Sticking the landing: A simple reduced-variance gradient for ADVI

Geoffrey Roeder, Yuhuai Wu, and David Duvenaud
{roeder, ywu, duvenaud}@cs.toronto.edu

Main Idea

- We give an estimator of the reparameterized ELBO gradient with lower variance when the variational approximation is close to truth
- Bigger improvement for flexible families like normalizing flows, IWAE, Hamiltonian variational inference
- Simple to implement

Three forms of the ELBO:

$$\mathcal{L}(\phi) = \mathbb{E}_{\mathbf{z}\sim q}[\log p(x|z)] - KL(q_\phi(z|x)||p(z)) \qquad \text{(exact KL)}$$
$$= \mathbb{E}_{\mathbf{z}\sim q}[\log p(x|z) + \log p(z))] + \mathbb{H}[q_\phi] \qquad \text{(exact entropy)}$$
$$= \mathbb{E}_{\mathbf{z}\sim q}[\log p(x|z) + \log p(z) - \log q_\phi(z|x)] \qquad \text{(Monte Carlo)}$$

- KL seems lowest variance, because it analytically integrates out some terms

Monte Carlo variance goes to zero!

... as the approximate posterior gets close to the true posterior

- If $q(z|x) = p(z|x)$, then a fully Monte Carlo estimator has *0* variance, since

$$\hat{\mathcal{L}}_{MC}(\phi) = \log p(x|z_i) + \log p(z_i) - \log q_\phi(z_i|x) \qquad z_i \sim_{\text{iid}} q(z) \qquad (1)$$
$$= \log p(z_i|x) + \log p(x) - \log p(z_i|x) \qquad (\text{using } q(z|x) = p(z|x)\ ) \qquad (2)$$
$$= \log p(x) \qquad (3)$$

a **constant** w.r.t. z!

But what about the gradient of the ELBO?

- It turns out that the naive fully Monte Carlo gradient estimator *doesn't* go to zero. Why not?

