

## Appendix to Generalized Doubly-Reparameterized Gradient Estimators

### Contents of the Appendix

A	Derivation of the GDREGs identity . . . . .	1
B	Derivation of the GDREGs estimator for the IWAE objective . . . . .	3
C	Derivation of the DREGs and GDREGs estimator for IWAE objectives of hierarchical VAEs . . . . .	5
D	Worked example for a 2-layer hierarchical VAE . . . . .	8
E	Surrogate losses to implement the DREGs and GDREGs estimators for IWAE objectives . . . . .	9
F	Implementation details . . . . .	11
G	Experimental details and additional results . . . . .	13
H	The cross-entropy for Gaussian distributions . . . . .	17

### A. Derivation of the GDREGs identity

Here, we detail the derivation of our main result, the GDREGs identity (Eq. (16)), which we restate here.

$$\mathbb{E}_{z \sim q_\phi(z)} [g_{\phi, \theta}(z) \nabla_{\theta} \log p_{\theta}(z)] = \mathbb{E}_{z \sim q_\phi(z)} \left[ \left( g_{\phi, \theta}(z) \nabla_z^{\text{TD}} \log \frac{q_\phi(z)}{p_\theta(z)} + \nabla_z^{\text{TD}} g_{\phi, \theta}(z) \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) \Big|_{\tilde{\epsilon} = \mathcal{T}_p^{-1}(z, \theta)} \right] \quad (16)$$

We can derive this identity in two ways: (1) through a reweighting correction as presented in the main paper, see [App. A.1](#); (2) through a flow-like transformation of  $z \sim q_\phi(z|x)$ , see [App. A.2](#).

#### A.1. Derivation via reweighting

In the main paper ([Sec. 4](#)), we explained that we need to make the sampling path of  $z$  depend on the parameters  $\theta$  to replace the score function  $\nabla_{\theta} \log p_{\theta}(z)$  with a pathwise derivative. At the same time, we evaluate the objective on samples  $z$  from the variational posterior  $q_\phi(z|x)$  during training. The derivation consists of the following three steps:

- ① *Temporarily* change the path such that it depends on  $\theta$ . We change the path by first using importance sampling reweighting to temporarily re-write the expectation,  $\mathbb{E}_{q_\phi(z)}[*] = \mathbb{E}_{p_\theta(z)} \left[ \frac{q_\phi(z)}{p_\theta(z)} * \right]$  (step ①a), and then by employing reparameterization on  $p_\theta(z)$ :  $z = \mathcal{T}_p(\tilde{\epsilon}; \theta)$  with  $\tilde{\epsilon} \sim p(\tilde{\epsilon}) = \mathcal{N}(0, \mathbb{I})$  (step ①b).
- ② Perform the gradient computation and collect all the terms.
- ③ Change the path back by un-doing the reparameterization and the reweighting so we can use samples  $z \sim q_\phi(z|x)$  to estimate the expectation.

$$\nabla_{\theta}^{\text{TD}} \mathbb{E}_{q_\phi(z)} [g_{\phi, \theta}(z)] \stackrel{\text{①a}}{=} \nabla_{\theta}^{\text{TD}} \mathbb{E}_{p_\theta(z)} \left[ \frac{q_\phi(z)}{p_\theta(z)} g_{\phi, \theta}(z) \right] \stackrel{\text{①b}}{=} \nabla_{\theta}^{\text{TD}} \mathbb{E}_{p(\tilde{\epsilon})} \left[ \frac{q_\phi(\mathcal{T}_p(\tilde{\epsilon}; \theta))}{p_\theta(\mathcal{T}_p(\tilde{\epsilon}; \theta))} g_{\phi, \theta}(\mathcal{T}_p(\tilde{\epsilon}; \theta)) \right] \quad (A.1)$$

$$\stackrel{\text{②}}{=} \mathbb{E}_{p(\tilde{\epsilon})} \left[ \nabla_z^{\text{TD}} \left( \frac{q_\phi(z)}{p_\theta(z)} g_{\phi, \theta}(z) \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) + \frac{q_\phi(z)}{p_\theta(z)} (\nabla_{\theta} g_{\phi, \theta}(z) - g_{\phi, \theta}(z) \nabla_{\theta} \log p_{\theta}(z)) \right]_{z = \mathcal{T}_p(\tilde{\epsilon}; \theta)} \quad (A.2)$$

$$\stackrel{\text{③}}{=} \mathbb{E}_{q_\phi(z)} \left[ \left( g_{\phi, \theta}(z) \nabla_z^{\text{TD}} \log \frac{q_\phi(z)}{p_\theta(z)} + \nabla_z^{\text{TD}} g_{\phi, \theta}(z) \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) \Big|_{\tilde{\epsilon} = \mathcal{T}_p^{-1}(z, \theta)} + \nabla_{\theta} g_{\phi, \theta}(z) - g_{\phi, \theta}(z) \nabla_{\theta} \log p_{\theta}(z) \right] \quad (A.3)$$

In the derivation we have used the identity  $x \nabla_* \log x = \nabla_* x$  repeatedly. By noting that  $\nabla_{\theta}^{\text{TD}} \mathbb{E}_{q_{\phi}(z)} [g_{\phi, \theta}(z)] = \mathbb{E}_{q_{\phi}(z)} [\nabla_{\theta} g_{\phi, \theta}(z)]$ , we can cancel these terms on the left hand side of Eq. (A.1) and right hand side of Eq. (A.3). By moving  $-\mathbb{E}_{q_{\phi}(z)} [g_{\phi, \theta}(z) \nabla_{\theta} \log p_{\theta}(z)]$  to the other side we obtain the desired result.

## A.2. Derivation via flow-like transformation

Here we provide a second derivation of the GDREGs identity (Eq. (16)) that does not explicitly use a re-weighting of the expectation as in App. A.1, but uses flow-like transformations instead (Rezende & Mohamed, 2015). We already noted in the main text that we can interpret the change of paths as a flow  $z \rightarrow \tilde{\epsilon} \rightarrow z$ , where  $\tilde{\epsilon}$  follows a more complicated distribution than the original  $\epsilon$  that might have been used to reparameterize samples  $z$  from  $q_{\phi}(z|x)$ . Here, we make this connection to normalizing flows more explicit.

We make use of the usual change of probability density formula for flows (see, e.g., Rezende & Mohamed (2015)):

$$y = f(x); x \sim p_x(x) \Rightarrow p_y(y) = p_x(x) |\nabla_x f(x)|^{-1}. \quad (\text{A.4})$$

Our main insight is that we can reparameterize  $z \sim q_{\phi}(z|x)$  in many different ways. Usually we sample an independent noise variable from a simple base distribution, for example,  $\epsilon \sim q(\epsilon) = \mathcal{N}(0, \mathbb{I})$ , and use it to reparameterize  $z$  as  $z = \mathcal{T}_q(\epsilon; \phi)$ . However, we can equally use a different base distribution  $\tilde{q}(\tilde{\epsilon})$  and reparameterize  $z$  as  $z = \mathcal{T}_p(\tilde{\epsilon}; \theta)$ ,  $\tilde{\epsilon} \sim \tilde{q}(\tilde{\epsilon})$ . Note that we reparameterize using  $p_{\theta}(z)$  in this case. In order for  $z$  still to be distributed according to  $q_{\phi}(z|x)$ , we have to choose  $\tilde{q}(\tilde{\epsilon})$  to be itself given by the normalizing flow  $\tilde{\epsilon} = \mathcal{T}_p^{-1}(z; \theta)$ ,  $z \sim q_{\phi}(z)$ , such that  $\tilde{q}(\tilde{\epsilon}) = \tilde{q}(\tilde{\epsilon}; \phi, \theta)$ . It may appear counter-intuitive to transform  $z$  as  $z \rightarrow \tilde{\epsilon} \rightarrow z$  because we are doing and then un-doing a transformation; however, it allows us to make the computation path depend on  $\theta$  when computing the gradients. We then separate the forward flow  $z = \mathcal{T}_p(\tilde{\epsilon}; \theta)$  and the backward flow  $\tilde{\epsilon} = \mathcal{T}_p^{-1}(z; \theta)$  by moving the forward flow into the path through reparameterization and the backward flow into the integration measure, see Eq. (A.7) below. To move the backward flow into the integration measure, we express  $\tilde{q}(\tilde{\epsilon})$  as a change of density:

$$\tilde{\epsilon} = \mathcal{T}_p^{-1}(z; \theta); z \sim q_{\phi}(z|x) \Rightarrow \tilde{q}(\tilde{\epsilon}) = q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \quad (\text{A.5})$$

Note that  $\tilde{q}(\tilde{\epsilon})$  is different from  $q(\epsilon) = \mathcal{N}(0, \mathbb{I})$  typically used to reparameterize samples of the approximate posterior  $z = \mathcal{T}_q(\epsilon; \phi)$ ;  $\epsilon \sim q(\epsilon)$ ! Using the change of density in Eq. (A.5) together with reparameterization as described above, we can re-write the following expectation over  $z$  as an integral over  $\tilde{\epsilon}$ :

$$\int q_{\phi}(z) g_{\phi, \theta}(z) dz = \int \tilde{q}(\tilde{\epsilon}) g_{\phi, \theta}(\mathcal{T}_p(\tilde{\epsilon}; \theta)) d\tilde{\epsilon} \quad \text{reparameterization} \quad (\text{A.6})$$

$$= \int \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \right]_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} g_{\phi, \theta}(\mathcal{T}_p(\tilde{\epsilon}; \theta)) d\tilde{\epsilon} \quad \text{change of density} \quad (\text{A.7})$$

In these integrals,  $\tilde{\epsilon}$  is an *independent* variable; its dependence on  $\theta$  has been moved into the path ( $g_{\phi, \theta}(\mathcal{T}_p(\tilde{\epsilon}; \theta))$ ) as well as the change of density (Eq. (A.5)). This is one way to understand where the `stop_gradient` in Figs. 2 and 3 comes from.

