

# Learning Proposals for Probabilistic Programs with Inference Combinators

Sam Stites\*<sup>1</sup>

Heiko Zimmermann\*<sup>2</sup>

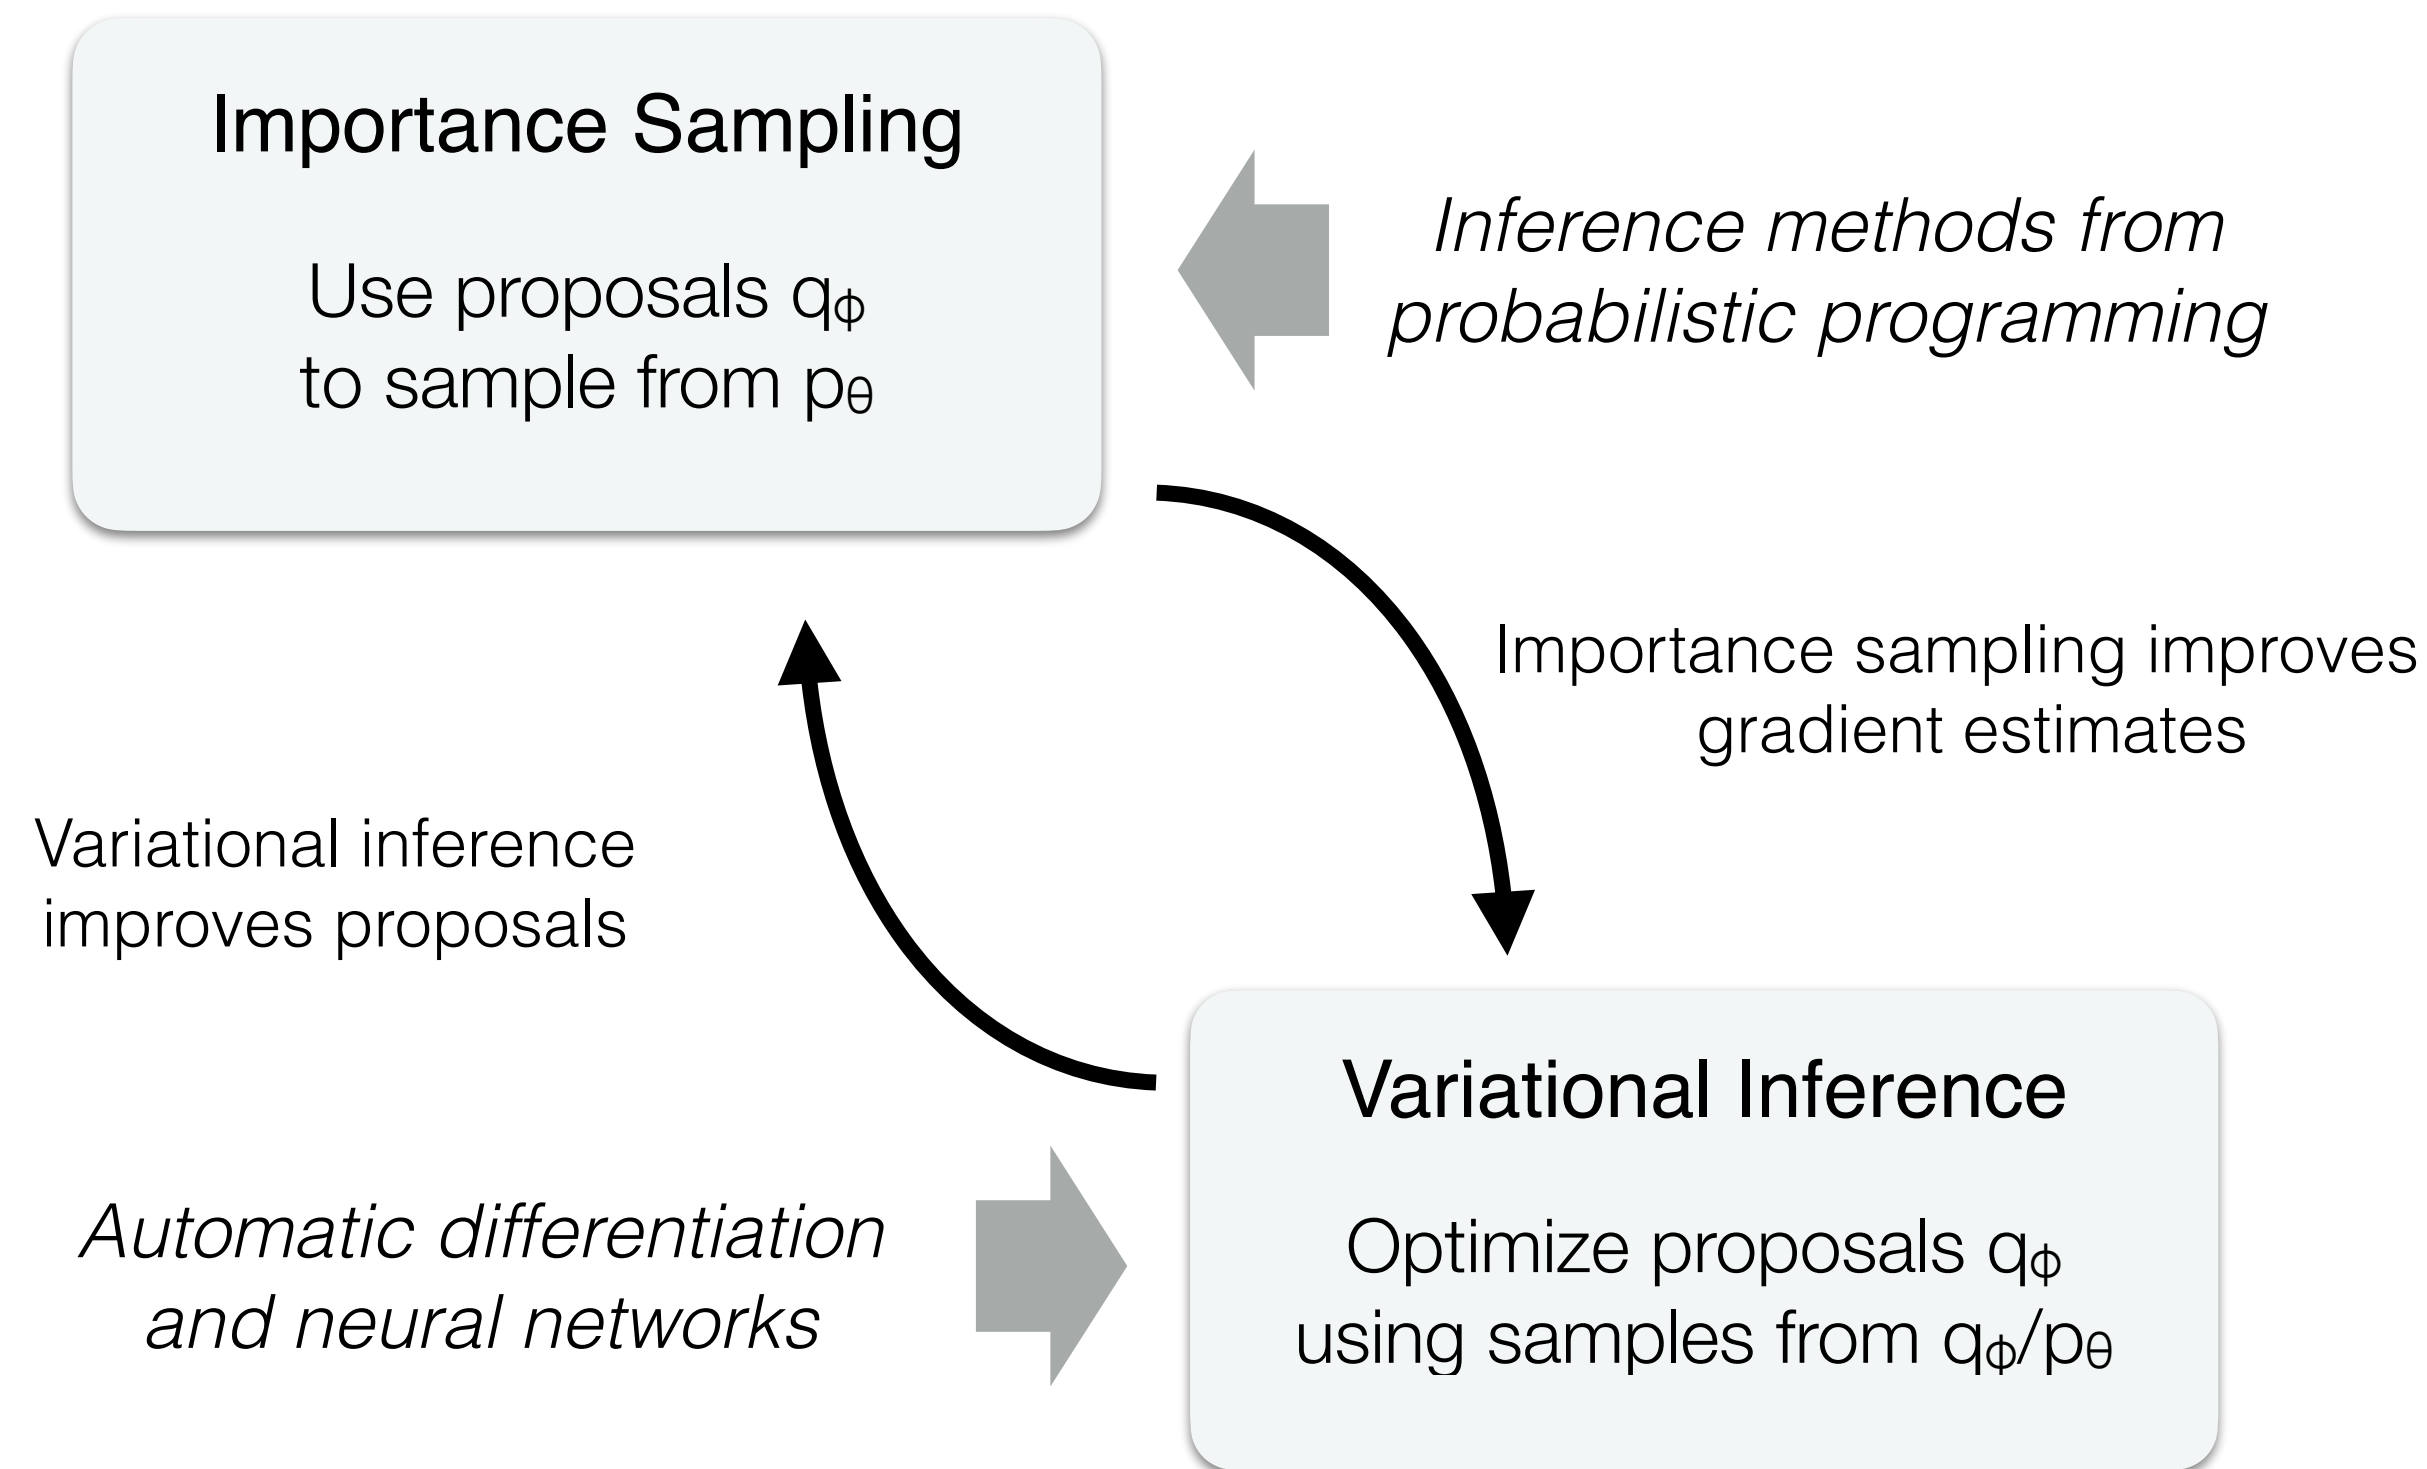
Hao Wu<sup>1</sup>

Eli Sennesh<sup>1</sup>

Jan-Willem van de Meent<sup>1,2</sup>

<sup>1</sup>Khoury College of Computer Sciences, Northeastern University, Massachusetts, USA <sup>2</sup>AMLab, University of Amsterdam, Amsterdam, The Netherlands

