

Amortized Rejection Sampling in Universal Probabilistic Programming

Saeid Naderiparizi, Adam Ścibior, Andreas Munk, Mehrdad Ghadiri, Atılım Güneş Baydin,
Bradley Gram-Hansen, Christian Schroeder de Witt, Robert Zinkov, Philip Torr, Tom Rainforth, Yee Whye Teh, Frank Wood

TL;DR

- Rejection sampling is widely used in implementing complex generative models.
- Inference in probabilistic programs including unbounded loops (e.g. rejection sampling) is hard.
- We address the problem of efficient amortized importance-sampling-based inference, in particular Inference Compilation (IC) [4], in such models.
- We show naive application of IC can produce importance weights with unbounded variance.
- We propose Amortized Rejection Sampling (ARS), an importance sampling procedure that produces finite variance weights and unbiased expectations for programs that include rejection sampling loops.
- We implement ARS in pyprob [1; 2] in a way that requires minimal modifications to user code.

1: $x \sim p(x)$	$x \sim q(x y)$
2:	$w \leftarrow \frac{p(x)}{q(x y)}$
3: for $k \in \mathbb{N}^+$ do	for $k \in \mathbb{N}^+$ do
4: $\mathbf{z}^k \sim p(\mathbf{z} x)$	$\mathbf{z}^k \sim q(\mathbf{z} x, y)$
5:	$w^k \leftarrow \frac{p(\mathbf{z}^k x)}{q(\mathbf{z}^k x, y)}$
6:	$w \leftarrow w w^k$
7: if $c(x, \mathbf{z}^k)$ then	if $c(x, \mathbf{z}^k)$ then
8: $\mathbf{z} = \mathbf{z}^k$	$\mathbf{z} = \mathbf{z}^k$
9: break	break
10: observe ($y, p(y \mathbf{z}, x)$)	$w_{IC} \leftarrow w p(y \mathbf{z}, x)$
(a) Original program	(b) Inference compilation
1: $x \sim p(x)$	$x \sim q(x y)$
2:	$w \leftarrow \frac{p(x)}{q(x y)}$
3: $\mathbf{z} \sim p(\mathbf{z} x, c(x, \mathbf{z}))$	$\mathbf{z} \sim q(\mathbf{z} x, y, c(x, \mathbf{z}))$
4:	$w \leftarrow w \frac{p(\mathbf{z} x, c(x, \mathbf{z}))}{q(\mathbf{z} x, y, c(x, \mathbf{z}))}$
5: observe ($y, p(y \mathbf{z}, x)$)	$w_C \leftarrow w p(y \mathbf{z}, x)$
(c) Equivalent to above	(d) ARS

IC weights

$$w_{IC} = \frac{p(x)}{q(x|y)} p(y|x, z) \prod_{k=1}^L w^k$$

Theorem: Under some mild conditions if the following holds then the variance of w_{IC} is infinite.

$$\mathbb{E}_{z \sim q(z|x,y)} \left[\frac{p(z|x)^2}{q(z|x,y)^2} (1 - p(A|x, z)) \right] \geq 1$$

where A is the event of $c(x, z)$ being satisfied.

Collapsed weights

$$w_C = \frac{p(x)}{q(x|y)} \frac{p(z|x, A)}{q(z|x, y, A)} p(y|x, z)$$

- $\mathbb{E}[w_{IC}] = \mathbb{E}[w_C]$ but these weights do not cause infinite variance importance sampling estimates.
- Unfortunately, we cannot directly compute w_C

Amortized Rejection Sampling (ARS)


$$w_C = \frac{p(x)}{q(x|y)} \frac{p(z|x)}{q(z|x, y)} p(y|x, z) \frac{q(A|x, y)}{p(A|x)}$$

- $q(A|x, y)$ is the probability of exiting the rejection sampling loop under the proposal.
- $p(A|x)$ is the probability of exiting the rejection sampling loop in the original probabilistic program.
- We use Monte Carlo to get unbiased estimates of $q(A|x, y)$ and $\frac{1}{p(A|x)}$.

