

Path Finding under Uncertainty through Probabilistic Inference

David Tolpin, Brooks Paige, Jan Willem van de Meent, Frank Wood

University of Oxford

{dtolpin,brooks,jwvdm,fwood}@robots.ox.ac.uk

Abstract

We introduce a new approach to solving path-finding problems under uncertainty by representing them as probabilistic models and applying domain-independent inference algorithms to the models. This approach separates problem representation from the inference algorithm and provides a framework for efficient learning of path-finding policies. We evaluate the new approach on the Canadian Traveller Problem, which we formulate as a probabilistic model, and show how probabilistic inference allows high performance stochastic policies to be obtained for this problem.

Introduction

In planning under uncertainty the objective is to find the optimal policy — a policy that maximizes the expected reward. In most interesting cases the optimal policy cannot be found exactly, and approximation schemes are used to discover the policy, either represented explicitly or as an implicit property of the planning algorithm, through *reinforcement learning*. Approximation schemes include value/policy iteration, Q-learning, policy gradient methods (Sutton and Barto 1998), as well as methods based on heuristic search (Bonet and Geffner 2001) and Monte Carlo sampling such as MCTS (Kocsis and Szepesvári 2006; Browne et al. 2012).

Domain-independent planning algorithms (Bonet and Geffner 2001; Haslum, Bonet, and Geffner 2005; Helmert 2006) can be applied to different domains with little modification, however for many applications domain-dependant techniques are still critical in order to obtain a high performance policy, and put the burden of implementation on the domain expert formulating the planning problem.

The framework of probabilistic inference (Pearl 1988) proposes solutions to a wide range of Artificial Intelligence problems by representing them as *probabilistic models*. Efficient domain-independent algorithms are available for several classes of representations, in particular for graphical models (Lauritzen 1996), where inference can be performed either exactly and approximately. However, graphical models typically require that the full graph of the model to be represented explicitly, and are not powerful enough for

problems where the state space is exponential in the problem size, such as the generative models common in planning (Szörényi, Kedenburg, and Munos 2014).

Probabilistic programs (Goodman et al. 2008; Mansinghka, Selsam, and Perov 2014; Wood, van de Meent, and Mansinghka 2014) can represent arbitrary probabilistic models, efficient approximate inference algorithms have recently emerged (Wingate, Stuhlmüller, and Goodman 2011; Wood, van de Meent, and Mansinghka 2014; Paige et al. 2014). In addition to expressive power, probabilistic programming separates modeling and inference, allowing the problem to be specified in a simple language which does not assume any particular inference technique.

In this paper, we show a connection between probabilistic inference and path finding, which allows many path-finding problems to be cast as inference problems using probabilistic programs. Based on this connection, we provide a generic scheme for expressing a path-finding problem as a probabilistic program that infers the path-finding policy. We illustrate this generic scheme by its application to the Canadian Traveller Problem (Papadimitriou and Yannakakis 1991; Bar-Noy and Schieber 1991; Nikolova and Karger 2008). In the empirical evaluation, we show that high performance stochastic policies can be obtained using domain-independent inference techniques. In the concluding section, we discuss other possible areas of application of probabilistic programming in planning, as well as possible difficulties.

Preliminaries

Probabilistic Programming

Probabilistic programs are regular programs extended by two constructs (Gordon et al. 2014):

- The ability to draw random values from probability distributions.
- The ability to condition values computed in the programs on probability distributions.

A probabilistic program implicitly defines a probability distribution over the program’s output. Formally, we define a probabilistic program as a stateful deterministic computation \mathcal{P} with the following properties:

- Initially, \mathcal{P} expects no arguments.

- On every call, \mathcal{P} returns either a distribution F , a distribution and a value (G, y) , a value z , or \perp .
- Upon returning F , \mathcal{P} expects a value x drawn from F as the argument to continue.
- Upon returning (G, y) or z , \mathcal{P} is invoked again without arguments.
- Upon returning \perp , \mathcal{P} terminates.

A program is run by calling \mathcal{P} repeatedly until termination. Every run of the program implicitly produces a sequence of pairs (F_i, x_i) of distributions and drawn from them values of latent random variables. We call this sequence a *trace* and denote it by \mathbf{x} . A trace induces a sequence of pairs (G_j, y_j) of distributions and values of observed random variables. We call this sequence an *image* and denote it by \mathbf{y} . We call a sequence of values z_k an *output* of the program and denote it by \mathbf{z} . Program output is deterministic given the trace.