To derive the GDREGs identity, we take the total derivative of Eq. (A.7) w.r.t.  $\theta$  and apply the chain- and product rule

$$\nabla_{\theta}^{\text{TD}} \int q_{\phi}(z) g_{\phi, \theta}(z) dz \stackrel{\text{A.7}}{=} \nabla_{\theta}^{\text{TD}} \int \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \right]_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} g_{\phi, \theta}(\mathcal{T}_p(\tilde{\epsilon}; \theta)) d\tilde{\epsilon} \quad (\text{A.8})$$

$$= \int \nabla_{\theta}^{\text{TD}} \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} g_{\phi, \theta}(z) \right]_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} d\tilde{\epsilon} \quad (\text{A.9})$$

$$= \int \nabla_z^{\text{TD}} \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} g_{\phi, \theta}(z) \right]_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) + \nabla_{\theta} \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} g_{\phi, \theta}(z) \right]_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} d\tilde{\epsilon} \quad (\text{A.10})$$

$$= \int \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \right] \left( g_{\phi, \theta}(z) \nabla_z^{\text{TD}} \log \left( q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \right) + \nabla_z^{\text{TD}} g_{\phi, \theta}(z) \right) \Big|_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) + \left[ q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} \right] \left( g_{\phi, \theta}(z) \nabla_{\theta} \log \left( q_{\phi}(z) |\nabla_z \mathcal{T}_p^{-1}(z; \theta)|^{-1} + \nabla_{\theta} g_{\phi, \theta}(z) \right) \right) \Big|_{z=\mathcal{T}_p(\tilde{\epsilon}; \theta)} d\tilde{\epsilon} \quad (\text{A.11})$$

where we have separated out the derivatives and used  $x\nabla_* \log x = \nabla_* x$ . We can now further separate terms and undo the change of density to replace  $q_\phi(\mathbf{z}) |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})|^{-1} = \tilde{q}(\tilde{\boldsymbol{\epsilon}})$  (Eq. (A.5)) after taking the derivatives. We obtain

$$\nabla_{\boldsymbol{\theta}}^{\text{TD}} \int q_\phi(\mathbf{z}) g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) d\mathbf{z} = \quad (\text{A.12})$$

$$\begin{aligned} &= \int \tilde{q}(\tilde{\boldsymbol{\epsilon}}) \left\{ \nabla_{\mathbf{z}}^{\text{TD}} \left( \log q_\phi(\mathbf{z}) + \log |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})|^{-1} \right) g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \Big|_{\mathbf{z}=\mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} \mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta}) + \right. \\ &\quad \left. + \nabla_{\mathbf{z}}^{\text{TD}} g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \Big|_{\mathbf{z}=\mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} \mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \Big|_{\mathbf{z}=\mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta})} + \right. \\ &\quad \left. - g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \Big|_{\mathbf{z}=\mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} \log |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})| \Big|_{\mathbf{z}=\mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta})} \right\} d\tilde{\boldsymbol{\epsilon}}. \end{aligned} \quad (\text{A.13})$$

To evaluate the gradients of the log Jacobians,  $\nabla_* \log |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})|$ , we can combine the log Jacobians with a simple base distribution  $q(\tilde{\boldsymbol{\epsilon}})$  to obtain  $p_\theta(\mathbf{z})$  because

$$\mathbf{z} = \mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta}); \tilde{\boldsymbol{\epsilon}} \sim q(\tilde{\boldsymbol{\epsilon}}) = \mathcal{N}(0, \mathbb{I}) \quad \Rightarrow \quad p_\theta(\mathbf{z}) = q(\tilde{\boldsymbol{\epsilon}}) |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})| \quad (\text{A.14})$$

through reparameterization of  $p_\theta(\mathbf{z})$  and by noting that

$$\nabla_* \log(f(x) \cdot c) = \nabla_* \log(f(x)) \quad \text{if } c \text{ is constant w.r.t. } * \quad (\text{A.15})$$

and that the simple base distribution  $q(\tilde{\boldsymbol{\epsilon}})$  is constant w.r.t. all gradients:

$$\nabla_* \log |\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})| \stackrel{(\text{A.15})}{=} \nabla_* \log (|\nabla_{\mathbf{z}} \mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})| q(\tilde{\boldsymbol{\epsilon}})) \stackrel{(\text{A.14})}{=} \nabla_* \log p_\theta(\mathbf{z}). \quad (\text{A.16})$$

This allows us to simplify Eq. (A.13) as

$$\begin{aligned} \nabla_{\boldsymbol{\theta}}^{\text{TD}} \int q_\phi(\mathbf{z}) g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) d\mathbf{z} &= \int q_\phi(\mathbf{z}) \left\{ \left( g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \nabla_{\mathbf{z}}^{\text{TD}} \log \frac{q_\phi(\mathbf{z})}{p_\theta(\mathbf{z})} + \nabla_{\mathbf{z}}^{\text{TD}} g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \right) \nabla_{\boldsymbol{\theta}} \mathcal{T}_p(\tilde{\boldsymbol{\epsilon}}; \boldsymbol{\theta}) \Big|_{\tilde{\boldsymbol{\epsilon}}=\mathcal{T}_p^{-1}(\mathbf{z}; \boldsymbol{\theta})} + \right. \\ &\quad \left. + \nabla_{\boldsymbol{\theta}} g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) - g_{\phi, \boldsymbol{\theta}}(\mathbf{z}) \nabla_{\boldsymbol{\theta}} \log p_\theta(\mathbf{z}) \right\} d\mathbf{z}, \end{aligned} \quad (\text{A.17})$$

which is identical to Eq. (A.3) and yields the GDREGs identity as explained above.

## B. Derivation of the GDREGs estimator for the IWAE objective

In this section we apply the GDREGs identity derived above to derive the GDREGs estimator for the IWAE objective, Eq. (22) in the main paper.

### B.1. Preliminaries on the IWAE objective

The importance weighted autoencoder (IWAE) objective is given by

$$\mathcal{L}_{\phi, \boldsymbol{\theta}}^{\text{IWAE}} = \mathbb{E}_{\mathbf{z}_{1:K} \sim q_\phi(\mathbf{z}_k | \mathbf{x})} \left[ \log \left( \frac{1}{K} \sum_{k=1}^K w_k \right) \right] \quad w_k = \frac{p_\theta(\mathbf{z}) p(\mathbf{x} | \mathbf{z}_k)}{q_\phi(\mathbf{z}_k | \mathbf{x})} \quad (\text{B.1})$$

where  $w_k$  are the importance weights (Burda et al., 2016).

Due to the structure of the IWAE objective, any gradient w.r.t. any of its parameters can be written as

$$\nabla_*^{\text{TD}} \mathcal{L}_{\phi, \boldsymbol{\theta}}^{\text{IWAE}} = \mathbb{E}_{\boldsymbol{\epsilon}_{1:K} \sim q(\boldsymbol{\epsilon})} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_*^{\text{TD}} \log w_k \right]; \quad \tilde{w}_k = \frac{w_k}{\sum_j w_j} \quad (\text{B.2})$$

using the chain rule and  $\nabla_* w_k = w_k \nabla_* \log w_k$ .  $\tilde{w}_k$  are the normalized importance weights, and we have reparameterized  $\mathbf{z}_k$  as  $\mathcal{T}_q(\boldsymbol{\epsilon}_k; \phi)$ . Typically, the derivatives we are interested in are w.r.t. the parameters  $\phi$  and  $\boldsymbol{\theta}$ .

We also note the following identity that we use in the derivation of the doubly reparameterized estimators,

$$\nabla_{\mathbf{z}}^{\text{TD}} \tilde{w}_k = (\tilde{w}_k - \tilde{w}_k^2) \nabla_{\mathbf{z}}^{\text{TD}} \log w_k \quad (\text{B.3})$$

which follows from applying the chain-rule and using  $\nabla_* w_k = w_k \nabla_* \log w_k$ .

Tucker et al. (2019) derive the DREGs identity (Eq. (9)) and use it to derive the following doubly-reparameterized gradient estimator (DREGs) w.r.t. the approximate posterior parameters  $\phi$  as:

$$\hat{\nabla}_{\phi}^{\text{DREGs}} \mathcal{L}_{\text{IWAE}} = \sum_{k=1}^K \tilde{w}_k^2 \nabla_{\mathbf{z}_k}^{\text{TD}} \log w_k \nabla_{\phi}^{\text{TD}} \mathcal{T}_q(\epsilon_k; \phi). \quad \epsilon_{1:K} \sim q(\epsilon) \quad (\text{B.4})$$

## B.2. Derivation of the GDREGs estimator

Similarly, we can derive a generalized doubly-reparameterized gradient (GDREGs) estimator w.r.t. the prior parameters  $\theta$ . We use the GDREGs identity (Eq. (16)) derived above with  $g_{\phi, \theta}(\mathbf{z}) = \tilde{w}_k$  and note that the reweighting term  $\log \frac{q_{\phi}(\mathbf{z})}{p_{\theta}(\mathbf{z})}$  looks like a log importance weight except for the missing likelihood:

$$\nabla_{\theta}^{\text{TD}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\theta}^{\text{TD}} \log w_k \right] = \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\theta}^{\text{TD}} \log p_{\theta}(\mathbf{z}) \right] \quad (\text{B.5})$$

$$\stackrel{(16)}{=} \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})} \left[ \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_k}^{\text{TD}} \log \frac{q_{\phi}(\mathbf{z}_k | \mathbf{x})}{p_{\theta}(\mathbf{z}_k)} + \nabla_{\mathbf{z}_k}^{\text{TD}} \tilde{w}_k \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \Big|_{\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)} \right] \quad (\text{B.6})$$

$$\stackrel{(B.3)}{=} \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})} \left[ \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_k}^{\text{TD}} \log \frac{q_{\phi}(\mathbf{z}_k | \mathbf{x})}{p_{\theta}(\mathbf{z}_k)} + (\tilde{w}_k - \tilde{w}_k^2) \nabla_{\mathbf{z}_k}^{\text{TD}} \log w_k \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \Big|_{\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)} \right] \quad (\text{B.7})$$

$$= \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})} \left[ \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_k}^{\text{TD}} \log p(\mathbf{x} | \mathbf{z}_k) - \tilde{w}_k^2 \nabla_{\mathbf{z}_k}^{\text{TD}} \log w_k \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \Big|_{\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)} \right]. \quad (\text{B.8})$$

Thus, the GDREGs estimator is given by:

$$\hat{\nabla}_{\theta}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_k}^{\text{TD}} \log p(\mathbf{x} | \mathbf{z}_k) - \tilde{w}_k^2 \nabla_{\mathbf{z}_k}^{\text{TD}} \log w_k \right) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \Big|_{\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)} \quad \mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x}). \quad (22)$$

Note that the  $\mathbf{z}_k$  are sampled from  $q_{\phi}(\mathbf{z}_k | \mathbf{x})$  but re-expressed as if they came from  $p_{\theta}(\mathbf{z})$ .

We can rewrite the importance weights as

$$w_k = \frac{p_{\theta}(\mathbf{z}_k) p(\mathbf{x} | \mathbf{z}_k)}{q_{\phi}(\mathbf{z}_k | \mathbf{x})} = \frac{p_{\theta}(\mathbf{z}_k | \mathbf{x}) p_{\theta}(\mathbf{x})}{q_{\phi}(\mathbf{z}_k | \mathbf{x})}. \quad (\text{B.9})$$

Thus, if the variational posterior  $q_{\phi}(\mathbf{z}_k | \mathbf{x})$  is equal to the true posterior  $p_{\theta}(\mathbf{z}_k | \mathbf{x})$ , all weights  $w_k$  become equal to  $p_{\theta}(\mathbf{x})$  and thus constant w.r.t.  $\mathbf{z}_k$ . In that case the second term in the GDREGs estimator Eq. (22) vanishes and the overall expression simplifies to

$$\hat{\nabla}_{\theta}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K \tilde{w}_k \nabla_{\mathbf{z}_k}^{\text{TD}} \log p(\mathbf{x} | \mathbf{z}_k) \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \Big|_{\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)}. \quad \mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x}) \quad (\text{B.10})$$

In contrast, the usual IWAE gradient involves the score function for  $p_{\theta}(\mathbf{z}_k)$ :

$$\hat{\nabla}_{\theta}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{i=1}^K \tilde{w}_i \nabla_{\theta} \log p_{\theta}(\mathbf{z}_i), \quad \mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x}). \quad (\text{B.11})$$

## C. Derivation of the DREGs and GDREGs estimator for IWAE objectives of hierarchical VAEs

In this section we derivations of and further details on the extension of DREGs and GDREGs to hierarchical VAEs with the IWAE objective.

