Reparameterized Variational Rejection Sampling

Martin Jankowiak
Generate Biomedicines

Du Phan
Google Research

AISTATS 2024

Consider a latent variable model:

$$p_{\boldsymbol{\theta}}(\underset{\substack{\nearrow\\ \text{observations}}}{\mathbf{x}}, \underset{\substack{\nearrow\\ \text{latent variable}}}{\mathbf{z}}) = \underset{\text{likelihood}}{p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})}\,\underset{\text{prior}}{p_{\boldsymbol{\theta}}(\mathbf{z})}$$

How do we approximate the posterior?

$$p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$$

How do we compute the evidence?

$$p_\theta(\mathbf{x}) = \int d\mathbf{z}\; p_\theta(\mathbf{x}, \mathbf{z})$$

Variational inference casts Bayesian inference as an **optimization problem**

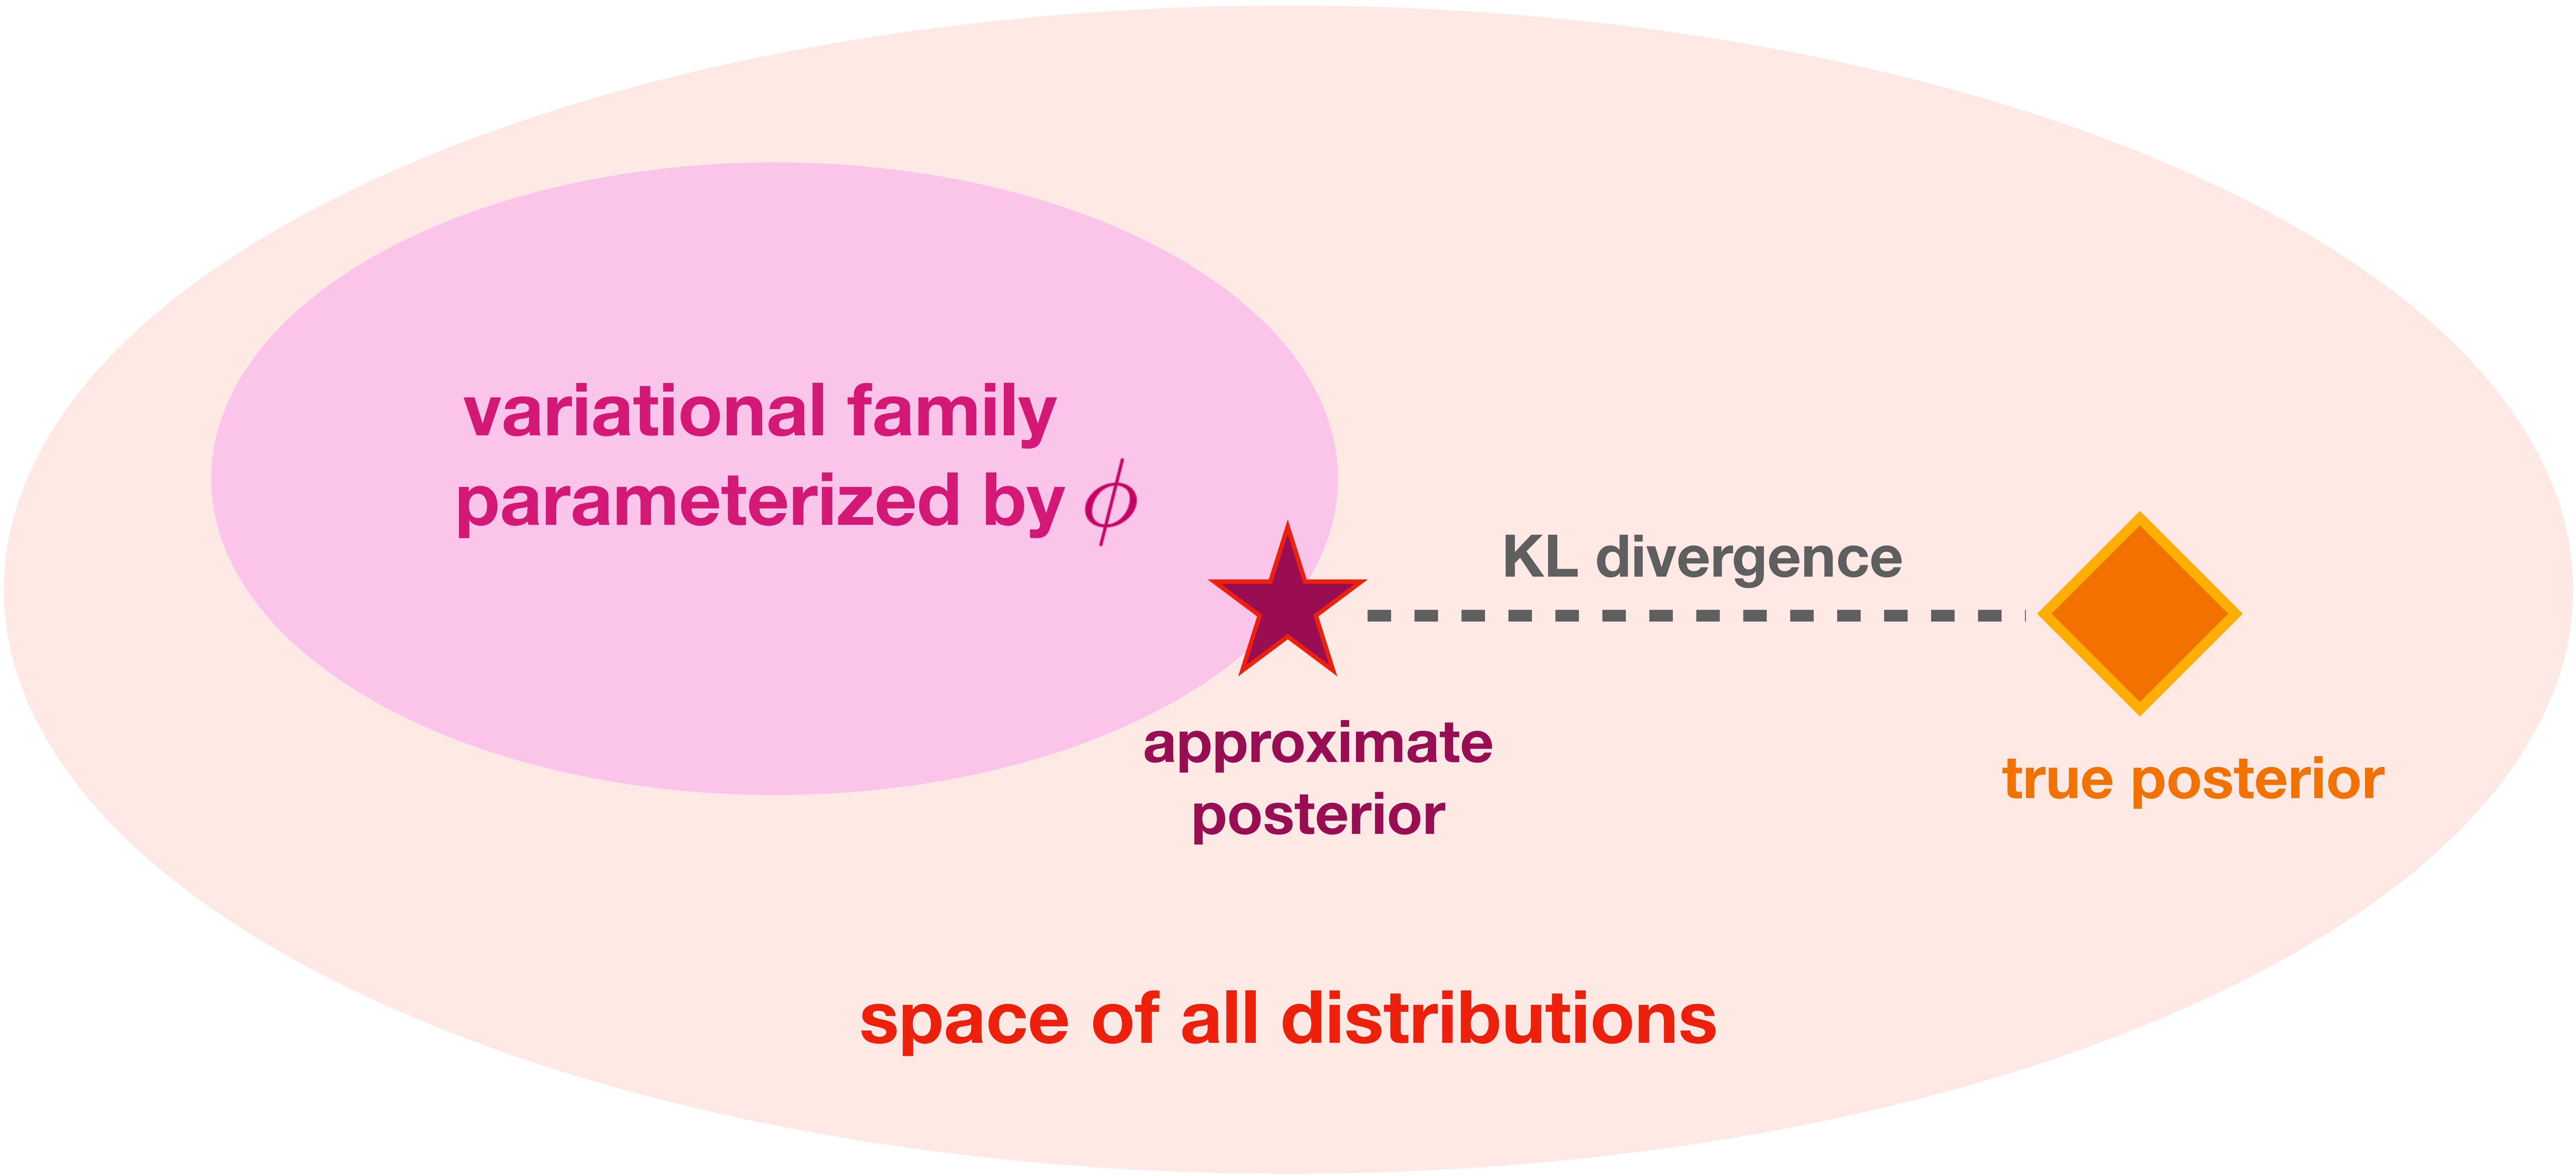

Variational inference casts Bayesian inference as an **optimization problem**

Concretely we optimize the ELBO variational objective:

$$\mathcal{L} \equiv \mathbb{E}_{q_{\boldsymbol{\phi}}(\mathbf{z})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) - \log q_{\boldsymbol{\phi}}(\mathbf{z})] \leq \log p_{\boldsymbol{\theta}}(\mathbf{x})$$

variational distribution

lower bound to log evidence

Parametric Variational Inference		**MCMC**	
Asymptotically **biased**	😢	Asymptotically **unbiased**	😍
Does estimate **evidence**	😍	Does **not** estimate **evidence**	😢
Does support **amortization**	😍	Does **not** support **amortization**	😢
Does support **subsampling**	😍	Does **not** support **subsampling**	😢
Does support **model learning**	😍	Does **not** support **model learning**	😢

How can we get the best of both worlds?

Varieties of variational inference

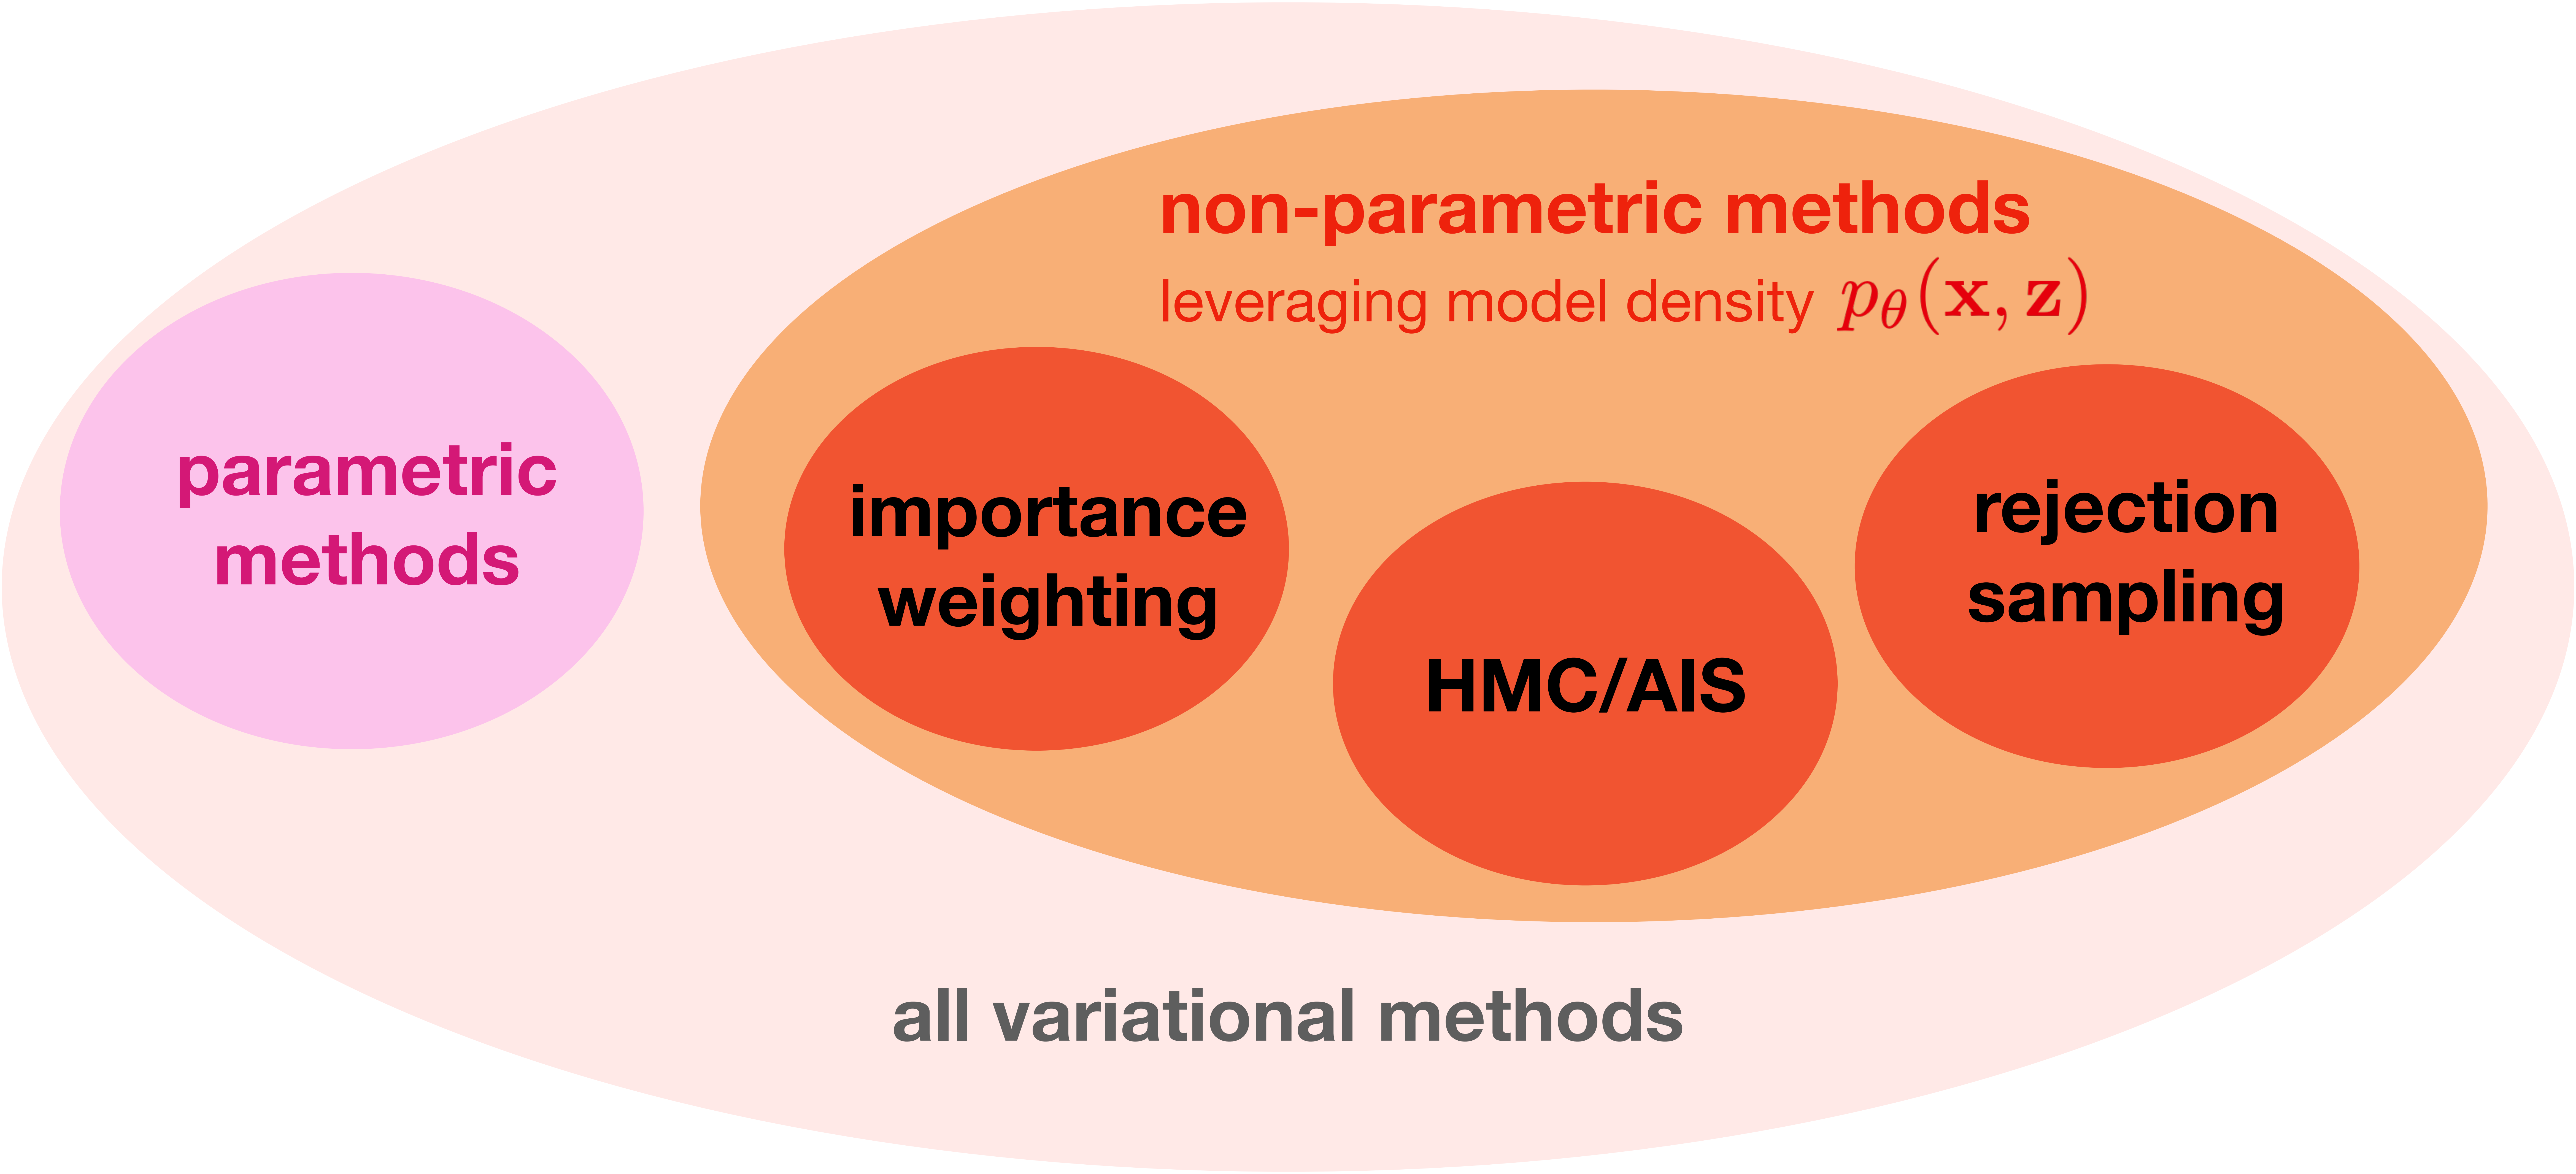

non-parametric methods

leveraging model density $p_\theta(\mathbf{x}, \mathbf{z})$

parametric methods

importance weighting

HMC/AIS

rejection sampling

all variational methods

Variational Rejection Sampling

Take a **parametric proposal distribution** $q_{\phi}(\mathbf{z})$
and **warp** it towards the posterior $p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$
via a smoothed form of **rejection sampling**:

variational distribution

$$r_{\phi,\boldsymbol{\theta}}(\mathbf{z}) \equiv \frac{q_{\phi}(\mathbf{z})a_{\phi,\boldsymbol{\theta}}(\mathbf{z})}{\mathcal{Z}_r} \text{ with } \mathcal{Z}_r \equiv \int d\mathbf{z}\, q_{\phi}(\mathbf{z})a_{\phi,\boldsymbol{\theta}}(\mathbf{z})$$

acceptance probability

$$a_{\phi,\boldsymbol{\theta}}(\mathbf{z}) \equiv \sigma(\log p_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{z}) - \log q_{\phi}(\mathbf{z}) + T)$$

↖ logistic function

↖ tunable threshold

Reference: *"Variational Rejection Sampling,"* A. Grover, R. Gummadi, M. Lazaro-Gredilla, D. Schuurmans, & S. Ermon, AISTATS 2018.

Aside: stochastic gradient estimators

score function gradient estimator

$$\nabla_{\phi}\mathbb{E}_{q_{\phi}(\mathbf{z})}\left[f(\mathbf{z})\right] = \mathbb{E}_{q_{\phi}(\mathbf{z})}\left[f(\mathbf{z})\nabla_{\phi}\log q_{\phi}(\mathbf{z})\right]$$

test function

pathwise gradient estimator

$$\nabla_{\phi}\mathbb{E}_{q_{\phi}(\mathbf{z})}\left[f(\mathbf{z})\right] = \mathbb{E}_{q_{\phi}(\mathbf{z})}\left[\frac{\partial f(\mathbf{z})}{\partial \mathbf{z}} \cdot \nabla_{\phi}\mathbf{z}\right]$$

utilizes gradient of test function

velocity field

Variational Rejection Sampling

VRS uses a **high-variance gradient estimator** for the parameters of the proposal distribution ϕ

score function

$$\nabla_{\phi}\mathcal{L} = \mathrm{COV}_{r_{\phi,\boldsymbol{\theta}}(\mathbf{z})}\left[\mathcal{A}(\mathbf{z}), a_{\phi,\boldsymbol{\theta}}(\mathbf{z})\boxed{\nabla_{\phi}\log q_{\phi}(\mathbf{z})}\right]$$

$$\mathcal{A}(\mathbf{z}) \equiv \log p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) - \log q_{\phi}(\mathbf{z}) - \log a_{\phi,\boldsymbol{\theta}}(\mathbf{z})$$