By definition, the probability of a trace is proportional to the product of the probability of all random choices \mathbf{x} and the likelihood of all observations \mathbf{y} :

$$p_{\mathcal{P}}(\mathbf{x}|\mathbf{y}) \propto \prod_{i=1}^{|\mathbf{x}|} p_{F_i}(x_i) \prod_{j=1}^{|\mathbf{y}|} p_{G_j}(y_j) \quad (1)$$

The objective of inference in probabilistic program \mathcal{P} is to discover the distribution of \mathbf{z} .

Several implementations of general probabilistic programming languages are available (Goodman et al. 2008; Mansinghka, Selsam, and Perov 2014; Wood, van de Meent, and Mansinghka 2014). Inference is usually performed using Monte Carlo sampling algorithms for probabilistic programs (Wingate, Stuhlmüller, and Goodman 2011; Wood, van de Meent, and Mansinghka 2014; Paige et al. 2014). While some algorithms are better suited for certain inference types, most can be used with any valid probabilistic program.

Canadian Traveller Problem

Canadian Traveller Problem (CTP) was introduced in (Papadimitriou and Yannakakis 1991) as a problem of finding the shortest travel distance in a graph where some edges may be blocked. There are several variants of CTP (Bar-Noy and Schieber 1991; Nikolova and Karger 2008; Bnaya, Felner, and Shimony 2009); here we consider the *stochastic* version. In the stochastic CTP we are given

- Undirected weighted graph $G = (V, E)$.
- The initial and the final location nodes s and t .
- Edge weights $w : E \rightarrow \mathcal{R}$.
- Traversability probabilities: $p_o : E \rightarrow (0, 1]$.

The actual state of each edge is fixed for every problem instance but becomes known only upon reaching one of the edge vertices. The goal is to find a *policy* that minimizes the expected *travel distance* from s to t . The travel distance is the sum of weights of all traversed edges during the travel, where traversing in each direction is counted separately.




Figure 1: A path in the graph of a probabilistic program

CTP problem is PSPACE-hard (Fried et al. 2013), however a number of heuristic algorithms were proposed, including high-quality policies based on Monte Carlo methods (Eyerich, Keller, and Helmert 2010). Policies are empirically compared by averaging the distance travel over multiple instantiations of the actual states of the edges (open or blocked) according to the traversal probabilities. Since the travel distance is defined only for instance where a path between s and t exists, instantiations in which t cannot be reached from s are ignored.

A trivial travel policy is realized by traversing the problem graph in a depth-first order until the final location is reached. The expected travel distance of the policy is bounded from above by the sum of weights of all edges in the graph by noticing that every edge is traversed at most once in each direction, and at most half of the edges are traversed on average.

Duality between Path Finding and Probabilistic Inference

We shall now show a connection between path finding and probabilistic inference. This connection was noticed earlier (Shimony and Charniak 1991) and was used to search for the *maximum a-posteriori* probability (MAP) assignment in graphical models using a best-first search algorithm. Here we further extend the analogy and establish a bilateral correspondence between inferring the distribution defined by a probabilistic model and learning the optimal policy in a path-finding problem.

We proceed in two steps. First, following earlier work, we establish a connection between a MAP assignment and the shortest path. Then, based on this analogy, we explain how discovering the optimal policy in a generative model can be translated into inferring the output distribution of a probabilistic program and vice versa.

Inference on probabilistic programs computes a representation of distribution (1). An equivalent form of (1) is obtained by taking logarithm of both sides:

$$\log p_{\mathcal{P}}(\mathbf{x}) = \sum_{i=1}^{|\mathbf{x}|} \log p_{F_i}(x_i) + \sum_{j=1}^{|\mathbf{y}|} \log p_{G_j}(y_j) + C \quad (2)$$

where C is a constant that does not depend on \mathbf{x} . To find the MAP assignment \mathbf{x}_{MAP} , one must maximize $\log p(\mathbf{x})$. One can view \mathbf{x} as a specification of a path in a graph where each node corresponds to either (F_i, x_i) or (G_j, y_j) , and the costs of edges entering (F_i, x_i) or (G_j, y_j) is $-\log p_{F_i}(x_i)$ or $-\log p_{G_j}(y_j)$, correspondingly (Figure 1). Then, finding

the MAP assignment is tantamount to finding the trace \mathbf{x} that produces the shortest path. (Shimony and Charniak 1991; Sun, Druzdel, and Yuan 2007) use this correspondence in their MAP algorithms for graphical models.