### C.1. Preliminaries and notation for the hierarchical IWAE objective

For a hierarchically structured model with  $L$  stochastic layers the IWAE objective is still given by Eq. (B.1) but with importance weights  $w_k$  given by

$$w_k = \frac{p_{\lambda}(\mathbf{x} | \mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) p_{\theta}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL})}{q_{\phi}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL} | \mathbf{x})}. \quad (\text{C.1})$$

Here,  $\mathbf{z}_{kl}$  denotes the  $k$ th importance sample ( $k \in \{1, \dots, K\}$ ) for the  $l$ th layer ( $l \in \{1, \dots, L\}$ ). Both the variational posterior and the prior distribution factorize according to their respective hierarchical structure. While the prior factorizes top-down in most cases, the variational posterior can have many different structures. In order for the distributions to be valid in the context of a VAE, we require the individual dependency graphs for the prior (generative path) and the variational posterior (inference path) to be directed acyclic graphs. Cycles would mean that a latent variable conditionally dependent on itself. To keep the dependency structure general, we write the factorization of the variational posterior and prior as follows:

$$q_{\phi}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL} | \mathbf{x}) = \prod_{l=1}^L q_{\phi_l}(\mathbf{z}_{kl} | \text{pa}_{\alpha}(l), \mathbf{x}) = \prod_{l=1}^L q_{\alpha_l(\text{pa}_{\alpha}(l); \phi_l)}(\mathbf{z}_{kl}) \quad (\text{C.2})$$

$$p_{\theta}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) = \prod_{l=1}^L p_{\theta_l}(\mathbf{z}_{kl} | \text{pa}_{\beta}(l)) = \prod_{l=1}^L p_{\beta_l(\text{pa}_{\beta}(l); \theta_l)}(\mathbf{z}_{kl}) \quad (\text{C.3})$$

Here,  $\alpha_l(\cdot; \phi_l)$  and  $\beta_l(\cdot; \theta_l)$  are the distribution parameters of the variational posterior and prior distribution in the  $l$ th layer, respectively, and we have made the dependencies of the conditional distributions explicit;  $\text{pa}_{\alpha}(l)$  denotes the ‘‘parents’’ of the latent variable  $\mathbf{z}_{kl}$  according to the dependency graph of the inference path (the factorization of the posterior); similarly,  $\text{pa}_{\beta}(l)$  denotes the latent variables that  $\mathbf{z}_{kl}$  directly depends on according to the factorization of the prior  $p_{\theta}$ . Typically, the prior is assumed to factorize top-down, such that  $\text{pa}_{\beta}(l) = \mathbf{z}_{k(l+1)}$  for all but the top-most layer.

The samples  $\mathbf{z}_{kl}$  are drawn from the variational posterior and can be expressed through reparameterization as  $\mathbf{z}_{kl} = \mathcal{T}_{q_l}(\epsilon_{kl}; \alpha_l(\text{pa}_{\alpha}(l), \phi_l))$ , where  $\epsilon_{kl}$  is an independent noise variable per importance sample and layer.

We note that it is these dependencies of the distribution parameters  $\alpha_l$  and  $\beta_l$  on  $\text{pa}_{\alpha}(l)$  and  $\text{pa}_{\beta}(l)$ , respectively, that give rise to the indirect score functions as discussed in Sec. 3.

### C.2. Derivation of the hierarchical DREGs estimator for IWAE

With notation fully set up we consider the reparameterized gradients of the IWAE objective w.r.t. the variational parameters in a particular stochastic layer  $\phi_l$ :

$$\nabla_{\phi_l}^{\text{TD}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \mathbb{E}_{\epsilon_{1:K} \sim q(\epsilon)} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\phi_l}^{\text{TD}} \log w_k \right] \quad (\text{C.4})$$

$$= \mathbb{E}_{\epsilon_{1:K} \sim q(\epsilon)} \left[ \sum_{k=1}^K \tilde{w}_k \left( \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log w_k \nabla_{\phi_l} \mathcal{T}_{q_l}(\epsilon_{kl}; \alpha_l(\text{pa}_{\alpha}(l), \phi_l)) + \nabla_{\phi_l} \log w_k \right) \right] \quad (\text{C.5})$$

where we have used the chain-rule to arrive at Eq. (C.5); the first term contains both the (true) pathwise gradients as well as the indirect score functions; the second term only contains a direct score function as we only take the partial derivative w.r.t.  $\phi_l$ .

We can rewrite this direct score function gradient because only one term in the (log-)importance weight directly depends on  $\phi_l$ ,

$$\nabla_{\phi_l} \log w_k = -\nabla_{\phi_l} \log q_{\alpha_l(\text{pa}_{\alpha}(l); \phi_l)}(\mathbf{z}_{kl}). \quad (\text{C.6})$$

Applying the DREGs identity to this term and using Eq. (B.3) yields:

$$\mathbb{E}_{\epsilon_{1:K} \sim q(\epsilon)} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\phi_l} \log w_k \right] = -\mathbb{E}_{\epsilon_{1:K} \sim q(\epsilon)} \left[ \sum_{k=1}^K (\tilde{w}_k - \tilde{w}_k^2) \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log w_k \nabla_{\phi_l} \mathcal{T}_{q_l}(\epsilon_{kl}; \boldsymbol{\alpha}_l(\text{pa}_{\boldsymbol{\alpha}}(l), \phi_l)) \right] \quad (\text{C.7})$$

which agrees with the first term in Eq. (C.5) up to the prefactor. Thus, both the true pathwise gradients as well as the indirect score functions appear twice and the prefactors partly cancel to give rise to the DREGs estimator for hierarchical IWAE objectives:

$$\widehat{\nabla}_{\phi_l}^{\text{DREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K \tilde{w}_k^2 \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log w_k \nabla_{\phi_l} \mathcal{T}_{q_l}(\epsilon_{kl}; \boldsymbol{\alpha}_l(\text{pa}_{\boldsymbol{\alpha}}(l), \phi_l)) \quad \epsilon_{1:K} \sim q(\epsilon) \quad (\text{C.8})$$

where  $\mathbf{z}_{kl} = \mathcal{T}_{q_l}(\epsilon_{kl}; \boldsymbol{\alpha}_l(\text{pa}_{\boldsymbol{\alpha}}(l), \phi_l))$ ,  $\forall l \in \{1, \dots, L\}$ ,  $\forall k \in \{1, \dots, K\}$  through reparameterization.

We emphasize that the total derivative w.r.t.  $\mathbf{z}_{kl}$  contains pathwise gradients as well as indirect score functions for both the variational posterior as well as for the prior. The hierarchical DREGs estimator otherwise looks very similar to the DREGs estimator in the single layer case (Tucker et al., 2019).

In App. E.1 we explain how to implement this estimator effectively and in a structure-agnostic way. That is, we do *not* have to derive a new estimator for each new dependency graph of the variational posterior or the prior.

### C.3. Derivation of the hierarchical GDREGs estimator for IWAE

Next, we derive the expression for the GDREGs estimator for hierarchical VAEs with IWAE objective.

Applying the GDREGs identity entails re-expressing the samples  $\mathbf{z}_{kl}$  from the variational posterior as if they were sampled from the prior. Starting from a sample  $(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) \sim q_{\phi}(\mathbf{z}_1, \dots, \mathbf{z}_L | \mathbf{x})$ , we use the inverse flow of  $p_{\theta}$  to obtain new noise variables for each layer,  $(\tilde{\epsilon}_{k1}, \dots, \tilde{\epsilon}_{kL})$ . We then use the forward flow of  $p_{\theta}$  to obtain back  $(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL})$  but with the gradient path now depending on  $\theta$  as discussed in App. A and Sec. 4.

More precisely, we find that

$$\mathbf{z}_{kl}^{(q)} = \mathcal{T}_{q_l}(\epsilon_{kl}; \boldsymbol{\alpha}_l(\text{pa}_{\boldsymbol{\alpha}}(l), \phi_l)) \quad \text{original sampling of } (\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) \sim q_{\phi}(\mathbf{z}_1, \dots, \mathbf{z}_L | \mathbf{x}) \quad (\text{C.9})$$

$$\tilde{\epsilon}_{kl} = \mathcal{T}_{p_l}^{-1}(\mathbf{z}_{kl}^{(q)}; \boldsymbol{\beta}_l(\text{pa}_{\boldsymbol{\beta}}(l), \theta_l)) \quad \text{inverse prior flow to obtain new “noise” variables} \quad (\text{C.10})$$

$$\mathbf{z}_{kl} = \mathcal{T}_{p_l}(\tilde{\epsilon}_{kl}; \boldsymbol{\beta}_l(\text{pa}_{\boldsymbol{\beta}}(l), \theta_l)) \quad \text{forward prior flow to re-express the } \mathbf{z}_{kl} \quad (\text{C.11})$$

where  $\epsilon_{kl} \sim q(\epsilon)$  follows a simple distribution that is different from the more complicated distribution of  $\tilde{\epsilon}_{kl}$ . Note how the initial reparameterization of a sample  $\mathbf{z}_{kl}$  depends on the dependency structure of the variational posterior (through  $\text{pa}_{\boldsymbol{\alpha}}(\cdot)$ ), while the other transformations depend on the dependency structure of the prior ( $\text{pa}_{\boldsymbol{\beta}}(\cdot)$ ).

As for DREGs, we note that only one term in the log importance weight directly depends on the variable  $\theta_l$ ,

$$\nabla_{\theta_l} \log w_k = \nabla_{\theta_l} \log p_{\boldsymbol{\beta}_l(\text{pa}_{\boldsymbol{\beta}}(l); \theta_l)}(\mathbf{z}_{kl}). \quad (\text{C.12})$$

With these prerequisites, we can compute the GDREGs estimator for parameters  $\theta_l$  of the  $l$ th stochastic layer.

$$\nabla_{\theta_l}^{\text{TD}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z} | \mathbf{x})} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\theta_l} \log w_k \right] \quad (\text{C.13})$$

$$\stackrel{\text{C.12}}{=} \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z} | \mathbf{x})} \left[ \sum_{k=1}^K \tilde{w}_k \nabla_{\theta_l} \log p_{\boldsymbol{\beta}_l(\text{pa}_{\boldsymbol{\beta}}(l); \theta_l)}(\mathbf{z}_{kl}) \right] \quad (\text{C.14})$$

$$\stackrel{\text{16}}{=} \mathbb{E}_{\mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z} | \mathbf{x})} \left[ \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log \frac{q_{\phi}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL} | \mathbf{x})}{p_{\theta}(\mathbf{z}_{k1}, \dots, \mathbf{z}_{kL})} + \right. \right. \\ \left. \left. + (\tilde{w}_k - \tilde{w}_k^2) \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log w_k \right) \nabla_{\theta_l} \mathcal{T}_{p_l}(\tilde{\epsilon}_{kl}; \theta_l) \Big|_{\tilde{\epsilon}_{kl} = \mathcal{T}_{p_l}^{-1}(\mathbf{z}_{kl}; \theta_l)} \right] \quad (\text{C.15})$$

$$\begin{aligned}
 \widehat{\nabla}_{\theta_l}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} &= \\
 &= \sum_{k=1}^K \left( \tilde{w}_k \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log p_{\lambda}(\mathbf{x} | \mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) - \tilde{w}_k^2 \nabla_{\mathbf{z}_{kl}}^{\text{TD}} \log w_k \right) \nabla_{\theta_l} \mathcal{T}_{p_l}(\tilde{\epsilon}_{kl}; \theta_l) \Big|_{\tilde{\epsilon}_{kl} = \mathcal{T}_{p_l}^{-1}(\mathbf{z}_{kl}; \theta_l)}; \quad \mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})
 \end{aligned} \tag{C.16}$$

where we suppressed dependencies on  $\text{pa}_*(l)$  where they are not necessary to simplify notation.

