

---

# Doubly Reparameterized Gradient Estimators for Monte Carlo Objectives

---

**George Tucker**  
Google Brain  
gjt@google.com

**Dieterich Lawson**  
New York University  
jdl404@nyu.edu

**Shixiang Gu**  
Google Brain  
shanegu@google.com

**Chris J. Maddison**  
University of Oxford, DeepMind  
cmaddis@stats.ox.ac.uk

## 1 Introduction

Following the influential work by (Kingma & Welling, 2013; Rezende et al., 2014), deep generative models with latent variables have been widely used to model data such as natural images (Rezende & Mohamed, 2015; Kingma et al., 2016; Chen et al., 2016; Gulrajani et al., 2016), speech and music time-series (Chung et al., 2015; Fraccaro et al., 2016; Krishnan et al., 2015), and video (Babaeizadeh et al., 2017; Ha & Schmidhuber, 2018; Denton & Fergus, 2018). Generally, marginalizing the latent variables is intractable, so instead of directly maximizing the marginal likelihood, a common approach is to maximize a tractable lower bound on the likelihood such as the variational evidence lower bound (ELBO) (Jordan et al., 1999; Blei et al., 2017).

Burda et al. (2015) introduced a multi-sample bound, IWAE, that is at least as tight as the ELBO and becomes increasingly tight as the number of samples increases. Counterintuitively, although the bound is tighter, Rainforth et al. (2018) theoretically and empirically showed that the standard inference network gradient estimator for the IWAE bound performs poorly as the number of samples increases due to a diminishing signal-to-noise ratio (SNR).

Roeder et al. (2017) proposed a lower-variance estimator of the gradient of the IWAE bound. We show that their estimator is biased, but that it is possible to construct an unbiased estimator with a second application of the reparameterization trick which we call the IWAE *doubly reparameterized gradient* (DReG) estimator. Our estimator is an unbiased, computationally efficient drop-in replacement, and does not suffer as the number of samples increases, resolving the counterintuitive behavior from previous work (Rainforth et al., 2018). Furthermore, our insight is applicable to alternative multi-sample training techniques for latent variable models: reweighted wake-sleep (RWS) (Bornschein & Bengio, 2014) and jackknife variational inference (JVI) (Nowozin, 2018).

In this work, we derive DReG estimators for IWAE, RWS, and JVI and demonstrate improved scaling with the number of samples on a simple example. Then, we evaluate DReG estimators on MNIST generative modeling, Omniglot generative modeling, and MNIST structured prediction tasks. In all cases, we demonstrate substantial unbiased variance reduction, which translates to improved performance over the original estimators.

## 2 Background

Our goal is to learn a latent variable generative model  $p_\theta(x, z) = p_\theta(z)p_\theta(x|z)$  where  $x$  are observed data and  $z$  are continuous latent variables. We maximize a variational lower bound on  $\log p_\theta(x)$  such as the ELBO

$$\log p_\theta(x) = \log \mathbb{E}_{p_\theta(z)}[p_\theta(x|z)] \geq \mathbb{E}_{q(z|x)} \left[ \log \frac{p_\theta(x, z)}{q(z|x)} \right], \quad (1)$$

---

Implementation of DReG estimators and code to reproduce experiments: [sites.google.com/view/dregs](https://sites.google.com/view/dregs).

where  $q(z|x)$  is a variational distribution. Following the influential work by (Kingma & Welling, 2013; Rezende et al., 2014), we consider the amortized inference setting with an inference network  $q_\phi(z|x)$ . Burda et al. (2015) introduced the importance weighted autoencoder (IWAE) bound

$$\mathbb{E}_{z_{1:K}} \left[ \log \left( \frac{1}{K} \sum_{i=1}^K \frac{p_\theta(x, z_i)}{q_\phi(z_i|x)} \right) \right] \leq \log p_\theta(x), \quad (2)$$

with  $z_{1:K} \sim \prod_i q_\phi(z_i|x)$ . The IWAE bound reduces to the ELBO when  $K = 1$ , is non-decreasing as  $K$  increases, and converges to  $\log p_\theta(x)$  as  $K \rightarrow \infty$  under mild conditions (Burda et al., 2015). When  $q_\phi$  is reparameterizable<sup>1</sup>, the standard gradient estimator of the IWAE bound is

$$\nabla_{\theta, \phi} \mathbb{E}_{z_{1:K}} \left[ \log \left( \frac{1}{K} \sum_{i=1}^K w_i \right) \right] = \nabla_{\theta, \phi} \mathbb{E}_{\epsilon_{1:K}} \left[ \log \left( \frac{1}{K} \sum_{i=1}^K w_i \right) \right] = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{w_i}{\sum_j w_j} \nabla_{\theta, \phi} \log w_i \right]$$

where  $w_i = p_\theta(x, z_i)/q_\phi(z_i|x)$ . A single sample estimator of this expectation is typically used as the gradient estimator.

As  $K$  increases, the bound becomes increasingly tight, however, Rainforth et al. (2018) show that the signal-to-noise ratio (SNR) of the inference network gradient estimator goes to 0. This does not happen for the model parameters ( $\theta$ ). Following up on this work, Le et al. (2018) demonstrate that this deteriorates the performance of learned models on practical problems. This motivates the search for lower variance inference network gradient estimators.

