

Maximum a Posteriori Estimation by Search in Probabilistic Programs

David Tolpin, Frank Wood

University of Oxford

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

Abstract

We introduce an approximate search algorithm for fast *maximum a posteriori* probability estimation in probabilistic programs, which we call Bayesian ascent Monte Carlo (BaMC). Probabilistic programs represent probabilistic models with varying number of mutually dependent finite, countable, and continuous random variables. BaMC is an anytime MAP search algorithm applicable to any combination of random variables and dependencies. We compare BaMC to other MAP estimation algorithms and show that BaMC is faster and more robust on a range of probabilistic models.

Introduction

Many Artificial Intelligence problems, such as approximate planning in MDP and POMDP, probabilistic abductive reasoning (Raghavan 2011), or utility-based recommendation (Shani and Gunawardana 2009), can be formulated as MAP estimation problems. 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; Wood, van de Meent, and Mansinghka 2014) can represent arbitrary probabilistic models. 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. Recent success in PMCMC methods enables efficient sampling from posterior distributions with few restrictions on the structure of the models (Wood, van de Meent, and Mansinghka 2014; Paige et al. 2014).

However, an efficient sampling scheme for finding a MAP estimate would be different from the scheme for inferring

the posterior distribution: only a single instantiation of model’s variables, rather than their joint distribution, must be found. This difference reminds of the difference between simple and cumulative reward optimization in many settings, for example, in Multi-armed bandits (Stoltz, Bubeck, and Munos 2011): when all samples contribute to the total reward, the algorithms are said to optimize the *cumulative reward*, which is the classical Multi-armed bandit settings. Alternatively, when only the quality of the final choice matters, the algorithms are said to optimize the *simple reward*. This setting is often called a *search problem*. Previous research demonstrated that different sampling schemes work better for either cumulative or simple reward, and algorithms which are optimal in one setting can be suboptimal in the other (Hay et al. 2012).

In this paper, we introduce a sampling-based search algorithm for fast MAP estimation in probabilistic programs, Bayesian ascent Monte Carlo (BaMC), which can be used with any combination of finite, countable and continuous random variables and any dependency structure. We empirically compare BaMC to other feasible MAP estimation algorithms, showing that BaMC is faster and more robust.

Preliminaries

Probabilistic Programming

Probabilistic programs are regular programs extended by two constructs (Gordon et al. 2014): a) the ability to draw random values from probability distributions, and b) the ability to condition values computed in the programs on probability distributions. A probabilistic program implicitly defines a probability distribution over program state. Formally, a probabilistic program is a *stateful deterministic computation* \mathcal{P} with the following properties:

- Initially, \mathcal{P} expects no arguments.
- On every invocation, \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 values drawn from them. We call this sequence a *trace* and denote it by \mathbf{x} . 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 program output.

Several implementations of general probabilistic programming languages are available (Goodman et al. 2008; 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 problem types, most can be used with any valid probabilistic program.

Maximum a Posteriori Probability Inference

Maximum *a posteriori* probability (MAP) inference is the problem of finding an assignment to the variables of a probabilistic model that maximizes their joint posterior probability (Murphy 2012). Sometimes, a more general problem of marginal MAP inference estimation is solved, when the distribution is marginalized over some of the variables (Doucet, Godsill, and Robert 2002; Mauá and de Campos 2012). In this paper we consider the simpler setting of MAP estimation, where assignment for all variables is sought, however the proposed algorithms can be extended to marginal MAP inference.

For certain graphical models the MAP assignment can be found exactly (Park and Darwiche 2003; Sun, Druzdzel, and Yuan 2007). However, in most advanced cases, e.g. models expressed by probabilistic programs, MAP inference is intractable, and approximate algorithms such as Stochastic Expectation-Maximization (Wei and Tanner 1990) or Simulated Annealing (Andrieu and Doucet 2000) are used.