Reparameterized Variational Rejection Sampling

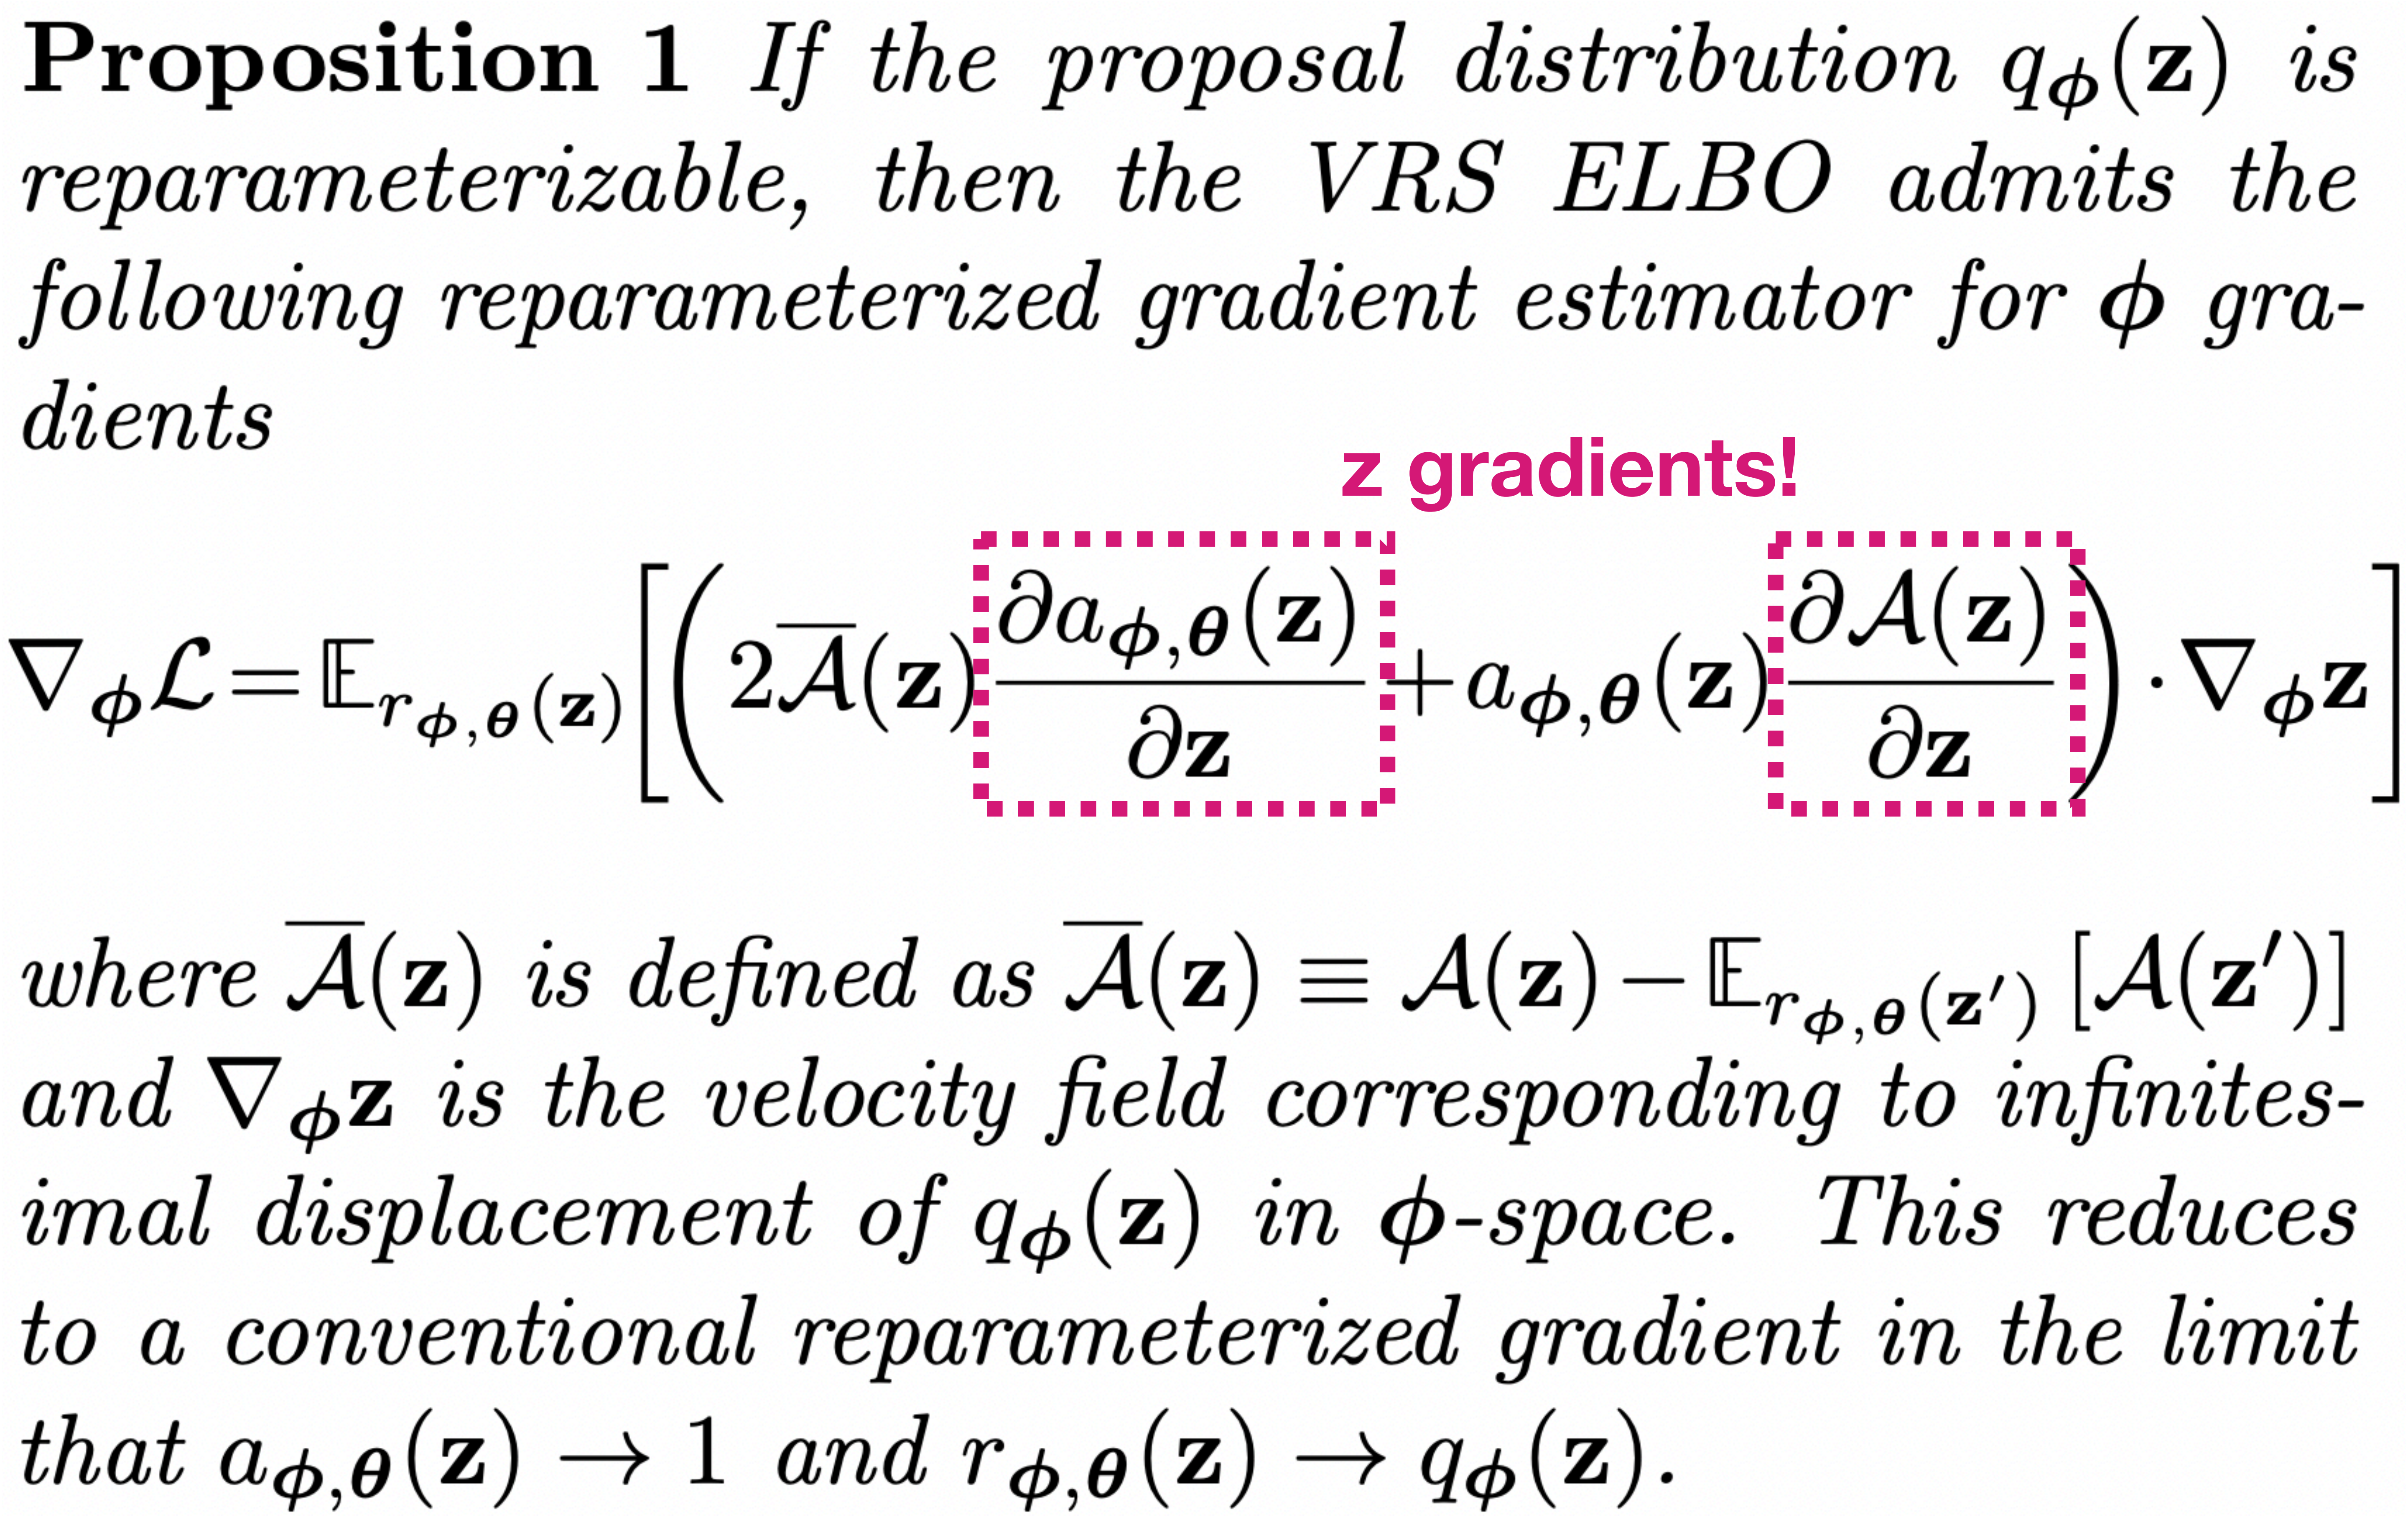

Proposition 1 *If the proposal distribution $q_{\boldsymbol{\phi}}(\mathbf{z})$ is reparameterizable, then the VRS ELBO admits the following reparameterized gradient estimator for $\boldsymbol{\phi}$ gradients*

z gradients!

$$\nabla_{\boldsymbol{\phi}}\mathcal{L}=\mathbb{E}_{r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})}\left[\left(2\overline{\mathcal{A}}(\mathbf{z})\boxed{\frac{\partial a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})}{\partial \mathbf{z}}}+a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})\boxed{\frac{\partial \mathcal{A}(\mathbf{z})}{\partial \mathbf{z}}}\right)\cdot\nabla_{\boldsymbol{\phi}}\mathbf{z}\right]$$

← no score function!

where $\overline{\mathcal{A}}(\mathbf{z})$ is defined as $\overline{\mathcal{A}}(\mathbf{z}) \equiv \mathcal{A}(\mathbf{z}) - \mathbb{E}_{r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}')}\left[\mathcal{A}(\mathbf{z}')\right]$ and $\nabla_{\boldsymbol{\phi}}\mathbf{z}$ is the velocity field corresponding to infinitesimal displacement of $q_{\boldsymbol{\phi}}(\mathbf{z})$ in $\boldsymbol{\phi}$-space. This reduces to a conventional reparameterized gradient in the limit that $a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \rightarrow 1$ and $r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \rightarrow q_{\boldsymbol{\phi}}(\mathbf{z})$.