To derive improved gradient estimators for  $\phi$ , it is informative to expand the total derivative<sup>2</sup> of the IWAE bound with respect to  $\phi$

$$\mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{w_i}{\sum_{j=1}^K w_j} \left( -\frac{\partial}{\partial \phi} \log q_\phi(z_i|x) + \frac{\partial \log w_i}{\partial z_i} \frac{dz_i}{d\phi} \right) \right]. \quad (3)$$

Previously, Roeder et al. (2017) found that the first term within the parentheses of Eq. 3 can contribute significant variance to the gradient estimator. When  $K = 1$ , this term analytically vanishes in expectation, so when  $K > 1$  they suggested dropping it. Below, we abbreviate this estimator as STL. As we show in Section 4.1, the STL estimator introduces bias when  $K > 1$ .

### 3 Doubly Reparameterized Gradient Estimators (DReGs)

Our insight is that we can estimate the first term within the parentheses of Eq. 3 efficiently with a second application of the reparameterization trick. To see this, first note that it suffices to focus on one of the  $K$  terms in the sum. Because the derivative is a partial derivative  $\frac{\partial}{\partial \phi}$ , it treats  $z_i = z(\epsilon_i, \phi)$  as a constant, so we can freely change the random variable that the expectation is over to  $z_{1:K}$ . Now,

$$\mathbb{E}_{z_{1:K}} \left[ \frac{w_i}{\sum_j w_j} \frac{\partial}{\partial \phi} \log q_\phi(z_i|x) \right] = \mathbb{E}_{z_{-i}} \mathbb{E}_{z_i} \left[ \frac{w_i}{\sum_j w_j} \frac{\partial}{\partial \phi} \log q_\phi(z_i|x) \right], \quad (4)$$

where  $z_{-i} = z_{1:i-1, i+1:K}$  is the set of  $z_{1:K}$  without  $z_i$ . The inner expectation resembles a REINFORCE gradient term (Williams, 1992), where we interpret  $\frac{w_i}{\sum_j w_j}$  as the “reward”. Now, we can use the following well-known equivalence between the REINFORCE gradient and the reparameterization trick gradient (See Appendix B for a derivation)

$$\mathbb{E}_{q_\phi(z|x)} \left[ f(z) \frac{\partial}{\partial \phi} \log q_\phi(z|x) \right] = \mathbb{E}_\epsilon \left[ \frac{\partial f(z)}{\partial z} \frac{\partial z(\epsilon, \phi)}{\partial \phi} \right]. \quad (5)$$

<sup>1</sup>Meaning that we can express  $z_i$  as  $z(\epsilon_i, \phi)$ , where  $z$  is a deterministic, differentiable function and  $p(\epsilon_i)$  does not depend on  $\theta$  or  $\phi$ . This allows gradients to be estimated using the reparameterization trick (Kingma & Welling, 2013; Rezende et al., 2014).

<sup>2</sup> $\log w_i$  depends on  $\phi$  in two ways:  $\phi$  and  $z_i = z(\epsilon_i, \phi)$ . The total derivative accounts for both sources of dependence and the partial derivative  $\frac{\partial}{\partial \phi}$  considers  $z_i$  as a constant.

This holds even when  $f$  depends on  $\phi$ . Typically, the reparameterization gradient estimator has lower variance than the REINFORCE gradient estimator, which suggests using it will improve performance. Applying the identity from Eq. 5 to the right hand side of Eq. 4 gives

$$\begin{aligned} \mathbb{E}_{z_i} \left[ \frac{w_i}{\sum_j w_j} \frac{\partial}{\partial \phi} \log q_\phi(z_i|x) \right] &= \mathbb{E}_{\epsilon_i} \left[ \frac{\partial}{\partial z_i} \left( \frac{w_i}{\sum_j w_j} \right) \frac{\partial z_i}{\partial \phi} \right] \\ &= \mathbb{E}_{\epsilon_i} \left[ \left( \frac{1}{\sum_j w_j} - \frac{w_i}{(\sum_j w_j)^2} \right) \frac{\partial w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi} \right] = \mathbb{E}_{\epsilon_i} \left[ \left( \frac{w_i}{\sum_j w_j} - \frac{w_i^2}{(\sum_j w_j)^2} \right) \frac{\partial \log w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi} \right]. \end{aligned} \quad (6)$$

This last expression can be efficiently estimated with a single Monte Carlo sample. When  $K = 1$ , this term vanishes exactly and we recover the estimator proposed in (Roeder et al., 2017) for the ELBO.

Substituting Eq. 6 into Eq. 3, we obtain a simplification due to cancellation of terms

$$\nabla_\phi \mathbb{E}_{z_{1:K}} \left[ \log \left( \frac{1}{K} \sum_{i=1}^K w_i \right) \right] = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \left( \frac{w_i}{\sum_j w_j} \right)^2 \frac{\partial \log w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi} \right]. \quad (7)$$

We call the algorithm that uses the single sample Monte Carlo estimator of this expression for the inference network gradient the IWAE doubly reparameterized gradient estimator (IWAE-DReG). In Appendix C, we show that in contrast to the standard IWAE gradient estimator, the SNR of the IWAE-DReG estimator exhibits the same scaling behaviour of  $\mathcal{O}(\sqrt{K})$  for both the generative and inference network gradients (i.e., improving in  $K$ ). Finally, we derive doubly reparameterized gradient estimators for RWS and JVI in Appendix D.