Simulated Annealing (SA) for MAP inference constitutes a universal approach which is based on Monte Carlo sampling. Simulated Annealing is a non-homogeneous version of Metropolis-Hastings algorithm where the acceptance probability is gradually changed in analogy with the physical process of annealing (Kirkpatrick, Gelatt, and Vecchi 1983). Convergence of Simulated Annealing algorithms depends on the properties of the *annealing schedule* — the rate with which the acceptance probability changes in the course of the algorithm (Lundy and Mees 1986). When the rate is too low, the SA algorithm may take too many iterations to find the global maximum. When the rate is too high, the algorithm may fail to find the global maximum at all and get stuck in a local maximum instead. Tuning the annealing schedule is necessary to achieve reasonable performance

with SA, and the best schedule depends on both the problem domain and model parameters.

Bayesian Ascent Monte Carlo

We introduce here an approximate search algorithm for fast MAP estimation in probabilistic programs, Bayesian ascent Monte Carlo (BaMC). The algorithm draws inspiration from Monte Carlo Tree Search (Kocsis and Szepesvári 2006). Unlike Simulated Annealing, BaMC uses the information about the probability of every sample to propose assignments in future samples, a kind of *adaptive proposal* in Monte Carlo inference. BaMC differs from known realizations of MCTS in a number of ways.

- The first difference between BaMC and MCTS as commonly implemented in online planning or game playing follows from the nature of inference in probabilistic models. In online planning and games, the search is performed with the root of the search corresponding to the current state of the agent. After a certain number of iterations, MCTS commits to an action, and restarts the search for the action to take in the next state. In probabilistic program inference assignment to all variables must be determined simultaneously, hence the sampling is always performed for all variables in the model.
- Additionally, probabilistic programs often involve a combination of finite, infinite countable, and infinite continuous random variables. Variants of MCTS for continuous variables were developed (Couëtoux et al. 2011), however mixing variables of different types in the same search is still an open problem. BaMC uses *open randomized probability matching*, also introduced here, to handle all variables in a unified way independently of variable type.
- Finally, BaMC is an any-time algorithm. Since BaMC searches for an estimate of the maximum of the posterior probability, every sample with a greater posterior probability than that of all previous samples is an improved estimate. BaMC outputs all such samples. As sampling goes on, the quality of solution improves, however any currently available solution is a MAP estimate of the model with increasing quality.

BaMC (Algorithm 1) maintains beliefs about probability distribution of *log weight* (the logarithm of unnormalized probability defined by Equation 1) of the trace for each value of each random variable in the probabilistic program. At every iteration (Algorithm 1) the algorithm runs the probabilistic program (lines 4–18) and computes the log weight of the trace. If the log weight of the trace is greater than the previous maximum log weight, the maximum log weight is updated, and the trace is output as a new MAP estimate (lines 19–21). Finally, the beliefs are updated from the log weight of the sample (lines 22–23).

The ability of BaMC to discover new, improved MAP estimates depends on the way values are selected for random variables (line 7). On one hand, new values should be drawn to explore the domain of the random variable. On the other hand, values which were tried previously and resulted in a

Algorithm 1 Monte Carlo search for MAP assignment.

```
1: max-log-weight  $\leftarrow -\infty$ 
2: loop
3:   trace  $\leftarrow ()$ , log-weight  $\leftarrow 0$ 
4:   result  $\leftarrow \mathcal{P}()$  /* probabilistic program */
5:   loop
6:     if result is  $F_i$  then
7:        $x_i \leftarrow \text{SELECTVALUE}(i, F_i)$ 
8:       log-weight  $\leftarrow$  log-weight +  $\log p_{F_i}(x_i)$ 
9:       log-weight $i$   $\leftarrow$  log-weight
10:      PUSH(trace, ( $F_i, x_i$ ))
11:      result  $\leftarrow \mathcal{P}(x_i)$ 
12:     else if result is ( $G_j, y_j$ ) then
13:       log-weight  $\leftarrow$  log-weight +  $\log p_{G_j}(y_j)$ 
14:       result  $\leftarrow \mathcal{P}()$ 
15:     else if result is  $z_k$  then
16:       OUTPUT( $z_k$ )
17:       result  $\leftarrow \mathcal{P}()$ 
18:     else break
19:   if log-weight > max-log-weight then
20:     OUTPUT(trace)
21:     max-log-weight  $\leftarrow$  log-weight
22:   for  $i$  in |trace| downto 1 do
23:     UPDATE( $i$ , log-weight - log-weight $i$ )
```

high-probability trace should be re-selected sufficiently often to discover high-probability assignments conditioned on these values.

