

AUXILIARY VARIABLES IN PROBABILISTIC PROGRAMS

EKANSH SHARMA AND DANIEL M. ROY

ABSTRACT. We extend the first-order meta language of Staton et al. (2016) with a new language construct, **slice-let**. This new construct allows a user to define a random variable x , with a potentially intractable distribution, by introducing an auxiliary random variable, u , and specifying only the conditional distributions of x given u and of u given x . In effect, the distribution of x need only be defined up to some level of approximation, determined by the value of u . We outline the denotational semantics for **slice-let** and give some example programs. Finally, we briefly discuss an approximate representation of the program that can be used by an inference algorithm.

1. INTRODUCTION

Most probabilistic programming languages work by allowing users to define prior distributions by specifying exact samplers (see, e.g., (Goodman et al., 2008; Staton et al., 2016)). This tact is limiting; for example, it is known that certain conditional independences cannot be represented by exact samplers (Ackerman et al., 2017). In this paper, we consider the problem of representing distributions approximately. The setup is as follows: We want to introduce a random variable, x , whose distribution is potentially intractable. Normally, one would have to specify an exact sampler to use this random variable within a PPL. Instead, the programmer introduces an auxiliary variable, u , usually by specifying a simple conditional distribution for u given x . Instead of the user specifying the distribution of x , the user specifies the conditional distribution of x given u . The semantics of our construct assign the expression the distribution of x . Implementations, however, can use the two conditional distributions to generate approximate or exact-approximate inference algorithms. These types of algorithms are known as slice samplers. In effect, the distribution of x given u treats u as a level of approximation and one need only sample x up to that level of approximation. Potential applications include exact-approximate representation of Dirichlet Processes (Papaspiliopoulos and Roberts, 2008), Indian Buffet Processes (Teh, Görür, and Ghahramani, 2007).

2. SYNTAX AND SEMANTICS

We extend the first-order meta-language described in Staton et al. (2016), the syntax of which is shown in Fig. 1. This language is an idealized PPL with infinitary type and term constructor, and a constant term for all measurable functions between measurable spaces. It also contains **sample** and **observe** primitives that facilitate construction of Bayesian probabilistic models. We augment this language by adding **slice-let** $u = \text{sample}(\mu_0)$ with $x = P_1$ **update** $u = P_2$. The typing rule for **slice-let** is

$$\frac{\Gamma \vdash_d \mu_0 : \mathbb{P}(\mathbb{R}_+) \quad \Gamma, u : \mathbb{R}_+ \vdash_p P_1 : \mathbb{A} \quad \Gamma, x : \mathbb{A} \vdash_p P_2 : \mathbb{R}_+}{\Gamma \vdash_p \text{slice-let } u = \text{sample}(\mu_0) \text{ with } x = P_1 \text{ update } u = P_2 : \mathbb{A} \times \mathbb{R}_+}$$

The construct **slice-let** generalizes the probabilistic term **let**: instead of providing sequential dependence between u and x , user provides an initial (potentially incorrect) distribution for u , $\llbracket \mu_0 \rrbracket(\gamma) : \Sigma_{\llbracket \mathbb{R}_+ \rrbracket} \rightarrow [0, 1]$ and two Markov transition kernels $\llbracket P_1 \rrbracket(\gamma) : \llbracket \mathbb{R}_+ \rrbracket \times \Sigma_{\llbracket \mathbb{A} \rrbracket} \rightarrow [0, 1]$ and $\llbracket P_2 \rrbracket(\gamma) : \llbracket \mathbb{A} \rrbracket \times \Sigma_{\llbracket \mathbb{R}_+ \rrbracket} \rightarrow [0, 1]$. Informally, the semantics of **slice-let** construct is the stationary distribution on $\llbracket \mathbb{A} \rrbracket \times \llbracket \mathbb{R}_+ \rrbracket$ of the Markov chain corresponding to the kernel $P' = P'_1 P'_2$, where $P'_i = \llbracket P_i \rrbracket(\gamma)$. More formally, $\llbracket \text{slice-let } u = \text{sample}(\mu_0) \text{ with } x = P_1 \text{ update } u = P_2 \rrbracket(\gamma)$ is the unique measure $\mu : \Sigma_{\llbracket \mathbb{A} \rrbracket} \times \Sigma_{\llbracket \mathbb{R}_+ \rrbracket} \rightarrow [0, 1]$ such that, for all bounded measurable functions f , $\int f(x', u) \mu(dx, du) P'_1(u, dx') = \int f(x, u') \mu(dx, du) P'_2(x, du') = \int f(x, u) \mu(dx, du)$.

