center manifold reduction, where one tries to find a subsystem (invariant under the full system flow) which contains the core of the bifurcation under consideration. The Lyapunov–Schmidt method takes a rather different point of view, by concentrating purely on the objects under study, such as equilibria or periodic solutions, setting up an equation for these objects in an appropriate space, and then performing a reduction on this equation without any further reference to the original dynamics. This is done by splitting both the unknown and the equation in two parts, solving one of the resulting partial equations and substituting the solutions in the second equation. The outcome is typically a low-dimensional set of algebraic equations – usually called bifurcation equations – the solutions of which correspond in a one-to-one way with the chosen solutions of the original system. The analysis of these bifurcation equations (using appropriate techniques) allows to prove existence and multiplicity results and to describe bifurcation scenarios.

The Liapunov–Schmidt reduction method has a long history, going back to the early 20th century. It has been used in many different forms, and sometimes under different names, to a large variety of problems both within and outside the strict domain of dynamical systems and bifurcation theory: continuous and discrete systems, ordinary and partial differential equations, functional differential equations, integral equations, variational problems, and so on. Within the theory of dynamical systems the method has been widely used (and remains to be used) for the study of bifurcations of equilibria and periodic orbits.

Introduction

The Lyapunov–Schmidt method is, as its name suggests, a method, a technique or procedure which can be applied to a wide class of problems and equations, ranging from local bifurcation problems for equilibria and periodic solutions in finite-dimensional systems to global existence and multiplicity results for solutions of infinite-dimensional problems. The common theme in each of these applications is that of a *reduction*: the original equation is partly solved, and substitution of the partial solution in the remaining equations gives a *reduced equation* of a lower dimension than the original one. In many cases – but there are notorious exceptions – the reduced problem is finite-dimensional, with low dimension (one, in the best case).

Lyapunov–Schmidt Method – An Overview

The basic idea behind the Lyapunov–Schmidt reduction method is a rather simple splitting algorithm which can in very general terms be described as follows. The starting point is an equation

$$F_\lambda(x) = 0 \tag{1}$$

for an unknown x belonging to some space X, and depending on some parameter λ. The mapping F_λ takes the space X into a space Y; in simple cases $Y = X$, and typically X is a subspace of Y (although this is not necessary for the method to work). For example, $X = Y$ may be the phase space, and (1) the condition for $x \in X$ to be an equilibrium of a given dynamical system on X. Or X and Y may be spaces of periodic mappings, while (1) represents the condition for such periodic mapping to be a solution of a given system.

The main ingredients for the splitting are two projections: a projection P of the space X onto a subspace U, and a projection Q of the space Y onto a subspace W; in most applications the subspaces U and W are finite-dimensional. The projection P is used to split the variable x:

$$x = Px + (I - P)x = u + v\,,$$

with $u = Px \in U$ and $v = (I - P)x \in V = (I - P)(X)$. The projection Q allows to split the Eq. (1) into a set of two equations, $(I - Q)F_\lambda(x) = 0$ and $QF_\lambda(x) = 0$. Bringing the two splittings together shows that (1) is equivalent to the system

$$(I - Q)F_\lambda(u + v) = 0\,, \quad QF_\lambda(u + v) = 0\,. \tag{2}$$

Of course these equations heavily depend on the choices for the projections P and Q. The main hypothesis which allows to perform a reduction is that the projections P and Q should be such that the first equation in (2) has for each (u, λ) in a certain domain a unique solution $v = v^*_\lambda(u)$. Depending on the particular problem this solution $v^*_\lambda(u)$ will be obtained by different techniques, but mostly some form of the contraction principle or the implicit function theorem is used. Substituting the solution of the first equation of (2) into the second equation results in the following

reduced equation, usually called the *bifurcation* or *determining equation*:

$$G_\lambda(u) = QF_\lambda\big(u + v^*_\lambda(u)\big) = 0\,. \tag{3}$$

Under the condition that both U and W are finite-dimensional this represents a finite set of equations for a finite number of unknowns. Every solution (u,λ) of (3) corresponds to a solution $(x,\lambda) = (u + v^*_\lambda(u),\lambda)$ of (1), and conversely, every solution (x,λ) of (1) (in the appropriate domain) can be written as $x = u + v^*_\lambda(u)$ for some solution (u,λ) of (3).

The main difficulty in analyzing the Eq. (3) consists in the fact that typically the solution $v^*_\lambda(u)$ of the first equation of (2) is not explicitly known, and therefore also the mapping G_λ (which takes U into W) is not explicitly known. It is therefore important that the set-up (i. e. the projections P and Q used in the reduction) should be such that at least some qualitative properties of v^*_λ and G_λ, such as continuity and differentiability, can be obtained; when the method is used to study local bifurcation problems also a good approximation of these mappings, such as for example a few terms in their Taylor expansion, is required to obtain relevant conclusions.

When the method is used to study (local) bifurcation phenomena one generally assumes that a particular solution (x_0,λ_0) of (1) is known, and one wants to obtain the full solution set of (1) close to this given solution. In such case a convenient choice for the projections P and Q is related to the linearization $L_0 = DF_{\lambda_0}(x_0)$ of the mapping F_λ: the projection P should be such that $U = \mathrm{N}(L_0) = \{x \in X \mid L_0 x = 0\}$, while Q should be a projection of Y onto a subspace W which is complementary to the range $\mathrm{R}(L_0) = \{A_0 x \mid x \in X\}$ of L_0, i. e. $(I - Q)(Y) = \mathrm{R}(L_0)$. With such choice, and under appropriate smoothness conditions for the mapping F_λ, the existence of a unique solution $v = v^*_\lambda(u)$ of the first equation in (2) is guaranteed by the implicit function theorem for all (u,λ) near $(u_0,\lambda_0) = (Px_0,\lambda_0)$. Also the approximation of $v^*_\lambda(u)$ and $G_\lambda(u)$ by Taylor expansions works quite well in such case.

Lyapunov–Schmidt Method – An Example

As a simple but illustrative example on how the method works in practice consider the second order scalar ordinary differential equation

$$\ddot{y} + \alpha\dot{y} + y + g(y,\dot{y}) = 0\,, \tag{4}$$

with $\alpha \in \mathbb{R}$ a parameter and $g(y,\dot{y})$ a smooth function of two variables such that $g(y,\dot{y}) = \mathrm{O}((|y| + |\dot{y}|)^2)$ as $(y,\dot{y}) \to (0,0)$. This equation has the equilibrium solution $y(t) \equiv 0$ for all values of the parameter; the linearization at this equilibrium,

$$\ddot{y} + \alpha\dot{y} + y = 0\,, \tag{5}$$

has for $\alpha = 0$ a two-dimensional space of periodic solutions given by $\{y(t) = a\cos t + b\sin t \mid (a,b) \in \mathbb{R}^2\}$. One may then ask which of these periodic solutions survive when the linear Eq. (5) is replaced by the non-linear Eq. (4). Since (5) is an approximation of (4) only for small $(y,\dot{y})$, and since (5) has periodic solutions only for $\alpha = 0$, the question can be stated more precisely as follows: determine, for all small values of α, all small periodic solutions of (4). An immediate problem which arises is that different periodic solutions may have different periods; indeed, the determination of the period is part of the problem. In the Lyapunov–Schmidt approach this is handled by setting $y(t) = x((1+\tau)t)$, resulting in the new equation

$$(1+\tau)^2\ddot{x} + \alpha(1+\tau)\dot{x} + x + g(x,(1+\tau)\dot{x}) = 0\,; \tag{6}$$

2π-periodic solutions of (6) correspond to $2\pi\,(1+\tau)^{-1}$-periodic solutions of (4). The new problem is then to determine, for all small values of (τ,α), all small 2π-periodic solutions of (6). Observe that this problem is symmetric with respect to phase shifts: if $x(t)$ is a 2π-periodic solution, then so is $x(t+\phi)$, for each ϕ.

To solve this problem a possible 2π-periodic solution $x(t)$ is written in the form

$$x(t) = a\cos t + b\sin t + v(t) = \rho\cos(t+\phi) + v(t)\,,$$

with $v(t)$ a twice continuously differentiable function such that

$$\int_0^{2\pi} v(t)\cos t\,\mathrm{d}t = \int_0^{2\pi} v(t)\sin t\,\mathrm{d}t = 0\,. \tag{7}$$

Taking the symmetry with respect to phase shifts into account one can set $\phi = 0$, such that $x(t)$ takes the form

$$x(t) = \rho\cos t + v(t)\,, \quad \text{with } v \in V^{(2)}\,; \tag{8}$$

here $V^{(2)}$ denotes the space of 2π-periodic C^2-functions satisfying (7), while $V^{(0)}$ will denote the space of continuous such functions. It is an easy exercise to verify that for each $\tilde{v} \in V^{(0)}$ the non-homogeneous linear equation

$$\ddot{v} + v = \tilde{v}(t) \tag{9}$$

has a unique solution $v \in V^{(2)}$, given by

$$v(t) = \int_0^t \sin(t-s)\tilde{v}(s)\mathrm{d}s + \frac{1}{2\pi}\int_0^{2\pi} \sin(t-s)s\tilde{v}(s)\mathrm{d}s\,.$$

For each $(\rho, \nu, \tau) \in \mathbb{R} \times V^{(2)} \times \mathbb{R}$ the term $g(\rho \cos t + \nu(t), (1+\tau)(-\rho \sin t + \dot{\nu}(t)))$ appearing in (6) is decomposed as

$$A(\rho, \nu, \tau) \cos t + B(\rho, \nu, \tau) \sin t + \tilde{g}(\rho, \nu, \tau)(t) \,,$$

with

$$A(\rho, \nu, \tau) = \frac{1}{\pi} \int_0^{2\pi} g(\rho \cos t + \nu(t), (1+\tau)(-\rho \sin t + \dot{\nu}(t))) \cos t \mathrm{d}t \,,$$

$$B(\rho, \nu, \tau) = \frac{1}{\pi} \int_0^{2\pi} g(\rho \cos t + \nu(t), (1+\tau)(-\rho \sin t + \dot{\nu}(t))) \sin t \mathrm{d}t \,,$$

and $\tilde{g}(\rho, \nu, \tau) \in V^{(0)}$. The Eq. (6) then splits into two equations, namely

$$[-(2\tau + \tau^2)\rho + A(\rho, \nu, \tau)] \cos t + \big[\alpha(1+\tau)\rho + B(\rho, \nu, \tau)\big] \sin t = 0 \quad (10)$$

and

$$(1+\tau)^2 \ddot{\nu} + \alpha(1+\tau)\dot{\nu} + \nu + \tilde{g}(\rho, \nu, \tau) = 0 \,. \quad (11)$$

Both equations are satisfied for $(\rho, \nu) = (0, 0)$. Using the solvability property of (9) mentioned above in combination with the implicit function theorem proves that the Eq. (11) has for each sufficiently small (ρ, τ, α) a unique small solution $\nu = \nu^*(\rho, \tau, \alpha) \in V^{(2)}$, with $\nu^*(0, \tau, \alpha) = 0$. Bringing this solution in (10) gives two scalar equations

$$\begin{aligned} &-(2\tau + \tau^2)\rho + A(\rho, \nu^*(\rho, \tau, \alpha), \tau) = 0 \\ \text{and} \quad &\alpha(1+\tau)\rho + B(\rho, \nu^*(\rho, \tau, \alpha), \tau) = 0 \,. \end{aligned} \quad (12)$$

It is easy to verify that $A(\rho, \nu^*(\rho, \tau, \alpha), \tau)$ and $B(\rho, \nu^*(\rho, \tau, \alpha), \tau)$ are of the order $\mathrm{O}(\rho^2)$ as $\rho \to 0$, such that

$$\begin{aligned} &A(\rho, \nu^*(\rho, \tau, \alpha), \tau) = \rho \tilde{A}(\rho, \tau, \alpha) \\ \text{and} \quad &B(\rho, \nu^*(\rho, \tau, \alpha), \tau) = \rho \tilde{B}(\rho, \tau, \alpha) \,, \end{aligned}$$

with $\tilde{A}(0, \tau, \alpha) = \tilde{B}(0, \tau, \alpha) = 0$, and (12) reduces for non-trivial solutions (i. e. solutions with $\rho \neq 0$) to

$$\begin{aligned} &-(2\tau + \tau^2) + \tilde{A}(\rho, \tau, \alpha) = 0 \\ \text{and} \quad &\alpha(1+\tau) + \tilde{B}(\rho, \tau, \alpha) = 0 \,. \end{aligned} \quad (13)$$