The estimator looks very similar to the GDREGs estimator for a single layer IWAE model Eq. (22). Note that just like above for hierarchical DREGs, the total gradients w.r.t.  $\mathbf{z}_{kl}$  give rise to both (true) pathwise gradients as well as indirect score functions through the hierarchical dependencies of the variational posterior and prior.

In App. E.2 we show how to implement the hierarchical GDREGs estimator Eq. (C.16) effectively and in a way that is agnostic to the structure of the model. That is, we do not have to derive a separate estimator for every dependency graph of the variational posterior and prior.

#### C.4. Double reparameterization and indirect score functions

In principle, we could apply double-reparameterization to the indirect score functions as well. However, as we explain now, we often cannot doubly-reparameterize *all* indirect score functions; moreover, even in cases where this is possible, it is still impractical, as the corresponding estimator depends on the exact model structure and would require adaptation to each dependency graph of the prior and variational posterior.

Double reparameterization of indirect score functions works in the same way as for the direct score functions except that  $g_{\phi, \theta}(\mathbf{z})$  is given by  $\tilde{w}_k^2$  instead of  $\tilde{w}_k$  in this case. The derivatives of  $\tilde{w}_k^2$  have a similar reproducing property as we observed in Eq. (B.3):

$$\nabla_{\mathbf{z}}^{\text{TD}} \tilde{w}_k^2 = 2(\tilde{w}_k^2 - \tilde{w}_k^3) \nabla_{\mathbf{z}}^{\text{TD}} \log w_k. \tag{C.17}$$

Thus, double reparameterization of the indirect score functions similarly gives rise to further indirect score functions. We note that these indirect score functions only appear for the “children” of the current stochastic layer, that is, stochastic variables in those layers that depend on the current layer. In this context, “children” refers to *all* children w.r.t. the dependency structure of both, the variational posterior *and* the prior. For a particular layer  $l$  we obtain indirect score functions from double reparameterization of all of its (direct or indirect) parent nodes. Following the dependency structure, we could collect all of these terms and reparameterize them to obtain pathwise gradients only.

However, a problem arises, because we need to account for dependencies of both the variational posterior *and* the prior. Reparameterization of a score function gives rise to indirect score functions in all its “children” layers for both the variational posterior and the prior. For general hierarchical structures, this leads to cycles, in that some of the children of one dependency tree (the variational posterior) are the parents in the other (the prior) and/or vice versa. In this case we are never able to collect all the terms and fully reparameterize all the score functions.

Moreover, even if the joint dependency graph of the variational posterior and the prior were acyclic, this derivation would be structure-specific and would need to be repeated for each hierarchical structure. We therefore do not doubly reparameterize the indirect score functions.

## D. Worked example for a 2-layer hierarchical VAE

In this section we show an example of using the IWAE objective with a 2-layer VAE model consisting of a prior  $p_{\theta_2}(z_2)p_{\theta_1}(z_1|z_2)$  and likelihood  $p_{\lambda}(x|z_1, z_2)$ . The inference network is bottom-up:  $q_{\phi_1}(z_1|x)q_{\phi_2}(z_2|x, z_1)$ . Therefore  $\text{pa}_{\alpha}(2) = z_1$  and  $\text{pa}_{\beta}(1) = z_2$ .

We hierarchically sample  $z_{k1}$  and  $z_{k2}$  from the approximate posterior and abbreviate:

$$z_{k1}(\phi_1) \equiv \mathcal{T}_{q_1}(\epsilon_{k1}; \alpha_1(\phi_1)) \quad (\text{D.1})$$

$$z_{k2}(\phi_1, \phi_2) \equiv \mathcal{T}_{q_{2|1}}(\epsilon_1; \alpha_{2|1}(x, z_{k1}(\phi_1), \phi_2)). \quad (\text{D.2})$$

We also explicitly distinguish between the distribution parameters  $\alpha_i$  and the network parameters  $\phi_i$  for the posterior as well as the distribution parameters  $\beta_i$  and the network parameters  $\theta_i$  for the prior. A single importance sample  $w_k$  with all of its functional dependencies is given by:

$$w_k = \frac{p_{\lambda}(x|z_{k1}(\phi_1), z_{k2}(\phi_1, \phi_2)) p_{\beta_2}(\theta_2)(z_{k2}(\phi_1, \phi_2)) p_{\beta_{1|2}}(z_{k2}(\phi_1, \phi_2), \theta_1)(z_{k1}(\phi_1))}{q_{\alpha_1(x, \phi_1)}(z_{k1}(\phi_1)) q_{\alpha_{2|1}(x, z_{k1}(\phi_1), \phi_2)}(z_{k2}(\phi_1, \phi_2))}. \quad (\text{D.3})$$

The DREGs estimator for the variational parameters of the lower latent layer,  $\phi_1$ , is then given by:

$$\widehat{\nabla}_{\phi_1}^{\text{DREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K \tilde{w}_k^2 \nabla_{z_{k1}}^{\text{TD}} \log w_k \nabla_{\phi_1} \mathcal{T}_{q_1}(\epsilon_{k1}; \alpha_1(x, \phi_1)); \quad \epsilon_{1:K} \sim q(\epsilon) \quad (\text{D.4})$$

$$\begin{aligned} \nabla_{z_{k1}}^{\text{TD}} \log w_k &= \nabla_{z_{k1}} \log w_k - \nabla_{\alpha_{2|1}} \log q_{\alpha_{2|1}(x, z_1(\phi_1), \phi_2)}(z_{k2}(\phi_1, \phi_2)) \nabla_{z_{k1}} \alpha_{2|1}(x, z_{k1}(\phi_1), \phi_2) \\ &\quad + \nabla_{\beta_{1|2}} \log p_{\beta_{1|2}}(z_{k2}(\phi_1, \phi_2), \theta_1)(z_{k1}(\phi_1)) \nabla_{z_{k1}} \beta_{1|2}(z_{k2}(\phi_1, \phi_2), \theta_1) \end{aligned} \quad (\text{D.5})$$

where we have expanded the total derivative w.r.t.  $z_{k1}$  into the (true) pathwise gradients and two indirect score functions.

Similarly, we can compute the DREGs estimator for the variational parameters of the upper level,  $\phi_2$ :

$$\widehat{\nabla}_{\phi_2}^{\text{DREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K \tilde{w}_k^2 \nabla_{z_{k2}}^{\text{TD}} \log w_k \nabla_{\phi_2} \mathcal{T}_{q_{2|1}}(\epsilon_{k2}; \alpha_2(x, z_{k1}(\phi_1), \phi_2)); \quad \epsilon_{1:K} \sim q(\epsilon) \quad (\text{D.6})$$

$$\nabla_{z_{k2}}^{\text{TD}} \log w_k = \nabla_{z_{k2}} \log w_k + \nabla_{\beta_{1|2}} \log p_{\beta_{1|2}}(z_2(\phi_1, \phi_2), \theta_1)[z_1(\phi_1)] \nabla_{z_{k2}} \beta_{1|2}(z_2(\phi_1, \phi_2), \theta_1) \quad (\text{D.7})$$

For this model structure of the prior and variational posterior, there is only one indirect score function for this gradient. Note that the indirect score functions are computed automatically in our surrogate losses that we introduce in the following section, such that we do not need to compute them manually; the hierarchical DREGs estimator can be implemented without having to trace the dependency structure of the model manually.

To compute the GDREGs estimator, we first have to re-express the samples  $z_1$  and  $z_2$  as if they were sampled from  $p_{\theta_1, \theta_2}(z_1, z_2)$ . We write this reparameterization as

$$z_{k2}(\theta_2) \equiv \mathcal{T}_{p_2}(\tilde{\epsilon}_{k2}; \beta_2(\theta_2)) \quad (\text{D.8})$$

$$z_{k1}(\theta_1, \theta_2) \equiv \mathcal{T}_{p_{2|1}}(\tilde{\epsilon}_{k1}; \beta_{1|2}(z_{k2}(\theta_2), \theta_1)) \quad (\text{D.9})$$

where  $\tilde{\epsilon}_1$  and  $\tilde{\epsilon}_2$  are noise variables drawn from  $\tilde{q}(\tilde{\epsilon})$ , which is given by the inverse prior flow of the samples drawn from  $q_{\phi_1, \phi_2}(x_1, x_2|x)$ . The full functional dependency of a single importance sample is given by:

$$w_k = \frac{p_{\lambda}(x|z_{k1}(\theta_1, \theta_2), z_{k2}(\theta_2)) p_{\beta_2}(\theta_2)(z_{k2}(\theta_2)) p_{\beta_{1|2}}(z_{k2}(\theta_2), \theta_1)(z_{k1}(\theta_1, \theta_2))}{q_{\alpha_1(x, \phi_1)}(z_{k1}(\theta_1, \theta_2)) q_{\alpha_{2|1}(x, z_{k1}(\theta_1, \theta_2), \phi_2)}(z_{k2}(\theta_2))}. \quad (\text{D.10})$$

The GDREGs estimator w.r.t. the prior parameters of the lower stochastic layer,  $\theta_1$ , is given by:

$$\widehat{\nabla}_{\theta_1}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{IWAE}} = \sum_{k=1}^K (\tilde{w}_k \nabla_{z_{k1}}^{\text{TD}} \log p_{\lambda}(x|z_{k1}, z_{k2}) - \tilde{w}_k^2 \nabla_{z_{k1}}^{\text{TD}} \log w_k) \nabla_{\theta_1} \mathcal{T}_{p_1}(\tilde{\epsilon}_{k1}; \theta_1)|_{\tilde{\epsilon}_{k1} = \mathcal{T}_{p_1}^{-1}(z_{k1}; \theta_1)}; \quad z_{1:K} \sim q_{\phi}(z_k|x)$$



$$\nabla_{\mathbf{z}_{k1}}^{\text{TD}} \log w_k = \nabla_{\mathbf{z}_{k1}} \log w_k - \nabla_{\alpha_{2|1}} \log q_{\alpha_{2|1}}(\mathbf{x}, \mathbf{z}_{k1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2), \phi_2)(\mathbf{z}_{k2}(\boldsymbol{\theta}_2)) \nabla_{\mathbf{z}_{k1}} \alpha_{2|1}(\mathbf{x}, \mathbf{z}_{k1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2), \phi_2) \quad (\text{D.11})$$