Open Randomized Probability Matching

Randomized probability matching (RPM), also called Thompson sampling (Thompson 1933), is used in many contexts where choices are made based on empirically determined choice rewards. It is a selection scheme that maintains beliefs about reward distributions of every choice, selects a choice with the probability that the average reward of the choice is the highest one, and revises beliefs based on observed rewards. Bayesian belief revision is usually used with randomized probability matching. Selection can be implemented efficiently by drawing a single sample from the belief distribution of average belief for every choice, and selecting the choice with the highest sample value (Scott 2010; Agrawal and Goyal 2012).

Here we extend randomized probability matching to domains of infinite or unknown size. We call this generalized version *open randomized probability matching* (ORPM) (Algorithm 2). ORPM is given a choice distribution, and selects choices from the distribution to maximize the total reward. ORPM does not know or assume the type of the choice distribution, but rather handles all distribution types in a unified way. Like RPM, ORPM maintains beliefs about the rewards of every choice. First, ORPM uses RPM to guess the reward distribution of a randomly drawn choice (lines 6–12). Then, ORPM uses RPM again to select a choice based on the beliefs of each choices, *including a randomly drawn choice* (lines 4–23). If the selected choice is a randomly drawn choice, the choice is drawn from the

Algorithm 2 Open randomized probability matching.

```
1: choices  $\leftarrow ()$ 
2: loop
3:   /*  $x_i \leftarrow \text{SELECTVALUE}(i, F_i)$  */
4:   if choices = () then best-choice  $\leftarrow$  random-choice
5:   else
6:     best-reward  $\leftarrow -\infty$ 
7:     best-belief  $\leftarrow \perp$ 
8:     for choice in choices do
9:       reward  $\leftarrow$  DRAW(RewardBelief(choice))
10:      if reward  $\geq$  best-reward then
11:        best-reward  $\leftarrow$  reward
12:        best-belief  $\leftarrow$  MeanRewardBelief(choice)
13:      best-reward  $\leftarrow$  DRAW(best-belief)
14:      best-choice  $\leftarrow$  random-choice
15:      for choice in choices do
16:        reward  $\leftarrow$  DRAW(MeanRewardBelief(choice))
17:        if reward  $\geq$  best-reward then
18:          best-reward  $\leftarrow$  reward
19:          best-choice  $\leftarrow$  choice
20:      if best-choice = random-choice then
21:        best-choice  $\leftarrow$  DRAW(ChoiceDistribution)
22:        RewardBelief(best-choice)  $\leftarrow$  PriorBelief
23:        choices  $\leftarrow$  APPEND(choices, best-choice)
24:      /* result  $\leftarrow \mathcal{P}()$  */
25:      reward  $\leftarrow$  EXECUTE(best-choice)
26:      /* UPDATE( $i$ , log-weight - log-weight $i$ ) */
27:      UPDATEREWARDBELIEF(best-choice, reward)
```

choice distribution (line 21) and added to the set of choices. Finally, an action is executed based on the choice (line 25) and the reward belief of the choice is updated based on the reward updated from the action.

The final form of Bayesian Ascent Monte Carlo is obtained by combining Monte Carlo search for MAP assignment (Algorithm 1) and open randomized probability matching (Algorithm 2). Selecting a value in line 7 corresponds to lines 4–23 of Algorithm 2. Line 27 of Algorithm 2 is performed at every iteration of the loop in lines 22–23 of Algorithm 1. The reliance on randomized probability matching allows to implement BaMC efficiently without any tunable parameters.

Belief Maintenance

In probabilistic models with continuous random choices $\log p_{\mathcal{P}}(\mathbf{x}|\mathbf{y})$ (Equation 1) may be both positive and negative, and is in general unbounded on either side, therefore we opted for the normal distribution to represent beliefs about choice rewards. Since parameters of reward distributions vary in wide bounds, and reasonable initial estimates are hard to guess, we used an uninformative prior belief, which is in practice equivalent to maintaining sample mean and variance estimates for each choice, and using the estimates as the parameters of the belief distribution. Let us denote by r_{ij} a random variable corresponding to the reward attributed to random choice j at selection point i . Then the reward




Figure 1: HMM with 16 observed states and unknown transition probabilities.




