1. Younes and Littman: *PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects* [YL04]
   We desribe a variation of the planning domain definition language, PDDL, that permits the modeling of probabilistic planning problems with rewards. This language, PPDDL1.0, was used as the input language for the probabilistic track of the 4th International Planning Competition. We provide the complete syntax for PPDDL1.0 and give a semantics of PPDDL1.0 planning problems in terms of Markov decision processes.

2. Marinescu and Coles: "Heuristic Guidance for Forward-Chaining Planning with Numeric Uncertainty" [MC16]
   Uncertainty hinders many interesting applications of planning – it may come in the form of sensor noise, unpredictable environments, or known limitations in problem models. In this paper we explore heuristic guidance for forward-chaining planning with continuous random variables, while ensuring a probability of plan success. We extend the Metric Relaxed Planning Graph heuristic to capture a model of uncertainty, providing better guidance in terms of heuristic estimates and dead-end detection. By tracking the accumulated error on nu- meric values, our heuristic is able to check if preconditions in the planning graph are achievable with a sufficient degree of confidence; it is also able to consider acting to reduce the accumulated error. Results indicate that our approach offers improvements in performance compared to prior work where a less-informed relaxation was used.

3. Scala et al.: "Interval-Based Relaxation for General Numeric Planning" [Sca+16]
   We generalise the interval-based relaxation to sequential numeric planning problems with non-linear conditions and effects, and cyclic dependencies. This effectively removes all the limi-tations on the problem placed in previous work on numeric planning heuristics, and even allows us to extend the planning language with a wider set of mathematical functions. Heuristics obtained from the generalised relaxation are pruning-safe. We derive one such heuristic and use it to solve discrete-time control-like planning problems with autonomous processes. Few planners can solve such problems, and search with our new heuristic compares favourably with them.

4. Yoo, Fitch, and Sukkarieh: "Provably-correct stochastic motion planning with safety constraints" [YFS13]
   Formal methods based on the Markov decision process formalism, such as probabilistic computation tree logic (PCTL), can be used to analyse and synthesise control policies that maximise the probability of mission success. In this paper, we consider a different objective. We wish to minimise time-to-completion while satisfying a given probabilistic threshold of success. This important problem naturally arises in motion planning for outdoor robots, where high quality mobility prediction methods are available but stochastic path planning typically relies on an arbitrary weighted cost function that attempts to balance the opposing goals of finding safe paths (minimising risk) while making progress towards the goal (maximising reward). We propose novel algorithms for model checking and policy synthesis in PCTL that (1) provide a quantitative measure of safety and completion time for a given policy, and (2) synthesise policies that minimise completion time with respect to a given safety threshold. We provide simulation results in a stochastic outdoor navigation domain that illustrate policies with varying levels of risk.

5. Camacho, Muise, and McIlraith: "From FOND to Robust Probabilistic Planning : Computing Compact Policies that Bypass Avoidable Deadends" [CMM16]
   We address the class of probabilistic planning problems where the objective is to maximize the probability of reaching a prescribed goal. The complexity of probabilistic planning problems makes it difficult to compute high quality solutions for large instances, and existing algorithms either do not scale, or do so at the expense of the solution quality. We leverage core similarities between probabilistic and fully observable non-deterministic (FOND) planning to construct a sound, offline probabilistic planner, ProbPRP, that exploits algorithmic advances from state-of-the-art FOND planner, PRP, to compute compact policies that are guaranteed to bypass avoidable deadends. We evaluate ProbPRP on a selection of benchmarks used in past probabilistic planning competitions. The results show that ProbPRP, in many cases, outperforms the state of the art, computing substantially more robust policies and at times doing so orders of magnitude faster.

6. Mombourquette, Muise, and McIlraith: "Logical Filtering and Smoothing: State Estimation in Partially Observable Domains" [MMM17]
State estimation is the task of estimating the state of a partially observable dynamical system given a sequence of executed actions and observations. In logical settings, state estimation can be realized via logical filtering, which is exact but can be intractable. We propose logical smoothing, a form of backwards reasoning that works in concert with approximated logical filtering to refine past beliefs in light of new observations. We characterize the notion of logical smoothing together with an algorithm for backwards-forwards state estimation. We also present an approximation of our smoothing algorithm that is space efficient. We prove properties of our algorithms, and experimentally demonstrate their behaviour, contrasting them with state estimation methods for planning. Smoothing and backwards-forwards reasoning are important techniques for reasoning about partially observable dynamical systems, introducing the logical analogue of effective techniques from control theory and dynamic programming.