Marsaglia [5; 6]



```
1: for  $k \in \mathbb{N}^+$  do
2:    $a_k \sim \text{Uniform}(-1, 1)$ 
3:    $b_k \sim \text{Uniform}(-1, 1)$ 
4:    $s = a_k^2 + b_k^2$ 
5:   if  $s < 1$  then
6:      $a = a_k$ 
7:      $b = b_k$ 
8:     break
```

```
10:  $\mu = a \sqrt{\frac{-2 \log(s)}{s}}$ 
11: observe( $y_1, \mathcal{N}(\mu, \sigma^2)$ )
12: observe( $y_2, \mathcal{N}(\mu, \sigma^2)$ )
```



Mini-SHERPA

Mini-SHERPA is a simplified model of high-energy reactions of particles [3]. It uses rejection sampling extensively to simulate a particle decay event and the energy deposited by the resulting particles in a simplified detector.



Algorithm

```
1:  $x \sim q(x|y)$ 
2:  $w \leftarrow \frac{p(x)}{q(x|y)}$ 
3: for  $k \in \mathbb{N}^+$  do
4:    $\mathbf{z}^k \sim q(\mathbf{z}|x, y)$ 
5:   if  $c(x, \mathbf{z}^k)$  then
6:      $\mathbf{z} = \mathbf{z}^k$ 
7:     break
8:    $w \leftarrow w \frac{p(\mathbf{z}|x)}{q(\mathbf{z}|x, y)}$ 
9:    $K \leftarrow 0$ 
10:  for  $i \in 1, \dots, N$  do
11:     $\mathbf{z}'_i \leftarrow q(\mathbf{z}|x, y)$ 
12:     $K \leftarrow K + c(\mathbf{z}, x)$ 
```

```
13: for  $j \in 1, \dots, M$  do
14:   for  $l \in \mathbb{N}^+$  do
15:      $\mathbf{z}''_{j,l} \leftarrow q(\mathbf{z}|x, y)$ 
16:     if  $c(x, \mathbf{z}''_{j,l})$  then
17:        $T_j \leftarrow l$ 
18:       break
19:      $T \leftarrow \frac{1}{M} \sum_{j=1}^M T_j$ 
20:      $w \leftarrow w \frac{K T}{N}$ 
21:      $w \leftarrow w p(y|\mathbf{z}, x)$ 
```

We introduce two new functions to tag the beginning and end of rejection sampling loops.

Original

```
x = sample(P_x)
while True:
  rs_start()
  z = sample(P_z(x))
  if c(x, z):
    rs_end()
  break
observe(P_y(x, z), y)
return x, z
```

Annotated

```
x = sample(P_x)
while True:
  rs_start()
  z = sample(P_z(x))
  if c(x, z):
    rs_end()
  break
observe(P_y(x, z), y)
return x, z
```

References

- [1] A. G. Baydin, L. Heinrich, W. Bhimji, B. Gram-Hansen, G. Louppe, L. Shao, K. Cranmer, F. Wood, et al. Efficient probabilistic inference in the quest for physics beyond the standard model. In *Thirty-second Conference on Neural Information Processing Systems (NeurIPS)*. 2019.
- [2] A. G. Baydin, L. Shao, W. Bhimji, L. Heinrich, L. Meadows, J. Liu, A. Munk, S. Naderiparizi, B. Gram-Hansen, G. Louppe, et al. Etalumis: Bringing probabilistic programming to scientific simulators at scale. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '19)*. 2019.
- [3] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter. Event generation with sherpa 1.1. *Journal of High Energy Physics*, 2009(02):007, 2009.
- [4] T. A. Le, A. G. Baydin, and F. Wood. Inference compilation and universal probabilistic programming. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1338–1348, Fort Lauderdale, FL, USA, 2017. PMLR.
- [5] G. Marsaglia and T. A. Bray. A convenient method for generating normal variables. *SIAM review*, 6(3):260–264, 1964.
- [6] F. Wood, J. W. Meent, and V. Mansingka. A new approach to probabilistic programming inference. In *Artificial Intelligence and Statistics*, pages 1024–1032, 2014.