Figure 2: Probabilistic Deterministic Infinite Automata on the first chapter of Alice’s Adventures in Wonderland.

belief distribution is

$$r_{ij} \sim \text{Bel}(r_{ij}) \triangleq \mathcal{N}(E(r_{ij}), \text{Var}(r_{ij})), \quad (2)$$

where $E(r_{ij})$ and $\text{Var}(r_{ij})$ are the sample mean and variance of the reward, correspondingly.

In the same uninformative prior setting, the *mean reward belief* r_{ij}^{mean} , used by randomized probability matching, is

$$r_{ij}^{\text{mean}} \sim \text{Bel}(r_{ij}^{\text{mean}}) \triangleq \mathcal{N}\left(E(r_{ij}), \frac{\text{Var}(r_{ij})}{n}\right) \quad (3)$$

where n is the sample size of r_{ij} (Gelman et al. 2003).

These beliefs can be computed efficiently, and provide sufficient information to guide MAP search. Informative priors on reward distributions can be imposed to improve convergence when available, such as in the approach described in (Bai, Wu, and Chen 2013).

Empirical Evaluation

We present here empirical results for MAP estimation on two problems, Hidden Markov Model with 16 observable states and unknown transition probabilities (Figures 1, 3) and Probabilistic Infinite Deterministic Automata (Pfau, Bartlett, and Wood 2010), applied to the first chapter of “Alice’s Adventures in Wonderland” as the training data. Both represent, for purposes of MAP estimation, probabilistic models of reasonable size, with a mix of discrete and continuous random variables.

For both problems, we compared BaMC, Simulated Annealing with exponential schedule and the schedule used




Figure 3: A single run of BaMC on HMM.

in (Lundy and Mees 1986) which we customarily call Lundy-Mees schedule, as well a lightweight implementation of Metropolis-Hastings (Wingate, Stuhlmüller, and Goodman 2011) adapted for MAP search, as a base line for the comparison. In Figures 1, 2 the solid lines correspond to the medians, and the dashed lines to the 25% and 75% quantiles of MAP estimates produced by each of the algorithms over 50 runs for 4000 iterations. For each annealing schedule of SA we kept only lines corresponding to the empirically best annealing rate (choosing the rate out of the list 0.8, 0.85, 0.9, 0.95). In both case studies, BaMC consistently outperformed other algorithms, finding high probability MAP estimates faster and with less variability between runs.

In addition, Figure 3 visualizes a single run of BaMC on HMM. The solid black line shows the produced MAP estimates, the light-blue lines are weights of individual samples, and the bright blue line is the smoothed median of individual sample weights. One can see that the smoothed median approaches the MAP estimate as the sampling goes on, reflecting the fact the BaMC samples gradually converge to a small set of high probability assignments.

Discussion

In this paper, we introduced BaMC, a search algorithm for fast MAP estimation. The algorithm is based on MCTS but differs in a number of important ways. In particular, the algorithm can search for MAP in models with any combination of variable types, and does not have any parameters that has to be tuned for a particular problem domain or configuration. As a part of BaMC we introduced open randomized probability matching, an extension of randomized probability matching to arbitrary variable types.

BaMC is simple and straightforward to implement both for MAP estimation in probabilistic programs, and for stochastic optimization in general. Empirical evaluation showed that BaMC outperforms Simulated Annealing despite the ability to tune the annealing schedule and rate in the latter. BaMC coped well with cases of both finite and infinite continuous variables present in the same problem.

Full analysis of algorithm properties and convergence is still a subject of ongoing work. Conditions under which BaMC converges, as well as the convergence rate, in particular in the continuous case, still need to be established, and may shed light on the boundaries of applicability of the algorithm. On the other hand, techniques used in BaMC, in particular open randomized probability matching, span beyond MAP estimation and stochastic optimization, and may

constitute a base for a more powerful search approach in continuous and mixed spaces in general.