For small (ρ, τ, α) these equations can be solved by the implicit function theorem for $(\tau, \alpha) = (\tau^*(\rho), \alpha^*(\rho))$, with $(\tau^*(0), \alpha^*(0)) = (0, 0)$. This means that for each small ρ one can find a parameter value $\alpha = \alpha^*(\rho)$ such that (4) has a periodic solution with "amplitude" equal to $|\rho|$ and given by

$$x^*(\rho)(t) = \rho \cos((1+\tau^*(\rho))t) + \nu^*(\rho, \tau^*(\rho), \alpha^*(\rho))((1+\tau^*(\rho))t) \,;$$

the minimal period of this solution is $2\pi(1+\tau^*(\rho))^{-1}$. In fact, the periodic orbits corresponding to ρ and $-\rho$ coincide, as can be seen from the following. The phase shift symmetry together with $-\rho \cos t = \rho \cos(t + \pi)$ and the uniqueness of the solution $\nu^*(\rho, \tau, \alpha)$ of (11) implies that $\nu^*(-\rho, \tau, \alpha)(t) = \nu^*(\rho, \tau, \alpha)(t + \pi)$, from which one can easily deduce that $A(-\rho, \nu^*(-\rho, \tau, \alpha), \tau) = -A(\rho, \nu^*(\rho, \tau, \alpha), \tau)$, $B(-\rho, \nu^*(-\rho, \tau, \alpha), \tau) = -B(\rho, \nu^*(\rho, \tau, \alpha), \tau)$ and $(\tau^*(-\rho), \alpha^*(-\rho)) = (\tau^*(\rho), \alpha^*(\rho))$. Hence

$$x^*(-\rho)(t) = x^*(\rho)(t + \pi(1+\tau^*(\rho))^{-1}) \,.$$

A further consequence is that $\tau^*(\rho)$ and $\alpha^*(\rho)$ are in fact functions of ρ^2, i. e. $\tau^*(\rho) = \hat{\tau}(\rho^2) = c\rho^2 + \text{h.o.t.}$ and $\alpha^*(\rho) = \hat{\alpha}(\rho^2) = d\rho^2 + \text{h.o.t.}$

It is interesting to notice the difference in treatment between the *auxiliary equation* (11) and the *bifurcation equations* (12). The auxiliary equation (11) is solved by a direct application of the implicit function theorem, while the same implicit function theorem can only be used to solve the bifurcation equations after some manipulation, which in this example simply consists in dividing by ρ to "eliminate" the trivial solution branch.

A further remark is that when the nonlinearity g in (4) does not depend on $\dot{y}$ (i. e. $g = g(y)$) then $\alpha^*(\rho) \equiv 0$ and we have what is usually called a *vertical bifurcation*. For the example this means that the equation $\ddot{y} + \alpha\dot{y} + y + g(y) = 0$ can have periodic solutions only for $\alpha = 0$, and that all small solutions of $\ddot{y} + y + g(y) = 0$ are indeed periodic. To see how this well known fact fits into the Lyapunov–Schmidt reduction observe that the equation $\ddot{y} + y + g(y) = 0$ is *reversible*, in the sense that if $y(t)$ is a solution then so is $y(-t)$. Using this reversibility one can show that $\nu^*(\rho, \tau, 0)(-t) = \nu^*(\rho, \tau, 0)(t)$; therefore $g(\rho \cos t + \nu^*(\rho, \tau, 0)(t))$ is an even function of t, and $B(\rho, \nu^*(\rho, \tau, 0), \tau) = 0$ for all (ρ, τ). This immediately implies that $\alpha^*(\rho) \equiv 0$, as claimed.

The foregoing example forms a very simple case of a *Hopf bifurcation*; this bifurcation will be treated in a more general context in Subsect. "Hopf Bifurcation".

Lyapunov–Schmidt Method – Historical Background

The Lyapunov–Schmidt method has been initiated in the early 20th century by A.M. Lyapunov [48] and E. Schmidt [55] in their (independent) studies of nonlinear integral equations. Later the method was used in diverse

forms in the work of (among many others) Cacciopoli [8], Cesari [9,10,11], Shimizu [56], Cronin [19], Bartle [5], Hale [26], Lewis [46], Vainberg and Trenogin [59], Antosiewicz [2], Rabinowitz [51], Bancroft, Hale and Sweet [4], Mawhin [49], Hall [33,34], Crandall and Rabinowitz [17], etc. Expositions on the method were given by Friedrichs [21], Nirenberg [50], Krasnosel'skii [43,44], Hale [27,28,30], Vainberg and Trenogin [60], and others. In particular the work of Cesari and Hale, who used the name *alternative method*, has put the method on a strong functional analytic basis. Apparently the name *Lyapunov–Schmidt method* was first introduced in the paper [59] of Vainberg and Trenogin; it has been used systematically for this type of approach since the seventies. In the seventies and the eighties the method was used extensively in a wide variety of bifurcation studies, not only in finite-dimensional systems but also in delay equations [29], reaction-diffusion equations [18,20], boundary value problems [17,52], and so on. In particular it forms together with group representation theory and singularity theory the main cornerstone for equivariant bifurcation theory, the study of bifurcation problems in the presence of symmetry (see e. g. the monographs by Sattinger [54], Vanderbauwhede [61], Golubitsky, Stewart and Schaeffer [23,24], and Chossat and Lauterbach [13], and the more recent theory of bifurcations in cell networks initiated by Golubitsky, Pivato and Stewart [25]). Later the method was extended and adapted to study the bifurcation of periodic orbits in Hamiltonian systems [63], reversible systems [41,42] and mappings [15,62], and there have been several numerical implementations [3,6,38]. Lin's method to study bifurcations near homoclinic and heteroclinic connections, based on the work of Lin [47] and developed by Sandstede [53] and Knobloch [40], is maybe from a technical point of view not a pure example of a Lyapunov–Schmidt reduction, but it certainly has all the flavors from it. In a different direction the method was used in combination with topological methods, in particular degree theory, to prove all kind of existence and multiplicity results and to obtain more global bifurcation results; a few references where one can read more about this are [22,37] and [39]. A recent account of the use of the Lyapunov–Schmidt method by mainly Russian authors is given in [57]. In recent years the Lyapunov–Schmidt method has been used to study (among many other topics) the existence of transition layers in singularly perturbed reaction-diffusion equations [31,32,58], and, in combination with a Newton scheme to solve small divisor problems, the existence of periodic and quasi-periodic solutions of nonlinear Schrödinger equations in one or more space dimensions [7,16].

The subsequent sections give a more detailed account of the method as it is used in the local bifurcation theory for equilibria and periodic orbits of finite-dimensional systems.

Lyapunov–Schmidt Method for Equilibria

The easiest example to illustrate the Lyapunov–Schmidt method is to consider the equilibria of a finite-dimensional parameter-dependent system of the form

$$\dot{x} = f(x, \lambda)\,, \tag{14}$$

with $x \in \mathbb{R}^n$ and $f\colon \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^n$ a smooth mapping. These equilibria are given by the solutions of the equation

$$f(x, \lambda) = 0\,. \tag{15}$$

Let (x_0, λ_0) be any solution of this equation; by a translation one can without loss of generality assume that $(x_0, \lambda_0) = (0, 0)$, and hence $f(0, 0) = 0$. The problem is to describe *all* solutions of (15) which are close to the given solution $(0, 0)$.

Let A_0 be the Jacobian $(n \times n)$-matrix $D_x f(0, 0)$. If A_0 is invertible, that is, if the nullspace $\mathrm{N}(A_0)$ is trivial, then by the implicit function theorem the Eq. (15) has for each small parameter value λ a unique small solution $x = x^*(\lambda) \in \mathbb{R}^n$. This means that the equilibrium $x = 0$ for $\lambda = 0$ can be continued for all nearby parameter values λ. Therefore one is mainly interested in the case where A_0 is not invertible and $\mathrm{N}(A_0)$ nontrivial, and that is when the Lyapunov–Schmidt method can be used to reduce the Eq. (15), as follows.

Let P be the orthogonal projection in $\mathbb{R}^n$ onto $\mathrm{N}(A_0)$. This projection allows to split a vector $x \in \mathbb{R}^n$ as $x = Px + (I - P)x = u + v$, with $u = Px \in \mathrm{N}(A_0)$ and $v = (I - P)x \in V = \mathrm{N}(A_0)^\perp$ (the orthogonal complement of $\mathrm{N}(A_0)$). It is known from elementary algebra that $\mathrm{R}(A_0) = \mathrm{N}(A_0^T)^\perp$, where A_0^T is the transpose of A_0, and that $\dim \mathrm{N}(A_0) = \dim \mathrm{N}(A_0^T)$. Hence, if Q is the orthogonal projection of $\mathbb{R}^n$ onto $W = \mathrm{N}(A_0^T)$, then $(I - Q)(\mathbb{R}^n) = \mathrm{R}(A_0)$ and $\mathbb{R}^n = W \oplus \mathrm{R}(A_0)$. Bringing the splitting $x = u + v$ into (15) and using Q to split (15) into $Qf(x, \lambda) = 0$ and $(I - Q)f(x, \lambda)$ results in the following system of two equations, equivalent to (15):

$$\begin{aligned} &\tilde{f}(u, v, \lambda) = Qf(u + v, \lambda) = 0 \\ \text{and} \quad &\hat{f}(u, v, \lambda) = (I - Q)f(u + v, \lambda) = 0\,. \end{aligned} \tag{16}$$

Clearly $\tilde{f}(0, 0, 0) = 0$ and $\hat{f}(0, 0, 0) = 0$; also, $D_v \hat{f}(0, 0, 0)$ is equal to the restriction of $(I - Q)A_0$ to V, which is an invertible linear operator from V onto $\mathrm{R}(A_0)$. From the implicit function theorem one can then conclude

that the *auxiliary equation* $\hat{f}(u, v, \lambda) = 0$ has, for each sufficiently small $(u, \lambda) \in U \times \mathbb{R}^k$, a unique small solution $v = v^*(u, \lambda) \in V$, with $v^* \colon U \times \mathbb{R}^k \to V$ a smooth mapping such that $v^*(0, 0) = 0$. Moreover, differentiating the identity $\hat{f}(u, v^*(u, \lambda), \lambda) = 0$ and using the fact that $A_0 u = 0$ for all $u \in U$ gives $(I - Q)A_0 D_u v^*(0, 0) = 0$, and hence $D_u v^*(0, 0) = 0$. Substitution of the solution $v = v^*(u, \lambda)$ in the first equation of (16) gives the bifurcation equation

$$g(u, \lambda) = \tilde{f}(u, v^*(u, \lambda), \lambda) = Q f(u + v^*(u, \lambda), \lambda) = 0 \,. \tag{17}$$

The bifurcation mapping $g \colon U \times \mathbb{R}^k \to W$ is smooth, and direct calculation shows that $g(0, 0) = 0$ and $D_u g(0, 0) = 0$ (just use the fact that $QA_0 x = 0$ for all $x \in \mathbb{R}^n$).

In order to come to conclusions about the equilibria of (14) one has now to study the bifurcation Eq. (17). Two remarks about this equation are in order here. First, since the solution $v = v^*(u, \lambda)$ is only implicitly defined, the mapping $v^*(u, \lambda)$ and hence also $g(u, \lambda)$ are usually not explicitly known; however, it is in principle possible to obtain the Taylor expansions of these mappings around the point $(u, \lambda) = (0, 0)$. Second, since $D_u g(0, 0) = 0$ it is not possible to use the implicit function theorem to solve (17) for u (or even part of u) as a function of the parameter λ.