The GDREGs estimator w.r.t. the prior parameters of the upper stochastic layer,  $\boldsymbol{\theta}_2$ , is given by:

$$\begin{aligned} \widehat{\nabla}_{\boldsymbol{\theta}_2}^{\text{GDREGs}} \mathcal{L}_{\phi, \boldsymbol{\theta}}^{\text{IWAE}} &= \sum_{k=1}^K (\tilde{w}_k \nabla_{\mathbf{z}_{k2}}^{\text{TD}} \log p_{\lambda}(\mathbf{x} | \mathbf{z}_{k1}, \mathbf{z}_{k2}) - \tilde{w}_k^2 \nabla_{\mathbf{z}_{k2}}^{\text{TD}} \log w_k) \nabla_{\boldsymbol{\theta}_2} \mathcal{T}_{p_2}(\tilde{\boldsymbol{\epsilon}}_{k2}; \boldsymbol{\theta}_2) |_{\tilde{\boldsymbol{\epsilon}}_{k2} = \mathcal{T}_{p_2}^{-1}(\mathbf{z}_{k2}; \boldsymbol{\theta}_2)}; \quad \mathbf{z}_{1:K} \sim q_{\phi}(\mathbf{z}_k | \mathbf{x}) \\ \nabla_{\mathbf{z}_{k2}}^{\text{TD}} \log w_k &= \nabla_{\mathbf{z}_{k2}} \log w_k - \nabla_{\alpha_{2|1}} \log q_{\alpha_{2|1}}(\mathbf{x}, \mathbf{z}_{k1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2), \phi_2)(\mathbf{z}_{k2}(\boldsymbol{\theta}_2)) \nabla_{\mathbf{z}_{k2}} \alpha_{2|1}(\mathbf{x}, \mathbf{z}_{k1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2), \phi_2) \\ &\quad + \nabla_{\beta_{1|2}} \log p_{\beta_{1|2}}(\mathbf{z}_2(\boldsymbol{\theta}_2), \boldsymbol{\theta}_1)(\mathbf{z}_{k1}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)) \nabla_{\mathbf{z}_{k2}} \beta_{1|2}(\mathbf{z}_{k2}(\boldsymbol{\theta}_2), \boldsymbol{\theta}_1) \end{aligned} \quad (\text{D.12})$$

Note how the GDREGs estimator for  $\boldsymbol{\theta}_2$  has two indirect score functions for the upper layer where the DREGs estimator for  $\phi_2$  only has one. This is due to the opposite factorization (opposite hierarchical dependency structure) of the variational posterior and the prior. This cyclic dependence is also the reason why we cannot replace all indirect score functions with DREGs and GDREGs gradients. Double reparameterization of one of the indirect score functions, leads to another indirect score function, whose double-reparameterization in turn leads back to the first indirect score function but with a different pre-factor.

Again, note that the indirect score functions are computed automatically in our surrogate losses, [App. E](#), and we do not need to manually trace the dependency structure or derive them.

## E. Surrogate losses to implement the DREGs and GDREGs estimators for IWAE objectives

As we discussed in [Sec. 4.1](#) and similar to [Tucker et al. \(2019\)](#), we use surrogate loss functions to compute the gradients w.r.t. the likelihood, proposal, and prior parameters. That is, we use different losses, such that backpropagation results in the respective gradient estimator. While [Tucker et al. \(2019\)](#) use a single surrogate loss to compute the gradient estimators for all parts of the objective, we choose to use separate surrogate losses for each of the three parameter groups (likelihood, variational posterior, prior). In principle, we could combine them into a single loss, but in order to keep presentation simple we keep them separate. Computationally this does not make a difference as modern deep learning frameworks avoid duplicate computation.

For the likelihood parameters, we use the regular (negative) IWAE objective [Eq. \(3\)](#) as a loss. That is, the gradient estimator for the likelihood parameters is given by the gradient of the negative IWAE objective.

To construct the other surrogate losses we need to stop the gradients at various points in the computation graph. In the following, we use the shorthand notation  $\underbrace{\phantom{x}}_{\text{wavy red X}}$  to indicate that we stop gradients into the underlined parts of an expression. Where it might be ambiguous, or to highlight where we do *not* stop gradients, we use the shorthand  $\overrightarrow{\phantom{x}}$  to indicate that gradients flow. For example,  $f(\underbrace{\phi}_{\text{wavy red X}}, \overrightarrow{\theta})$  means that we backpropagate gradients into  $\phi$  but not into  $\theta$ .

### E.1. DREGs for variational posterior parameters $\phi$

#### E.1.1. SINGLE STOCHASTIC LAYER

Here we reproduce part of the surrogate loss for the variational parameters  $\phi$  by [Tucker et al. \(2019\)](#) for the single stochastic layer case:

$$\begin{aligned} L_{\text{DREGs}}(\phi) &= \sum_{k=1}^K \underbrace{\tilde{w}_k^2}_{\text{wavy red X}} \left( \log p_{\lambda}(\mathbf{x} | \overrightarrow{\mathbf{z}_k}) + \log p_{\beta(\boldsymbol{\theta})}(\overrightarrow{\mathbf{z}_k}) - \log q_{\alpha(\phi)}(\overrightarrow{\mathbf{z}_k}) \right) \\ \overrightarrow{\mathbf{z}_k} &= \mathcal{T}_q(\boldsymbol{\epsilon}_k; \overrightarrow{\phi}) \quad \boldsymbol{\epsilon}_k \sim q(\boldsymbol{\epsilon}_k) \end{aligned} \quad (\text{E.1})$$

That is, we sample  $\mathbf{z}_k \sim q_{\phi}(\mathbf{z}_k | \mathbf{x})$  as usual (by reparameterizing independent noise variables  $\boldsymbol{\epsilon}_k$ ) but stop the gradients of the parameters that parameterize the distributions when evaluating their densities,  $\log q_{\alpha(\phi)}(\overrightarrow{\mathbf{z}_k})$ . In addition we stop the gradients around the normalized importance weights  $\tilde{w}_k$ . Differentiating  $L_{\text{DREGs}}$  w.r.t. the proposal parameters  $\phi$  yields the DREGs estimator [Eq. \(11\)](#). Note that we do not explicitly stop gradients into  $\lambda$  or  $\boldsymbol{\theta}$  because we use separate surrogate losses for those parameter groups. If we were to use a combined loss, we would potentially have to stop gradients into these parameters as well, depending on the estimator used.

To practically implement this surrogate loss, we use two copies of the variational posterior distribution. An unaltered one

(no stopped gradients) to sample  $\mathbf{z}$  and one with gradients into the proposal parameters stopped to evaluate the log densities. The stopped gradient makes sure that we do not obtain a direct score function as we have doubly-reparameterized it.

Note that for single-stochastic-layer models we could also stop the gradients of the distribution parameters  $\alpha$  instead as they only depend on  $\phi$ . We emphasize that this is not possible for hierarchical models as this would eliminate the indirect score functions and thus produce potentially biased gradients.

### E.1.2. MULTIPLE STOCHASTIC LAYERS

For multiple layers, the surrogate loss for the DREGs estimator Eq. (C.8) is given by:

$$\begin{aligned}
 L_{\text{DREGs}}(\phi) &= \sum_{k=1}^K \tilde{w}_k^2 \log w_k \\
 \log w_k &= \log p_{\lambda}(\mathbf{x} | \mathbf{z}_{k1}, \dots, \mathbf{z}_{kL}) + \sum_{l=1}^L \log p_{\beta_l}(\text{pa}_{\beta}(l); \theta_l)(\mathbf{z}_{kl}) - \sum_{l=1}^L \log q_{\alpha_l}(\text{pa}_{\alpha}(l); \phi_l)(\mathbf{z}_{kl}) \\
 \mathbf{z}_{kl} &= \mathcal{T}_{q_l}(\epsilon_{kl}; \alpha_l(\text{pa}_{\alpha}(l); \phi_l)) \quad \epsilon_{kl} \sim q(\epsilon_{kl})
 \end{aligned} \tag{E.2}$$

Again, we do not explicitly stop gradients into  $\lambda$  or  $\theta_l$  as we only take gradients w.r.t.  $\phi_l$ .

The indirect score functions arise due to the indirect dependence of the distribution parameters  $\alpha_l(\text{pa}_{\alpha}(l); \phi_l)$  and  $\beta_l(\text{pa}_{\beta}(l); \theta_l)$  on the parent latent variables  $\text{pa}_{\alpha}(l)$  and  $\text{pa}_{\beta}(l)$ , respectively. Note how the former depends on the hierarchical structure of the variational posterior, whereas the latter depends on the hierarchical structure of the prior.

To implement this surrogate loss effectively, we again use two copies of the variational posterior distribution. One un-altered one (without stopped gradients) from which we sample the individual reparameterized  $\mathbf{z}_{kl}$  and through which gradients can flow; we use these samples to evaluate densities at and to parameterize the distribution parameters at subsequent layers. Derivatives w.r.t.  $\phi_l$  will then give rise to pathwise gradients and indirect score functions. We use the second copy of the variational posterior, where we have stopped the parameters  $\phi_l$ , to evaluate the density at for the log importance weights in the last summand of Eq. (E.2).

## E.2. GDREGs for prior parameter $\theta$

### E.2.1. SINGLE STOCHASTIC LAYER

$$\begin{aligned}
 L_{\text{GDREGs}}(\theta) &= \sum_{k=1}^K \tilde{w}_k \log p_{\lambda}(\mathbf{x} | \mathbf{z}_k) - \tilde{w}_k^2 \left( \log p_{\lambda}(\mathbf{x} | \mathbf{z}_k) + \log p_{\beta}(\theta)(\mathbf{z}_k) - \log q_{\alpha}(\phi)(\mathbf{z}_k) \right) \\
 \mathbf{z}_k &= \mathcal{T}_p(\tilde{\epsilon}_k; \theta) \\
 \tilde{\epsilon}_k &= \mathcal{T}_p^{-1}(\mathcal{T}_q(\epsilon_k; \phi); \theta) \quad \epsilon_k \sim q(\epsilon_k)
 \end{aligned} \tag{E.3}$$

Taking the derivative of Eq. (E.3) w.r.t.  $\theta$  gives rise to the GDREGs estimator for the single stochastic layer IWAE objective. As explained in Sec. 4, we need to re-express  $\mathbf{z}_k$  such that its path depends on  $\theta$ . In effect, we first sample  $\mathbf{z}_k = \mathcal{T}_q(\epsilon_k; \phi)$ , then compute the new noise variable  $\tilde{\epsilon}_k = \mathcal{T}_p^{-1}(\mathbf{z}_k; \theta)$ , and re-compute  $\mathbf{z}_k = \mathcal{T}_p(\tilde{\epsilon}_k; \theta)$ . Note that we have to stop gradients into the noise variables  $\tilde{\epsilon}_k$  to obtain the correct gradient estimator. This explains the `stop_grad` in Fig. 2.

As above, we do not explicitly stop gradients into  $\lambda$  and  $\phi$  as we use separate losses for these parameter groups and only compute gradients of Eq. (E.3) w.r.t.  $\theta$ .

To effectively implement this loss, we use two copies of the prior distribution. One that we implement as a normalizing flow and a second one with stopped gradients into the parameters. We then proceed as follows:

- Compute the new noise variables  $\tilde{\epsilon}_k$  by using the inverse flow  $\mathcal{T}_p^{-1}$  on the samples  $\mathbf{z}_k$  from the variational posterior.
- Stop the gradients into  $\tilde{\epsilon}_k$ .
- Use the forward flow  $\mathcal{T}_p(\tilde{\epsilon}_k; \theta)$  to re-compute  $\mathbf{z}_k$  but with path dependent on  $\theta$ . These samples when derived w.r.t.  $\theta$  will give rise to the pathwise gradients.

