

Output-Sensitive Adaptive Metropolis-Hastings for Probabilistic Programs

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

May 11th, 2015

Paper: <http://arxiv.org/abs/1501.05677>
Slides: <http://offtopia.net/almh-slides.pdf>

Outline

Probabilistic Programming

Inference

Output-sensitive Adaptive Metropolis-Hastings

Empirical Evaluation

Summary

Intuition

Probabilistic program:

- ▶ A program with random computations.
- ▶ Distributions are conditioned by ‘observations’.
- ▶ Values of certain expressions are ‘predicted’ — **the output**.

Can be written in any language (extended by `sample` and `observe`).

Example: Model Selection

```
1  (let [;; Model
2      dist (sample (categorical [[normal 1/4] [gamma 1/4]
3                                [uniform-discrete 1/4]
4                                [uniform-continuous 1/4]]))
5      a (sample (gamma 1 1))
6      b (sample (gamma 1 1))
7      d (dist a b)]
8
9  ;; Observations
10 (observe d 1)
11 (observe d 2)
12 (observe d 4)
13 (observe d 7)
14
15 ;; Explanation
16 (predict :d (type d))
17 (predict :a a)
18 (predict :b b)))
```

Definition

A **probabilistic program** is a stateful deterministic computation \mathcal{P} :

- ▶ Initially, \mathcal{P} expects no arguments.
- ▶ On every call, \mathcal{P} returns
 - ▶ a distribution F ,
 - ▶ a distribution and a value (G, y) ,
 - ▶ a value z ,
 - ▶ or \perp .
- ▶ Upon returning F , \mathcal{P} expects $x \sim F$.
- ▶ Upon returning \perp , \mathcal{P} terminates.

A program is run by calling \mathcal{P} repeatedly until termination.

The probability of each **trace** is $\propto \prod_{i=1}^{|\mathbf{x}|} p_{F_i}(x_i) \prod_{j=1}^{|\mathbf{y}|} p_{G_j}(y_j)$.

Outline

Probabilistic Programming

Inference

Output-sensitive Adaptive Metropolis-Hastings

Empirical Evaluation

Summary




Inference Objective

- ▶ Suggest **most probable explanation** (MPE) - most likely assignment for all non-evidence variable given evidence.
- ▶ Approximately **compute integral** of the form

$$\Phi = \int_{-\infty}^{\infty} \varphi(x)p(x)dx$$

- ▶ Continuously and **infinitely generate a sequence of samples** drawn from the distribution of the output expression — so that someone else puts it in good use (vague but common). ✓

Example: Inference Results



Importance Sampling

loop

Run \mathcal{P} , computing weight $w = \prod_{j=1}^{|\mathbf{y}|} p_{G_j}(y_j)$.
output \mathbf{z} , w .

end loop

- ▶ Simple — good.
- ▶ Slow convergence (unless one knows the answer) — bad.

Can we do better?

Lightweight Metropolis-Hastings (LMH)

Run \mathcal{P} once, remember \mathbf{x}, \mathbf{z} .

loop

Uniformly select x_i .

Propose a value for x_i .

Run \mathcal{P} , remember \mathbf{x}', \mathbf{z}' .

Accept $(\mathbf{x}, \mathbf{z} = \mathbf{x}', \mathbf{z}')$ or reject with MH probability.

Output \mathbf{z} .

end loop

Can we do better?

Adaptive MCMC

Adaptation opportunities:

1. Random walk (instead of proposing from priors).
2. Adapting random walk parameters.
3. Selecting each x_i with different probability. ✓

Outline

Probabilistic Programming

Inference

Output-sensitive Adaptive Metropolis-Hastings

Empirical Evaluation

Summary

Lightweight MH with Adaptive Scheduling

Maintains vector of weights \mathbf{W} of random choices:

- 1: Initialize \mathbf{W}^0 to a constant.
- 2: Run \mathcal{P} once.
- 3: **for** $t = 1 \dots \infty$ **do**
- 4: Select x_i^t with probability $\alpha_i^t = W_i^t / \sum_{i=1}^{|\mathbf{x}^t|} W_i^t$.
- 5: Propose a value for x_i^t .
- 6: Run \mathcal{P} , accept or reject with MH probability.
- 7: **if** accepted **then**
- 8: Compute \mathbf{W}^{t+1} based on the *program output*.
- 9: **else**
- 10: $\mathbf{W}^{t+1} \leftarrow \mathbf{W}^t$
- 11: **end if**
- 12: **end for**

Quantifying Influence of Program Output

- ▶ Objective: faster convergence of *program output* \mathbf{z} .
- ▶ Adaptation parameter: probabilities of selecting random choices for modification.
- ▶ Optimization target: maximize the **change** in the program output:

$$R^t = \frac{1}{|\mathbf{z}^t|} \sum_{k=1}^{|\mathbf{z}^t|} \mathbf{1}(z_k^t \neq z_k^{t-1}).$$

W_i reflects the anticipated change in \mathbf{z} from modifying x_i .

Delayed Changes

Modifying x2 affects the output ...

```
1      (let [x1 (sample (normal 1 10))
2            x2 (sample (normal x1 1))]
3          (observe (normal x2 1) 2)
4          (predict x1))
```

... but only when x1 is also modified.

Backpropagating rewards

- ▶ For each x_i , reward r_i and count c_i are kept.
- ▶ A *history* of modified random choices is attached to every z_j .

When modification of x_k accepted:

- 1: Append x_k to the history.
- 2: **if** $\mathbf{z}^{t+1} \neq \mathbf{z}^t$ **then**
- 3: $w \leftarrow \frac{1}{|\text{history}|}$
- 4: **for** x_m **in** history **do**
- 5: $\bar{r}_m \leftarrow r_m + w, c_m \leftarrow c_m + w$
- 6: **end for**
- 7: Flush the history.
- 8: **else**
- 9: $c_k \leftarrow c_k + 1$
- 10: **end if**

Outline

Probabilistic Programming

Inference

Output-sensitive Adaptive Metropolis-Hastings




Empirical Evaluation

Summary



Convergence — GP hyperparameter estimation

$$f \sim \mathcal{GP}(m, k),$$



$$\text{where } m(x) = ax^2 + bx + c, \quad k(x, x') = de^{\frac{(x-x')^2}{2g}}.$$



1000 samples



10000 samples





100000 samples

Sample Size — Kalman Smoother

$$\mathbf{x}_t \sim \text{Norm}(\mathbf{A} \cdot \mathbf{x}_{t-1}, \mathbf{Q}), \quad \mathbf{y}_t \sim \text{Norm}(\mathbf{C} \cdot \mathbf{x}_t, \mathbf{R}).$$

$$\mathbf{A} = \begin{bmatrix} \cos \omega & -\sin \omega \\ \sin \omega & \cos \omega \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} q & 0 \\ 0 & q \end{bmatrix}.$$



100 16-dimensional observations,
500 samples after 10 000 samples of burn-in.

Outline

Probabilistic Programming

Inference

Output-sensitive Adaptive Metropolis-Hastings

Empirical Evaluation

Summary

Summary

- ▶ A scheme of rewarding random choices based on program output.
- ▶ An approach to propagation of choice rewards to MH proposal scheduling parameters.
- ▶ An application of this approach to LMH, where the probabilities of selecting each variable for modification are adjusted.

Thank You