We shall turn now to a more general case when the MAP assignment of a part of the trace \mathbf{x}^θ is inferred. In a probabilistic program, this is expressed by selecting \mathbf{x}^θ as the program output, $\mathbf{z} \leftarrow \mathbf{x}^\theta$. The distribution of \mathbf{z} is marginalized over the rest of the trace $\mathbf{x}^{-\theta} = \mathbf{x} \setminus \mathbf{x}^\theta$, and finding the MAP assignment for \mathbf{x}^θ corresponds to finding the mode of the output distribution:

$$\begin{aligned} \mathbf{x}_{MAP}^\theta &= \arg \max p_{\mathcal{P}}(\mathbf{z}) \\ &\propto \arg \max_{\mathbf{x}^{-\theta}} \int \left[\prod_{i=1}^{|\mathbf{x}^\theta|} p_{F_i^\theta}(\mathbf{x}_i^\theta) \prod_{j=1}^{|\mathbf{y}|} p_{G_j}(y_j) \right] p_{\mathcal{P}}(\mathbf{x}^{-\theta}) d\mathbf{x}^{-\theta}, \end{aligned} \quad (3)$$

where the integrand in equation (3) depends on $\mathbf{x}^{-\theta}$. Just like in the case of MAP assignment to all random variables, equation (3) corresponds to a path finding problem: \mathbf{x}^θ can be viewed as a policy, and determining \mathbf{x}_{MAP}^θ corresponds to learning a policy which minimizes the expected path length

$$\mathbb{E}_{\mathbf{x}^{-\theta}} \left[- \sum_{i=1}^{|\mathbf{x}^\theta|} \log p_{F_i^\theta}(x_i^\theta) - \sum_{j=1}^{|\mathbf{y}|} \log p_{G_j}(y_j) \right] \quad (4)$$

While in principle policy learning algorithms could be used for MAP estimation, a greater potential lies, in our opinion, in casting planning problems as probabilistic programs and learning the optimal policies by estimating the modes of the programs' distributions. We suggest to adopt the Bayesian approach, according to which prior beliefs are imposed on policy parameters, and the optimal policy is learned through inferring posterior beliefs by conditioning the beliefs on observations. We explore this approach in the next section.

Stochastic Policy Learning through Probabilistic Inference

We have shown that in order to infer the optimum policy, a probabilistic program for policy learning should run the agent on the distribution of problem instance and policies, and compute probability of each execution such that the logarithm of the probability is equal to the negated travel cost. The generic program shown in Algorithm 1 achieves this by

Algorithm 1 Policy learning through probabilistic inference.

Require: *agent, Instances, Policies*
1: *instance* \leftarrow DRAW(*Instances*)
2: *policy* \leftarrow DRAW(*Policies*)
3: *cost* \leftarrow RUN(*agent, instance, policy*)
4: OBSERVE(1, Bernoulli(e^{-cost}))
5: PRINT(*policy*)

randomly drawing problem instances and policies from their

distributions supplied as program arguments (lines 1 and 2) and updating the log probability of the sample (line 4) by calling OBSERVE. OBSERVE adds the log probability of its first argument, the value, with respect to its second argument, the distribution. Consequently, the log probability of the output policy

$$\begin{aligned} \log p_{\mathcal{P}}(\text{policy}) &= \\ &= \log p_{Policies}(\text{policy}) + \log e^{-cost(\text{policy})} + C \\ &= \underline{-cost(\text{policy})} + \log p_{Policies}(\text{policy}) + C \end{aligned} \quad (5)$$

When policies are drawn from their distribution uniformly, $\log p_{Policies}(\text{policy})$ is the same for any policy, and does not affect the distribution of policies specified by the probabilistic program:

$$\log p_{\mathcal{P}}(\text{policy}) = -cost(\text{policy}) + C' \quad (6)$$

In practice, this is achieved by using a uniform distribution on policy parameters, such as the uniform continuous or discrete distribution for scalars, the categorical distribution with equal choice probabilities for discrete choices, or the symmetric Dirichlet distribution with parameter 1 for real vectors. Alternatively, if different policies have different probabilities with respect to the distribution *Policies* from which the policies are drawn, their log probabilities (taken with the opposite sign) have the interpretation of the costs of the corresponding policies and provide a means for specifying preferences of the model designer with respect to different policies. In either case, the optimal policy is approximated by estimating the mode of the program output.