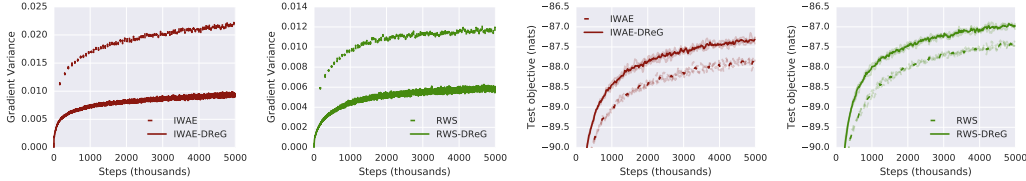


Figure 1: MNIST generative modeling trained according to IWAE and RWS. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent.  $K = 64$ .

## 4 Experiments

### 4.1 Toy Gaussian

We reimplemented the Gaussian example from (Rainforth et al., 2018). Consider the generative model with  $z \sim N(\theta, I)$  and  $x|z \sim N(z, I)$  and inference network  $q_\phi(z|x) \sim N(Ax + b, \frac{2}{3}I)$ , where  $\phi = \{A, b\}$ . As in (Rainforth et al., 2018), we sample a set of parameters for the model and inference network close to the optimal parameters (perturbed by zero-mean Gaussian noise with standard deviation 0.01), then estimate the gradient of the inference network parameters for increasing number of samples ( $K$ ). We plot the SNR, the squared bias, and the variance of the gradient estimators in Appendix Fig. 2. IWAE-DReG is unbiased, its SNR increases with  $K$ , and it has the lowest variance. Furthermore, we can see the bias present in the STL estimator.

### 4.2 Generative Modeling and Structured Prediction

Training generative models of the binarized MNIST digits dataset is a standard benchmark task for latent variable models. We used the single latent layer architecture from (Burda et al., 2015) with additional details in Appendix A. We trained models with the IWAE gradient, the RWS wake update, and with the JVI estimator. In all three cases, the doubly reparameterized gradient estimator reduced variance and as a result substantially improved performance (Fig. 1 and Appendix Figs. 3,4).

Finally, we performed generative modeling with the Omniglot dataset and structured prediction of the bottom half of an MNIST digit using the top half as the context. Again, we found that the doubly reparameterized gradient estimator reduced variance and as a result improved test performance (Appendix Figs. 5,6,7, and 8).

## Acknowledgments

We thank Ben Poole and Diederik P. Kingma for helpful discussion and comments on drafts of this paper. We thank Sergey Levine and Jascha Sohl-Dickstein for insightful discussion.

## References

- Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H Campbell, and Sergey Levine. Stochastic variational video prediction. *International Conference on Learning Representations*, 2017.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 2017.
- Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. *International Conference on Learning Representations*, 2014.
- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *International Conference on Learning Representations*, 2015.
- George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. *International Conference on Learning Representations*, 2016.
- Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pp. 2980–2988, 2015.
- Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. *International Conference on Machine Learning*, 2018.
- Marco Fraccaro, Søren Kaae Sønderby, Ulrich Paquet, and Ole Winther. Sequential neural models with stochastic layers. In *Advances in neural information processing systems*, pp. 2199–2207, 2016.
- Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville. Pixelvae: A latent variable model for natural images. *International Conference on Learning Representations*, 2016.
- David Ha and Jürgen Schmidhuber. World models. *Advances in neural information processing systems*, 2018.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2013.
- Diederik P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pp. 4743–4751, 2016.
- Rahul G Krishnan, Uri Shalit, and David Sontag. Deep kalman filters. *arXiv preprint arXiv:1511.05121*, 2015.
- Tuan Anh Le, Adam R Kosior, N Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep. *arXiv preprint arXiv:1805.10469*, 2018.
- Sebastian Nowozin. Debiasing evidence approximations: On importance-weighted autoencoders and jackknife variational inference. *International Conference on Learning Representations*, 2018.

- Tom Rainforth, Adam R Kosiorek, Tuan Anh Le, Chris J Maddison, Maximilian Igl, Frank Wood, and Yee Whye Teh. Tighter variational bounds are not necessarily better. *International Conference on Machine Learning*, 2018.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538, 2015.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286, 2014.
- Geoffrey Roeder, Yuhuai Wu, and David Duvenaud. Sticking the landing: An asymptotically zero-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 2017.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

## Appendix

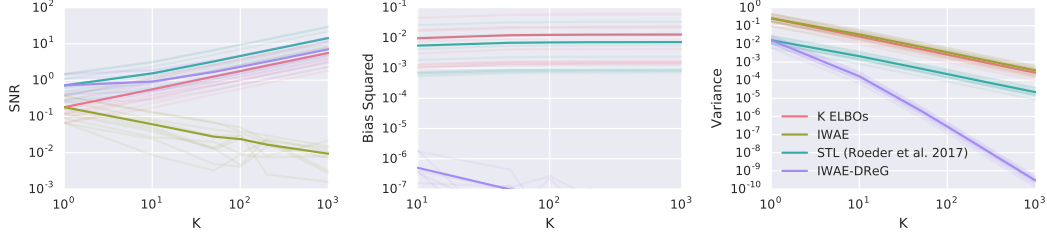


Figure 2: Signal-to-noise ratios (SNR), bias squared, and variance of gradient estimators with increasing  $K$  over 10 random trials with 1000 measurement samples per trial (mean in bold). All dimensions of  $\phi$  behaved qualitatively similarly, so for clarity, we show curves for a single randomly chosen dimension of  $\phi$ . The observed “bias” for IWAE-DReG is not statistically significant under a paired t-test (as expected because IWAE-DReG is unbiased). IWAE-DReG is unbiased, its SNR increases with  $K$ , and it has the lowest variance of the estimators considered here.