Types:

$$\mathbb{A}, \mathbb{B} ::= \mathbb{R} \mid 1 \mid \mathbb{P}(\mathbb{A}) \mid \mathbb{A} \times \mathbb{B} \mid \mathbb{A} + \mathbb{B}$$

Expressions:

$$\begin{aligned} t, u, v ::= & x \mid * \mid (t, u) \mid (i, t) \mid \pi_j(t) \mid \text{case } t \text{ of } \{(i, x) \Rightarrow u_i\}_{i \in I} \\ & \mid f(t) \mid \text{sample}(t) \mid \text{observe}(t, u) \\ & \mid \text{slice-let } u = \text{sample}(\mu_0) \text{ with } x = P_1 \text{ update } u = P_2 \\ & \mid \text{let } x = t \text{ in } u \mid \text{return}(t) \end{aligned}$$

FIGURE 1. Syntax for Language

$$\begin{aligned} \text{slice-let } u = \text{sample}(\delta_{u_0}) \text{ with } x = \text{sample}(\text{Uniform}(\text{normal-inv}(u, \mu, \sigma))) \\ \text{update } u = \text{sample}(\text{Uniform}[0, \text{density-normal}(x, \mu, \sigma)]) \end{aligned}$$

FIGURE 2. Slice-sampling a normal distribution

$$\begin{aligned} \text{let truncated-sample} = \text{slice-let } u = \text{sample}(\delta_{u_0}) \text{ with } (x, n) = \text{truncated-sample} \\ \text{update } u = \text{sample}(\text{Uniform}[0, 2^{-n}]) \end{aligned}$$

FIGURE 3. Bernoulli Process w/ decaying weights

3. EXAMPLES

Fig. 2 gives a program for slice-sampling a normal distribution (Neal, 2003). Let \mathcal{P} represent the program in Fig. 2. In this toy example, define $\text{density-normal}(x, \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma}} \exp(-\frac{(x-\mu)^2}{2\sigma^2})$, define $\text{normal-inv}(u, \mu, \sigma) = \{x \mid \text{density-normal}(x, \mu, \sigma) > u\}$ and let $u_0 \in [0, \text{normal-density}(\mu, \mu, \sigma)]$. Denotationally, the following equation holds:

$$\llbracket \mathcal{P} \rrbracket(\gamma) = \llbracket \text{sample}(\mathcal{N}(\mu, \sigma)) \rrbracket(\gamma).$$

A similar approach can be used to sample from any univariate absolutely continuous distribution (Neal, 2003).

In the second example, let $Z = (z_1, z_2, \dots)$, be an infinite dimensional vector such that $z_i \stackrel{\text{ind}}{\sim} \text{Bernoulli}(2^{-i})$. An exact representation for such infinite dimensional object requires an infinite loop. Now, we write a function $\text{truncated-sample} : (0, 1) \rightarrow \{0, 1\}^\infty \times \mathbb{N}$ in the language. Define the terms $\Gamma \vdash_d t_n : \{0, 1\}^n$ as follows:

$$\begin{aligned} t_{n+1} &:= \text{let } x = \text{flip}(2^{-(n+1)}) \text{ in return}(t_n, x) \\ t_1 &:= \text{flip}(2^{-1}) \end{aligned}$$

Also, define the terms $\Gamma, x : \{0, 1\}^n \vdash_d \text{last-active}_n : \mathbb{N}$ as follows:

$$\begin{aligned} \text{last-active}_{n+1} &:= \text{if } (\pi_{n+1}(x) == 1) \text{ return}(n+1) \\ &\quad \text{else let } x = (\pi_1(x), \dots, \pi_n(x)) \text{ in last-active}_n \\ \text{last-active}_1 &:= \text{if } (\pi_1(x) == 1) \text{ return}(1) \text{ else return}(0). \end{aligned}$$

Then, we can write the a term $\Gamma, u : \mathbb{R}_+ \vdash_p \text{truncated-sample} : \{0, 1\}^\infty \times \mathbb{N}$ as follows:

$$\begin{aligned} \text{truncated-sample} &:= \text{case } (\lfloor -\log_2(u) \rfloor, ()) \text{ of} \\ &\quad (n, ()) \implies \text{let } x = t_n \text{ in return } ((x, 0, 0, \dots), \text{last-active}_n) \end{aligned}$$

Given these definitions, the program in Fig. 3 represents exactly the distribution of interest on $\{0, 1\}^\infty$.

4. INFERENCE

There is an obvious program transformation, based on unfolding, for a language with `slice-let` that can be used as an approximate representation for an inference algorithm. In general, sampling based inference algorithms in probabilistic programming languages requires exact samples from the prior distribution. Each sample is weighted by the score of the trace. For a language with `slice-let`, we allow the user to pre-select the level of approximation they want, say k . Given this approximation level, we can sample from the k -step Markov transition kernel $(P')^k$ defined earlier. As $k \rightarrow \infty$, the distribution of the sample converges in several metrics. We are also devising a MCMC algorithm for the construct.

REFERENCES

- Ackerman, N. L., J. Avigad, C. E. Freer, D. M. Roy, and J. M. Rute (2017). “On Computable Representations of Exchangeable Data”. In: *PPS Workshop, Principles of Programming Languages (POPL 2017)*.
- Goodman, N. D., V. K. Mansinghka, D. Roy, K. Bonawitz, and J. B. Tenenbaum (2008). “Church: A language for generative models”. In: *In Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI 2008)*.
- Neal, R. M. (2003). “Slice sampling”. *Annals of statistics*, pp. 705–741.
- Papaspiliopoulos, O. and G. O. Roberts (2008). “Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models”. *Biometrika* 95.1, pp. 169–186.
- Staton, S., H. Yang, F. Wood, C. Heunen, and O. Kammar (2016). “Semantics for Probabilistic Programming: Higher-order Functions, Continuous Distributions, and Soft Constraints”. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS '16. New York, NY, USA: ACM, pp. 525–534. ISBN: 978-1-4503-4391-6.
- Teh, Y. W., D. Görür, and Z. Ghahramani (2007). “Stick-breaking construction for the Indian buffet process”. In: *Artificial Intelligence and Statistics*, pp. 556–563.