When policies are drawn uniformly, the scale of the travel cost does not affect the choice of optimal policy. However, as follows from equation (6), the shape of the probability density (or probability mass for discrete distributions) depends on the cost scale — the higher the cost, the sharper the shape. Thus, by altering the cost scale we can affect the performance of the inference algorithm: on one hand, the mode estimate of a sharper function can be computed with higher accuracy, on the other hand, when $p_{\mathcal{P}}(\text{policy})$ changes too fast with its argument in the high probability region, approximate inference algorithms converge slowly. The right scale depends on the probabilistic program, and finding the most appropriate scale is a parameter optimization problem.

Note that the probabilistic program for policy learning is independent of the inference algorithm which would be used to obtain the results. The programmer does not need to make any assumptions about the way the mode of the output distribution is estimated, and even whether approximate or exact inference (if appropriate) is performed.

Case Study: Canadian Traveller Problem

We evaluated the proposed policy learning scheme on the Canadian Traveller Problem (Algorithm 2). The algorithm draws CTP problem instances from a given graph with traversability of each edge randomly selected according to the probabilities p , and learns a stochastic policy based on depth-first search — the policy is specified by a vector of probabilities of selecting each of the adjacent edges in every

node. When the policy is realized, the selection probabilities are conditioned such that only open unexplored edges are selected, in accordance with the base depth-first search traversal. $\text{Dirichlet}(\mathbf{1}^{\deg(v)})$ is a uniform distribution, hence the

Algorithm 2 Learning stochastic policy for the Canadian traveller problem

Require: $\text{CTP}(G, s, t, w, p)$
1: $instance \leftarrow \text{DRAW}(\text{CTP}(G, w, p))$
2: **for** $v \in V$ **do**
3: $policy(v) \leftarrow \text{DRAW}(\text{Dirichlet}(\mathbf{1}^{\deg(v)}))$
4: **end for**
5: **repeat**
6: $(reached, distance) \leftarrow \text{STDFS}(instance, policy)$
7: **until** $reached$
8: $\text{OBSERVE}(1, \text{Bernoulli}(e^{-distance}))$
9: $\text{PRINT}(policy)$

log probability of a trace is equal to the path cost taken with the opposite sign. STDFS (line 6) is a flavour of depth-first search which enumerates node children in a random order according to the policy for the current node. An optimal policy is expected to assign a higher probability to edges leading to shorter paths having lower probability to become blocked.

To assess the quality of learned policies we generated several CTP problem specifications by triangulating a randomly drawn set of either 50 or 20 nodes from Poisson-distributed points on a unit square. The average DFS travel cost in fully traversable instances was 7.9 for 50 node instances, and 5.7 for 20 node instances. The same traversal probability in the range $[0.35, 1.0]$ is assigned to every edge of the graph (the bond percolation threshold for Delaunay triangulation is ≈ 0.33 (Becker and Ziff 2009), hence instances with $p < 0.3$ are disconnected with high probability). A 50 node instance is shown in Figure 2. The s and t nodes are marked by the red circles, and edge weights are equal to the Euclidean distances between the nodes.

Lightweight Metropolis-Hastings (Wingate, Stuhlmüller, and Goodman 2011) was used for inference. We learned a policy for each problem specification by running the inference algorithm for 10,000 iterations. Then, we evaluated policies returned at different numbers of iterations on 1,000 randomly drawn instances to estimate the average travel cost. The average computation time of learning and evaluation per instance was ≈ 80 s on Intel Core i5 CPU.

The results are shown in Figure 3, where the solid lines correspond to the average travel cost over the set of problems of the corresponding size, and dashed lines to 95% confidence intervals. For both 50 and 20 node problems, the policy mostly converged after ≈ 1000 iterations, achieving 50–80% improvement compared to the uniform stochastic policy. While a further refinement of the policy is possible, a different type of policy should be learned to obtain significantly better results, for example, a deterministic policy which takes online information into account. This, however, would complicate the probabilistic program which we chose to keep as simple as possible — the actual implementation




Figure 2: An instance of CTP with 50 nodes. Initial (1) and final (25) locations are marked by red circles; edge weights are Euclidean distances between edge vertices.