Logistic Regression Experiment

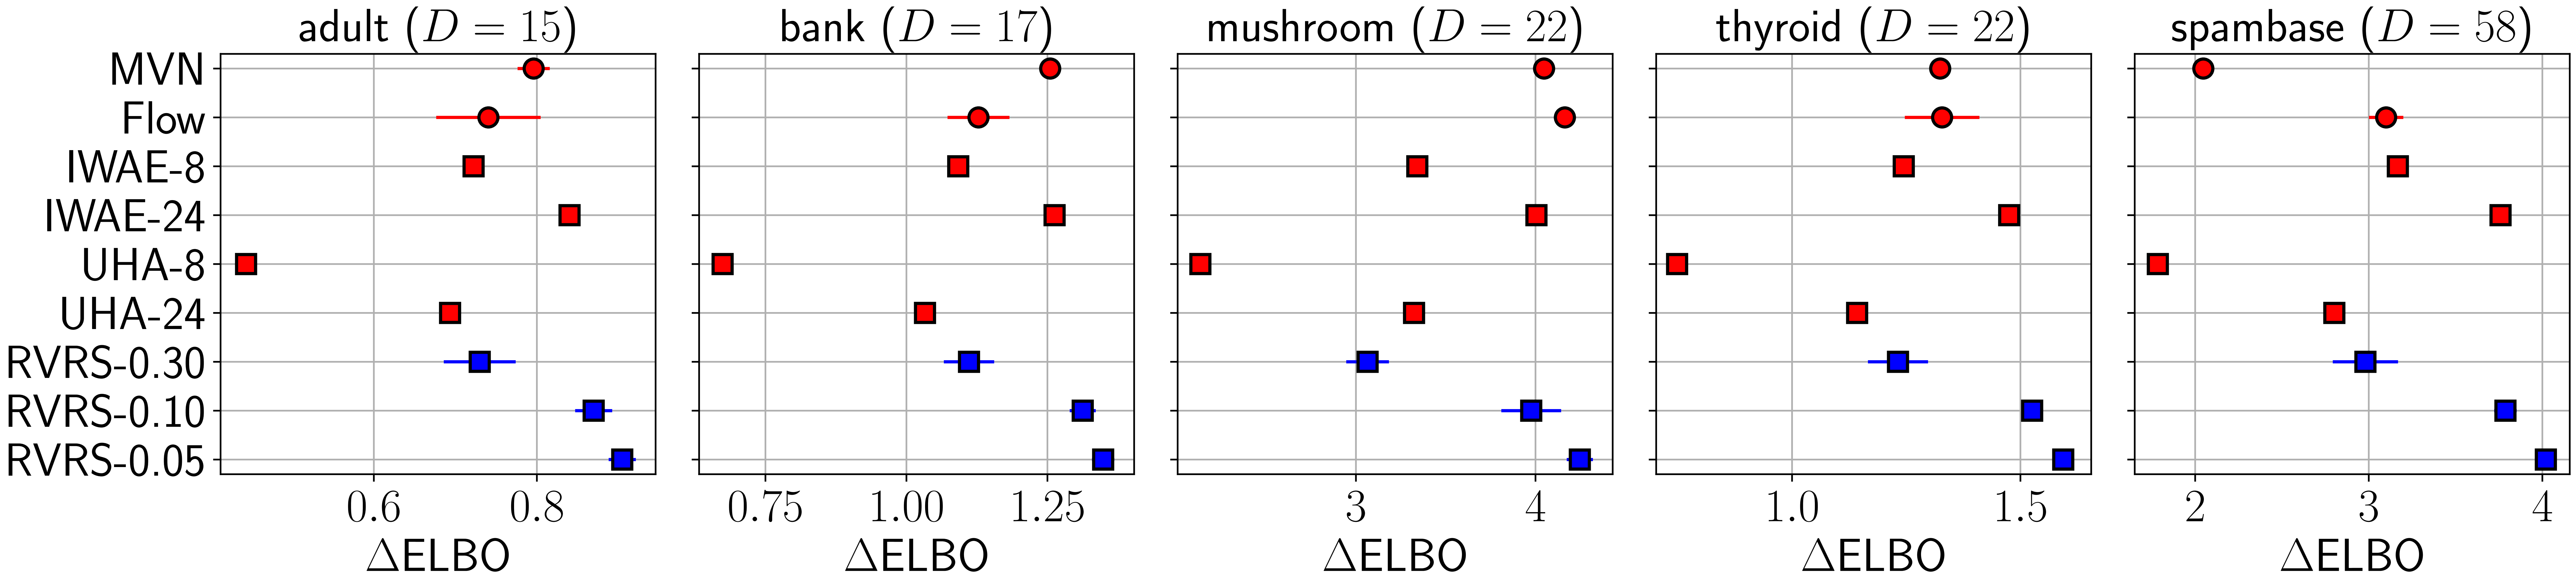

Summary & Outlook

- RVRS works surprisingly well given its relative simplicity
- Offers an attractive trade-off between computational cost and inference fidelity
- Can also be applied to local latent variable models: see paper
- The design space for further hybrid variational inference algorithms remains open!