Lyapunov–Schmidt Reduction at a Simple Zero Eigenvalue

In the case of just one single parameter (i. e. $k = 1$ and $\lambda \in \mathbb{R}$) one can argue as follows. If at some parameter value one has an equilibrium at which the linearization is invertible, then, as explained before, this equilibrium can be continued for nearby parameter values. This leads to one-parameter branches of equilibria of the form $\{(x^*(\lambda), \lambda) \mid \lambda \in I\}$, with $I \subset \mathbb{R}$ some interval. This continuation stops at parameter values $\lambda = \lambda_0$ where the linearization $A_\lambda = D_x f(x^*(\lambda), \lambda)$ becomes singular, i. e. when zero is an eigenvalue of A_{λ_0}. For one-parameter families of matrices, such as $\{A_\lambda \mid \lambda \in I\}$, this typically happens in such a way that zero will be a *simple eigenvalue* of A_{λ_0}, which means that $\dim \mathrm{N}(A_{\lambda_0}) = 1$ and $\mathbb{R}^n = \mathrm{N}(A_{\lambda_0}) \oplus \mathrm{R}(A_{\lambda_0})$. Shifting $(x^*(\lambda_0), \lambda_0)$ to the origin and using an appropriate linear transformation this motivates the following hypothesis:

$$f(0, 0) = 0 \quad \text{and} \quad A_0 x = (0, \hat{A}_0 v), \quad \forall x = (u, v) \in \mathbb{R} \times \mathbb{R}^{n-1} = \mathbb{R}^n \,, \tag{18}$$

where (as before) $A_0 = D_x f(0, 0)$ and where $\hat{A}_0$ is an *invertible* $(n - 1) \times (n - 1)$-matrix. With this set-up the Lyapunov–Schmidt reduction is straightforward: both P and Q are the projections in $\mathbb{R}^n$ onto the first component, the decomposition $x = (u, v)$ used in (18) is precisely the one needed to apply the reduction, and the resulting bifurcation equation is a single scalar equation

$$g(u, \lambda) = f_1(u, v^*(u, \lambda), \lambda) = 0 \,. \tag{19}$$

As in the general case $g(0, 0) = 0$ and $\partial g / \partial u(0, 0) = 0$. Depending on some further hypotheses this bifurcation equation leads to different bifurcation scenario's which are discussed in the next subsections. It should be noted that under appropriate conditions the bifurcation function $g(u, \lambda)$ also contains some stability information for the equilibria of (14) corresponding to solutions (u, λ) of (19) (see e. g. [23]): such equilibrium is stable if $\partial g / \partial u(u, \lambda) < 0$, and unstable if $\partial g / \partial u(u, \lambda) > 0$.

Saddle-Node Bifurcation

Just keeping the lowest order terms in u and λ the bifurcation function $g(u, \lambda)$ can be written as

$$g(u, \lambda) = \alpha u^2 + \beta \lambda + \text{h.o.t.} \,, \qquad \alpha = \frac{1}{2} \frac{\partial^2 f_1}{\partial u^2}(0, 0, 0) \,, \quad \beta = \frac{\partial f_1}{\partial \lambda}(0, 0, 0) \,. \tag{20}$$

Assuming that $\beta \neq 0$ one can solve (19) for λ as a function of u to obtain

$$\lambda = \lambda^*(u) = -\alpha \beta^{-1} u^2 + \text{h.o.t.} \tag{21}$$

If also $\alpha \neq 0$ then one has, in case $\alpha\beta > 0$, for $\lambda < 0$ two equilibria, one stable and one unstable, which merge for $\lambda = 0$ and then disappear for $\lambda > 0$; in case $\alpha\beta < 0$ one has the opposite scenario. This is called a *saddle-node* or *limit point* bifurcation.

Pitchfork Bifurcation

Suppose that the vectorfield $f(x, \lambda)$ appearing in (14) has some symmetry properties, for example $f(-x, \lambda) = -f(x, \lambda)$. Then clearly one has for each parameter value $\lambda \in \mathbb{R}$ the *trivial* equilibrium $x^*(\lambda) = 0$; setting (as before) $A_\lambda = D_x f(0, \lambda)$ one will typically find isolated parameter values where zero is a simple eigenvalue of A_λ. Assuming that $\lambda = 0$ is such a critical parameter value, and working through the Lyapunov–Schmidt reduction, it is not hard to see that $v^*(-u, \lambda) = -v^*(u, \lambda)$, $g(-u, \lambda) = -g(u, \lambda)$ and $g(0, \lambda) = 0$. Then clearly both coefficients α and β appearing in (20) will be zero, and the conditions for a saddle-node bifurcation are not satisfied. Instead, one

can write the bifurcation function $g(u,\lambda)$ in the form

$$g(u,\lambda) = u\tilde{g}(u,\lambda)\,,$$
$$\text{with}\ \ \tilde{g}(-u,\lambda) = \tilde{g}(u,\lambda)\ \ \text{and}\ \ \tilde{g}(0,0) = 0\,; \quad (22)$$

hence the non-trivial solutions of (19) (these are solutions with $u \neq 0$) are obtained by solving the equation

$$\tilde{g}(u,\lambda) = 0\,. \quad (23)$$

Further calculations show that

$$\tilde{g}(u,\lambda) = \gamma u^2 + \delta\lambda + \text{h.o.t.}\,,$$
$$\gamma = \frac{1}{6}\frac{\partial^3 f_1}{\partial u^3}(0,0,0)\,, \quad \delta = \frac{\partial^2 f_1}{\partial u \partial\lambda}(0,0,0)\,. \quad (24)$$

Under the condition $\delta \neq 0$ it is possible to solve (23) for λ as a function of u, giving

$$\lambda = \lambda^*(u)\,, \quad \text{with}\ \ \lambda^*(0) = 0, \lambda^*(-u) = \lambda^*(u)$$
$$\text{and}\ \ \lambda^*(u) = -\gamma\delta^{-1}u^2 + \text{h.o.t.} \quad (25)$$

When also $\gamma \neq 0$ this results in the following bifurcation scenario when $\delta > 0$ and $\gamma < 0$: the zero equilibrium is stable for $\lambda < 0$ and becomes unstable for $\lambda > 0$; at $\lambda = 0$ two branches of symmetrically related non-trivial equilibria bifurcate from the trivial branch of equilibria; these nontrivial equilibria exist for $\lambda > 0$ and are stable. For other signs of γ and δ similar scenario's hold. This type of bifurcation is known as a *pitchfork bifurcation*.

Transcritical Bifurcation

The arguments used to obtain the pitchfork bifurcation remain partly valid in the following situation which is not necessarily related to symmetry. Suppose that the system (14) has a *given* equilibrium $x^*(\lambda)$ for each $\lambda \in \mathbb{R}$; such equilibrium may be a consequence of symmetry properties (see the case of a pitchfork bifurcation) or of some other property of the system. For example, in population models one always has the zero equilibrium; and in oscillation systems there is usually some basic equilibrium which persists for all parameter values. By a parameter dependent translation one can always arrange that $x^*(\lambda) = 0$ and hence $f(0,\lambda) = 0$ for all λ. Assuming as before that zero is a simple eigenvalue of $A_0 = D_x f(0,0)$ and applying the Lyapunov–Schmidt method one finds that

$$v^*(0,\lambda) = 0, \quad g(0,\lambda) = 0\,,$$
$$\forall\lambda \in \mathbb{R} \Rightarrow g(u,\lambda) = u\tilde{g}(u,\lambda)\,, \quad (26)$$

with $\tilde{g}(u,\lambda)$ a smooth function such that $\tilde{g}(0,0) = 0$. So again non-trivial equilibria are obtained by solving the Eq. (23), in which now the function $\tilde{g}(u,\lambda)$ can be written as

$$\tilde{g}(u,\lambda) = \eta u + \delta\lambda + \text{h.o.t.}\,,$$
$$\eta = \frac{1}{2}\frac{\partial^2 f_1}{\partial u^2}(0,0,0)\,, \quad \delta = \frac{\partial^2 f_1}{\partial u\partial\lambda}(0,0,0)\,. \quad (27)$$

Assuming that $\delta \neq 0$ one can solve (23) for λ as a function of u, resulting in

$$\lambda = \lambda^*(u) = -\eta\delta^{-1}u + \text{h.o.t.} \quad (28)$$

When also $\eta \neq 0$ this represents a *transcritical bifurcation*: for each $\lambda \neq 0$ there exists next to the trivial equilibrium at zero also a non-trivial equilibrium which approaches zero as λ tends to zero. For the same parameter value $\lambda \neq 0$ the trivial and the nontrivial equilibria have opposite stability properties, and there is an *exchange of stability* at $\lambda = 0$. For example, if $\delta > 0$ then the trivial equilibrium is stable for $\lambda < 0$ and unstable for $\lambda > 0$, while the non-trivial equilibrium is unstable for $\lambda < 0$ and stable for $\lambda > 0$.

A remark about the condition $\delta \neq 0$ used to obtain the pitchfork and the transcritical bifurcation is in order here. In both cases it was assumed that zero is a simple eigenvalue of $A_0 = A_{\lambda=0}$, where $A_\lambda = D_x f(0,\lambda)$. This simple eigenvalue can be continued for small values of λ, i. e. the matrix A_λ has, for sufficiently small values of λ, a simple real eigenvalue $\mu(\lambda)$ which depends smoothly on λ and is such that $\mu(0) = 0$. It is then not very hard to prove that the condition $\delta \neq 0$ is equivalent to the *transversality condition* $\mu'(0) \neq 0$ which states that the simple eigenvalue $\mu(\lambda)$ crosses zero with non-zero speed when λ crosses zero. Similar transversality conditions are needed to obtain most of the elementary bifurcations which appear in continuous systems such as (14) or in discrete systems such as considered in the next section.

A further remark has to do with the condition that zero should be a simple eigenvalue of $A_0 = D_x f(0,0)$. In certain cases this condition can be satisfied by restricting a larger system to an appropriate invariant subspace. This happens in particular when one studies bifurcations in equivariant systems; in many such systems the simple zero eigenvalue can be obtained by restricting the full system to invariant subsystems consisting of solutions which are invariant under certain isotropy subgroups. Usually the symmetry also forces the origin to be an equilibrium. Pitchfork or transcritical bifurcations observed for such invariant subsystems form then special cases of the *Equivariant Branching Lemma* [12,61]; see [61] and [24] for more details.

Lyapunov–Schmidt Method in Discrete Systems

The discrete analogue of the continuous system (14) has the form

$$x_{j+1} = f(x_j, \lambda)\,, \quad j = 0, 1, 2, \ldots\,, \tag{29}$$

again with $x_j \in \mathbb{R}^n$ $(j \in \mathbb{N})$ and $f\colon \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^n$ smooth. For a given value λ of the parameter the equilibria of (29) are given by the *fixed points* of $f_\lambda = f(\cdot, \lambda)$, i. e. by the solutions of the equation

$$f(x, \lambda) = x\,. \tag{30}$$

The continuation and bifurcation of such fixed points under a change of the parameter can be studied by the same techniques as the ones used in Sect. "Lyapunov–Schmidt Method for Equilibria": it is sufficient to replace $f(x, \lambda)$ by $F(x, \lambda) = f(x, \lambda) - x$. At a given solution $(x_0, \lambda_0) \in \mathbb{R}^n \times \mathbb{R}^k$ of (29) the linearization $L_0 = D_x F(x_0, \lambda_0)$ of $F_{\lambda_0} = F(\cdot, \lambda_0)$ is given by $L_0 = A_0 - I_n$, where $A_0 = D_x f(x_0, \lambda_0)$ and I_n is the identity matrix on $\mathbb{R}^n$. Therefore, if 1 is not an eigenvalue of A_0 then the fixed point x_0 can be continued for all λ sufficiently close to λ_0, and if 1 is a simple eigenvalue of A_0 (and $k = 1$) then, depending on some further conditions, one may have at $\lambda = \lambda_0$ a saddle-node, a pitchfork bifurcation or a transcritical bifurcation of fixed points.

Discrete systems such as (29) arise for example frequently in mathematical biology and mathematical economics. A special situation arises when the mapping f in (29) is a Poincaré map constructed near a given periodic orbit Γ_0 of a continuous parameter-dependent autonomous $(n+1)$-dimensional system (see e. g. [45]); in that case a fixed point of f_λ corresponds to a periodic orbit of the original continuous system, usually with the origin corresponding to Γ_0. The periodic orbit Γ_0 has always 1 as a Floquet multiplier; if this multiplier is *simple*, then 1 is not an eigenvalue of the linearization at the origin of the Poincaré map, and the periodic orbit Γ_0 can be continued for nearby parameter values. If $k = 1$ (a single parameter) and the multiplier 1 of Γ_0 has multiplicity two, then 1 will be a simple eigenvalue of the linearization of the Poincaré map, and one will typically obtain a saddle-node of fixed points of the Poincaré map, corresponding to a saddle-node of periodic orbits in the original continuous system. In particular cases such as the presence of symmetry one may also see a transcritical or pitchfork bifurcation of periodic orbits.