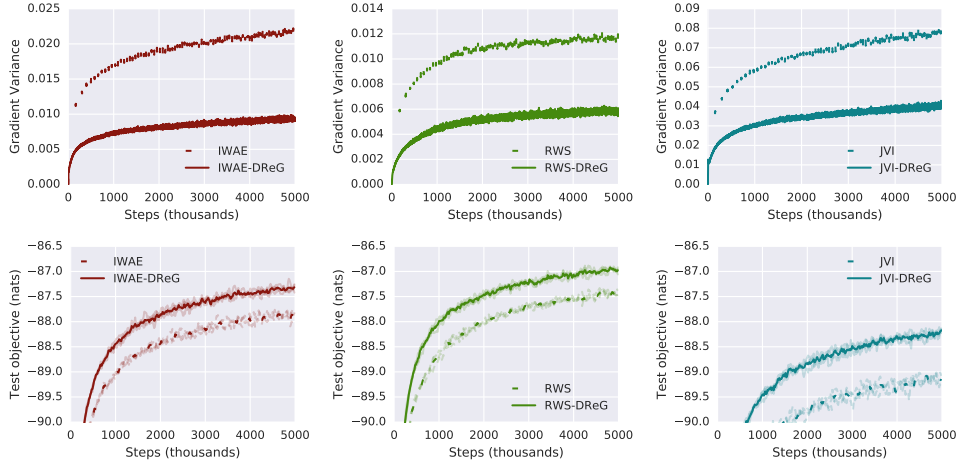


Figure 3: MNIST generative modeling trained according to IWAE (left), RWS (middle), and JVI (right). The top row compares the variance of the original gradient estimator (dashed) with the variance of the doubly reparameterized gradient estimator (solid). The bottom row compares test performance. The left and middle plots show the IWAE (stochastic) lower bound on the test set. The right plot shows the JVI estimator (which is not a bound) on the test set. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent. All methods used  $K = 64$ .

## A Experiment Details

### A.1 Generative modeling

The generative model used 50 Gaussian latent variables with an isotropic prior and passed  $z$  through two deterministic layers of 200 tanh units to parameterize factorized Bernoulli outputs. The inference network passed  $x$  through two deterministic layers of 200 tanh units to parameterize a factorized Gaussian distribution over  $z$ . Because our interest was in improved gradient estimators and optimization performance, we used the dynamically binarized MNIST dataset, which minimally suffers from overfitting. We used the standard split of MNIST into train, validation, and test sets.

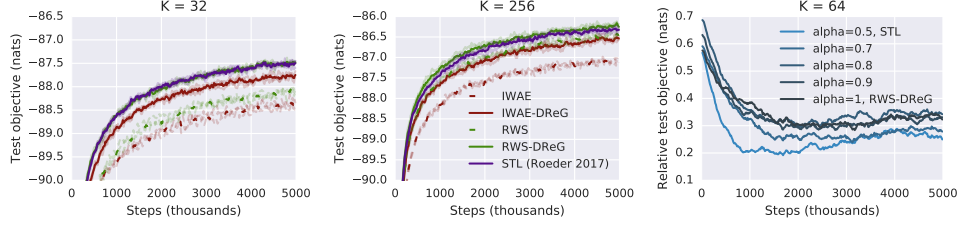


Figure 4: Log-likelihood lower bounds for generative modeling on MNIST. The left and middle plots compare performance with different number of samples  $K = 32, 256$ . The bold lines are the average over three trials, and individual trials are displayed as semi-transparent). The right plot compares performance as the convex combination between IWAE-DReG and RWS-DReG is varied (Eq. 10). To highlight differences, we plot the difference between the test IWAE bound and the test IWAE bound IWAE-DReG achieved at that step.

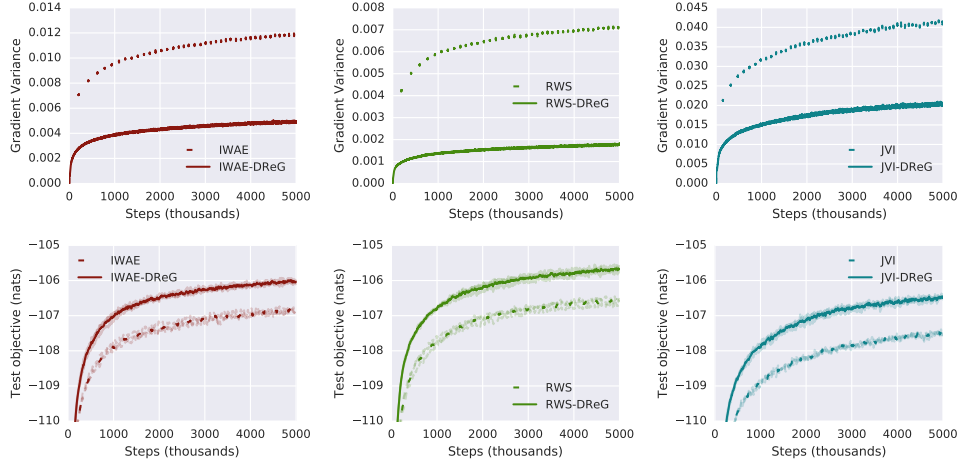


Figure 5: Omniglot generative modeling trained according to IWAE (left), RWS (middle), and JVI (right). The top row compares the variance of the original gradient estimator (dashed) with the variance of the doubly reparameterized gradient estimator (solid). The bottom row compares test performance. The left and middle plots show the IWAE (stochastic) lower bound on the test set. The right plot shows the JVI estimator (which is not a bound) on the test set. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent. All methods used  $K = 64$ .