BACKUP
SLIDES

Comparing VRS and RVRS gradient variance for a mean-field proposal distribution

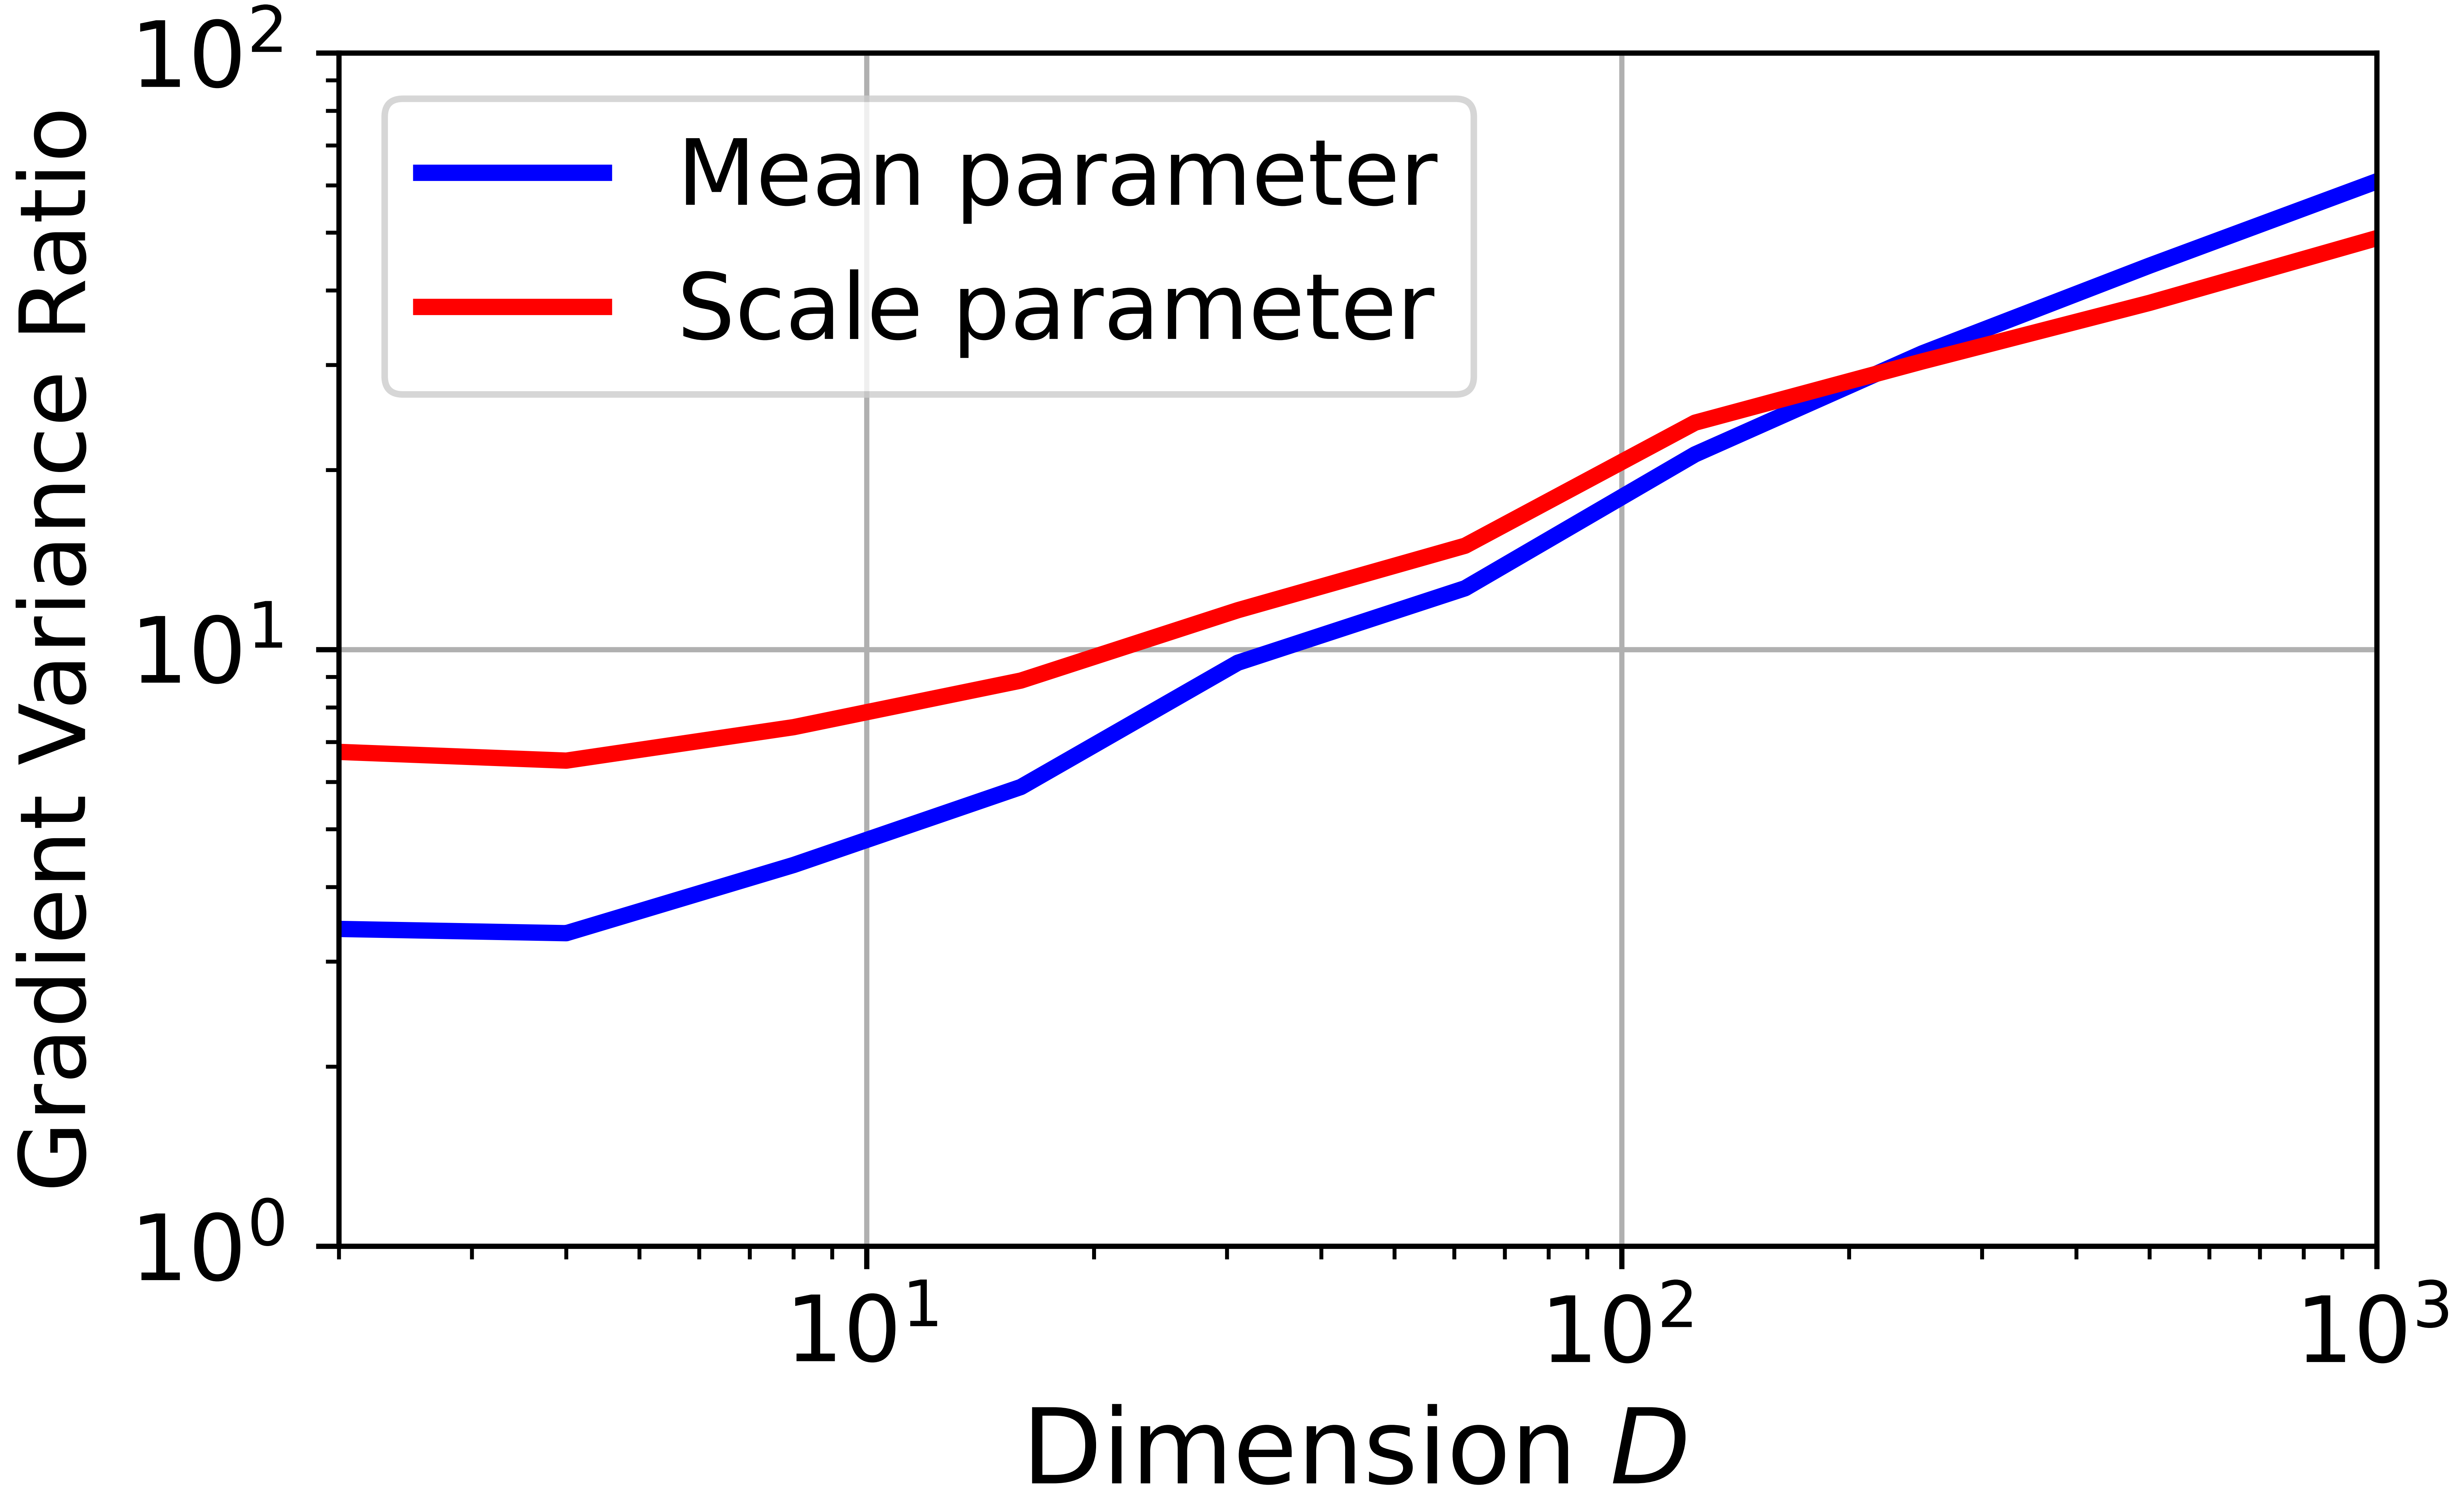