- Use the second copy of the prior (with stopped gradients into its parameters) to evaluate the log density at the samples  $z_k$ . The stopped gradients make sure that we do not obtain the direct score function.

### E.2.2. MULTIPLE STOCHASTIC LAYERS

For multiple stochastic layers the surrogate loss that gives rise to the GDREGs estimator Eq. (C.16) is given by:

$$\begin{aligned}
 L_{\text{GDREGs}}(\theta) &= \sum_{k=1}^K \tilde{w}_k \log p_\lambda(x | z_{k1}, \dots, z_{kL}) - \tilde{w}_k^2 \log w_k \\
 \log w_k &= \log p_\lambda(x | z_{k1}, \dots, z_{kL}) + \sum_{l=1}^L \log p_{\beta_l}(\text{pa}_{\beta}(l); \theta_l)(z_{kl}) - \sum_{l=1}^L \log q_{\alpha_l}(\text{pa}_{\alpha}(l); \phi_l)(z_{kl}) \\
 z_{kl} &= \mathcal{T}_{p_l}(\tilde{\epsilon}_{kl}; \beta_l(\text{pa}_{\beta}(l), \theta_l)) \\
 \tilde{\epsilon}_{kl} &= \mathcal{T}_{p_l}^{-1}(z_{kl}^{(q)}; \beta_l(\text{pa}_{\beta}(l), \theta_l)) \\
 z_{kl}^{(q)} &= \mathcal{T}_{q_l}(\epsilon_{kl}; \alpha_l(\text{pa}_{\alpha}(l), \phi_l)) \quad \epsilon_{kl} \sim q(\epsilon_{kl})
 \end{aligned} \tag{E.4}$$

As for the single layer case, we need to re-express variational posterior samples  $z_{kl}$  as if they were sampled from the prior. To obtain the correct gradients, we again have to stop gradients into the new noise variables  $\tilde{\epsilon}_{kl}$ , also see Fig. 3.

As for hierarchical DREGs, the indirect score functions stem from the second and third term of  $\log w_k$  and arise because the distribution parameters  $\alpha_l$  and  $\beta_l$  depend on the “parent” stochastic layers.

As before we use two copies of the prior distribution, one with regular gradients that is set up as a flow, and a second with stopped gradients into the parameters. This allows us to implement the GDREGs estimator regardless of the model structure.

## F. Implementation details

In our implementations we use NumPy (Harris et al., 2020), JAX (Bradbury et al., 2018), Haiku (Hennigan et al., 2020), as well as tensorflow probability and tensorflow distributions (Dillon et al., 2017).

**Stopping gradients.** All major frameworks allow for gradients to be stopped or interrupted. For example, in TensorFlow (Abadi et al., 2016) we can use `tf.stop_gradient` and in JAX we can use `jax.lax.stop_gradient`. To implement stopped gradients w.r.t. the parameters of a distribution we use `haiku.experimental.custom_getter` contexts, which allow us to manipulate the parameters before they are used to construct the respective networks; in this case we use the context to stop gradients.

**Re-expressing samples.** To re-express variational posterior samples as if they came from the prior, we directly implement the computation flow as it is described in e.g. Figs. 2 and 3 and detailed in App. E.2.

In the code listings Listings 1 and 2 we provide (pseudo-)code for a simple implementation of the surrogate objectives. Listing 1 contains the import statements as well as the function and class definitions to create parameterized distributions that allow for

1. stopping the gradients into their parameters, and
2. re-expressing samples using the bijector interface in tensorflow probability.

In Listing 2 we implement surrogate objectives for the naive and the GDREGs estimators of the cross-entropy. More specifically, we wish to estimate:

$$\nabla_{\theta}^{\text{TD}} \mathcal{L}^{ce} = \nabla_{\theta}^{\text{TD}} \mathbb{E}_{z \sim q_{\phi}(z|x)} [\log p_{\theta}(z)]. \tag{F.1}$$

The naive estimator (using a single Monte Carlo sample) is given by

$$\widehat{\nabla}_{\theta}^{\text{naive}} \mathcal{L}^{ce} = \nabla_{\theta} \log p_{\theta}(z) \quad z \sim q_{\phi}(z). \tag{F.2}$$

The corresponding GDREGs estimator (again using a single MC sample) is given by Eq. (21):

$$\widehat{\nabla}_{\theta}^{\text{GDREGs}} \mathcal{L}^{ce} = \nabla_{\mathbf{z}} \log \frac{q_{\phi}(\mathbf{z})}{p_{\theta}(\mathbf{z})} \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) \Big|_{\tilde{\epsilon}=\mathcal{T}_p^{-1}(\mathbf{z}, \theta)} \quad \mathbf{z} \sim q_{\phi}(\mathbf{z}). \quad (21)$$

Listing 1: Function and class definitions necessary to define the surrogate objectives for the DREGs and GDREGs estimators.

```

1 from typing import List
2 import haiku as hk
3 import jax
4 import jax.lax as lax
5 import jax.numpy as jnp
6 from tensorflow_probability.substrates import jax as jtfd
7 jtfd = jtfd.distributions
8
9
10 def stop_grad_getter(next_getter, value, _):
11     """A custom getter that stops gradients of parameters."""
12     return lax.stop_gradient(next_getter(value))
13
14
15 def reparameterize_as_if_from(p: jtfd.TransformedDistribution,
16                               z_q: jnp.ndarray) -> jnp.ndarray:
17     """Reparameterize samples z_q as if they were sampled from p.
18
19     Transforms: z_q -> stop_gradient(noise) -> z_q_as_if_from_p
20     This transformation leaves the numerical value unchanged but alters the
21     gradients.
22
23     Args:
24         p: a TransformedDistribution object
25         z_q: Samples to be reparameterized.
26
27     Returns:
28         Reparameterized samples.
29     """
30
31     eps = p.bijector.inverse(z_q)
32     eps = jax.lax.stop_gradient(eps)
33     return p.bijector.forward(eps)
34
35
36 class ConditionalNormal(hk.Module):
37     """A Normal distribution that is conditioned through an MLP."""
38
39     def __init__(
40         self,
41         output_size: int,
42         hidden_layer_sizes: List[int],
43         name: str = "conditional_normal"):
44         """Creates a conditional Normal distribution.
45
46         Args:
47             output_size: The dimension of the random variable.
48             hidden_layer_sizes: The sizes of the hidden layers of the fully connected
49                 network used to condition the distribution on the inputs.
50             name: The name of this distribution.
51         """
52         super(ConditionalNormal, self).__init__(name=name)
53         self.name = name
54         self.fcnet = hk.nets.MLP(
55             output_sizes=hidden_layer_sizes + [2 * output_size],
56             activation=jnp.tanh,

```

```

57     activate_final=False,
58     with_bias=True,
59     name=name + "_fcnet")
60
61 def condition(self, inputs):
62     """Computes the parameters of a normal distribution based on the inputs."""
63     outs = self.fcnet(inputs)
64     mu, sigma = jnp.split(outs, 2, axis=-1)
65     sigma = jax.nn.softplus(sigma)
66     return mu, sigma
67
68 def __call__(self, inputs, **kwargs):
69     """Creates a normal distribution conditioned on the inputs."""
70     # Optional 'stop_gradient_params' argument stops the parameters
71     # of the distribution.
72     if kwargs.get("stop_gradient_params", False):
73         with hk.experimental.custom_getter(stop_grad_getter):
74             mu, sigma = self.condition(inputs)
75     else:
76         mu, sigma = self.condition(inputs)
77
78     # Optional 'as_flow' argument parameterizes the distribution as a flow
79     # to have access to 'Bijector.inverse' and 'Bijector.forward'
80     # to use with the function 'reparameterize_as_if_from'
81     if kwargs.get("as_flow", False):
82         bijector = jtfp.bijectors.Chain(
83             [jtfp.bijectors.Shift(shift=mu), jtfp.bijectors.Scale(scale=sigma)])
84         base = jtfd.Normal(loc=jnp.zeros_like(mu), scale=jnp.ones_like(sigma))
85         return jtfd.TransformedDistribution(
86             distribution=base, bijector=bijector, name=self.name + "_flow")
87
88     return jtfd.Normal(loc=mu, scale=sigma)

```

Listing 2: Code to implement the surrogate objective for the naive estimator of the cross-entropy as well as for the GDREGs estimator Eq. (21). Computing derivatives w.r.t. the parameters of the prior  $p$  using automatic differentiation gives rise to the correct expressions for the estimators.

```

1 # Create distributions
2 # Inputs 'x' and context 'c'
3 q = ConditionalNormal(x)
4 p = ConditionalNormal(c)
5 q_stop = ConditionalNormal(x, stop_gradient_params=True)
6 p_stop = ConditionalNormal(c, stop_gradient_params=True)
7 p_flow = ConditionalNormal(c, as_flow=True)
8
9 # Sample from the variational posterior
10 z_q = q.sample(sample_shape=[num_samples], seed=hk.next_rng_key()) # [k, bs, z]
11
12 # Reparameterize the samples from q as if they were sampled from p
13 z_q_as_p = reparameterize_as_if_from(p_flow, z_q)
14
15 # Cross-entropy surrogate losses
16 cross_entropy_naive = p.log_prob(z_q)
17 cross_entropy_gdregs = q_stop.log_prob(z_q_as_p) - p_stop.log_prob(z_q_as_p)

```

## G. Experimental details and additional results

In this section we provide additional experimental details as well as additional results.

### G.1. Illustrative example

As discussed in the main text, we use a 2-layer linear VAE inspired by the single layer example of [Rainforth et al. \(2018\)](#); [Tucker et al. \(2019\)](#). We use a top-down generative model  $z_2 \rightarrow z_1 \rightarrow \mathbf{x}$ , with  $z_1, z_2, \mathbf{x} \in \mathbb{R}^D$  and  $D = 5$ . The hierarchical prior is given by  $z_2 \sim \mathcal{N}(0, \mathbb{I})$ ,  $z_1|z_2 \sim \mathcal{N}(\boldsymbol{\mu}_\theta(z_2), \boldsymbol{\sigma}_\theta^2(z_2))$ , and the likelihood is given by  $\mathbf{x}|z_1 \sim \mathcal{N}(z_1, \mathbb{I})$ . We choose a bottom-up variational posterior that factorizes as:  $q_{\phi_1}(z_1|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\phi_1}(\mathbf{x}), \boldsymbol{\sigma}_{\phi_1}^2(\mathbf{x}))$  and  $q_{\phi_2}(z_2|z_1) = \mathcal{N}(\boldsymbol{\mu}_{\phi_2}(z_1), \boldsymbol{\sigma}_{\phi_2}^2(z_1))$ . All functions  $\boldsymbol{\mu}_*$  and  $\boldsymbol{\sigma}_*$  are given by linear functions with weights and biases; the likelihood and the upper layer of the prior do not have any learnable parameters.

To generate data, we sample 512 datapoints from the model where we have set  $\boldsymbol{\mu}_\theta(z_2) = z_2$  and  $\boldsymbol{\sigma}_\theta(z_2) = 1$ .