7. Belle: "Probabilistic Planning by Probabilistic Programming" [Bel18]
Automated planning is a major topic of research in artificial intelligence, and enjoys a long and distinguished history. The classical paradigm assumes a distinguished initial state, comprised of a set of facts, and is defined over a set of actions which change that state in one way or another. Planning in many real-world settings, however, is much more involved: an agent's knowledge is almost never simply a set of facts that are true, and actions that the agent intends to execute never operate the way they are supposed to. Thus, probabilistic planning attempts to incorporate stochastic models directly into the planning process. In this article, we briefly report on probabilistic planning through the lens of probabilistic programming: a programming paradigm that aims to ease the specification of structured probability distributions. In particular, we provide an overview of the features of two systems, HYPE and ALLEGRO, which emphasise different strengths of probabilistic programming that are particularly useful for complex modelling issues raised in probabilistic planning. Among other things, with these systems, one can instantiate planning problems with growing and shrinking state spaces, discrete and continuous probability distributions, and non-unique prior distributions in a first-order setting.

8. Biscaia and Mateus: "A Temporal Logic for Planning under Uncertainty" [BM13]
   Dealing with uncertainty in the context of planning has been an active research subject in AI. Addressing the case when uncertainty evolves over time can be difficult. In this work, we provide a solution to this problem by proposing a temporal logic to reason about quantities and probability. For this logic, we provide a PSPACE SAT algorithm together with a complete calculus. The algorithm enables us to perform planning under uncertainty via SAT, extending a technique used for classic planning. We can show that any obtained plan will have certain properties (desired or undesired). The calculus can also be used to derive the impossibility of a plan, given a set of specification.

9. Mateus et al.: "Probabilistic Situation Calculus" [Mat+01]
   In this article we propose a Probabilistic Situation Calculus logical language to represent and reason with knowledge about dynamic worlds in which actions have uncertain effects. Uncertain effects are modeled by dividing an action into two subparts: a deterministic (agent produced) input and a probabilistic reaction (produced by nature). We assume that the probabilities of the reactions have known distributions. Our logical language is an extension to Situation Calculae in the style proposed by Raymond Reiter. There are three aspects to this work. First, we extend the language in order to accommodate the necessary distinctions (e.g., the separation of actions into inputs and reactions). Second, we develop the notion of Randomly Reactive Automata in order to specify the semantics of our Probabilistic Situation Calculus. Finally, we develop a reasoning system in MATHEMATICA capable of performing temporal projection in the Probabilistic Situation Calculus.

10. Dehnert et al.: "PROPhESY: A PRObabilistic ParamEter SYnthesis Tool" [Deh+10]
    We present PROPhESY, a tool for analyzing parametric Markov chains (MCs). It can compute a rational function (i.e., a fraction of two polynomials in the model parameters) for reachability and expected reward objectives. Our tool outperforms state-of-the-art tools and supports the novel feature of conditional probabilities. PROPhESY supports incremental automatic parameter synthesis (using SMT techniques) to

determine "safe" and "unsafe" regions of the parameter space. All values in these regions give rise to instantiated MCs satisfying or violating the (conditional) probability or expected reward objective. PROPhESY features a web front-end supporting visualization and user- guided parameter synthesis. Experimental results show that PROPhESY scales to MCs with millions of states and several parameters.