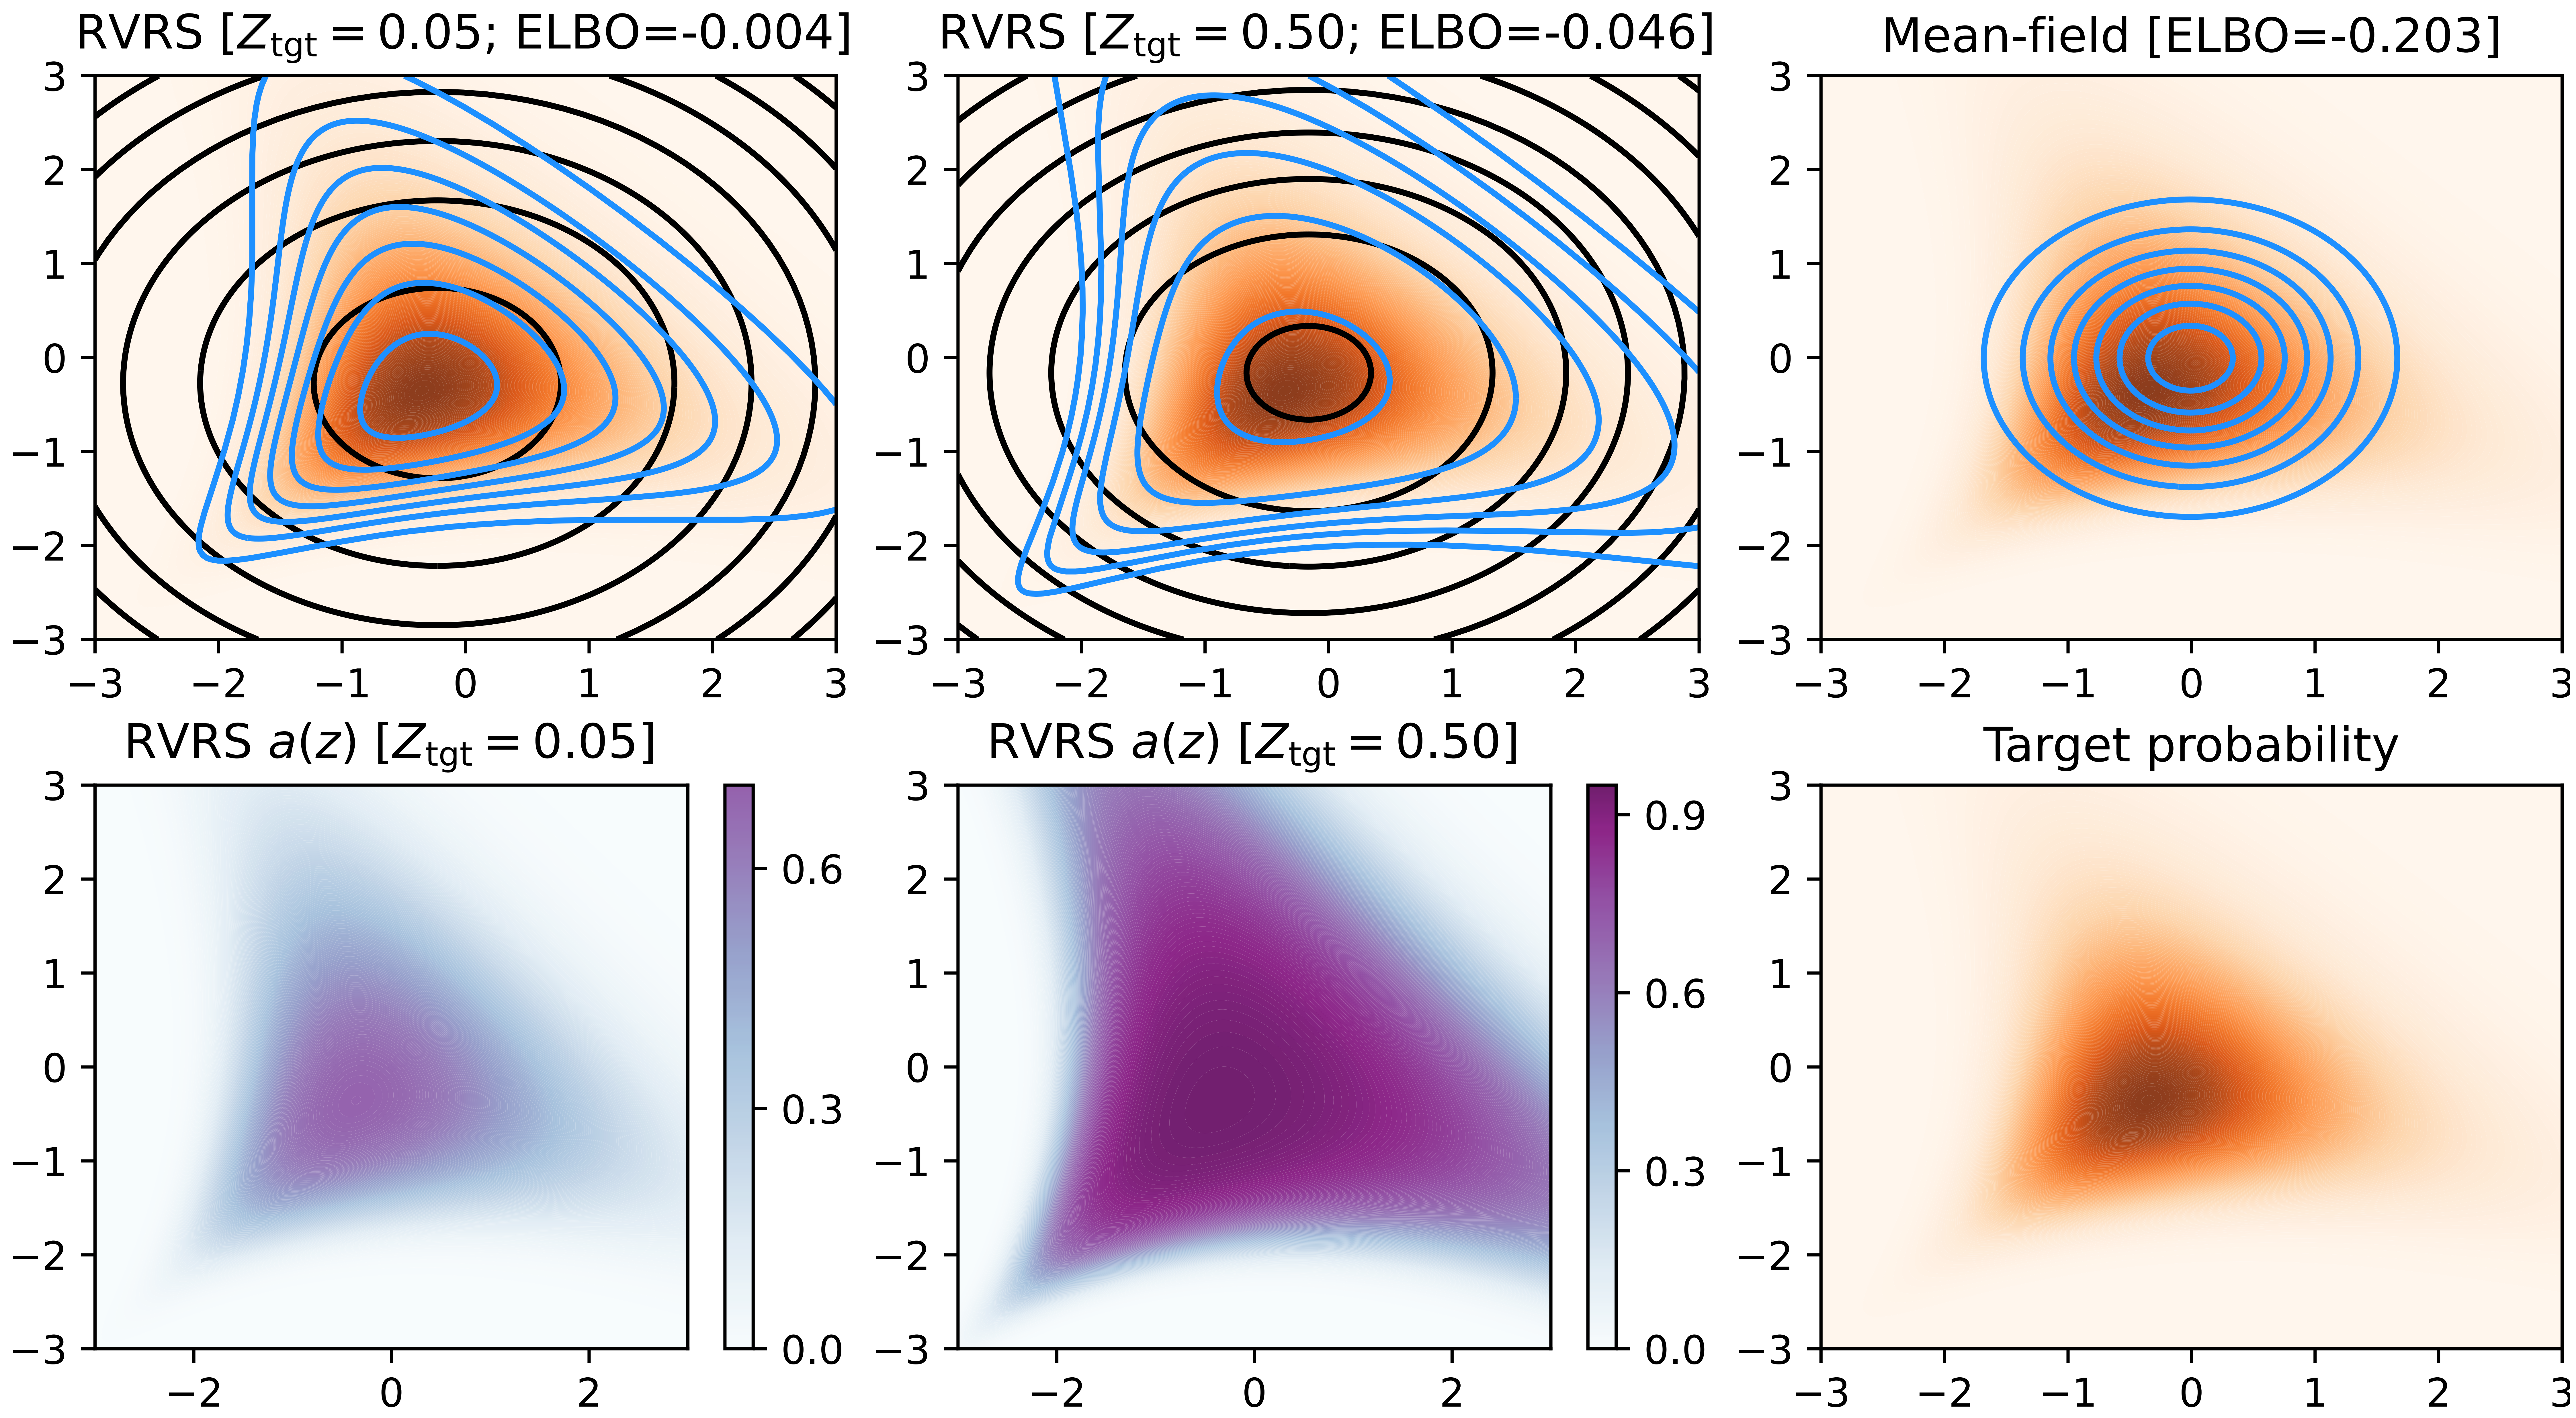
RVRS [$Z_{\text{tgt}} = 0.05$; ELBO=-0.004]
RVRS [$Z_{\text{tgt}} = 0.50$; ELBO=-0.046]
Mean-field [ELBO=-0.203]
RVRS $a(z)$ [$Z_{\text{tgt}} = 0.05$]
RVRS $a(z)$ [$Z_{\text{tgt}} = 0.50$]
Target probability
3
2
1
0
−1
−2
−3
0.0
0.3
0.6
0.9