Periodic Points in Discrete Systems

Next to the bifurcation of fixed points for (29) one can also study the bifurcation of *periodic points* from fixed points. For any $\lambda \in \mathbb{R}^k$ and any integer $q \geq 1$ a q-periodic point of $f_\lambda = f(\cdot, \lambda)$ is a point $x \in \mathbb{R}^n$ such that the orbit $\{x_j \mid j \in \mathbb{N}\}$ of (29) starting at $x_0 = x$ has the property that $x_{j+q} = x_j$ for all $j \in \mathbb{N}$. In case f_λ is a Poincaré map associated to a periodic solution with minimal period $T_0 > 0$ of a continuous system then such q-periodic point corresponds to a periodic orbit of the continuous system with minimal period near qT_0. Such periodic solution is called a *subharmonic solution* (when $q \geq 2$); the bifurcation of periodic points from fixed points for discrete systems such as (29) is therefore strongly related to subharmonic bifurcation in continuous systems.

The periodicity condition $x_{j+q} = x_j$ (for all $j \in \mathbb{N}$) is equivalent to the condition that $x_q = x_0 = x$, or to

$$f_\lambda^q(x) = x, \qquad f_\lambda^q = f_\lambda \circ f_\lambda \circ \cdots \circ f_\lambda \quad (q \text{ times}). \tag{31}$$

Hence a q-periodic point of f_λ is a fixed point of f_λ^q (the q-th iterate of f_λ), and therefore one could in principle study the bifurcation of q-periodic points of f_λ by studying the bifurcation of fixed points of f_λ^q. Observe that a fixed point of f_λ is automatically also a fixed point of f_λ^q, for any $q \geq 1$.

However, there is a better approach which takes explicitly into account an important property of (31), namely that if x is a solution then so are all the other points of the orbit $\{x_j = f_\lambda^j(x) \mid j \in \mathbb{N}\}$ generated by x. The basic idea is the following: instead of considering all orbits of (29) and then finding all those which are q-periodic, one considers all q-periodic sequences with elements in $\mathbb{R}^n$ and then tries to determine which ones are actually orbits of the system (29). To work this out one has to introduce (for the chosen value of q) the space

$$\mathcal{O}_q = \left\{z = (x_j)_{j\in\mathbb{Z}} \mid x_j \in \mathbb{R}^n,\ x_{j+q} = x_j,\ \forall j \in \mathbb{Z}\right\}. \tag{32}$$

This *orbit space* is finite-dimensional, with dimension qn. Any map $g\colon \mathbb{R}^n \to \mathbb{R}^n$ can be lifted to a map $\hat{g}\colon \mathcal{O}_q \to \mathcal{O}_q$ by setting

$$\hat{g}(z) = (g(x_j))_{j\in\mathbb{Z}}\,, \quad \forall z = (x_j)_{j\in\mathbb{Z}} \in \mathcal{O}_q\,.$$

The *shift operator* is the linear mapping $\sigma\colon \mathcal{O}_q \to \mathcal{O}_q$ defined by

$$(\sigma \cdot z)_j = x_{j+1}\,, \quad \forall j \in \mathbb{Z}\,, \quad \forall z = (x_j)_{j\in\mathbb{Z}} \in \mathcal{O}_q\,.$$

Observe that $\sigma^q = I_{\mathcal{O}_q}$, and hence σ generates a $\mathbb{Z}_q$-action on $\mathcal{O}_q$ ($\mathbb{Z}_q$ is the cyclic group of order q, i. e. $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$). Also $\hat{g} \circ \sigma = \sigma \circ \hat{g}$ for any $g\colon \mathbb{R}^n \to \mathbb{R}^n$. And finally, identifying $\mathcal{O}_q$ with $(\mathbb{R}^n)^q$ and using the standard scalar product it is easily seen that $\langle \sigma \cdot z_1, \sigma \cdot z_2\rangle = \langle z_1, z_2\rangle$ for all $z_1, z_2 \in \mathcal{O}_q$, i. e. σ is orthogonal and $\sigma^T = \sigma^{-1}$.

A q-periodic sequence $z = (x_j)_{j\in\mathbb{Z}} \in \mathcal{O}_q$ forms an orbit under the iteration of f_λ if and only if $f_\lambda(x_j) = x_{j+1}$ for all $j \in \mathbb{Z}$, that is, if and only if

$$\hat{f}_\lambda(z) = \sigma \cdot z\,. \tag{33}$$

The advantage of this equation over the equivalent Eq. (31) is that (33) is $\mathbb{Z}_q$-equivariant: both sides of the equation commute with σ, and if $(z,\lambda) \in \mathcal{O}_q \times \mathbb{R}^k$ is a solution, then so is $(\sigma^j \cdot z, \lambda)$ for any $j \in \mathbb{Z}$. This equivariance plays an important role in the bifurcation analysis for q-periodic points of (29) since it puts strong restrictions on the possible form of the bifurcation equation. The $\mathbb{Z}_q$-equivariance of (33) also fits nicely together with the Lyapunov–Schmidt reduction, as will be shown next.

Lyapunov–Schmidt Reduction for Periodic Points

Assume that (30) has a solution $(x_0, \lambda_0) \in \mathbb{R}^n \times \mathbb{R}^k$; as before one can without loss of generality translate this solution to the origin and assume that $f(0,0) = 0$. Consider then, for a given integer $q \geq 1$, the following problem:

(P_q) Determine, for all small values of the parameter λ, all small q-periodic points of f_λ.

The bifurcation equation for this problem can be obtained by a Lyapunov–Schmidt reduction near the solution $(z,\lambda) = (0,0)$ of (33), as follows. Linearization of (33) at $(0,0)$ gives the equation

$$\hat{A}_0 \cdot z = \sigma \cdot z\,, \quad A_0 = D_x f(0,0)\,, \tag{34}$$

and the corresponding linear operator $L_0 = \hat{A}_0 - \sigma\colon \mathcal{O}_q \to \mathcal{O}_q$. The nullspace $\mathrm{N}(L_0)$ consists of the q-periodic orbits of the linear discrete system $x_{j+1} = A_0 x_j$ and can be described as follows. Let

$$U = \mathrm{N}\big(A_0^q - I_n\big)\,, \quad S = A_0\big|_U\colon U \to U \quad \text{and} \quad \zeta\colon U \to \mathcal{O}_q,\ u \mapsto \zeta(u) = (S^j u)_{j\in\mathbb{Z}}\,. \tag{35}$$

Then $\mathrm{N}(L_0) = \zeta(U)$, and ζ is an isomorphism between U and $\mathrm{N}(L_0)$. Observe that the subspace U of $\mathbb{R}^n$ is the sum of the eigenspaces of A_0 corresponding to all eigenvalues $\mu \in \mathbb{C}$ of A_0 for which $\mu^q = 1$, i. e. all eigenvalues μ which are q-th roots of unity. A further observation is that the linear operator S generates a $\mathbb{Z}_q$-action on U (since $S^q = I_U$ by definition of U and S), and that

$$\zeta(Su) = \sigma \cdot \zeta(u)\,, \quad \forall u \in U\,. \tag{36}$$

The description of $\mathrm{R}(L_0) = \mathrm{N}(L_0^T)^\perp$ and the necessary projections can be somewhat simplified by assuming that all eigenvalues of A_0 which are q-th roots of unity are *semi-simple*, an assumption which implies that $\mathbb{R}^n = \mathrm{N}(A_0^q - I_n) \oplus \mathrm{R}(A_0^q - I_n) = U \oplus Y$; writing $x \in \mathbb{R}^n$ as $x = (u, y)$, with $u \in U$ and $y \in Y$, and constructing an appropriate basis of $\mathbb{R}^n$ one can then put A_0 in the following (diagonal) form:

$$A_0 x = A_0(u,y) = (Su, \tilde{A}_0 y)\,, \quad \forall x = (u,y) \in U \times Y = \mathbb{R}^n\,, \tag{37}$$

where S is orthogonal, i. e. $S^T S = I_U$, and where $\tilde{A}_0\colon Y \to Y$ is such that $\tilde{A}_0^q - I_Y$ is invertible. Under these assumptions (see [62] for the general case) it is easily seen that $\mathrm{N}(L_0^T) = \mathrm{N}(\widehat{A}_0^T - \sigma^T)$ is equal to $\mathrm{N}(L_0)$ (use the fact that $\sigma^T = \sigma^{-1}$ and $S^T = S^{-1}$). Therefore $\mathrm{R}(L_0)$ is the orthogonal complement in $\mathcal{O}_q = (\mathbb{R}^n)^q$ of $\mathrm{N}(L_0) = \zeta(U)$, and every $z \in \mathcal{O}_q$ has a unique decomposition of the form

$$z = \zeta(u) + v\,, \quad \text{with } u \in U \text{ and } v \in V = \mathrm{R}(L_0)\,. \tag{38}$$

In particular, going back to (33) and replacing x by $\zeta(u) + v$, one can write

$$\hat{f}_\lambda(\zeta(u) + v) = \zeta(g(u,v,\lambda)) + h(u,v,\lambda)\,,$$

with $g\colon U \times V \times \mathbb{R}^k \to U$ and $h\colon U \times V \times \mathbb{R}^k \to V$ smooth mappings such that $g(0,0,0) = 0$, $h(0,0,0) = 0$, $D_u g(0,0,0) = S$, $D_v g(0,0,0) = 0$, $D_u h(0,0,0) = 0$, $D_v h(0,0,0) = \hat{A}_0\big|_V$, and

$$\begin{aligned} g(Su, \sigma \cdot v, \lambda) &= S g(u,v,\lambda)\,, \\ h(Su, \sigma \cdot v, \lambda) &= \sigma \cdot h(u,v,\lambda)\,. \end{aligned} \tag{39}$$

This last property is a consequence of (36), the $\mathbb{Z}_q$-equivariance of $\hat{f}_\lambda$, and the fact that σ leaves the decomposition $\mathcal{O}_q = \zeta(U) \oplus V$ invariant since σ commutes with L_0.

It follows from the foregoing that the Eq. (33) can be split in the two equations

$$g(u,v,\lambda) = Su \quad \text{and} \quad h(u,v,\lambda) = \sigma \cdot v\,. \tag{40}$$

Since $h(0,0,0) = 0$, $D_v h(0,0,0) = \hat{A}_0\big|_V$ and $L_0 = \hat{A}_0 - \sigma$ is invertible on $V = \mathrm{R}(L_0)$, the second equation in (40) can, locally near $(u,v,\lambda) = (0,0,0)$, be solved by the implicit function theorem for $v = v^*(u,\lambda)$; the mapping $v^*\colon U \times \mathbb{R}^k$ is smooth, with $v^*(0,0) = 0$ and $D_u v^*(0,0) = 0$. Also, (39) and the uniqueness of solutions given by the implicit function theorem imply that

$$v^*(Su, \lambda) = \sigma \cdot v^*(u,\lambda)\,. \tag{41}$$

Substituting the solution $v = v^*(u,\lambda)$ into the first equation of (40) gives the bifurcation equation

$$g^*(u,\lambda) = Su\,, \quad \text{with } g^*(u,\lambda) = g(u, v^*(u,\lambda), \lambda)\,. \tag{42}$$

This is an equation on the subspace U of $\mathbb{R}^n$; the mapping $g^*\colon U \times \mathbb{R}^k \to U$ is smooth, with $g^*(0,0) = 0$, $D_u g^*(0,0) = S$ and

$$g^*(Su,\lambda) = Sg^*(u,\lambda)\,, \quad \forall (u,\lambda) \in U \times \mathbb{R}^k\,. \tag{43}$$

This last property ensures that the bifurcation equation (42) is $\mathbb{Z}_q$-equivariant, i. e. commutes with the $\mathbb{Z}_q$-action generated by S on U, and if $(u,\lambda) \in U \times \mathbb{R}^k$ is a solution then so is $(\tilde{u},\lambda)$ for any $\tilde{u}$ belonging to the $\mathbb{Z}_q$-orbit $\{S^j u \mid j \in \mathbb{Z}\}$ generated by u.