## 1 Importance Sampling and Variational Inference



## 2 Properly Weighted Program Evaluations

**Denotational Semantics.** We assume the existence of denotational semantics, which for a program  $f$ , define a prior and unnormalized density

$$\llbracket f(c') \rrbracket_\gamma(\tau) = \gamma_f(\tau; c'), \quad \llbracket f(c') \rrbracket_p(\tau) = p_f(\tau; c').$$

**Inference.** Given the unnormalized density, we want to approximate the corresponding normalized density


$$\pi_f(\tau; c') = \frac{\gamma_f(\tau; c')}{Z_f(c')}, \quad Z_f(c') = \int d\tau \gamma_f(\tau; c').$$

**Definition.** A evaluation  $c, \tau, \rho, w \leftarrow q(c')$  is strictly properly weighted for an unnormalized density  $\gamma_q \equiv Z_q \pi_q$ , if for all measurable functions  $h$

$$\mathbb{E}_{q(c')} [w h(\tau)] = Z_q(c') \mathbb{E}_{\pi_q(\cdot; c')} [h(\tau)].$$



## 3 A DSL for importance sampling



**Operational Semantics.** We propose a grammar for composing importance samplers and develop inference rules which guarantee that these samplers are valid by construction.

$f ::=$  A primitive program  
 $p ::= f \mid \text{extend}(p, f)$   
 $q ::= p \mid \text{resample}(q) \mid \text{compose}(q', q) \mid \text{propose}(p, q)$

**Theorem.** Evaluation of an inference program  $q(c)$  is strictly properly weighted for its unnormalized density  $\llbracket q(c) \rrbracket_\gamma$ .

$$\begin{aligned} c_1, \tau_1, \rho_1, w_1 &\leftarrow q(c_0) & c_2, \tau_2, \rho_2, w_2 &\leftarrow p(c_0)[\tau_1] \\ c_3, \tau_3, \rho_3, w_3 &\leftarrow \text{marginal}(p)(c_0)[\tau_2] \\ u_1 &= \prod_{\alpha \in \text{dom}(\rho_1) \setminus (\text{dom}(\tau_1) \cup \text{dom}(\tau_2))} \rho_1(\alpha) \\ c_3, \tau_3, \rho_3, w_2 \cdot w_1 / u_1 &\leftarrow \text{propose}(p, q)(c_0) \end{aligned}$$

$c_1, \tau_1, \rho_1, w_1 \leftarrow q(c_0) \quad c_2, \tau_2, \rho_2, w_2 \leftarrow p(c_0)$   
 $\text{dom}(\rho_1) \cap \text{dom}(\rho_2) = \emptyset \quad \text{dom}(\rho_2) = \text{dom}(\tau_2)$   
 $c_2, \tau_2, \rho_2 = \text{REINDEX}(\tau_1, \tau_1, \tau_1) \quad \bar{w}_2 = \text{MEAN}(\bar{w}_1)$   
 $c_3, \tau_3, \rho_3, w_1, w_2 \leftarrow \text{compose}(q_2, q_1)(c_0) \quad c_3, \tau_3, \rho_3, w_1, w_2 \leftarrow \text{extend}(p, f)(c_0)$   
 $\bar{c}_3, \bar{\tau}_3, \bar{\rho}_3, \bar{w}_2 \leftarrow \text{resample}(q)(c_0)$

## 4 Learning Neural Proposals

**Optimization.** We optimize the parameters of the neural proposals by optimizing a divergence or divergence-based stochastic bound  $\mathcal{D}$  at each level of nesting (propose statement).

$$\mathcal{D}(\theta, \phi) = D_1(p_1 \parallel \pi_1) + \sum_{k=2}^K D_k(\pi_{k-1} \parallel \pi_k)$$

Initial proposal program    Intermediate IS targets    Final IS targets

## 5 Example: Amortized Population Gibbs (APG) Samplers

**Writing inference programs with combinators.** Combinators allow us to implement otherwise complicated inference algorithms with comparative ease.