11. Möhring, Uetz, and Schulz: "Approximation in stochastic scheduling: the power of LP-based priority policies"
    We consider the problem to minimize the total weighted completion time of a set of jobs with individual release dates which have to be scheduled on identical parallel machines. Job processing times are not known in advance, they are realized on-line according to given probability distributions. The aim is to find a scheduling policy that minimizes the objective in expectation. Motivated by the success of LP-based approaches to deterministic scheduling, we present a polyhedral relaxation of the performance space of stochastic parallel machine scheduling. This relaxation extends earlier relaxations that have been used, among others, by Hall et al. [1997] in the deterministic setting. We then derive constant performance guarantees for priority policies which are guided by optimum LP solutions, and thereby generalize previous results from deterministic scheduling. In the absence of release dates, the LP-based analysis also yields an additive performance guarantee for the WSEPT rule which implies both a worst-case performance ratio and a result on its asymptotic optimality, thus complementing previous work by Weiss [1990]. The corresponding LP lower bound generalizes a previous lower bound from deterministic scheduling due to Eastman et al. [1964], and exhibits a relation between parallel machine problems and corresponding problems with only one fast single machine. Finally, we show that all employed LPs can be solved in polynomial time by purely combinatorial algorithms.

References

[Bel18] Vaishak Belle. "Probabilistic Planning by Probabilistic Programming". In: *AAAI Workshop on Planning and Inference*. 2018. arXiv: 1801.08365. URL: http://arxiv.org/abs/1801.08365.

[BM13] Manuel Biscaia and Paulo Mateus. "A Temporal Logic for Planning under Uncertainty". In: *Proceedings of the Twenty-Sixth International Florida Artificial Intelligence Research Society Conference*. 2013, pp. 1–11. URL: http://www.aaai.org/ocs/index.php/FLAIRS/FLAIRS13/paper/download/5873/6145.

[CMM16] Alberto Camacho, Christian Muise, and Sheila A. McIlraith. "From FOND to Robust Probabilistic Planning : Computing Compact Policies that Bypass Avoidable Deadends". In: *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2016)*. 2016, pp. 65–69. URL: http://www.aaai.org/ocs/index.php/ICAPS/ICAPS16/paper/download/13188/12662.

[Deh+10] Christian Dehnert et al. "PROPhESY: A PRObabilistic ParamEter SYnthesis Tool". In: *Proceedings of the 22nd International Conference on Compuer Aided Verification (CAV)*. 2010, pp. 214–231. URL: https://link.springer.com/chapter/10.1007/978-3-319-21690-4_13.

[Mat+01] Paulo Mateus et al. "Probabilistic Situation Calculus". In: *Annals of Mathematics and Artificial Intelligence* 32.1-4 (2001), pp. 393–431. ISSN: 10122443. DOI: 10.1023/A:1016738205696.

[MC16] Liana Marinescu and Andrew Coles. "Heuristic Guidance for Forward-Chaining Planning with Numeric Uncertainty". In: *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2016)*. Vol. 285. ICAPS. 2016, pp. 1694–1695. ISBN: 9781614996712. DOI: 10.3233/978-1-61499-672-9-1694. URL: http://www.aaai.org/ocs/index.php/ICAPS/ICAPS16/paper/download/13186/12682.

[MMM17] Brent Mombourquette, Christian Muise, and Sheila A. McIlraith. "Logical Filtering and Smoothing: State Estimation in Partially Observable Domains". In: *Proceedings of the 31st Conference on Artificial Intelligence (AAAI 2017)*. 2017, pp. 3613–3621. URL:

http://www.aaai.org/ocs/index.php/AAAI/AAAI17/paper/download/14958/14166.

[MUS99] Rolf H. Möhring, Marc Uetz, and Andreas S. Schulz. "Approximation in stochastic scheduling: the power of LP-based priority policies". In: *Journal of the ACM* 46.6 (1999), pp. 924–942. ISSN: 00045411. DOI: 10.1145/331524.331530.

[Sca+16] Enrico Scala et al. "Interval-Based Relaxation for General Numeric Planning". In: *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*. 2016. URL: http://users.cecs.anu.edu.au/~thiebaux/papers/ecai16.pdf.

[YFS13] C. Yoo, R. Fitch, and S. Sukkarieh. "Provably-correct stochastic motion planning with safety constraints". In: *2013 IEEE International Conference on Robotics and Automation*. May 2013, pp. 981–986. DOI: 10.1109/ICRA.2013.6630692.

[YL04] Hakan Younes and Michael Littman. *PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects*. Tech. rep. Pittsburgh, PA, USA: Carnegie Mellon University, 2004. URL: http://reports-archive.adm.cs.cmu.edu/anon/anon/home/ftp/usr0/ftp/2004/CMU-CS-04-167.pdf.