Period-Doubling Bifurcation

Suppose that $f(0,0) = 0$ and that $A_0 = D_x f(0,0)$ has -1 as a simple eigenvalue, while $+1$ is not an eigenvalue. Then, taking $q = 2$ in the foregoing, dim $U = 1$, $Su = -u$ for all $u \in U$, and the bifurcation equation (42) is a scalar equation:

$$g^*(u,\lambda) = -u\,, \quad \text{with } \frac{\partial g^*}{\partial u}(0,0) = -1 \quad \text{and } g^*(-u,\lambda) = -g^*(u,\lambda)\,. \tag{44}$$

Due to the symmetry the discussion of this bifurcation equations runs completely analogous to the one in Subsect. "Pitchfork Bifurcation" and leads, for $k = 1$ and under appropriate genericity conditions, to a pitchfork bifurcation. However, the interpretation of the bifurcation diagram is different in this case: the trivial solution $u = 0$ of (42) corresponds to a fixed point $x_0(\lambda)$ of f_λ which changes stability as λ crosses zero, while the symmetric pairs $\pm\hat{u}(\lambda)$ of nontrivial solutions which exist either for $\lambda < 0$ or for $\lambda > 0$ correspond to the two points $\hat{x}(\lambda)$ and $f_\lambda(\hat{x}(\lambda))$ of a period-two orbit of f_λ. Such bifurcation is known as a *period-doubling bifurcation* or a *flip bifurcation*. In case f_λ represents a Poincaré map near a given (nontrivial) periodic orbit with minimal period T_0 of a continuous system such period-doubling bifurcation means the following: the given periodic orbit persists for nearby parameter values, with minimal period T_λ near T_0, but when the parameter crosses the critical value there bifurcates a branch of periodic solutions with minimal period near $2T_0$; so also in the continuous system one sees a period-doubling.

Bifurcation of q-Periodic Points for $q \geq 3$

Next let $q \geq 3$ and assume that $f(0,0) = 0$ and that A_0 has a pair of simple eigenvalues

$$\exp(\pm i\theta_0)\,, \quad \text{with } \theta_0 = \frac{2\pi p}{q}, \quad \gcd(p,q) = 1\,. \tag{45}$$

Assume also that A_0 has no other eigenvalues μ with $\mu^q = 1$; in particular $+1$ and -1 are no eigenvalues. Then dim $U = 2$ and the linear operator S in (37) has the matrix form

$$S = \begin{pmatrix} \cos\theta_0 & -\sin\theta_0 \\ \sin\theta_0 & \cos\theta_0 \end{pmatrix}.$$

Identifying $u = (u_1,u_2) \in U = \mathbb{R}^2$ with $z = u_1 + \mathrm{i}z_2 \in \mathbb{C}$ the expression (37) for A_0 takes the more explicit form

$$A_0 x = A_0(z,y) = \left(\mathrm{e}^{\mathrm{i}\theta_0} z, \tilde{A}_0 y\right), \quad \forall x = (z,y) \in \mathbb{C} \times Y = \mathbb{R}^n\,. \tag{46}$$

The Lyapunov–Schmidt reduction for q-periodic points as explained in Subsect. "Lyapunov–Schmidt Reduction for Periodic Points" leads to a complex bifurcation equation

$$g^*(z,\lambda) = \mathrm{e}^{\mathrm{i}\theta_0} z\,, \quad \text{with } g^*(z,0) = \mathrm{e}^{\mathrm{i}\theta_0} z + \mathrm{O}(|z|^2) \quad \text{and } g^*\left(\mathrm{e}^{\mathrm{i}\theta_0} z,\lambda\right) = \mathrm{e}^{\mathrm{i}\theta_0} g^*(z,\lambda)\,. \tag{47}$$

It can be shown (see e. g. [24] or [62]) that the $\mathbb{Z}_q$-equivariance of g^* implies that

$$g^*(z,\lambda) = \left(e^{\mathrm{i}\theta_0} + \phi_1(|z|^2,\lambda)\right)z + \phi_2(|z|^2,\lambda)\bar{z}^{q-1} + \mathrm{O}(|z|^{q+1}), \tag{48}$$

with ϕ_1 and ϕ_2 complex valued functions and $\phi_1(0,0) = 0$. The bifurcation equation (47) has for each $\lambda \in \mathbb{R}^k$ the trivial solution $z = 0$, corresponding to the continuation of the fixed point $x = 0$ at $\lambda = 0$ (remember that it was assumed that $+1$ is not an eigenvalue of A_0). To obtain nontrivial solutions one can multiply (47) by $\bar{z}$, set $z = \rho e^{\mathrm{i}\theta}$ and divide by ρ^2; the resulting equation has the form

$$\tilde{g}(\rho,\theta,\lambda) = \phi_1(\rho^2,\lambda) + \phi_2(\rho^2,\lambda)\rho^{q-2}\mathrm{e}^{-\mathrm{i}q\theta} + \mathrm{O}(\rho^q) = 0\,, \tag{49}$$

with $\tilde{g}(\rho,\theta+\theta_0,\lambda) = \tilde{g}(\rho,\theta,\lambda)$ and $\tilde{g}(-\rho,\theta+\pi,\lambda) = \tilde{g}(\rho,\theta,\lambda)$. The standard way to handle (49) is to assume $k = 2$ (i. e. $\lambda = (\lambda_1,\lambda_2) \in \mathbb{R}^2$) together with the transversality condition

$$\frac{\partial(\Re\phi_1, \Im\phi_1)}{\partial(\lambda_1,\lambda_2)}(0,0) \neq 0\,. \tag{50}$$

Then (49) can be solved for $\lambda = \lambda_q^*(\rho,\theta)$, with $\lambda_q^*(0,\theta,\lambda) = 0$, $\lambda_q^*(\rho,\theta+\theta_0,\lambda) = \lambda_q^*(\rho,\theta,\lambda)$ and $\lambda_q^*(-\rho,\theta+\pi,\lambda) = \lambda_q^*(\rho,\theta,\lambda)$. Both the transversality condition and the solution set $\{(\rho,\theta,\lambda_q^*(\rho,\theta)) \mid (\rho,\theta) \in \mathbb{R}^2\}$ require some further explanation.

Since $\exp(\pm\theta_0)$ are simple eigenvalues of A_0 it follows that these eigenvalues can be smoothly continued for nearby parameter values, i. e. A_λ has for small λ a pair

of simple eigenvalues $\{\mu^*(\lambda), \bar{\mu}^*(\lambda)\}$ near $\exp(\pm\theta_0)$. The transversality condition (50) means that the mapping $\lambda \mapsto \mu^*(\lambda)$ is a local diffeomorphism near $\lambda = 0$. Under this condition the solution $\lambda = \lambda_q^*(\rho, \theta)$ of (49) implies that for each λ in the range

$$\mathcal{A}_q = \{\lambda_q^*(\rho, \theta) \mid (\rho, \theta) \in \mathbb{R}^2\}$$

the mapping f_λ has at least one q-periodic orbit near $x = 0$. For $q \geq 5$ the set $\mathcal{A}_q$ typically forms a horn-like domain in the λ-plane, with tip at the origin; higher values of q correspond to sharper tips. Indeed, it follows from (48) that

$$\lambda_q^*(\rho, \theta) = \rho^2\alpha(\rho^2) + \rho^{q-2}e^{-q\theta}\beta(\rho^2) + \mathrm{O}(\rho^q)\,,$$

for some complex valued functions $\alpha(\rho^2)$ and $\beta(\rho^2)$; if $\alpha(0) \neq 0$ this shows that the "thickness" of $\mathcal{A}_q$ at a distance d from the origin is of the order $d^{(q-2)/2}$ (for $q = 3$ and $q = 4$ the situation is different, see e. g. [62] or [36]). The horn-like regions $\mathcal{A}_q$ are known as *Arnol'd tongues*. For parameter values λ in the interior of $\mathcal{A}_q$ and close to $\lambda = 0$ the mapping f_λ has two q-periodic orbits which annihilate each other in a saddle-node bifurcation when the parameter crosses the boundary of $\mathcal{A}_q$. Moreover, the hypotheses imply that f_λ undergoes a *Neimark–Sacker bifurcation*: when λ crosses the curve defined by $|\mu^*(\lambda)| = 1$ the stability of the fixed point changes and an invariant curve bifurcates from the fixed point. The Arnol'd tongue $\mathcal{A}_q$ is contained in the region where an invariant curve exists, and the two periodic orbits are contained in this invariant curve; on the curve, one periodic orbit is stable, the other unstable.

When f_λ represents a Poincaré map the bifurcation picture as described in the foregoing can be lifted to the original continuous system, leading to the bifurcation of invariant tori and subharmonics. For more details see e. g. [36] or [45]. For a discussion of the case where the mappings f_λ are reversible see [15].

Lyapunov–Schmidt Method for Periodic Solutions

The ideas described in Sect. "Lyapunov–Schmidt Method in Discrete Systems and in particular in Subsect. "Lyapunov–Schmidt Reduction for Periodic Points" can be adapted for the treatment of the bifurcation of periodic orbits from equilibria in continuous systems (14); however, there are a few complications which arise for this continuous case. First, the orbit space consists of continuous periodic functions defined on the real line, hence this orbit space is infinite-dimensional; this necessitates a more careful handling of splittings and projections. Second, when looking for periodic solutions of (14) the (minimal) period is one of the unknowns of the problem. As will be seen in the treatment which follows the finite symmetry group $\mathbb{Z}_q$ which plays an important role in the bifurcation analysis of periodic orbits in discrete systems will be replaced by the continuous circle group S^1 in the case of continuous systems.

In order to get a more precise description of the problem, assume that $f(0, 0) = 0$ and consider the linearized equation

$$\dot{x} = A_0 x\,, \tag{51}$$

(with $A_0 = D_x f(0, 0)$ as before). The solutions have the form $x(t) = \mathrm{e}^{A_0 t}x_0$ with $x_0 \in \mathbb{R}^n$; such solution is T_0-periodic (for some $T_0 > 0$) if $x_0 \in \mathrm{N}(\mathrm{e}^{A_0 T_0} - I_n)$. In order to exclude a possible bifurcation of equilibria one should require that zero is not an eigenvalue of A_0, which entails that without further loss of generality one can also assume that $f(0, \lambda) = 0$ for all λ. By a time rescale one can put $T_0 = 2\pi$. The nullspace $\mathrm{N}(\mathrm{e}^{2\pi A_0} - I_n)$ is then given by the direct sum of the (real) eigenspaces corresponding to all pairs of eigenvalues of the form $\pm ik$ $(k = 1, 2, \ldots)$. A further simplification in the formulation of the Lyapunov–Schmidt reduction occurs if one assumes that all such eigenvalues are semisimple, leading to the following hypotheses about (14) which will be assumed further on:

(H) $f(0, \lambda) = 0$ for all λ, $A_0 = D_x f(0, 0)$ is invertible, and 1 is a semisimple eigenvalue of $\mathrm{e}^{2\pi A_0}$, i. e. $\mathbb{R}^n = \mathrm{N}(\mathrm{e}^{2\pi A_0} - I_n) \oplus \mathrm{R}(\mathrm{e}^{2\pi A_0} - I_n) = U \oplus Y$, where $U = \mathrm{N}(\mathrm{e}^{2\pi A_0} - I_n)$ and $Y = \mathrm{R}(\mathrm{e}^{2\pi A_0} - I_n)$.

By an appropriate choice of basis one can then more explicitly assume that

$$A_0 x = A_0(u, y) = (Su, \tilde{A}_0 y)\,, \qquad \forall x = (u, y) \in U \times Y = \mathbb{R}^n\,, \tag{52}$$

with $S\colon U \to U$ and $\tilde{A}_0\colon Y \to Y$ linear and invertible, S anti-symmetric ($S^T = -S$), $\mathrm{e}^{2\pi S} = I_U$ and $\mathrm{N}(\mathrm{e}^{2\pi \tilde{A}_0} - I_Y) = \{0\}$. The space of 2π-periodic solutions of (51) is then given by $\{\mathrm{e}^{St}u \mid u \in U\}$. The question one wants to answer is then which of these periodic solutions survive when we turn on the nonlinearities and change parameters. More precisely, the problem can be formulated as follows:

(P) Determine, for all $(T, \lambda) \in \mathbb{R} \times \mathbb{R}^k$ near $(2\pi, 0)$, all small T-periodic solutions of (14).