**APG algorithm block from paper:**

```

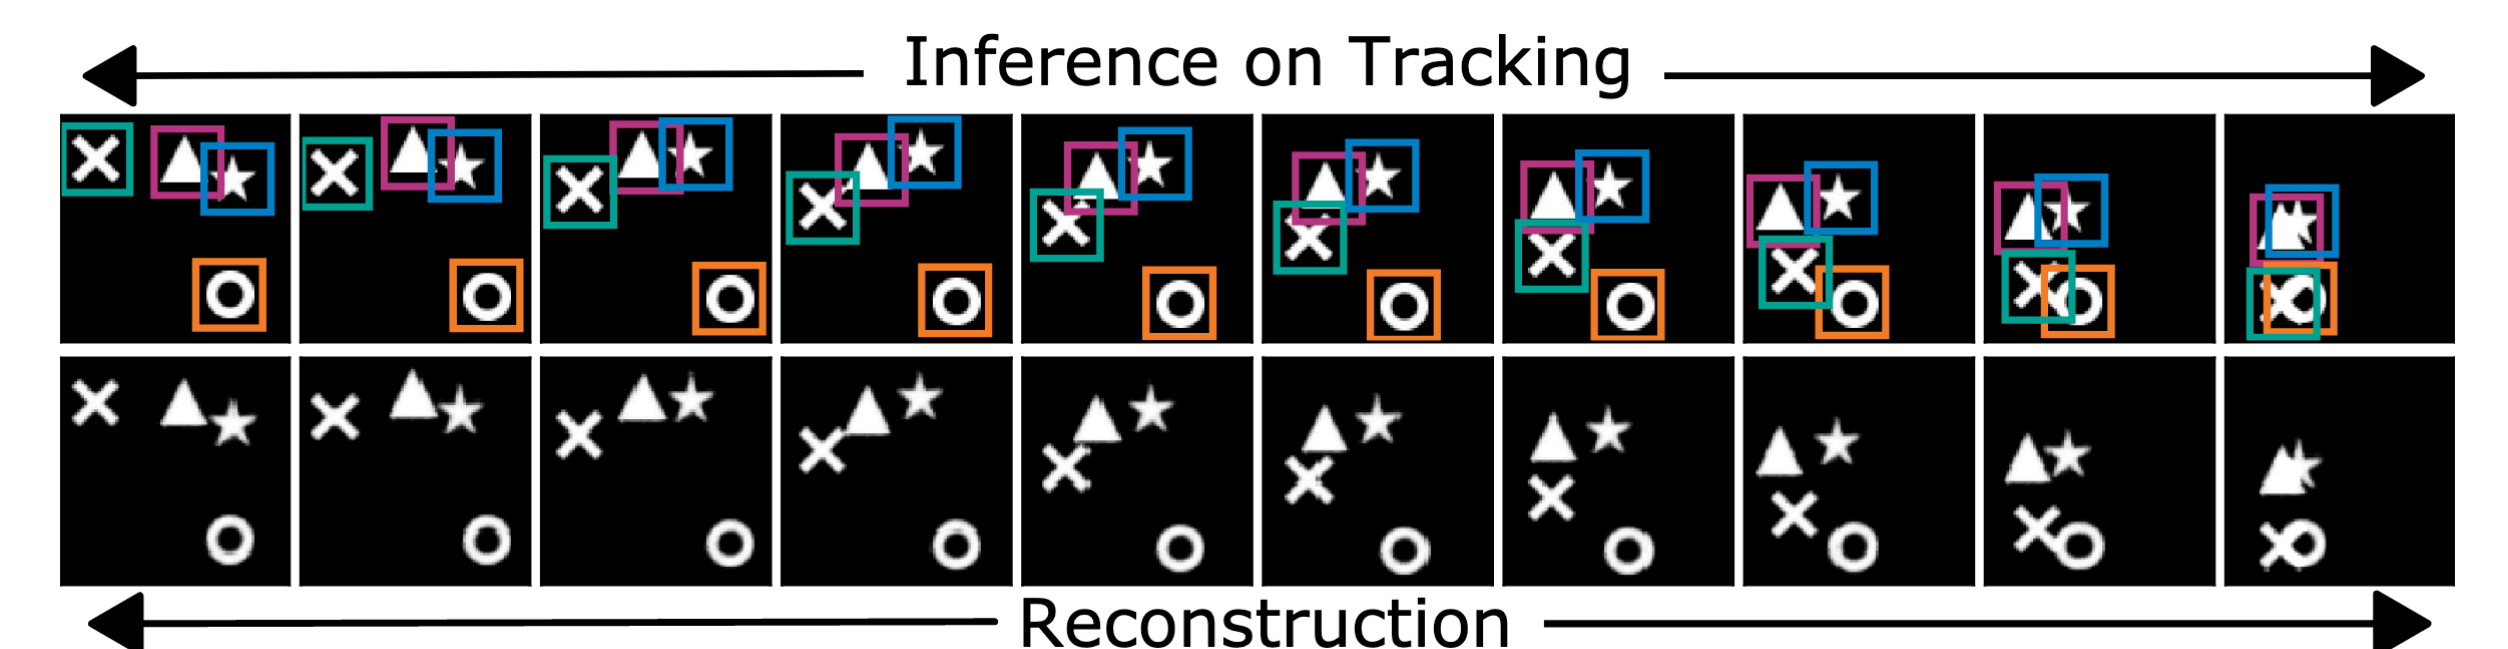
for n in 1, ..., N do
  G_phi = 0
  x^n ~ p^DATA(x)
  for l in 1, ..., L do
    z^{n,l} ~ q_phi(z | x^n)
    w^{n,l} ~ p_theta(x^n, z^{n,l}) / q_phi(z^{n,l})
  for k in 2, ..., K do
    z-tilde, w-tilde = z^{n,k-1}, w^{n,k-1}
    for b in 1, ..., B do
      z-tilde, w-tilde = RESAMPLE(z-tilde, w-tilde)
      for l in 1, ..., L do
        z_b^l ~ q_phi(. | x^n, z-tilde_b^l)
        w-tilde^l = (p_theta(x^n, z_b^l, z-tilde_b^l) q_phi(z_b^l | x^n, z-tilde_b^l)) / (p_theta(x^n, z-tilde_b^l, z-tilde_b^l) q_phi(z-tilde_b^l | x^n, z-tilde_b^l)) w-tilde^l
        z_b^l = z-tilde_b^l
      G_phi = G_phi + sum_{l=1}^L (w-tilde^l / sum_{l'} w-tilde^{l'}) * (d/dphi) log q_phi(z_b^l | x^n, z-tilde_b^l)
    z^{n,k}, w^{n,k} = z-tilde, w-tilde
  return G_phi, z, w
    
```

**APG code in combinators:**

```

def pop_gibbs(target, proposal, kernels, sweeps):
  q = propose(partial(target, suffix=0), partial(proposal, suffix=0))
  for s in range(sweeps):
    for k in kernels:
      q = propose(
        extend(partial(target, suffix=s+1), partial(k, suffix=s)),
        compose(partial(k, suffix=s+1), resample(q, dim=0)))
  return q
    
```

**Inference results:**



**Implementation.**

Combinators can be implemented on top of most current PPLs. We provide an implementation based on probtorch and an open design document for pyro at the following:

<https://github.com/probtorch/combinators>

<https://bit.ly/pyro-design>