Acknowledgments

This work is supported under DARPA PPAML through the U.S. AFRL under Cooperative Agreement number FA8750-14-2-0004.

References

- Agrawal, S., and Goyal, N. 2012. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proc. of the 25th Annual Conference on Learning Theory (COLT)*.
- Andrieu, C., and Doucet, A. 2000. Simulated annealing for maximum a posteriori parameter estimation of hidden Markov models. *Information Theory, IEEE Transactions on* 46(3):994–1004.
- Bai, A.; Wu, F.; and Chen, X. 2013. Bayesian mixture modelling and inference based Thompson sampling in Monte-Carlo tree search. In *Advances in Neural Information Processing Systems* 26. 1646–1654.
- Couëtoux, A.; Hooock, J.-B.; Sokolovska, N.; Teytaud, O.; and Bonnard, N. 2011. Continuous upper confidence trees. In *LION'11*, 433–445.
- Doucet, A.; Godsill, S.; and Robert, C. 2002. Marginal maximum a posteriori estimation using Markov chain Monte Carlo. *Statistics and Computing* 12(1):77–84.
- Gelman, A.; Carlin, J. B.; Stern, H. S.; and Rubin, D. B. 2003. *Bayesian Data Analysis, Second Edition (Chapman & Hall/CRC Texts in Statistical Science)*. Chapman and Hall/CRC.
- 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, A. D.; Henzinger, T. A.; Nori, A. V.; and Rajamani, S. K. 2014. Probabilistic programming. In *International Conference on Software Engineering (ICSE, FOSE track)*.
- Hay, N.; Russell, S. J.; Tolpin, D.; and Shimony, S. E. 2012. Selecting computations: Theory and applications. In *UAI*, 346–355.
- Kirkpatrick, S.; Gelatt, C. D.; and Vecchi, M. P. 1983. Optimization by simulated annealing. *Science* 220(4598):671–680.
- 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, S. 1996. *Graphical Models*. Clarendon Press.
- Lundy, M., and Mees, A. 1986. Convergence of an annealing algorithm. *Mathematical Programming* 34(1):111–124.
- Mauá, D. D., and de Campos, C. P. 2012. Anytime marginal map inference. In *Proc. of the 28th International Conference on Machine Learning (ICML)*, 1471–1478.
- Murphy, K. P. 2012. *Machine learning: a probabilistic perspective*.
- Paige, B.; Wood, F.; Doucet, A.; and Teh, Y. 2014. Asynchronous anytime sequential Monte Carlo. In *Advances in Neural Information Processing Systems*.
- Park, J. D., and Darwiche, A. 2003. Solving MAP exactly using systematic search. In *UAI*, 459–468.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Pfau, D.; Bartlett, N.; and Wood, F. 2010. Probabilistic deterministic infinite automata. In Lafferty, J.; Williams, C.; Shawe-Taylor, J.; Zemel, R.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 23*. Curran Associates, Inc. 1930–1938.
- Raghavan, S. 2011. Bayesian abductive logic programs: A probabilistic logic for abductive reasoning.
- Scott, S. L. 2010. A modern Bayesian look at the multi-armed bandit. *Appl. Stoch. Model. Bus. Ind.* 26(6):639–658.
- Shani, G., and Gunawardana, A. 2009. Evaluating recommender systems. Technical Report MSR-TR-2009-159, Microsoft Research.
- Stoltz, G.; Bubeck, S.; and Munos, R. 2011. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science* 412(19):1832–1852.
- Sun, X.; Druzdzal, 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.
- Szörényi, B.; Kedenburg, G.; and Munos, R. 2014. Optimistic planning in Markov decision processes using a generative model. In *Advances in Neural Information Processing Systems* 27. 1035–1043.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25:285–294.
- Wei, G. C. G., and Tanner, M. A. 1990. A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *Journal of the American Statistical Association* 85(411):699–704.
- 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, F.; van de Meent, J. W.; and Mansinghka, V. 2014. A new approach to probabilistic programming inference. In *Artificial Intelligence and Statistics*.