Lyapunov–Schmidt Reduction for the Problem (P)

Analogically to the treatment for discrete systems in Sect. "Lyapunov–Schmidt Method in Discrete Systems" the first step should be to reformulate (P) in an appropriate orbit space; to do so one has to fix the period, which is done as

follows. Let $\hat{x}(t)$ be a T-periodic solution of (14), define $\tau \in \mathbb{R}$ by $(1+\tau)T = 2\pi$, and set $\tilde{x}(t) = \hat{x}((1+\tau)^{-1}t)$. Then $\tilde{x}(t)$ is a 2π-periodic solution of the equation

$$(1+\tau)\dot{x} = f(x,\lambda)\,. \tag{53}$$

Conversely, if $\tilde{x}(t)$ is a 2π-periodic solution of (53) then $\hat{x}(t) = \tilde{x}((1+\tau)t)$ is a T-periodic solution of (14), with $T = 2\pi(1+\tau)^{-1}$. Therefore the problem (P) is equivalent to

(P*) Determine, for all small $(\tau,\lambda) \in \mathbb{R}\times\mathbb{R}^k$, all small 2π-periodic solutions of (53).

The orbit space for this problem is the space $C^0_{2\pi}$ of all continuous 2π-periodic mappings $x\colon \mathbb{R} \to \mathbb{R}^n$; this is a Banach space when one uses the norm $\|x\|^0_{2\pi} = \sup\{\|x(t)\| \mid t \in \mathbb{R}\}$. On $C^0_{2\pi}$ one can define a natural S^1-action (where $S^1 = \mathbb{R}/2\pi\mathbb{R}$ is the circle group) by the *phase shift operator*

$$\Sigma\colon S^1 \times C^0_{2\pi} \to C^0_2\pi,\ (\phi,x) \mapsto \Sigma_\phi(x)\,,$$
$$\text{with}\ \ \Sigma_\phi(x)(t) = x(t+\phi)\,. \tag{54}$$

The infinitesimal generator of this S^1-action is defined on the dense subspace $C^1_{2\pi}$ of continuously differentiable 2π-periodic mappings $x\colon \mathbb{R} \to \mathbb{R}^n$ and is given by

$$D_t\colon C^1_{2\pi} \to C^0_{2\pi},\ x \mapsto D_t x = \left.\frac{\mathrm{d}}{\mathrm{d}\phi}\right|_{\phi=0} \Sigma_\phi(x)\,,$$
$$(D_t x)(t) = \dot{x}(t),\ \forall (x,t) \in C^1_{2\pi}\times\mathbb{R}\,. \tag{55}$$

The subspace $C^1_{2\pi}$ is itself a Banach space when one uses the norm $\|x\|^1_{2\pi} = \|x\|^0_{2\pi} + \|D_t x\|^0_{2\pi}$; the space $C^1_{2\pi}$ is continuously imbedded in $C^0_{2\pi}$. Also, any continuous mapping $g\colon \mathbb{R}^n \to \mathbb{R}^n$ can be lifted to a mapping on $C^0_{2\pi}$ by defining

$$\hat{g}\colon C^0_{2\pi} \to C^0_{2\pi},\ x \mapsto \hat{g}(x)\,,$$
$$\text{with}\ \ \hat{g}(x)(t) = g(x(t))\,,\quad \forall (x,t) \in C^0_{2\pi}\times\mathbb{R}\,.$$

Moreover, such a lifting commutes with the shift operator: $\hat{g}\circ\Sigma_\phi = \Sigma_\phi \circ \hat{g}$ for all $\phi \in S^1$.

Using the foregoing definitions the problem (P*) translates into finding all small solutions $(x,\tau,\lambda) \in C^1_{2\pi}\times\mathbb{R}\times\mathbb{R}^k$ of the equation

$$\widehat{f}_\lambda(x) = (1+\tau)D_t x\,. \tag{56}$$

(Observe the analogy and the differences with (33)). This equation is S^1-equivariant: both sides commute with Σ_ϕ for all $\phi \in S^1$; therefore, if (x,τ,λ) is a solution, then so is $(\Sigma_\phi(x),\tau,\lambda)$ for all $\phi \in S^1$.

Next define $F\colon C^1_{2\pi}\times\mathbb{R}\times\mathbb{R}^k \to C^0_{2\pi}$ by $F(x,\tau,\lambda) = \widehat{f}_\lambda - (1+\tau)D_t$; then $F(0,\tau,\lambda) = 0$, and the linearization $L_0 = D_x F(0,0,0)\colon C^1_{2\pi} \to C^0_{2\pi}$ is given by $L_0 = \widehat{A}_0 - D_t$. The nullspace $\mathrm{N}(L_0) \subset C^1_{2\pi}$ is the space of all 2π-periodic solutions of (51) which was described before; hence there exists an isomorphism $\zeta\colon U \to \mathrm{N}(L_0)$ given by

$$\zeta(u)(t) = \mathrm{e}^{St}u\,,\quad \forall u \in U,\ \forall t \in \mathbb{R}\,. \tag{57}$$

It follows from the definition of U and $S = A_0\big|_U$ that S is the infinitesimal generator of an S^1-action on U given by

$$(\phi,u) \in S^1\times U \mapsto \mathrm{e}^{S\phi}u \in U\,.$$

Also, $\zeta(e^{S\phi}u) = \Sigma_\phi\zeta(u)$ and $\zeta(Su) = D_t\zeta(u)$ for all $(\phi,u) \in S^1\times U$.

In order to describe $\mathrm{R}(L_0)$ one has to find those $z \in C^0_{2\pi}$ for which the non-homogeneous linear 2π-periodic differential equation

$$\dot{x} = A_0 x + z(t) \tag{58}$$

has a 2π-periodic solution. The *Fredholm alternative* for this problem (see e. g. [28] or [61]) states that this will be the case if and only if

$$\int_0^{2\pi} \langle z(t), x(t)\rangle \mathrm{d}t = 0\,,\quad \forall x \in C^1_{2\pi}\colon\ -\dot{x} = A_0^T x\,.$$

Stated differently: if we introduce on $C^0_{2\pi}$ a scalar product $\langle\cdot,\cdot\rangle_{C^0_{2\pi}}$ by

$$\langle x,\hat{x}\rangle_{C^0_{2\pi}} = \int_0^{2\pi} \langle x(t),\hat{x}(t)\rangle \mathrm{d}t\,,\quad \forall x\,,\ \hat{x} \in C^0_{2\pi}\,,$$

and if we define the *adjoint operator* $L_0^*\colon C^1_{2\pi} \to C^0_{2\pi}$ by $L_0^* = \widehat{A}_0^T + D_t$, then $\mathrm{R}(L_0) = \mathrm{N}(L_0^*)^\perp$. It follows easily from (52) that $\mathrm{N}(L_0^*) = \mathrm{N}(L_0)$, and therefore $\mathrm{R}(L_0) = \mathrm{N}(L_0)^\perp$ and $C^0_{2\pi} = \mathrm{N}(L_0) \oplus \mathrm{R}(L_0)$. Moreover, $\mathrm{N}(L_0)$ is finite-dimensional and $\mathrm{R}(L_0)$ is a closed subspace of $C^0_{2\pi}$; these properties imply that $L_0\colon C^1_{2\pi} \to C^0_{2\pi}$ is a *Fredholm operator of index zero* (see the next section for the definition).

One can associate the orthogonal splitting $C^0_{2\pi} = \mathrm{N}(L_0)\oplus\mathrm{R}(L_0) = \zeta(U)\oplus V$ (with $V = \mathrm{R}(L_0)$) to a continuous projection $P\colon C^0_{2\pi} \to C^0_{2\pi}$ such that $\mathrm{R}(P) = \zeta(U)$ and $\mathrm{N}(P) = V$; explicitly $P = \zeta\circ\tilde{P}$, with $\tilde{P}\colon C^0_{2\pi} \to U$ given by

$$\tilde{P}(x) = \frac{1}{2\pi}\int_0^{2\pi} \mathrm{e}^{-Ss}u(s)\mathrm{d}s\,,$$
$$\forall x = x(\cdot) = (u(\cdot),y(\cdot)) \in C^0_{2\pi}\,. \tag{59}$$

Observe that $\tilde{P}\circ\Sigma_\phi = \mathrm{e}^{S\phi}\circ\tilde{P}$ and $P\circ\Sigma_\phi = \Sigma_\phi\circ P$ for every $\phi \in S^1$, and that the splitting $C^0_{2\pi} = \zeta(U)\oplus V$ is invariant under the S^1-action; as a consequence the infinitesimal generator D_t of the S^1-action maps $V\cap C^1_{2\pi}$ into V. Every $x \in C^1_{2\pi}$ can then be written as $x = \zeta(u)+v$, with $u = \tilde{P}(x) \in U$ and $v = (I-P)x \in V\cap C^1_{2\pi}$; in

a similar way $\widehat{f}_\lambda(\zeta(u)+v) = \zeta(g(u,v,\lambda)) + h(u,v,\lambda)$, with $g(u,v,\lambda) = \tilde{P}\hat{f}_\lambda(\zeta(u)+v) \in U$ and $h(u,v,\lambda) = (I-P)\hat{f}_\lambda(\zeta(u)+v) \in V$; the mappings $g\colon U\times V\times\mathbb{R}^k \to U$ and $h\colon U\times V\times\mathbb{R}^k \to V$ are smooth, with $g(0,0,\lambda)=0$, $h(0,0,\lambda)=0$, and satisfy the equivariance relations

$$g\big(e^{S\phi}u, \Sigma_\phi v, \lambda\big) = e^{S\phi}g(u,v,\lambda) \quad \text{and}$$
$$h\big(e^{S\phi}u, \Sigma_\phi v, \lambda\big) = \Sigma_\phi(h(u,v,\lambda)) , \quad \forall\phi\in S^1 .$$

Using the foregoing decompositions the Eq. (56) splits into two equations

$$\begin{aligned} &g(u,v,\lambda) = (1+\tau)Su \\ \text{and}\quad &h(u,v,\lambda) = (1+\tau)D_t v , \end{aligned} \tag{60}$$

to be solved for small $(u,v,\tau,\lambda)\in U\times(V\cap C^1_{2\pi})\times\mathbb{R}\times\mathbb{R}^k$.

Since $h(0,0,\lambda) = 0$ and the restriction of L_0 to $V\cap C^1_{2\pi}$ is an isomorphism of $V\cap C^1_{2\pi}$ onto V the second equation in (60) can be solved for $v = v^*(u,\tau,\lambda)$, with $v^*\colon U\times\mathbb{R}\times\mathbb{R}^k \to V\cap C^1_{2\pi}$ a smooth mapping, $v^*(0,\tau,\lambda)=0$, $D_u v^*(0,0,0)=0$ and $v^*(e^{S\phi}u,\tau,\lambda) = \Sigma_\phi(v^*(u,\tau,\lambda))$ for all $\phi\in S^1$. Bringing this solution into the first equation of (60) gives the bifurcation equation

$$\begin{aligned} &g^*(u,\tau,\lambda) = (1+\tau)Su , \\ &\text{with}\quad g^*(u,\tau,\lambda) = g(u,v^*(u,\tau,\lambda),\lambda) . \end{aligned} \tag{61}$$

The mapping $g^*\colon U\times\mathbb{R}\times\mathbb{R}^k \to U$ is smooth, with $g^*(0,\tau,\lambda)=0$, $D_u g^*(0,\tau,0) = S$ and

$$g^*(e^{S\phi}u,\tau,\lambda) = e^{S\phi}g^*(u,\tau,\lambda) , \quad \forall\phi\in S^1 . \tag{62}$$

This S^1-equivariance of g^* implies that for each solution $(u,\tau,\lambda)\in U\times\mathbb{R}\times\mathbb{R}^k$ of (62) also $(e^{S\phi}u,\tau,\lambda)$ is a solution, for each $\phi\in S^1$. Solving (61) near the origin and working back through the reduction allows to solve the problems (P*) and (P).