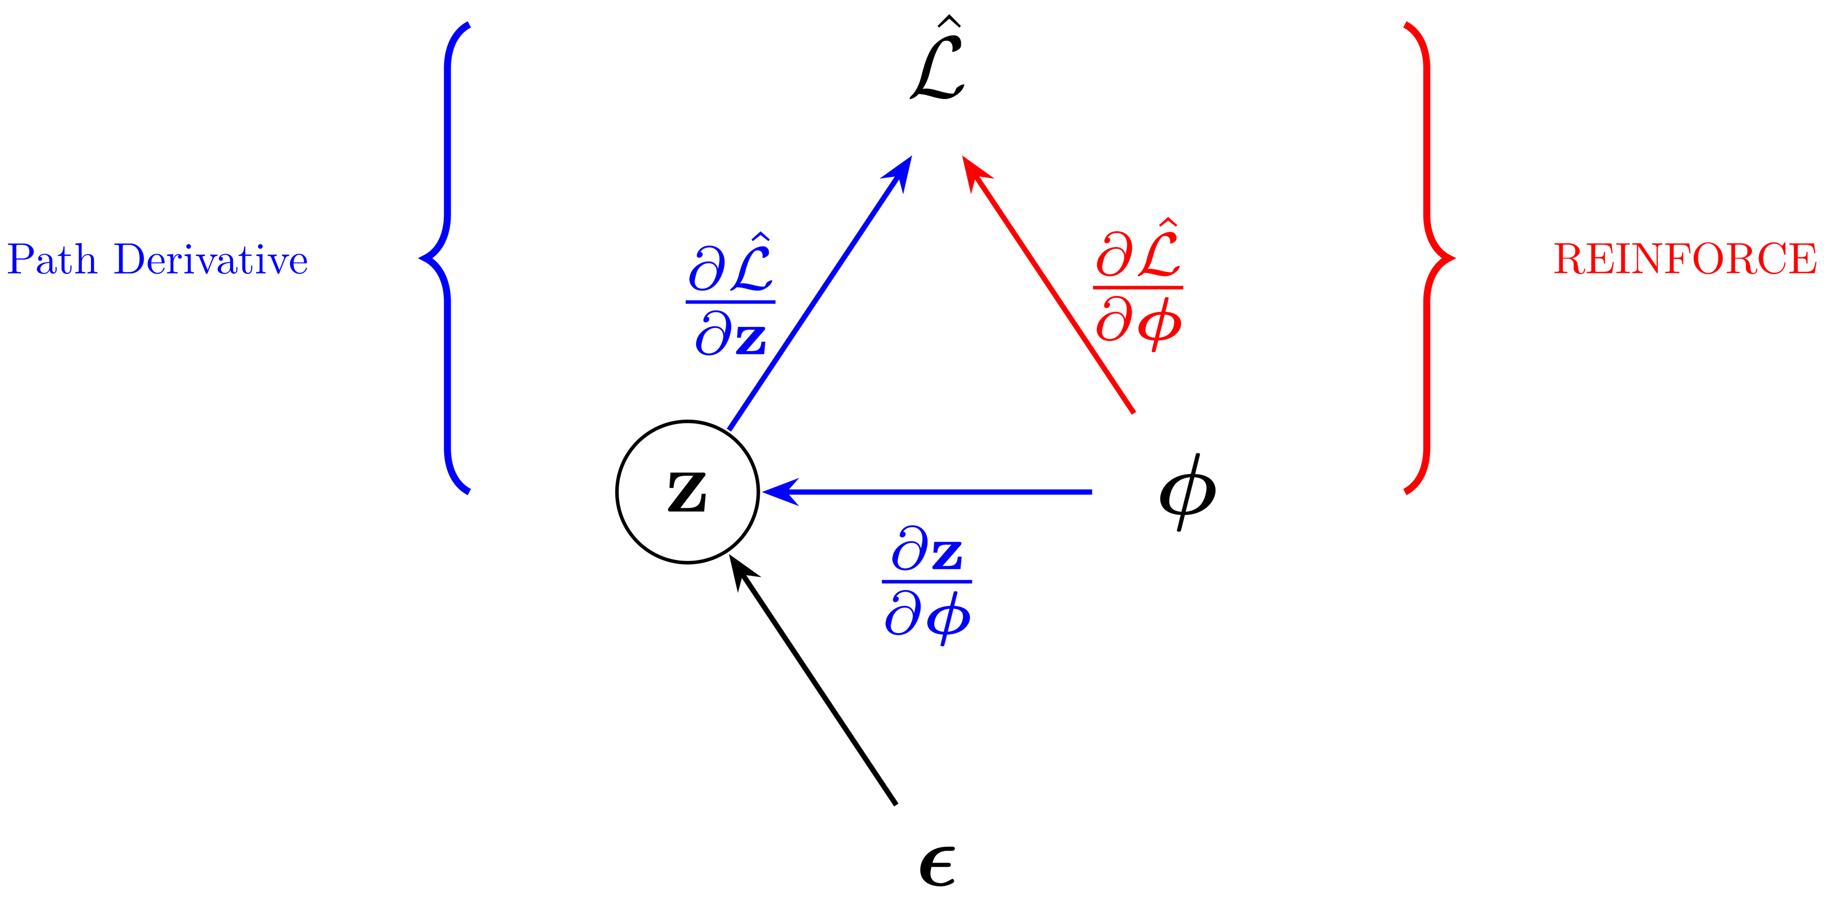

$$\hat{\nabla}_{\text{MC}} = \nabla_\phi \left[\log p(x|z_\phi) + \log p(z_\phi) - \log q_\phi(z_\phi|x)\right] \qquad \epsilon \sim_{\text{iid}} \mathcal{N}(0, I)$$

$$= \underbrace{\frac{\partial \log p(z_\phi|x)}{\partial z_\phi}\frac{\partial z_\phi}{\partial \phi} - \frac{\partial \log q(z_\phi|x)}{\partial z_\phi}\frac{\partial z_\phi}{\partial \phi}}_{\text{Path Derivative}} - \underbrace{\frac{\partial \log q_\phi(z|x)}{\partial \phi}}_{\text{REINFORCE}}.$$

- The Path Derivative component is analytically 0, but the REINFORCE gradient has variance (equal to Fisher information) even when $q(z|x) = p(z|x)$.