## A.2 Structured prediction on MNIST

Structured prediction is another common benchmark task for latent variable models. In this task, our goal is to model a complex observation  $x$  given a context  $c$  (i.e., model the conditional distribution  $p(x|c)$ ). We can use a conditional latent variable model  $p_\theta(x, z|c) = p_\theta(x|z, c)p_\theta(z|c)$ , however, as before, computing the marginal likelihood is generally intractable. It is straightforward to adapt the bounds and techniques from the previous section to this problem.

To evaluate our method in this context, we use the standard task of modeling the bottom half of a binarized MNIST digit from the top half. We use a similar architecture, but now learn a conditional prior distribution  $p_\theta(z|c)$  where  $c$  is the top half of the MNIST digit. The conditional prior feeds  $c$  to two deterministic layers of 200 tanh units to parameterize a factorized Gaussian distribution over  $z$ . To model the conditional distribution  $p_\theta(x|c, z)$ , we concatenate  $z$  with  $c$  and feed it to two deterministic layers of 200 tanh units to parameterize factorized Bernoulli outputs.

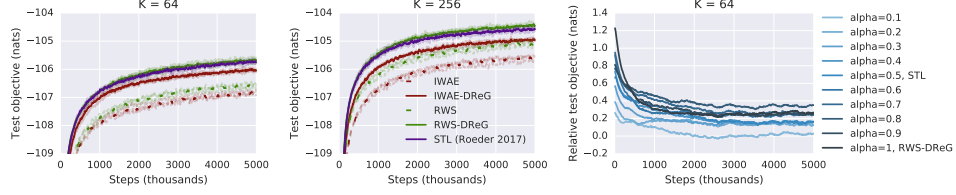


Figure 6: Log-likelihood lower bounds for structured prediction on Omniglot. The left plot uses  $K = 64$  samples and the right plot uses  $K = 256$  samples. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent). The right plot compares performance as the convex combination between IWAE-DReG and RWS-DReG is varied. To highlight differences, we plot the difference between the test IWAE bound and the test IWAE bound IWAE-DReG achieved at that step.

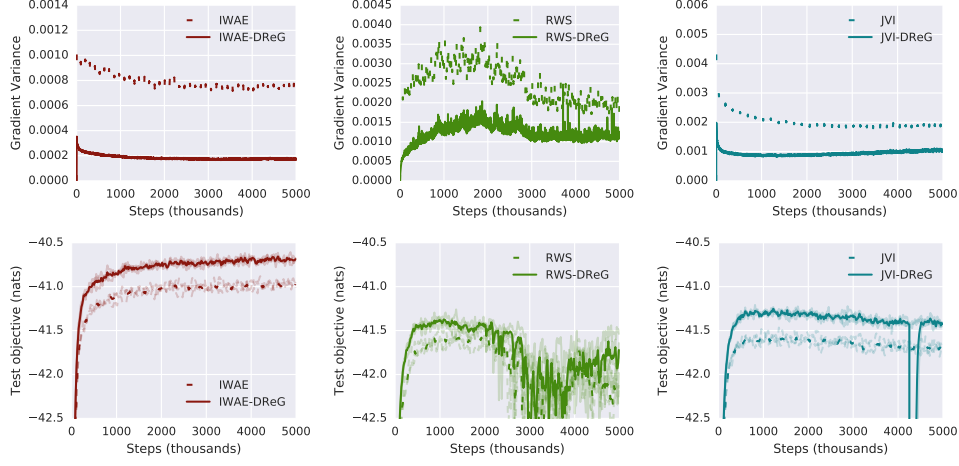


Figure 7: Structured prediction on MNIST according to IWAE (left), RWS (middle), and JVI (right). The top row compares the variance of the original gradient estimator (dashed) with the variance of the doubly reparameterized gradient estimator (solid). The bottom row compares test performance. The left and middle plots show the IWAE (stochastic) lower bound on the test set. The right plot shows the JVI estimator (which is not a bound) on the test set. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent. All methods used  $K = 64$ .

## B Equivalence between REINFORCE gradient and reparameterization trick gradient

Given a function  $f(z, \phi)$ , we have

$$\mathbb{E}_{q_\phi(z)} \left[ f(z, \phi) \frac{\partial \log q_\phi(z)}{\partial \phi} \right] = \mathbb{E}_\epsilon \left[ \frac{\partial f(z, \phi)}{\partial z} \frac{\partial z(\epsilon, \phi)}{\partial \phi} \right],$$

for a reparameterizable distribution  $q_\phi(z)$ . To see this, note that

$$\begin{aligned} \frac{d}{d\phi} \int_z q_\phi(z) f(z, \phi) dz &= \int_z \frac{\partial}{\partial \phi} q_\phi(z) f(z, \phi) dz = \int_z f(z, \phi) \frac{\partial}{\partial \phi} q_\phi(z) + q_\phi(z) \frac{\partial}{\partial \phi} f(z, \phi) dz \\ &= \int_z f(z, \phi) q_\phi(z) \frac{\partial \log q_\phi(z)}{\partial \phi} dz + \mathbb{E}_{q_\phi(z)} \left[ \frac{\partial f(z, \phi)}{\partial \phi} \right] \\ &= \mathbb{E}_{q_\phi(z)} \left[ f(z, \phi) \frac{\partial \log q_\phi(z)}{\partial \phi} \right] + \mathbb{E}_{q_\phi(z)} \left[ \frac{\partial f(z, \phi)}{\partial \phi} \right], \end{aligned}$$