The method outlined in this subsection can be adapted for special types of systems, such as Hamiltonian or reversible systems, and can also be generalized to the case where the linearization A_0 has some nilpotencies (non-semisimple eigenvalues); see e. g. [41,63] and [42] for more details. In these papers there is also established a link between the Lyapunov–Schmidt reduction and normal form theory.

Hopf Bifurcation

The simplest possible case in which the reduction of Subsect. "Lyapunov–Schmidt Reduction for the Problem (P)" can be applied is when A_0 has $\pm i$ as simple eigenvalues, while there are no other eigenvalues of the form $\pm ki$, with $k\in\mathbb{N}$ (this is called a *non-resonance condition*). Then dim $U=2$ and S can be brought into the matrix form

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} .$$

Identifying $u=(u_1,u_2)\in U=\mathbb{R}^2$ with $z=u_1+iu_2\in\mathbb{C}$ the linear operator A_0 gets the form (compare with (52))

$$\begin{aligned} A_0 x = A_0(z,y) = (iz, \tilde{A}_0 y) , \\ \forall x=(z,y)\in\mathbb{C}\times\mathbb{R}^{n-2} = \mathbb{R}^n . \end{aligned} \tag{63}$$

The bifurcation equation is then a single complex equation

$$g^*(z,\tau,\lambda) = (1+\tau)iz , \tag{64}$$

with $g^*\colon\mathbb{C}\times\mathbb{R}\times\mathbb{R}^k\to\mathbb{C}$ such that $g^*(0,\tau,\lambda)=0$, $g^*(z,\tau,0) = iz + O(|z|^2)$ and

$$g^*(e^{i\phi}z,\tau,\lambda) = e^{i\phi}g^*(z,\tau,\lambda) , \quad \forall\phi\in S^1 . \tag{65}$$

This equivariance property of g^* implies (see e. g. [24]) that one can write

$$g^*(z,\tau,\lambda) = \tilde{g}(|z|^2,\tau,\lambda)z , \quad \text{with } \tilde{g}(0,\tau,0) = i . \tag{66}$$

The Eq. (64) has for each (τ,λ) the trivial solution $z=0$, corresponding to the equilibrium $x=0$ of (14) Using (66) the nontrivial solutions of (64) are given by the solutions of the equation

$$\tilde{g}(|z|^2,\tau,\lambda) = (1+\tau)i . \tag{67}$$

The solutions of this equation come in circles $\{(e^{i\phi}\rho, \tau,\lambda) \mid \phi\in S^1\}$; periodic solutions of (14) corresponding to different points on such circle are related to each other by phase shift, and therefore each circle of solutions of (67) corresponds to a unique periodic orbit of (14).

The standard way to handle (67) is to assume $k=1$ (i. e. $\lambda\in\mathbb{R}$) and

$$\Re\,\frac{\partial\tilde{g}}{\partial\lambda}(0,0,0) \neq 0 . \tag{68}$$

Then (67) can be split in real and imaginary parts and solved by the implicit function theorem for $(\tau,\lambda) = (\tau^*(|z|^2),\lambda^*(|z|^2))$, with $(\tau^*(0),\lambda^*(0)) = (0,0)$. For each sufficiently $\rho>0$ there is a parameter value $\lambda=\lambda^*(\rho^2)$ such that the vectorfield f_λ has a small periodic orbit of "amplitude" ρ and with period $2\pi(1+\tau^*(\rho^2))^{-1}$. Such bifurcation of periodic orbits from an equilibrium is known as a *Hopf bifurcation*, or *Andronov–Hopf bifurcation*, or even *Poincaré–Andronov–Hopf bifurcation*. The papers [35] and [1] contain some of the original work on this topic; an extensive discussion on the history of Hopf bifurcation together with many more references can be found in Chap. 9 of [14].

Since the eigenvalues $\pm i$ of A_0 were assumed to be simple they can be continued to eigenvalues $\alpha(\lambda) \pm i\beta(\lambda)$ of $A_\lambda = D_x f_\lambda(0)$, for small values of λ. It is not hard to prove that the condition (68) is equivalent to

$$\alpha'(0) \neq 0\,; \tag{69}$$

this transversality condition means that as λ crosses zero the pair of eigenvalues $\alpha(\lambda) \pm i\beta(\lambda)$ of the linearization A_λ at the equilibrium $x = 0$ crosses the imaginary axis with non-zero transversal speed.

Similarly as for equilibria the requirement for Hopf bifurcation to have a pair of *simple* eigenvalues crossing the imaginary axis can sometimes be achieved by restricting (56) to periodic loops with some particular *space-time symmetries*. Working out this idea leads to the *Equivariant Hopf Bifurcation Theorem*. There exists a massive literature developing and applying this theorem; for more details see the basic reference [24].

Lyapunov–Schmidt Method in Infinite Dimensions

The Hopf bifurcation problem as explained in the preceding section leads to an Eq. (56) in the infinite-dimensional orbit space $C^0_{2\pi}$. Also bifurcation problems associated to boundary value problems for partial differential equations typically lead to equations in infinite-dimensional spaces (see e. g. [39] for some examples). The aim of this section is to briefly sketch a basic abstract setting which, under the appropriate conditions, allows to reduce these infinite-dimensional problems to finite-dimensional ones.

Consider an equation

$$F(x, \lambda) = 0\,, \tag{70}$$

with $F\colon X \times \mathbb{R}^k \to Y$ a smooth parameter-dependent mapping between two Banach spaces X and Y. Assuming that $x = 0$ is a solution for $\lambda = 0$, one can ask the question to what extend this solution persists for nearby values of λ, i. e. one wants to find all solutions (x, λ) of (70) near $(0, 0)$. The first step in solving this local problem consists in considering the linearized equation

$$A_0 x = 0\,, \quad \text{with } A_0 = D_x F(0, 0)\,, \tag{71}$$

and to determine the nullspace and range of A_0. The following hypothesis about A_0 allows to proceed with the Lyapunov–Schmidt reduction:

(F) The continuous linear operator $A_0\colon X \to Y$ is a *Fredholm operator*:

- (i) the nullspace $\mathrm{N}(A_0)$ is finite-dimensional;
- (ii) the range $\mathrm{R}(A_0)$ is closed in Y;
- (iii) the range $\mathrm{R}(A_0)$ has finite codimension in Y.

The *index* of such Fredholm operator is given by the number $\dim \mathrm{N}(A_0) - \mathrm{codim}\mathrm{R}(A_0)$; in many applications this index is zero, i. e. $\mathrm{codim}\mathrm{R}(A_0) = \dim \mathrm{N}(A_0)$.

Under the condition (F) there exist continuous linear projections $P\colon X \to X$ and $Q\colon Y \to Y$ such that

$$\mathrm{R}(P) = \mathrm{N}(A_0) \quad \text{and} \quad \mathrm{N}(Q) = \mathrm{R}(A_0)\,, \tag{72}$$

resulting in the topological decompositions $X = \mathrm{N}(A_0) \oplus \mathrm{N}(P) = U \oplus V$ and $Y = \mathrm{R}(Q) \oplus \mathrm{R}(A_0) = W \oplus \mathrm{R}(A_0)$, with $U = \mathrm{N}(A_0)$, $V = \mathrm{N}(P)$ and $W = \mathrm{R}(Q)$. The restriction of A_0 to V is an isomorphism from V onto $\mathrm{R}(A_0)$. Replacing $x \in X$ by $x = Px + (I-P)x = u + v$ in (70), and splitting the equation itself using Q results in the following system equivalent to (70):

$$QF(u + v, \lambda) = 0\,, \quad (I - Q)F(u + v, \lambda) = 0\,. \tag{73}$$

The second of these equations is sometimes called the *auxiliary equation*; the corresponding mapping $F_{\mathrm{aux}}\colon U \times V \times \mathbb{R}^k \to \mathrm{R}(A_0)$ defined by $F_{\mathrm{aux}}(u, v, \lambda) = (I-Q)F(u+v, \lambda)$ is such that $F_{\mathrm{aux}}(0, 0, 0) = 0$, $D_u F_{\mathrm{aux}}(0, 0, 0) = 0$ and $D_v F_{\mathrm{aux}}(0, 0, 0) = (I - Q)A_0\big|_V$. Since this last operator is an isomorphism from V onto $\mathrm{R}(A_0)$ one can apply the implicit function theorem to solve the auxiliary equation in the neighborhood of $(u, v, \lambda) = (0, 0, 0)$ for $v = v^*(u, \lambda)$. The mapping $v^*\colon U \times \mathbb{R}^k \to V$ is smooth, with $v^*(0, 0) = 0$ and $D_u v^*(0, 0) = 0$. Substituting $v = v^*(u, \lambda)$ into the first equation of (73) we obtain the bifurcation equation

$$G(u, \lambda) = QF(u + v^*(u, \lambda), \lambda) = 0\,. \tag{74}$$

The bifurcation mapping $G\colon U \times \mathbb{R}^k \to W$ is smooth, with $G(0, 0) = 0$, $D_u G(0, 0) = 0$ and $D_\lambda G(0, 0) = QD_\lambda F(0, 0)$. The further analysis of the bifurcation Eq. (74) will of course strongly depend on the particularities of the problem, such as the dimensions of U and W, possible symmetries, etc.

A well known and widely applied particular case is the by now classical theorem of Crandall and Rabinowitz [17] on bifurcation from simple eigenvalues. The hypotheses for this result are the following:

(CR) $F\colon X \times \mathbb{R} \to Y$ is smooth, with $F(0, \lambda) = 0$ for all $\lambda \in \mathbb{R}$, and:

- (i) the nullspace of $D_x F(0, 0)$ is one-dimensional;
- (ii) the range of $D_x F(0, 0)$ is closed and has codimension one in Y;
- (iii) $D_\lambda D_x F(0, 0) \cdot u_0 \notin \mathrm{R}(D_x F(0, 0))$, where $u_0 \neq 0$ belongs to $\mathrm{N}(D_x F(0, 0))$.

Under these hypotheses $\dim U = \dim W = 1$ in the foregoing, and therefore $U = \{\rho u_0 \mid \rho \in \mathbb{R}\}$ and $W = \{\eta w_0 \mid \eta \in \mathbb{R}\}$ for some non-zero elements $u_0 \in \mathrm{N}(D_x F(0, 0))$ and $w_0 \notin \mathrm{R}(D_x F(0, 0))$. Also $v^*(0, \lambda) = 0$ and $G(0, \lambda) = 0$

for all $\lambda \in \mathbb{R}$; writing $G(\rho u_0, \lambda) = g(\rho, \lambda) w_0$ it follows that also $g(0, \lambda) = 0$ for all λ. Hence

$$g(\rho, \lambda) = \rho \tilde{g}(\rho, \lambda), \quad \text{with } \tilde{g}(\rho, \lambda) = \int_0^1 \frac{\partial g}{\partial \rho}(s\rho, \lambda) \mathrm{d}s .$$

Clearly $\tilde{g}(0, 0) = 0$, since $\tilde{g}(0, 0) w_0 = \frac{\partial g}{\partial \rho}(0, 0) w_0 = D_u G(0, 0) \cdot u_0$ and $D_u G(0, 0) = 0$. Moreover,

$$\begin{aligned} \frac{\partial \tilde{g}}{\partial \lambda}(0, 0) w_0 &= \frac{\partial^2 g}{\partial \lambda \partial \rho}(0, 0) w_0 = D_\lambda D_u G(0, 0) \cdot u_0 \\ &= Q D_\lambda D_x F(0, 0) \cdot u_0 \neq 0 , \end{aligned}$$

by the hypothesis (CR)-(iii), and so $\frac{\partial \tilde{g}}{\partial \lambda}(0, 0) \neq 0$. Nonzero solutions of the bifurcation equation $g(\rho, \lambda) = 0$ are given by the solutions of $\tilde{g}(\rho, \lambda) = 0$, and using the implicit function theorem this last equation can be solved for $\lambda = \lambda^*(\rho)$, with $|\rho|$ sufficiently small and $\lambda^*(0) = 0$. This gives, under the hypotheses (CR), a unique branch

$$\begin{aligned} \{(x^*(\rho), \lambda^*(\rho)) =| \, 0 < |\rho| < \rho_0\} , \\ \text{with } x^*(\rho) = \rho u_0 + v^*(\rho u_0, \lambda^*(\rho)) , \end{aligned}$$

of nontrivial solutions of (70) bifurcation from the trivial branch $\{(0, \lambda) \mid \lambda \in \mathbb{R}\}$ at $\lambda = 0$.