Variational Rejection Sampling

Sampling is easy

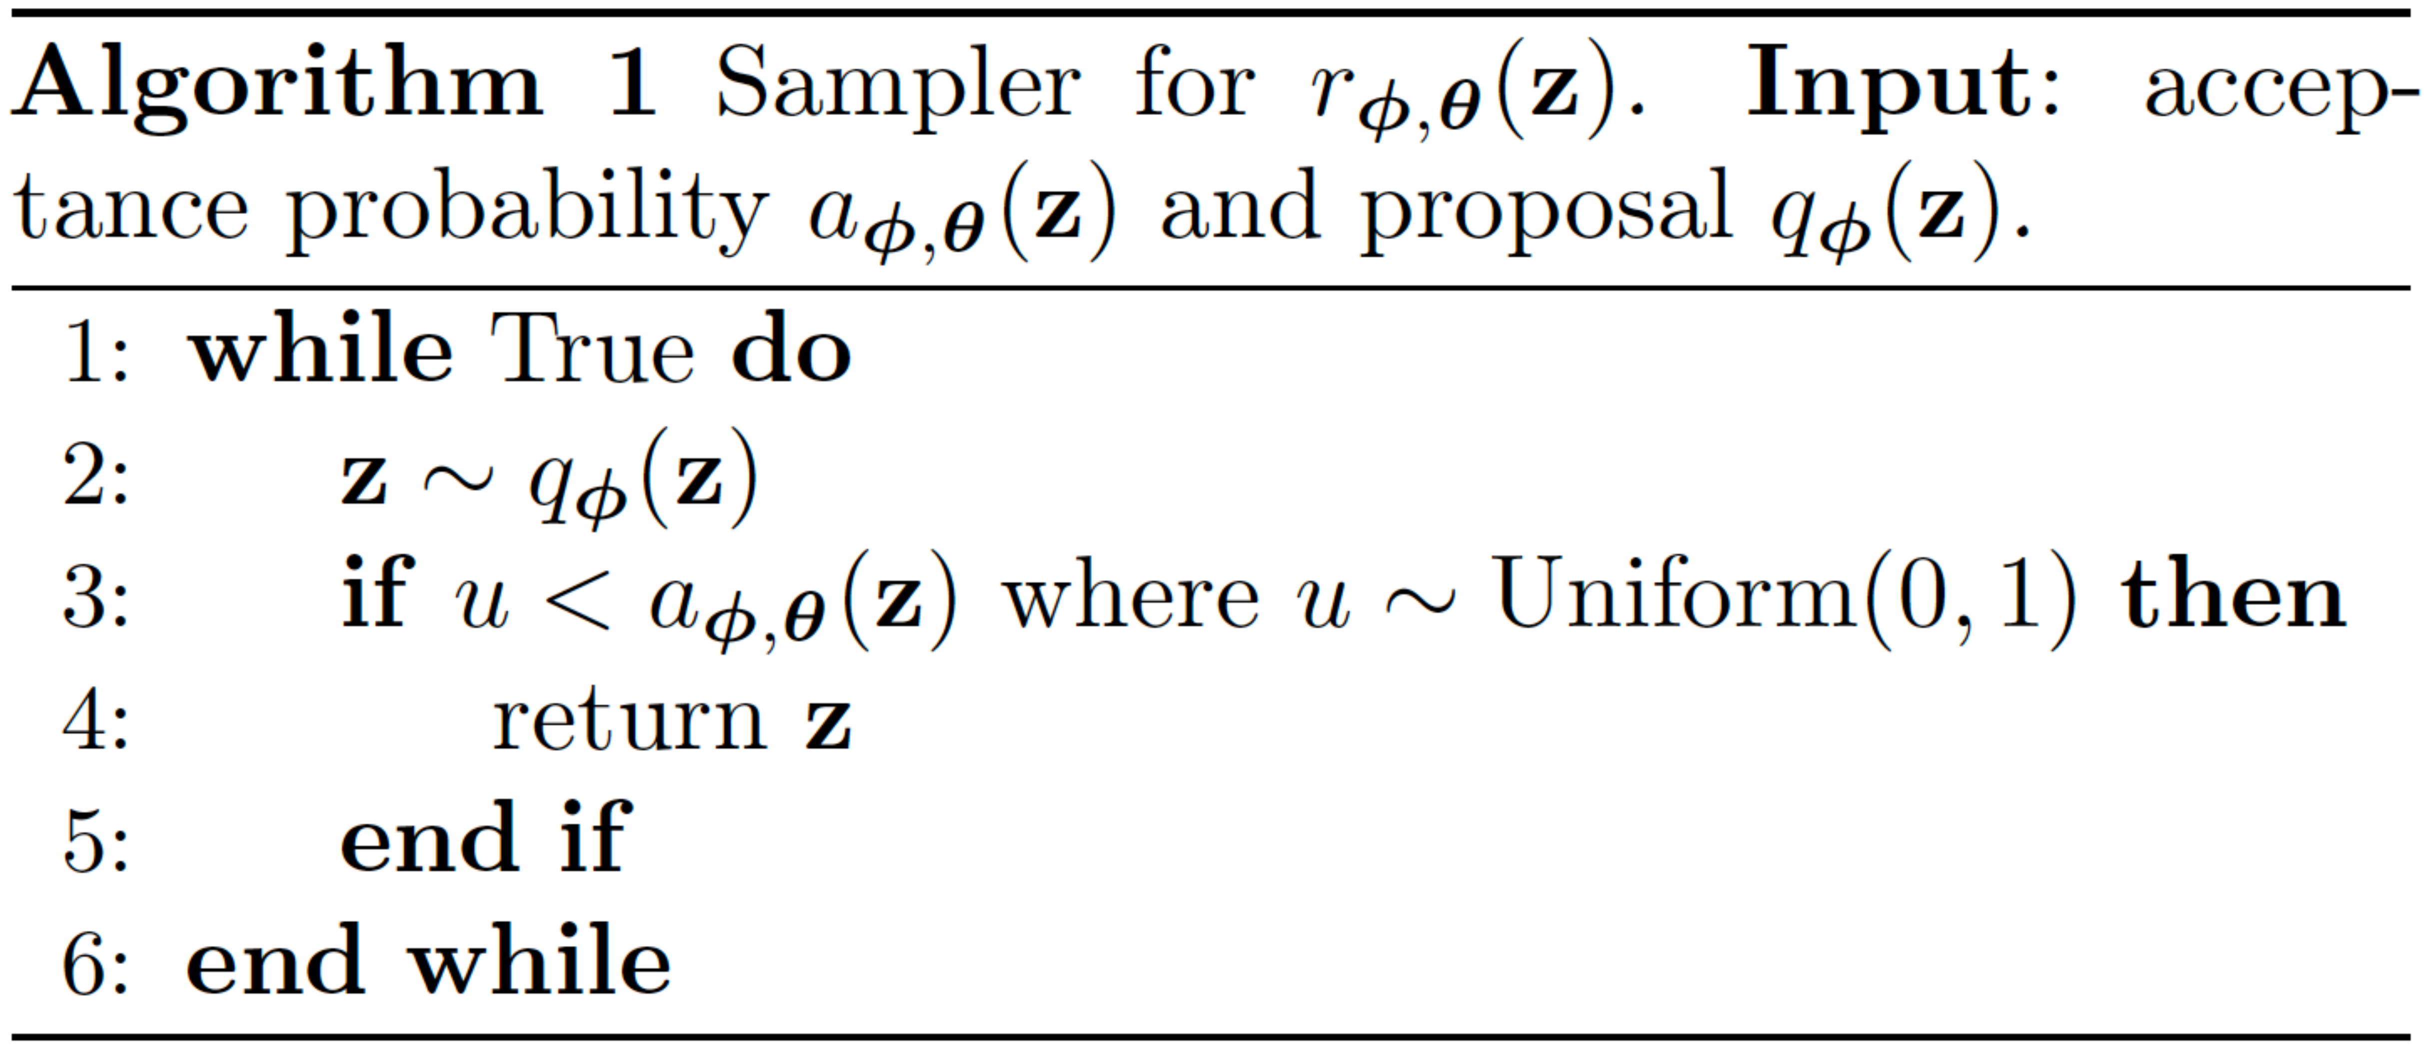

Algorithm 1 Sampler for $r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})$. **Input:** acceptance probability $a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})$ and proposal $q_{\boldsymbol{\phi}}(\mathbf{z})$.

```
while True do
    $\mathbf{z} \sim q_{\boldsymbol{\phi}}(\mathbf{z})$
    if $u < a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z})$ where $u \sim \mathrm{Uniform}(0,1)$ then
        return $\mathbf{z}$
    end if
end while
```

The frequency of rejected samples is controlled by T:

- As $T \to \infty$ we have $a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \to 1$ and $r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \to q_{\boldsymbol{\phi}}(\mathbf{z})$
- As $T \to -\infty$ we have $a_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \to 0$ and $r_{\boldsymbol{\phi},\boldsymbol{\theta}}(\mathbf{z}) \to p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$