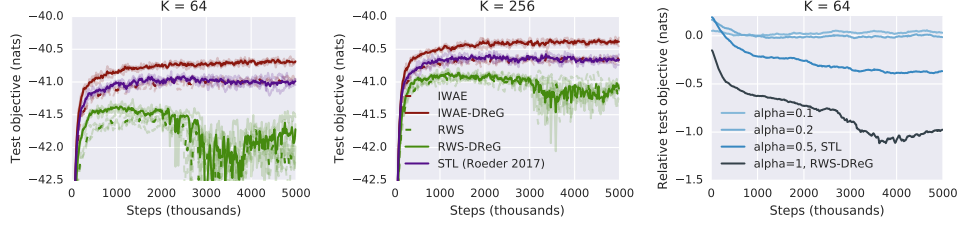


Figure 8: Log-likelihood lower bounds for structured prediction on MNIST. The left plot uses  $K = 64$  samples and the right plot uses  $K = 256$  samples. The bold lines are the average over three trials, and individual trials are displayed as semi-transparent). The right plot compares performance as the convex combination between IWAE-DReG and RWS-DReG is varied (Eq. 10). To highlight differences, we plot the difference between the test IWAE bound and the test IWAE bound IWAE-DReG achieved at that step.

via the REINFORCE gradient. On the other hand,

$$\begin{aligned} \frac{d}{d\phi} \int_z q_\phi(z) f(z, \phi) dz &= \frac{d}{d\phi} \mathbb{E}_{q_\phi(z)} [f(z, \phi)] = \frac{d}{d\phi} \mathbb{E}_\epsilon [f(z(\epsilon, \phi), \phi)] = \mathbb{E}_\epsilon \left[ \frac{d}{d\phi} f(z(\epsilon, \phi), \phi) \right] \\ &= \mathbb{E}_\epsilon \left[ \frac{\partial f(z, \phi)}{\partial z} \frac{\partial z(\epsilon, \phi)}{\partial \phi} \right] + \mathbb{E}_\epsilon \left[ \frac{\partial f(z, \phi)}{\partial \phi} \Big|_{z=z(\epsilon, \phi)} \right] \\ &= \mathbb{E}_\epsilon \left[ \frac{\partial f(z, \phi)}{\partial z} \frac{\partial z(\epsilon, \phi)}{\partial \phi} \right] + \mathbb{E}_{q_\phi(z)} \left[ \frac{\partial f(z, \phi)}{\partial \phi} \right], \end{aligned}$$

via the reparameterization trick. Thus, we conclude that

$$\mathbb{E}_{q_\phi(z)} \left[ f(z, \phi) \frac{\partial \log q_\phi(z)}{\partial \phi} \right] + \mathbb{E}_{q_\phi(z)} \left[ \frac{\partial f(z, \phi)}{\partial \phi} \right] = \mathbb{E}_\epsilon \left[ \frac{\partial f(z, \phi)}{\partial z} \frac{\partial z(\epsilon, \phi)}{\partial \phi} \right] + \mathbb{E}_{q_\phi(z)} \left[ \frac{\partial f(z, \phi)}{\partial \phi} \right],$$

from which the identity follows.

## C Informal asymptotic analysis

At a high level, Rainforth et al. (2018) show that the expected value of the IWAE gradient of the inference network collapses to zero with rate  $1/K$ , while its standard deviation is only shrinking at a rate of  $1/\sqrt{K}$ . This is the essence of the problem that results in the SNR (expectation divided by standard deviation) of the inference network gradients going to zero at a rate  $\mathcal{O}((1/K)/(1/\sqrt{K})) = \mathcal{O}(1/\sqrt{K})$ , worsening with  $K$ . In contrast, Rainforth et al. (2018) show that the generation network gradients scales like  $\mathcal{O}(\sqrt{K})$ , improving with  $K$ .

Because the IWAE-DReG estimator is unbiased, we cannot hope to change the scaling of the expected value in  $K$ , but we can hope to change the scaling of the variance. In particular, in this section, we provide an informal argument, via the delta method, that the standard deviation of IWAE-DReG scales like  $K^{-3/2}$ , which results in an overall scaling of  $\mathcal{O}(\sqrt{K})$  for the inference network gradient's SNR (i.e., increasing with  $K$ ). Thus, the SNR of the IWAE-DReG estimator improves similarly in  $K$  for both inference and generation networks.