Outlook

The Liapunov–Schmidt method is a well established reduction technique which can and has been used on a wide variety of problems. The preceding sections just give a small glimpse of its use in local bifurcation theory; the references given in the introduction will point the reader to a whole plethora of other applications. The basics of the method are rather simple, but depending on the mappings and the spaces one is using a number of technical conditions (sometimes easy, sometimes more difficult) need to be verified before the reduction can be carried out. In view of the steady increase in applications which the method has seen over the last half century one can safely say that the Lyapunov–Schmidt method will remain a valuable tool for future research in bifurcation theory, dynamical systems theory and nonlinear analysis.

Bibliography

Primary Literature

1. Andronov A, Leontovich FA (1938) Sur la théorie de la variation de la structure qualitative de la division du plan en trajectoires. Dokl Akad Nauk 21:427–430
2. Antosiewicz HA (1966) Boundary value problems for nonlinear ordinary differential equations. Pacific J Math 17:191–197
3. Ashwin P, Böhmer K, Mei Z (1993) A numerical Lyapunov–Schmidt method for finitely determined systems. In: Allgower E, Georg K, Miranda R (eds) Exploiting Symmetry in Applied and Numerical Analysis. Lectures in Appl Math, vol 29. AMS, Providence, pp 49–69
4. Bancroft S, Hale JK, Sweet D (1968) Alternative problems for nonlinear functional equations. J Differ Equ 4:40–56
5. Bartle R (1953) Singular points of functional equations. Trans Amer Math Soc 75:366–384
6. Böhmer K (1993) On a numerical Lyapunov–Schmidt method for operator equations. Computing 53:237–269
7. Bourgain J (1998) Quasi-periodic Solutions of Hamiltonian Perturbations of 2D Linear Schrödinger Equations. Ann Math 148(2):363–439
8. Cacciopoli R (1936) Sulle correspondenze funzionali inverse diramata: teoria generale e applicazioni ad alcune equazioni funzionali non lineari e al problema di Plateau. Atti Accad Naz Lincei Rend Cl Sci Fis Mat Natur Sez I 24:258–263, 416–421
9. Cesari L (1940) Sulla stabilità delle soluzioni dei sistemi di equazioni differenziali lineari a coefficienti periodici. Atti Accad Ital Mem Cl Fis Mat Nat 6(11):633–692
10. Cesari L (1963) Periodic solutions of hyperbolic partial differential equations. In: Intern. Symposium on Nonlinear Differential Equations and Nonlinear Mechanics. Academic Press, New York, pp 33–57
11. Cesari L (1964) Functional analysis and Galerkin's method. Michigan Math J 11:385–414
12. Cicogna C (1981) Symmetry breakdown from bifurcations. Lett Nuovo Cimento 31:600–602
13. Chossat P, Lauterbach R (2000) Methods in Equivariant Bifurcations and Dynamical Systems. Advanced Series in Nonlinear Dynamics, vol 15. World Scientific, Singapore
14. Chow S-N, Hale JK (1982) Methods of Bifurcation Theory. Grundlehren der Mathematischen Wissenschaften, vol 251. Springer, New York
15. Ciocco M-C, Vanderbauwhede A (2004) Bifurcation of Periodic Points in Reversible Diffeomorphisms. In: Elaydi S, Ladas G, Aulbach B (eds) New Progress in Difference Equations – Proceedings ICDEA2001, Augsburg. Chapman & Hall, CRC Press, pp 75–93
16. Craig W, Wayne CE (1993) Newton's method and periodic solutions of nonlinear wave equations. Comm Pure Appl Math 46:1409–1498
17. Crandall MG, Rabinowitz PH (1971) Bifurcation from simple eigenvalues. J Funct Anal 8:321–340
18. Crandall MG, Rabinowitz PH (1978) The Hopf bifurcation theorem in infinite dimensions. Arch Rat Mech Anal 67:53–72
19. Cronin J (1950) Branch points of solutions of equations in Banach spaces. Trans Amer Math Soc 69:208–231; (1954) 76:207–222
20. Da Prato G, Lunardi A (1986) Hopf bifurcation for fully nonlinear equations in Banach space. Ann Inst Henri Poincaré 3:315–329
21. Friedrichs KO (1953–54) Special Topics in Analysis. Lecture Notes. New York University, New York
22. Gaines R, Mawhin J (1977) Coincidence Degree, and Nonlinear Differential Equations. Lect Notes in Math, vol 568. Springer, Berlin
23. Golubitsky M, Schaeffer DG (1985) Singularities and Groups in Bifurcation Theory, vol 1. Appl Math Sci, vol 51. Springer, New York
24. Golubitsky M, Stewart IN, Schaeffer DG (1988) Singularities and

Groups in Bifurcation Theory, vol 2. Appl Math Sci, vol 69. Springer, New York
25. Golubitsky M, Pivato M, Stewart I (2004) Interior symmetry and local bifurcation in coupled cell networks. Dyn Syst 19:389–407
26. Hale JK (1954) Periodic solutions of non-linear systems of differential equations. Riv Mat Univ Parma 5:281–311
27. Hale JK (1963) Oscillations in nonlinear systems. McGraw-Hill, New York
28. Hale JK (1969) Ordinary Differential Equations. Wiley, New York
29. Hale JK (1977) Theory of Functional Differential Equations. Appl Math Sciences, vol 3. Springer, New York
30. Hale JK (1984) Introduction to dynamic bifurcation. In: Salvadori L (ed) Bifurcation Theory and Applications. Lecture Notes in Math, vol 1057. Springer, Berlin, pp 106–151
31. Hale JK, Sakamoto K (1988) Existence and stability of transition layers. Japan J Appl Math 5:367–405
32. Hale JK, Sakamoto K (2005) A Lyapunov–Schmidt method for transition layers in reaction-diffusion systems. Hiroshima Math J 35:205–249
33. Hall WS (1970) Periodic solutions of a class of weakly nonlinear evolution equations. Arch Rat Mech Anal 39:294–332
34. Hall WS (1970) On the existence of periodic solutions of the equation $D_{tt}u + (-1)^p D_x^{2p} u = \varepsilon f(t, x, u)$. J Differ Equ 7:509–526
35. Hopf E (1942) Abzweigung einer periodischen Lösung von einer stationären Lösung eines Differentialsystems. Ber Math Phys Kl Sachs Acad Wiss Leipzig 94:1–22. English translation in: Marsden J, McCracken M (1976) Hopf Bifurcation and its Applications. Springer, New York
36. Iooss G, Joseph D (1981) Elementary Stability and Bifurcation Theory. Springer, New York
37. Ize J (1976) Bifurcation theory for Fredholm operators. Mem AMS 174
38. Jepson A, Spence A (1989) On a reduction process for nonlinear equations. SIAM J Math Anal 20:39–56
39. Kielhöfer H (2004) Bifurcation Theory, An Introduction with Applications to PDEs. Appl Math Sciences, vol 156. Springer, New York
40. Knobloch J (2000) Lin's method for discrete dynamical systems. J Differ Eq Appl 6:577–623
41. Knobloch J, Vanderbauwhede A (1996) A general reduction method for periodic solutions in conservative and reversible systems. J Dyn Diff Eqn 8:71–102
42. Knobloch J, Vanderbauwhede A (1996) Hopf bifurcation at *k*-fold resonances in conservative systems. In: Broer HW, van Gils SA, Hoveijn I, Takens F (eds) Nonlinear Dynamical Systems and Chaos. Progress in Nonlinear Differential Equations and Their Applications, vol 19. Birkhäuser, Basel, pp 155–170
43. Krasnosel'skii MA (1963) Topological Methods in the Theory of Nonlinear Integral Equations. Pergamon Press, Oxford
44. Krasnosel'skii MA, Vainiko GM, Zabreiko PP, Rutitskii YB, Stesenko VY (1972) Approximate solutions of operator equations. Wolters-Noordhoff, Groningen
45. Kuznetsov YA (1998) Elements of Applied Bifurcation Theory, 2nd edn. Appl Math Sci, vol 112. Springer, New York
46. Lewis DC (1956) On the role of first integrals in the perturbations of periodic solutions. Ann Math 63:535–548
47. Lin XB (1990) Using Melnikov's method to solve Silnikov's problems. Proc Roy Soc Edinb 116A:295–325
48. Lyapunov AM (1906) Sur les figures d'équilibre peu différentes des ellipsoï, des d'une masse liquide homogène dotées d'un mouvement de rotation. Zap Akad Nauk St Petersb, pp 1–225; (1908), pp 1–175; (1912), pp 1–228; (1914), pp 1–112
49. Mawhin J (1969) Degré topologique et solutions périodiques des systèmes différentiels non linéaires. Bull Soc Roy Sci Liège 38:308–398
50. Nirenberg L (1960–61) Functional Analysis. Lecture Notes. New York University, New York
51. Rabinowitz PH (1967) Periodic solutions of nonlinear hyperbolic partial differential equations. Comm Pure Appl Math 20:145–206
52. Rabinowitz PH (1976) A Survey of Bifurcation Theory. In: Cesari L, Hale JK, LaSalle J (eds) Dynamical Systems, An International Symposium, vol 1. Academic Press, New York, pp 83–96
53. Sandstede B (1993) Verzweigungstheorie homokliner Verdopplungen. Ph D Thesis, University of Stuttgart, IAAS Report No. 7
54. Sattinger DH (1979) Group Theoretic Methods in Bifurcation Theory. Lecture Notes in Math, vol 762. Springer, Berlin
55. Schmidt E (1908) Zur Theorie des linearen und nichtlinearen Integralgleichungen, 3 Teil. Über die Auflösung der nichtlinearen Integralgleichungen und die Verzweigung ihrer Lösungen. Math Ann 65:370–399
56. Shimizu T (1948) Analytic operations and analytic operational equations. Math Japon 1:36–40
57. Sidorov N, Loginov B, Sinitsyn A, Falaleev M (2002) Lyapunov–Schmidt Methods in Nonlinear Analysis and Applications. Mathematics and Its Applications, vol 550. Kluwer, Dordrecht
58. Taniguchi M (1995) A Remark on Singular Perturbation Methods via the Lyapunov–Schmidt Reduction. Publ RIMS, Kyoto Univ 31:1001–1010
59. Vainberg MM, Trenogin VA (1962) The method of Lyapunov and Schmidt in the theory of nonlinear differential equations and their further development. Russ Math Surv 17:1–60
60. Vainberg MM, Trenogin VA (1974) Theory of Branching of Solutions of Non-Linear Equations. Noordhoff, Leyden
61. Vanderbauwhede A (1982) Local bifurcation and symmetry. Research Notes in Mathematics, vol 75. Pitman, London
62. Vanderbauwhede A (2000) Subharmonic bifurcation at multiple resonances. In: Elaydi S, Allen F, Elkhader A, Mughrabi T, Saleh M (eds) Proceedings of the Mathematics Conference, Birzeit, August 1998. World Scientific, Singapore, pp 254–276
63. Vanderbauwhede A, Van der Meer JC (1995) A general reduction method for periodic solutions near equilibria in Hamiltonian systems. In: Langford WF, Nagata W (eds) Normal Forms and Homoclinic Chaos. Fields Institute Comm. AMS, Providence, pp 273–294

Books and Reviews

Alligood KT, Sauer T, Yorke JA (1997) Chaos, An Introduction to Dynamical Systems. Springer, New York
Chicone C (2006) Ordinary Differential Equations with Applications, 2nd edn. Texts in Applied Math, vol 34. Springer, New York
Govaerts W (2000) Numerical Methods for Bifurcations of Dynamical Equilibria. SIAM Publications, Philadelphia
Hale JK, Koçak H (1991) Dynamics and Bifurcations. Texts in Applied Mathematics, vol 3. Springer, New York
Ma T, Wang S (2005) Bifurcation Theory and Applications. World Scientific Series on Nonlinear Science, vol A53. World Scientific, Singapore