Figure 3: Average travel cost vs. number of samples for problems with 50 and 20 nodes and traversability probabilities 0.85 and 0.5. The policies mostly converged after ≈ 1000 samples.

of the program is just above 100 lines of code, including the implementation of DFS.

A learned policy for a 50 node problem is visualized in Figure 4. Edge widths correspond to the confidence about the policy for the edge. Edges with higher precision (lower variance) of the policy are broader. Edge color is blue when a traversal through the edge is much more probable in one than in the other direction, and green when traversal in either direction has the same probability, with shades of green and blue reflecting how directed the edge is. As we would expect in a converged policy, edges in the center of the graph are thicker, that is, more explored, than at the periphery, where changes in the policy are less likely to affect the average




Figure 4: Visualization of policy learned for blocking probability 0.5 on instance in Figure 2. Broader edges correspond to more explored components of the policy.

travel cost. Bright blue (uni-directional) edges are mostly radial relative to the direction from the initial position (node 1) to the goal (node 25), and many well-explored tangential edges are green (bi-directional). This corresponds to an intuition about the policy — traversals through radial edges are mostly in the direction of the goal, and through the tangential edges in either direction to find an alternative route when the edge leading to the goal is blocked.

Discussion

We introduced a new approach to policy learning based on casting a policy learning task as a probabilistic program. The main contributions of the paper are:

- Discovery of bilateral correspondence between probabilistic inference and policy learning for path finding.
- A new approach to policy learning based on the established correspondence.
- A realization of the approach for the Canadian traveller problem, where improved policies were consistently learned by probabilistic program inference.

The proposed approach can be extended to many different planning problems, both in well-known path-finding applications and in other domains involving policy learning under uncertainty; Partially observable Markov Decision Processes and generalized Multi-armed bandit settings are just two examples. At the same time, the exposure of probabilistic programming tools to different domains and new applications is challenging. These tools were initially developed with certain applications in mind. Our limited experience shows that the probabilistic programming paradigm scales well to new applications and larger problems. However, as more problems are approached using the probabilistic programming methodology, apparent weaknesses and limita-

tions are uncovered, and a more powerful and flexible inference algorithm will have to be developed.

The policy learning algorithm presented here follows the offline learning scheme — the policy is selected before acting, and then used unmodified until the goal is reached. Although this is, indeed, the easiest way to cast policy learning as probabilistic inference, online learning can also be implemented so that when additional computation during acting is justified by the time cost, the policy is updated based on the information gathered online, as in some of state-of-the-art algorithms for CTP (Eyerich, Keller, and Helmert 2010). Moreover, the time cost of updating the policy incrementally based on the new evidence is lower than of inferring a new policy due to the any-time nature of Bayesian updating. Online inference is a subject of ongoing research in probabilistic programming.

By performing inference on a probabilistic program, we obtain a representation of *distribution of policies* rather than a single policy. We then use this distribution to select a policy. When the inference is performed approximately, which is a common case, the expected quality of the selected policy improves with more computation. In the most basic setting, a fixed threshold on the number of iterations of the inference algorithm can be imposed. In general, however, determining when to stop the inference and commit to a particular policy, whether in offline or online setting, is a rational metareasoning decision (Russell and Wefald 1991; Hay et al. 2012). Making this decision in an informed and systematic way is another topic for research.