Fix: remove the REINFORCE gradient!

- REINFORCE component (score function) has expectation 0
- **Still unbiased estimator of ELBO gradient**
- This can be interpreted as a control variate

In other words:

When the variational approximation is exact...

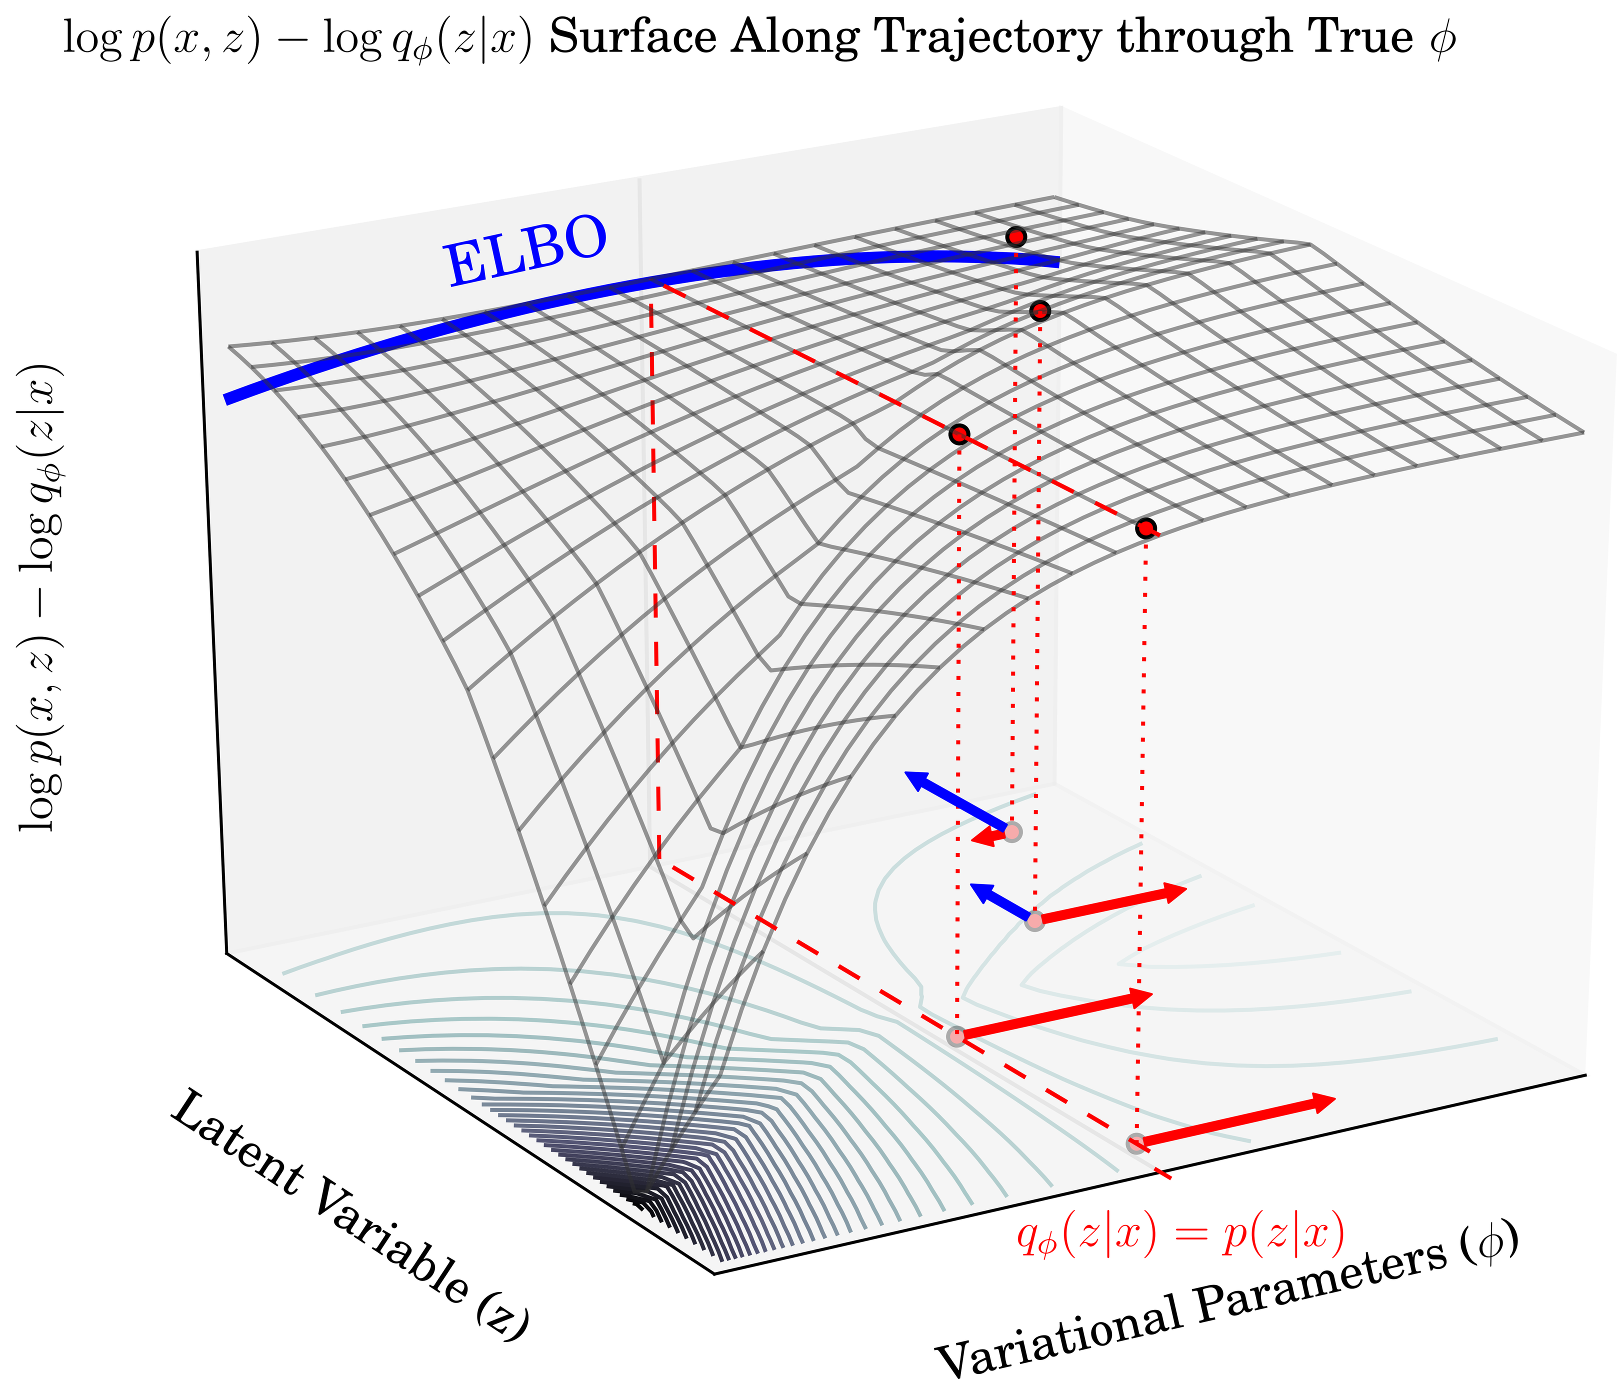

... the gradient w.r.t. the parameters is non-zero

The new gradient estimator fixes this:

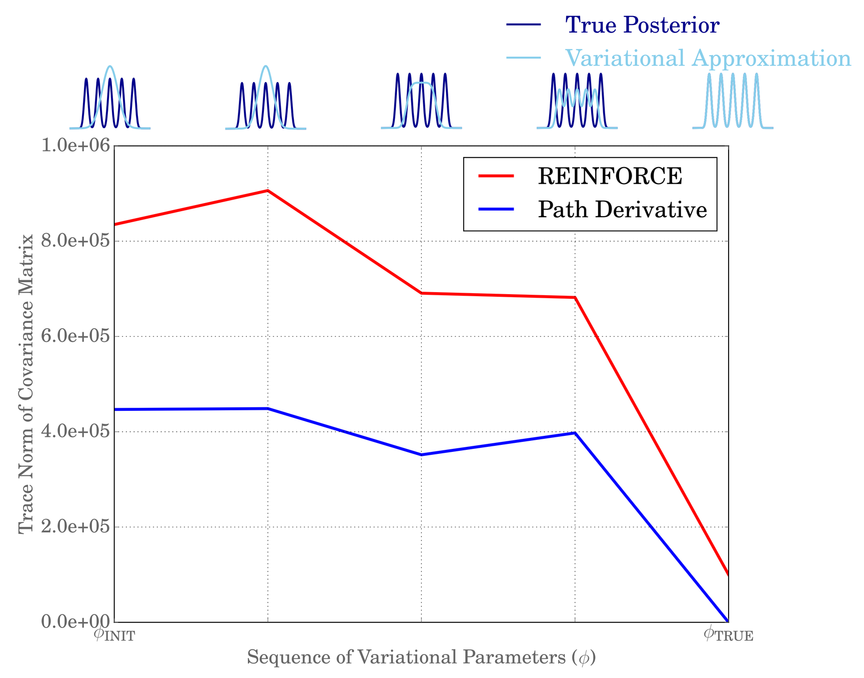

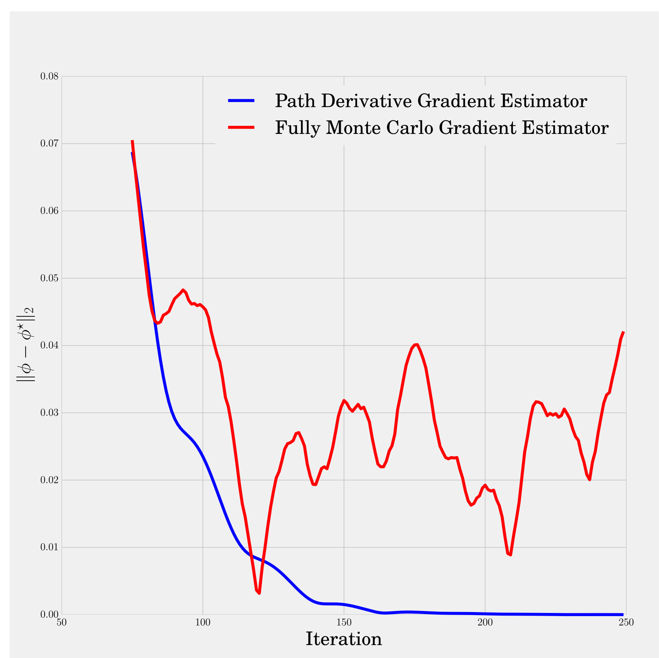

So the optimizer "sticks the landing"

Implementing this is easy

- Block gradient through variational params. Ex. for Gaussian:

```
logq = -SUM(SQUARE(z - mu_z))) / (2*EXP(log_sig_z)) - C - SUM(log_sig_z)
```

becomes

```
logq = -SUM(SQUARE(z - BLOCK(mu_z)) / (2*EXP(BLOCK(log_sig_z)) - C - 0.5*SUM(BLOCK(log_sig_z)
```

- Use `gradient_disconnected`, `stop_gradient` in `Theano`, `TensorFlow`
- In `autograd`, define custom gradient that only evaluates path derivative

Future work

- Unlikely that lower variance is maintained away from true posterior
- Control variates use an optimal scaling parameter: implement this!
- Empirical study of performance in multi-sample setting (IWAE)

Related work

- BBSVI (no reparameterization) uses a similar control variate
- Tan et al. 2016 note phenomenon in sparse precision Gaussian VI
- Han et al. 2015 note phenomenon in Gaussian copula VI models
- Our goal is to present unified analysis with easy implementation