We will appeal to the delta method on a two-variable function  $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ . Define the following notation for the partials of  $g$  evaluated at the mean of random variables  $X, Y$ ,

$$g_x(X, Y) = \left. \frac{\partial g(x, y)}{\partial x} \right|_{(x, y) = (\mathbb{E}(X), \mathbb{E}(Y))}$$

The delta method approximation of  $\text{Var}(g(X, Y))$  is given by (Section 5.5 of Casella & Berger),

$$\text{Var}(g(X, Y)) \approx g_x(X, Y)^2 \text{Var}(X) + 2g_x(X, Y)g_y(X, Y)\text{Cov}(X, Y) + g_y(X, Y)^2 \text{Var}(Y)$$

Now, assume without loss of generality that  $\phi$  is a single real-valued parameter. Let  $u_i = w_i^2 \frac{\partial \log w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi}$ ,  $X = \sum_{i=1}^K u_i$ , and  $Y = \sum_{i=1}^K w_i$ . Let  $g(X, Y) = X/Y^2$ , then  $g(X, Y)$  is the

IWAE-DReG estimator whose variance we seek to understand. Letting  $Z = \mathbb{E}(w_i)$  and  $U = \mathbb{E}(u_i)$  we get in this case after cancellations,

$$\text{Var}(g(X, Y)) \approx \frac{1}{Z^4} \frac{\text{Var}(X)}{K^4} - \frac{4U}{Z^5} \frac{\text{Cov}(X, Y)}{K^4} + \frac{4U^2}{Z^6} \frac{\text{Var}(Y)}{K^4}$$

Because  $w_i$  are all mutually independent, we get  $\text{Var}(Y) = K\text{Var}(w_i)$ . Similarly for  $\text{Var}(X)$  and  $u_i$ . Because the  $w_i$  and  $u_i$  are identically distributed and independent for  $i \neq j$ , we have  $\text{Cov}(X, Y) = K\text{Cov}(w_i, u_i)$ . All together we can see that  $\text{Var}(g(X, Y))$  scales like  $K^{-3}$ . Thus, the standard deviation scales like  $K^{-3/2}$ .

## D Alternative training algorithms

Now, we review alternative training algorithms for deep generative models and derive their doubly reparameterized versions.

### D.1 Reweighted Wake Sleep (RWS)

Bornschein & Bengio (2014) introduced RWS, an alternative multi-sample update for latent variable models that uses importance sampling. Computing the gradient of the log marginal likelihood

$$\nabla_{\theta} \log p_{\theta}(x) = \frac{\nabla_{\theta} \int_z p_{\theta}(x, z) dz}{p_{\theta}(x)} = \frac{\int_z p_{\theta}(x, z) \nabla_{\theta} \log p_{\theta}(x, z) dz}{p_{\theta}(x)} = \mathbb{E}_{p_{\theta}(z|x)} [\nabla_{\theta} \log p_{\theta}(x, z)],$$

requires samples from  $p_{\theta}(z|x)$ , which is generally intractable. We can approximate the gradient with a self-normalized importance sampling estimator

$$\mathbb{E}_{p_{\theta}(z|x)} [\nabla_{\theta} \log p_{\theta}(x, z)] \approx \mathbb{E}_{z_{1:K}} \left[ \sum_i \frac{w_i}{\sum_j w_j} \nabla_{\theta} \log p_{\theta}(x, z_i) \right],$$

where  $z_{1:K} \sim \prod_i q_{\phi}(z_i|x)$ . Interestingly, this is precisely the same as the IWAE gradient of  $\theta$ , so the RWS update for  $\theta$  can be interpreted as maximizing the IWAE lower bound in terms of  $\theta$ . Instead of optimizing a joint objective for  $p$  and  $q$ , RWS optimizes a separate objective for the inference network. (Bornschein & Bengio, 2014) propose a “wake” update and a “sleep” update for the inference network. Le et al. (2018) provide empirical support for solely using the wake update for the inference network, so we focus on that update.

The wake update approximately minimizes the KL divergence from  $p_{\theta}(z|x)$  to  $q_{\phi}(z|x)$ . The gradient of the KL term is

$$\nabla_{\phi} \mathbb{E}_{p_{\theta}(z|x)} [\log p_{\theta}(z|x) - \log q_{\phi}(z|x)] = -\mathbb{E}_{p_{\theta}(z|x)} \left[ \frac{\partial}{\partial \phi} \log q_{\phi}(z|x) \right].$$

The wake update of the inference network approximates the intractable expectation by self-normalized importance sampling

$$-\mathbb{E}_{p_{\theta}(z|x)} \left[ \frac{\partial}{\partial \phi} \log q_{\phi}(z|x) \right] \approx -\mathbb{E}_{z_{1:K}} \left[ \sum_i \frac{w_i}{\sum_j w_j} \frac{\partial}{\partial \phi} \log q_{\phi}(z_i|x) \right], \quad (8)$$

with  $z_i \sim q_{\phi}(z_i|x)$ . Le et al. (2018) note that this update does not suffer from diminishing SNR as  $K$  increases. However, a downside is that the updates for  $p$  and  $q$  are not gradients of a unified objective, so could potentially lead to instability or divergence.

### Doubly Reparameterized Reweighted Wake update

The wake update gradient for the inference network (Eq. 8) can be reparameterized

$$-\mathbb{E}_{z_{1:K}} \left[ \sum_i \frac{w_i}{\sum_j w_j} \frac{\partial}{\partial \phi} \log q_{\phi}(z_i|x) \right] = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_i \left( \frac{w_i^2}{(\sum_j w_j)^2} - \frac{w_i}{\sum_j w_j} \right) \frac{\partial \log w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi} \right]. \quad (9)$$

We call the algorithm that uses the single sample Monte Carlo estimator of this expression as the wake update for the inference network RWS-DReG.