We then train the parameters  $\phi$  and  $\theta$  in all linear layers using SGD on the IWAE objective til convergence. We then evaluate the gradient variance and gradient signal-to-noise ratio for each estimator. For the proposal parameters  $\phi$  we compare DREGs to the naive score function (labelled as IWAE) and to STL; for the prior parameters  $\theta$  we compare GDREGs to IWAE.

In [Fig. 4](#) in the main paper we show the average gradient variance and gradient signal-to-noise ratio (SNR). The average is taken over all parameters of either the variational posterior or the prior. The gradient variance and gradient SNR for individual parameters exhibit the same qualitative behaviour.

### G.2. Conditional and unconditional image modelling

For conditional and unconditional image modelling, we use VAEs with one or multiple stochastic layers, where the generative path is top-down and the inference path is bottom-up, as specified in [Eq. \(23\)](#), which we reproduce here for convenience:

$$\begin{aligned} q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c}) &= q_{\phi_1}(z_1|\mathbf{x}, \mathbf{c}) \prod_{l=2}^L q_{\phi_l}(z_l|z_{l-1}, \mathbf{x}, \mathbf{c}) \\ p_\theta(\mathbf{z}|\mathbf{c}) &= p_{\theta_L}(z_L|\mathbf{c}) \prod_{l=1}^{L-1} p_{\theta_l}(z_l|z_{l+1}, \mathbf{c}) \\ p_\lambda(\mathbf{x}|\mathbf{z}) &= p_\lambda(\mathbf{x}|z_1, \dots, z_L). \end{aligned} \tag{23}$$

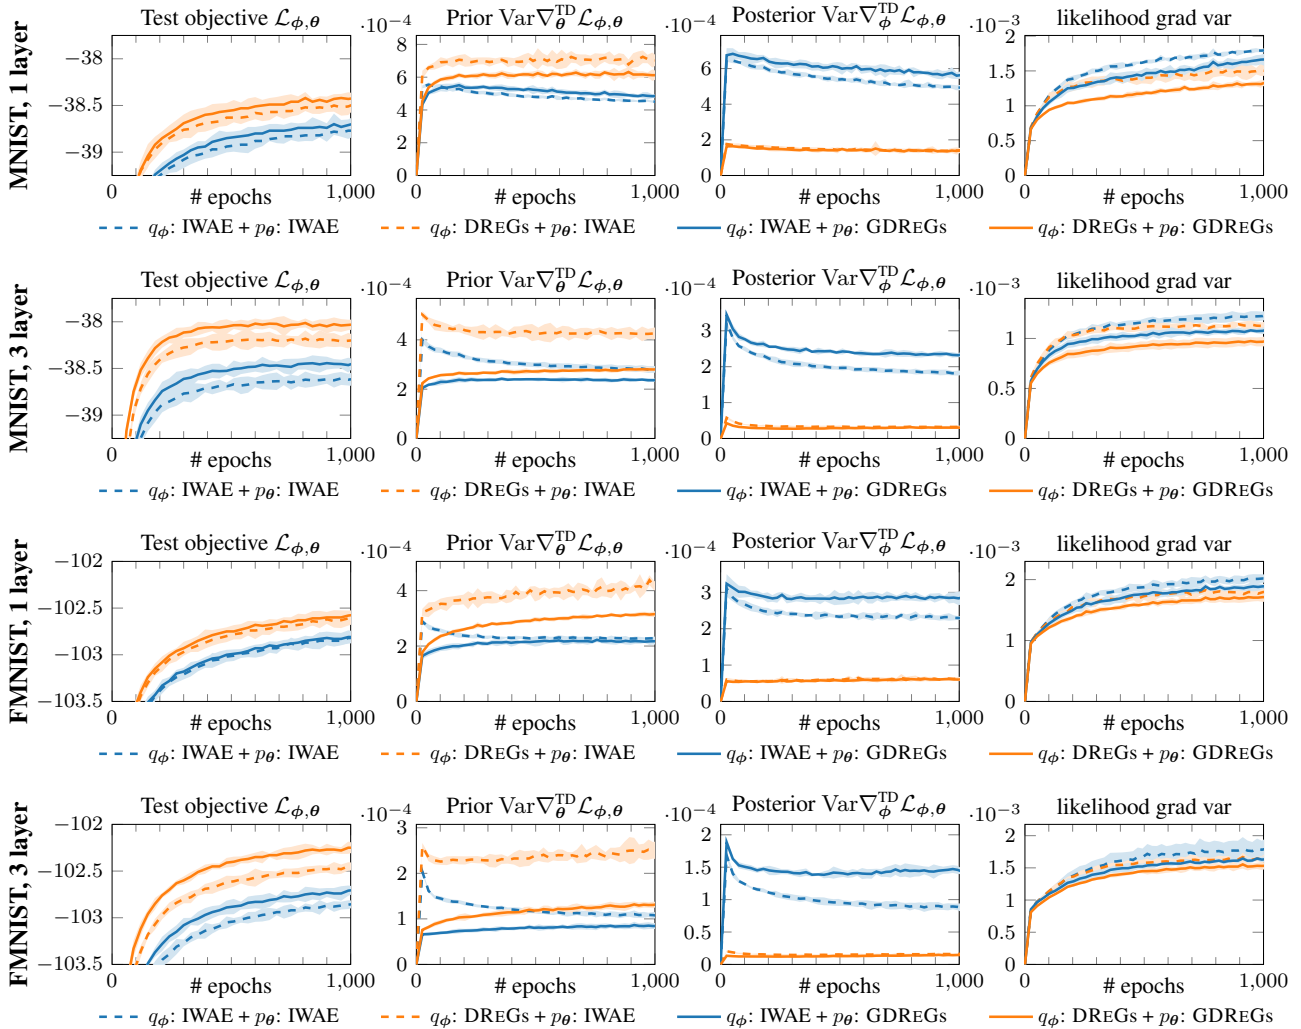
For conditional image modelling, we predict the bottom half of an image given its top half as in [Tucker et al. \(2019\)](#), providing the top half as an additional context input  $\mathbf{c}$  to the prior and variational posterior. Given the above model structure,  $\text{pa}_\alpha(l) = z_{k(l-1)}$  and  $\text{pa}_\beta(l) = z_{k(l+1)}$ .

Each conditional distribution in [Eq. \(23\)](#) is given by an MLP of 2 hidden layers of 300  $\tanh$  units each. If a distribution has multiple inputs, we concatenate them along the feature dimension. The prior and variational posterior are all given by diagonal Gaussian distributions, whereas the likelihood is given by a Bernoulli distribution. The unconditional prior distribution in the uppermost layer  $p_{\theta_L}(z_L)$  is given by a standard Normal distribution,  $p_{\theta_L}(z_L) = \mathcal{N}(0, \mathbb{I})$ . All latents  $z_l$  are 50 dimensional.

To avoid overfitting, we use dynamically binarized versions of the datasets. We use the Adam optimizer with default learning rate  $3 \cdot 10^{-4}$  and default parameters  $b_1 = 0.9$ ,  $b_2 = 0.999$ , and  $\epsilon = 10^{-8}$ . We use a batch size of 64 and  $K = 64$  importance samples for training and evaluation. Note that for testing we report test objective values rather than an estimate of  $\log p(\mathbf{x})$  by using a large number of importance samples. We do this as we are interested in the relative behaviour of the estimators.

In [Fig. G.1](#) we provide further plots that show the evolution of the test objective and the gradient variance of the prior, the variational posterior, as well as the likelihood throughout training on *conditional modelling* of MNIST and FashionMNIST (predict bottom half given top half). As in the main paper we find that using GDREGs for the prior instead of the naive estimator (denoted as IWAE) always improves performance on the test objective regardless of the estimator for the variational posterior. We also note that GDREGs always reduces gradient variance for the prior early in training and also typically throughout training, especially for deeper models and when combined with the DREGs estimator for the variational posterior.

Interestingly, using the GDREGs estimator for the prior leads to an increase in the gradient variance for the variational posterior when we use the naive estimator but similar or even lower posterior gradient variance when combined with the DREGs estimator (third column in [Fig. G.1](#)). It always leads to a slight improvement in gradient variance for the likelihood parameters (fourth column in [Fig. G.1](#)). We hypothesize that this is the case because lower gradient variance in for one set

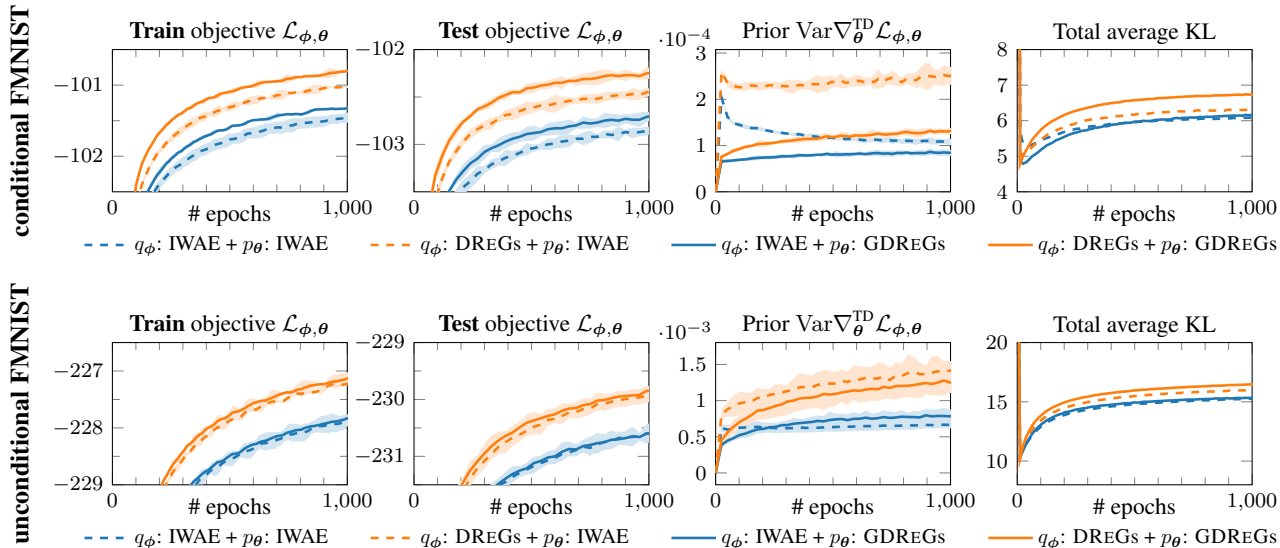


**Figure G.1:** Additional results for *conditional* image modelling with VAEs on MNIST and FashionMNIST (FMNIST) with 1 stochastic layer and 3 stochastic layers. For each task we show the test objective, the prior gradient variance, the variational posterior gradient variance, and the likelihood gradient variance. Means over 5 reruns; shaded areas denote  $\pm 1.96$  standard deviations  $\sigma$ .

of parameters makes it easier to estimate the gradients for another set as a secondary effect.

In [Fig. G.2](#) we show that the training objective behaves qualitatively similar to the test objective in that the applying the DREGs or GDREGs estimator results in improved objective values in the same way, regardless of whether we consider the training or test objective. That is, our hierarchical extension of DREGs results in a better training objective values for both conditional and unconditional tasks. GDREGs is particularly helpful for conditional tasks.