References

- [Bar-Noy and Schieber 1991] Bar-Noy, A., and Schieber, B. 1991. The Canadian traveller problem. In *Proc. of the Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '91, 261–270.
- [Becker and Ziff 2009] Becker, A. M., and Ziff, R. M. 2009. Percolation thresholds on two-dimensional voronoi networks and delaunay triangulations. *Phys. Rev. E* 80:041101.
- [Bnaya, Felner, and Shimony 2009] Bnaya, Z.; Felner, A.; and Shimony, S. E. 2009. Canadian traveler problem with remote sensing. In Boutilier, C., ed., *IJCAI*, 437–442.
- [Bonet and Geffner 2001] Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artif. Intell.* 129(1-2):5–33.
- [Browne et al. 2012] Browne, C.; Powley, E.; Whitehouse, D.; Lucas, S.; Cowling, P.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A survey of Monte Carlo tree search methods. *Computational Intelligence and AI in Games, IEEE Transactions on* 4(1):1–43.
- [Eyerich, Keller, and Helmert 2010] Eyerich, P.; Keller, T.; and Helmert, M. 2010. High-quality policies for the Canadian traveler’s problem. In *Proc. of the Twenty-Fourth AAAI Conference on Artificial Intelligence, Atlanta, Georgia, USA, July 11-15, 2010*.
- [Fried et al. 2013] Fried, D.; Shimony, S. E.; Benbassat, A.; and Wenner, C. 2013. Complexity of Canadian traveler problem variants. *Theor. Comput. Sci.* 487:1–16.
- [Goodman et al. 2008] Goodman, N. D.; Mansinghka, V. K.; Roy, D. M.; Bonawitz, K.; and Tenenbaum, J. B. 2008. Church: a language for generative models. In *Proc. of Uncertainty in Artificial Intelligence*.
- [Gordon et al. 2014] Gordon, A. D.; Henzinger, T. A.; Nori, A. V.; and Rajamani, S. K. 2014. Probabilistic programming. In *International Conference on Software Engineering (ICSE, FOSE track)*.
- [Haslum, Bonet, and Geffner 2005] Haslum, P.; Bonet, B.; and Geffner, H. 2005. New admissible heuristics for domain-independent planning. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA*, 1163–1168.
- [Hay et al. 2012] Hay, N.; Russell, S. J.; Tolpin, D.; and Shimony, S. E. 2012. Selecting computations: Theory and applications. In *UAI*, 346–355.
- [Helmert 2006] Helmert, M. 2006. The Fast Downward planning system. *J. Artif. Int. Res.* 26(1):191–246.
- [Kocsis and Szepesvári 2006] Kocsis, L., and Szepesvári, C. 2006. Bandit based Monte Carlo planning. In *Proc. of the 17th European Conference on Machine Learning, ECML'06*, 282–293.
- [Lauritzen 1996] Lauritzen, S. 1996. *Graphical Models*. Clarendon Press.
- [Mansinghka, Selsam, and Perov 2014] Mansinghka, V. K.; Selsam, D.; and Perov, Y. N. 2014. Venture: a higher-order probabilistic programming platform with programmable inference. *CoRR* abs/1404.0099.
- [Nikolova and Karger 2008] Nikolova, E., and Karger, D. R. 2008. Route planning under uncertainty: The Canadian traveller problem. In *Proc. of the 23rd National Conference on Artificial Intelligence - Volume 2, AAAI'08*, 969–974. AAAI Press.
- [Paige et al. 2014] Paige, B.; Wood, F.; Doucet, A.; and Teh, Y. 2014. Asynchronous anytime sequential Monte Carlo. In *Advances in Neural Information Processing Systems*.
- [Papadimitriou and Yannakakis 1991] Papadimitriou, C. H., and Yannakakis, M. 1991. Shortest paths without a map. *Theor. Comput. Sci.* 84(1):127–150.
- [Pearl 1988] Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [Russell and Wefald 1991] Russell, S., and Wefald, E. 1991. *Do the right thing: studies in limited rationality*. Cambridge, MA, USA: MIT Press.
- [Shimony and Charniak 1991] Shimony, S. E., and Charniak, E. 1991. A new algorithm for finding MAP assignments to belief networks. In *Proc. of the Sixth Annual Conference on Uncertainty in Artificial Intelligence, UAI '90*, 185–196. New York, NY, USA: Elsevier Science Inc.
- [Sun, Druzdzel, and Yuan 2007] Sun, X.; Druzdzel, M. J.; and Yuan, C. 2007. Dynamic weighting A* search-based MAP algorithm for bayesian networks. In *Proc. of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, 2385–2390. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [Sutton and Barto 1998] Sutton, R. S., and Barto, A. G. 1998. *Introduction to Reinforcement Learning*. Cambridge, MA, USA: MIT Press, 1st edition.
- [Szörényi, Kedenburg, and Munos 2014] Szörényi, B.; Kedenburg, G.; and Munos, R. 2014. Optimistic planning in markov decision processes using a generative model. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc. 1035–1043.
- [Wingate, Stuhlmüller, and Goodman 2011] Wingate, D.; Stuhlmüller, A.; and Goodman, N. D. 2011. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proc. of the 14th Artificial Intelligence and Statistics*.
- [Wood, van de Meent, and Mansinghka 2014] Wood, F.; van de Meent, J. W.; and Mansinghka, V. 2014. A new approach to probabilistic programming inference. In *Artificial Intelligence and Statistics*.