Interestingly, the inference network gradient estimator from (Roeder et al., 2017) can be seen as the sum of the IWAE gradient estimator and the wake update of the inference network (as the wake update minimizes, we add the negative of Eq. 9). Their positive results motivate further exploration of convex combinations of IWAE-DReG and RWS-DReG

$$\mathbb{E}_{\epsilon_{1:K}} \left[ \sum_i \left( \alpha \frac{w_i}{\sum_j w_j} + (1 - 2\alpha) \frac{w_i^2}{(\sum_j w_j)^2} \right) \frac{\partial \log w_i}{\partial z_i} \frac{\partial z_i}{\partial \phi} \right]. \quad (10)$$

We refer to the algorithm that uses the single sample Monte Carlo estimator of this expression as DReG( $\alpha$ ). When  $\alpha = 1$ , this reduces to RWS-DReG, when  $\alpha = 0$ , this reduces to IWAE-DReG and when  $\alpha = 0.5$ , this reduces STL.

## D.2 Jackknife Variational Inference (JVI)

Alternatively, Nowozin (2018) reinterprets the IWAE lower bound as a biased estimator for the log marginal likelihood. He analyzes the bias and introduces a novel family of estimators, Jackknife Variational Inference (JVI), which trade off reduction in bias for increased variance. This additional flexibility comes at the cost of no longer being a stochastic lower bound on the log marginal likelihood. The first-order JVI has significantly reduced bias compared to IWAE, which empirically results in a better estimate of the log marginal likelihood with fewer samples (Nowozin, 2018). For simplicity, we focus on the first-order JVI estimator

$$K \times \mathbb{E}_{z_1:K} \left[ \log \left( \frac{1}{K} \sum_i w_i \right) \right] - \frac{K-1}{K} \sum_i \mathbb{E}_{z_{-i}} \left[ \log \left( \frac{1}{K-1} \sum_{j \neq i} w_j \right) \right].$$

It is straightforward to apply our approach to higher order JVI estimators.

## Doubly Reparameterized Jackknife Variational Inference (JVI)

The JVI estimator is a linear combination of  $K$  and  $K - 1$  sample IWAE estimators, so we can use the doubly reparameterized gradient estimator (Eq. 7) for each term.

## E Unified Surrogate Objectives for Estimators

In the main text, we assumed that  $\theta$  and  $\phi$  were disjoint, however, it can be helpful to share parameters between  $p$  and  $q$  (e.g., (Fraccaro et al., 2016)). With the IWAE bound, we differentiate a single objective with respect to both the  $p$  and  $q$  parameters. Thus it is straightforward to adapt IWAE and IWAE-DReG to the shared parameter setting. In this section, we discuss how to deal with shared parameters in RWS.

Suppose that both  $p$  and  $q$  are parameterized by  $\theta$ . If we denote the unshared parameters of  $q$  by  $\phi$ , then we can restrict the RWS wake update to only  $\phi$ . Alternatively, with a modified RWS wake update, we can derive a single surrogate objective for each scenario such that taking the gradient with respect to  $\theta$  results in the proper update. For clarity, we introduce the following modifier notation for  $p_\theta(x, z_i)$ ,  $q_\theta(z_i|x)$ , and  $w_i$  which are functions of  $\theta$  and  $z_i = z(\theta, \epsilon_i)$ . We use  $\hat{X}$  to mean  $X$  with stopped gradients with respect to  $z_i$ ,  $\tilde{X}$  to mean  $X$  with stopped gradients with respect to  $\theta$  (but not  $\theta$  is not stopped in  $z(\theta, \epsilon_i)$ ), and  $\bar{X}$  to mean  $X$  with stopped gradients for all variables. Then, we can use the following surrogate objectives:

IWAE:

$$L_{IWAE}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} \log w_i \right] \quad (11)$$

DReG IWAE:

$$L_{DReG-IWAE}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} \log \tilde{p}_\theta(x, z_i) + \left( \frac{\bar{w}_i}{\sum_j \bar{w}_j} \right)^2 \log \hat{w}_i \right] \quad (12)$$

RWS:

$$L_{RWS}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} (\log \tilde{p}_\theta(x, z_i) + \log \tilde{q}_\theta(z_i|x)) \right] \quad (13)$$

DReG RWS:

$$L_{DReG-RWS}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} \log \tilde{p}_\theta(x, z_i) + \left( \frac{\bar{w}_i}{\sum_j \bar{w}_j} - \left( \frac{\bar{w}_i}{\sum_j \bar{w}_j} \right)^2 \right) \log \hat{w}_i \right] \quad (14)$$

STL:

$$L_{STL}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} (\log \tilde{p}_\theta(x, z_i) + \log \hat{w}_i) \right] \quad (15)$$

DReG( $\alpha$ ):

$$L_{DReG(\alpha)}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} \log \tilde{p}_\theta(x, z_i) + \left( \alpha \frac{\bar{w}_i}{\sum_j \bar{w}_j} + (1 - 2\alpha) \left( \frac{\bar{w}_i}{\sum_j \bar{w}_j} \right)^2 \right) \log \hat{w}_i \right] \quad (16)$$

The only subtle difference is that DReG( $\alpha = 0.5$ ) does not correspond exactly to STL due to the scaling between terms:

$$L_{DReG(\alpha=0.5)}(\theta) = \mathbb{E}_{\epsilon_{1:K}} \left[ \sum_{i=1}^K \frac{\bar{w}_i}{\sum_j \bar{w}_j} (\log \tilde{p}_\theta(x, z_i) + 0.5 \log \hat{w}_i) \right] \quad (17)$$