Moreover, in the main text we had hypothesized that GDREGs performs better on conditional image modelling tasks than unconditional tasks because access to the context makes the variational posterior and the prior more similar. In the case of analytically computed cross-entropy in [App. H](#), we derive that the GDREGs estimator outperforms a naive estimator of the score function in terms of gradient variance when the posterior and prior are similar. We hypothesize that this also holds more generally for the IWAE objective, where we cannot compute the gradient variance in closed form but only estimate it empirically. To investigate this, we computed the total average KL (rightmost column in [Fig. G.2](#)) over all latent variables and found that it is indeed lower for conditional than unconditional modelling, which indicates that the distributions are closer together in this case.




**Figure G.2:** Train objective (leftmost column) and total average KL (rightmost column) in addition to the test objective and prior gradient variance for conditional and unconditional FashionMNIST on a model with 3 stochastic layers.

### G.3. Offline evaluation of the DREGs and GDREGs estimator

In the image modelling experiments in Sec. 5.2 we discussed the gradient variance of the different estimators for the posterior and prior parameters. However, we only analyzed estimators *online* on their respective runs; that is, in Fig. 6 the prior and posterior gradient variance shown corresponds to the variance of the estimator also used during training.

Here, we provide an ablation study for conditional image modelling on MNIST with a 2-layer VAE where we evaluate the different estimators offline; that is, for each combination of estimators used for training we also show the gradient variance of the other estimators.




**Figure G.3:** Gradient variance of the **posterior** parameters  $\phi$  as estimated with the naive (IWAE) estimator (—) and the DREGs estimator (—) for different combinations of estimators used during training as indicated by the sub-figure title. For example, in the leftmost plot we compare the posterior gradient variance as estimated by the naive (IWAE) estimator to that of the DREGs estimator on an experiment where we trained both the posterior parameters  $\phi$  as well as the prior parameters  $\theta$  with the naive (IWAE) estimator.

In Fig. G.3 we consider the gradient variance w.r.t. the **variational posterior** parameters  $\phi$  and compare the naive (IWAE) estimator to the DREGs estimator. We find that, regardless which combination of estimators has been used during training, the DREGs estimator *always* results in a better (lower) gradient variance than the naive estimator.

Similarly, in Fig. G.4 we consider the gradient variance w.r.t. the **prior** parameters  $\theta$  and compare the naive (IWAE) estimator to the GDREGs estimator. We find that generally the GDREGs gradient estimates have lower variance than the naive (IWAE) estimates. However, when we use the naive estimator for the prior parameters during training, this reduction is smaller and may only be present in the beginning of training. However, consistently, the GDREGs gradient estimates





**Figure G.4:** Gradient variance of the **prior** parameters  $\theta$  as estimated with the naive (IWAE) estimator (—) and the GDREGs estimator (—) for different combinations of estimators used during training as indicated by the sub-figure title. For example, in the leftmost plot we compare the prior gradient variance as estimated by the naive (IWAE) estimator to that of the GDREGs estimator on an experiment where we trained both the posterior parameters  $\phi$  as well as the prior parameters  $\theta$  with the naive (IWAE) estimator.

have lower variance when we use the DREGs estimator during training to estimate the variational posterior parameters.

## H. The cross-entropy for Gaussian distributions

Here, we investigate the properties of the GDREGs estimator compared to the naive estimator in a setting where all quantities of interest can be computed in closed form, the cross-entropy of two Gaussian distributions,  $q_{\phi}(z) = \mathcal{N}(z; \mu_q, \sigma_q)$  and  $p_{\theta}(z) = \mathcal{N}(z; \mu_p, \sigma_p)$ .

The negative cross-entropy is given by

$$\mathcal{L}_{\phi, \theta}^{\text{ce}} = \mathbb{E}_{z \sim q_{\phi}(z)} [\log p_{\theta}(z)] = \frac{1}{2} \log(2\pi) + \log \sigma_p + \frac{\sigma_q^2 + (\mu_p - \mu_q)^2}{2\sigma_p^2} \quad (\text{H.1})$$

Note that in this analytic case, we can compute the gradients without having to sample. However, in the following we want to compare the naive (score function) estimator to the GDREGs estimator.

The naive estimator of this score function is given by:

$$\widehat{\nabla}_{\theta}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{ce}} = \nabla_{\theta} \log p_{\theta}(z); \quad z \sim q_{\phi}(z) \quad (\text{H.2})$$

while the GDREGs estimator is given by (see Eq. (21)):

$$\widehat{\nabla}_{\theta}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{ce}} = \nabla_z \log \frac{q_{\phi}(z)}{p_{\theta}(z)} \nabla_{\theta} \mathcal{T}_p(\tilde{\epsilon}; \theta) \Big|_{\tilde{\epsilon} = \mathcal{T}_p^{-1}(z, \theta)}; \quad z \sim q_{\phi}(z) \quad (21)$$

### H.1. Gradient variance of the estimators

In the case under consideration, the parameters  $\theta$  are given by the mean and variance of the prior,  $\theta = \{\mu_p, \sigma_p\}$ , and we can compute both the expectation as well as the variance of these gradient estimators in closed form.

**Naive estimator**

$$\mathbb{E}_{q_{\phi}(z)} \left[ \widehat{\nabla}_{\mu_p}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\mu_q - \mu_p}{\sigma_p^2} \quad (\text{H.3})$$

$$\text{Var}_{q_{\phi}(z)} \left[ \widehat{\nabla}_{\mu_p}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\sigma_q^2}{\sigma_p^4} \quad (\text{H.4})$$

$$\mathbb{E}_{q_{\phi}(z)} \left[ \widehat{\nabla}_{\sigma_p}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\sigma_q^2 - \sigma_p^2}{\sigma_p^3} + \frac{(\mu_q - \mu_p)^2}{\sigma_p^3} \quad (\text{H.5})$$

$$\text{Var}_{q_{\phi}(z)} \left[ \widehat{\nabla}_{\sigma_p}^{\text{naive}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = 2 \frac{\sigma_q^4}{\sigma_p^6} + 4 \sigma_q^2 \frac{(\mu_q - \mu_p)^2}{\sigma_p^6} \quad (\text{H.6})$$

Note that all operations are element-wise.

**The proposed GDREGs estimator**

$$\mathbb{E}_{q_\phi(z)} \left[ \widehat{\nabla}_{\mu_p}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\mu_q - \mu_p}{\sigma_p^2} \quad (\text{H.7})$$

$$\text{Var}_{q_\phi(z)} \left[ \widehat{\nabla}_{\mu_p}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\sigma_q^2 (\sigma_p^2 - \sigma_q^2)^2}{\sigma_p^4 \sigma_q^4} \quad (\text{H.8})$$

$$\mathbb{E}_{q_\phi(z)} \left[ \widehat{\nabla}_{\sigma_p}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{\sigma_q^2 - \sigma_p^2}{\sigma_p^3} + \frac{(\mu_q - \mu_p)^2}{\sigma_p^3} \quad (\text{H.9})$$

$$\text{Var}_{q_\phi(z)} \left[ \widehat{\nabla}_{\sigma_p}^{\text{GDREGs}} \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = 2 \frac{(\sigma_q^2 - \sigma_p^2)^2}{\sigma_p^6} + \frac{(\sigma_p^2 - 2\sigma_q^2)^2 (\mu_q - \mu_p)^2}{\sigma_q^2 \sigma_p^6} \quad (\text{H.10})$$

Note that all operations are element-wise.

Both estimators have equal expectation; this is because GDREGs is an unbiased estimator.

Comparing the variances in Eq. (H.4) and Eq. (H.8) we note that the GDREGs estimator has a lower gradient variance than the naive estimator for the mean parameters  $\mu_p$  if

$$\sigma_p^2 \leq 2\sigma_q^2. \quad (\text{H.11})$$

Similarly, comparing Eq. (H.6) and Eq. (H.10), we find that the GDREGs estimator has lower variance than the naive estimator for the variance parameters  $\sigma_p$  if

$$\sigma_p^2 \leq 4\sigma_q^2 \left( 1 - \frac{\sigma_q^2}{(\mu_p - \mu_q)^2 + 2\sigma_q^2} \right) \quad (\text{H.12})$$

In the case of  $\mu_p = \mu_q$ , this also reduces to  $\sigma_p^2 \leq 2\sigma_q^2$ .

Thus, we expect the GDREGs estimator to perform better than the naive estimator when  $p_\theta(z)$  and  $q_\phi(z)$  are close together.

**H.2. Constructing the optimal control variate**

Because the GDREGs estimators and the naive estimators have the same expectation, we can build a control variate out of their difference:

$$\left[ \widehat{\nabla}_\theta^{\text{naive}} + \alpha (\widehat{\nabla}_\theta^{\text{GDReG}} - \widehat{\nabla}_\theta^{\text{naive}}) \right] \mathcal{L}_{\phi, \theta}^{\text{ce}}. \quad (\text{H.13})$$

We can then compute its optimal strength  $\alpha^*$ , by minimizing its variance,

$$\alpha^* = \arg \min_{\alpha} \text{Var}_{q_\phi(z)} \left[ \left[ \widehat{\nabla}_\theta^{\text{naive}} + \alpha (\widehat{\nabla}_\theta^{\text{GDReG}} - \widehat{\nabla}_\theta^{\text{naive}}) \right] \mathcal{L}_{\phi, \theta}^{\text{ce}} \right]. \quad (\text{H.14})$$

We find that:

$$\alpha_\mu^* = \frac{\sigma_q^2}{\sigma_p^2} \quad (\text{H.15})$$

$$\alpha_\sigma^* = \frac{2\sigma_q^2 (\mu_p - \mu_q)^2 + \sigma_q^2}{\sigma_p^2 (\mu_p - \mu_q)^2 + 2\sigma_q^2} \quad (\text{H.16})$$

$$\text{Var}_{q_\phi(z)} \left[ \left[ \widehat{\nabla}_{\mu_p}^{\text{naive}} + \alpha_\mu^* (\widehat{\nabla}_{\mu_p}^{\text{GDReG}} - \widehat{\nabla}_{\mu_p}^{\text{naive}}) \right] \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = 0 \quad (\text{H.17})$$

$$\text{Var}_{q_\phi(z)} \left[ \left[ \widehat{\nabla}_{\sigma_p}^{\text{naive}} + \alpha_\sigma^* (\widehat{\nabla}_{\sigma_p}^{\text{GDReG}} - \widehat{\nabla}_{\sigma_p}^{\text{naive}}) \right] \mathcal{L}_{\phi, \theta}^{\text{ce}} \right] = \frac{2\sigma_q^4 (\mu_q - \mu_p)^2}{\sigma_p^6 (\mu_q - \mu_p)^2 + 2\sigma_q^2} \quad (\text{H.18})$$

Note that the expression for the optimal strength has different form for the mean  $\mu_p$  and variance  $\sigma_p$  parameters. Moreover, note that the analytic estimator has zero gradient variance whereas our estimator with control variate still has non-zero gradient variance for the variance parameters  $\sigma_p$ .

## Generalized Doubly-Reparameterized Gradient Estimators

---

This optimal estimator holds for a single layer VAE, where both the variational posterior as well as the prior are Gaussian. In a hierarchical model where both the prior and the posterior are factorized top-down, the same derivation holds for the lowest stochastic layer (similar to how semi-analytic approximations for the conditional KL can be derived in this case). Unfortunately, in this case the expectations cannot be computed in closed form anymore.