

Path Finding under Uncertainty through Probabilistic Inference

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

June 8th, 2015

Paper: <http://arxiv.org/abs/1502.07314>
Slides: <http://offtopia.net/ctp-pp-slides.pdf>

Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

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

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

Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

Summary

Inference Objective

- ▶ 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).

Inference Objective

- ▶ 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).
- ▶ Approximately **compute integral** of the form

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




Inference Objective

- ▶ 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).
- ▶ Approximately **compute integral** of the form

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

- ▶ Suggest **most probable explanation** (MPE) - most likely assignment for all non-evidence variables given evidence. ✓

Example: Inference Results



Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

Summary


Connection between MAP and Shortest Path

Maximizing the (logarithm of) trace probability

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

corresponds to finding the shortest path in a graph $G = (V, E)$:

- ▶ $V = \{(F_i, x_i)\} \cup \{(G_j, y_j)\}$.
- ▶ Edge costs are $-\log p_{F_i}(x_i)$ or $-\log p_{H_j}(y_j)$.



Marginal MAP as Policy Learning

In Marginal MAP, assignment of a *part* of the trace \mathbf{x}^θ is inferred.
In a probabilistic program:

- ▶ \mathbf{x}^θ becomes the program output \mathbf{z} .
- ▶ \mathbf{z} is marginalized over $\mathbf{x} \setminus \mathbf{x}^\theta$.
- ▶ $\mathbf{x}_{MAP}^\theta = \arg \max p_{\mathcal{P}}(\mathbf{z})$.

Determining \mathbf{x}_{MAP}^θ corresponds to learning a policy \mathbf{x}^θ which minimizes the *expected* path length

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

Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

Summary

Policy Learning through Probabilistic Inference

Require: *agent*, *Instances*, *Policies*

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

The log probability of the output policy is

$$\log p_{\mathcal{P}}(policy) = -cost(policy) + \log p_{Policies}(policy) + C$$

When policies are drawn uniformly

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

Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

Summary

Canadian Traveller Problem

CTP is a problem finding the shortest travel distance in a graph where some edges may be blocked.

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: $p_o : E \rightarrow (0, 1]$.

find the shortest travel distance from s to t — the sum of weights of all traversed edges.

The Simplest CTP Instance — Two Roads

Given

- ▶ two roads with probability being open p_1 and p_2 ,
- ▶ costs of each road c_1 and c_2 ,
- ▶ cost of bumping into a blocked road c_b ,

learn the optimum policy q .

```
1 (defquery tworoads
2   (loop []
3     (let [o1 (sample (flip p1))
4           o2 (sample (flip p2))]
5       (if (not (or o1 o2)) (recur)
6         (let [q (sample (uniform-continuous 0. 1.))
7               s (sample (flip (- 1 q)))]
8           (let [distance (if s (if o1 c1 (+ c2 cb))
9                             (if o2 c2 (+ c1 cb)))]
10              (observe +factor+ (- distance))
11              (predict :q q))))))))
```

Learning Stochastic Policy for CTP

Depth-first search based policy:

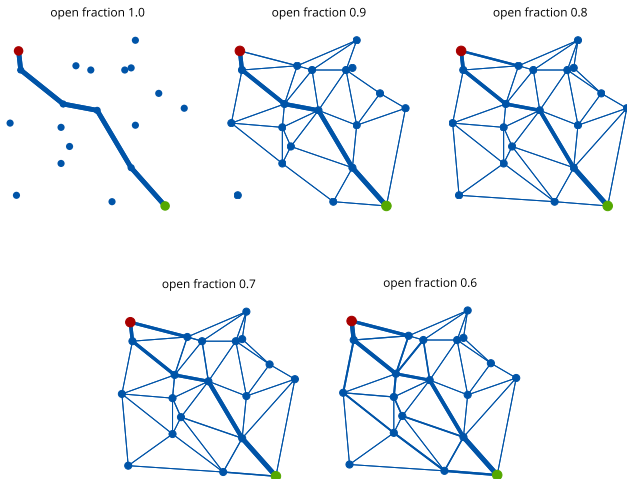
- ▶ the agent traverses G in depth-first order.
- ▶ the policy specifies the probabilities of selecting each adjacent edge in every node.

Require: $\text{CTP}(G, s, t, w, p)$

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

Inference Results — CTP Travel Graphs

Learned policies:



Line widths indicate the frequency of travelling each edge.

Outline

Probabilistic Programming

Inference

Path Finding and Probabilistic Inference

Stochastic Policy Learning

Case Study: Canadian Traveller Problem

Summary

Summary

- ▶ 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.

Thank You