
On Signal-to-Noise Ratio Issues in Variational Inference for Deep Gaussian Processes

Tim G. J. Rudner^{*1} Oscar Key^{*2} Yarin Gal¹ Tom Rainforth¹

Abstract

We show that the gradient estimates used in training Deep Gaussian Processes (DGPs) with importance-weighted variational inference are susceptible to signal-to-noise ratio (SNR) issues. Specifically, we show both theoretically and via an extensive empirical evaluation that the SNR of the gradient estimates for the latent variable’s variational parameters *decreases* as the number of importance samples *increases*. As a result, these gradient estimates degrade to pure noise if the number of importance samples is too large. To address this pathology, we show how doubly-reparameterized gradient estimators, originally proposed for training variational autoencoders, can be adapted to the DGP setting and that the resultant estimators completely remedy the SNR issue, thereby providing more reliable training. Finally, we demonstrate that our fix can lead to consistent improvements in the predictive performance of DGP models.

1. Introduction

Deep Gaussian Processes (DGPs) are a powerful class of probabilistic models for supervised learning tasks (Damianou & Lawrence, 2013; Bui et al., 2016; Salimbeni & Deisenroth, 2017). They are a multi-layer hierarchical generalization of conventional Gaussian processes (GPs, Rasmussen & Williams (2006)), themselves flexible non-parametric models that have seen a wide range of applications to a variety of machine learning problems (Mockus, 1994; Hensman et al., 2014; Wilson & Nickisch, 2015). DGPs aim to retain the advantages of GPs—such as their well-calibrated uncertainty estimates and robustness to overfitting—while also overcoming their limitations—such as

^{*}Equal contribution ¹Department of Computer Science, University of Oxford, Oxford, United Kingdom ²Computer Science Department, University College London, London, United Kingdom. Correspondence to: Tim G. J. Rudner <tim.rudner@cs.ox.ac.uk>.




Figure 1. Signal-to-noise ratio (SNRs) in IWVI for DGPs under the standard gradient estimators (REG) and our improved estimator (DREG) as a function of the number of importance samples (K). The different lines show SNRs for four different variational parameters. As K increases, the SNR of REG decreases and the SNR of DREG increases.

the restricted form of their predictive distribution. In essence, DGPs are an even more general and powerful class of models than traditional GPs (Damianou & Lawrence, 2013), allowing us to combine multi-layer function composition with a principled Bayesian approach to obtaining predictive uncertainties.

Unfortunately, this modeling power comes with a caveat: Exact inference in DGPs is intractable and one has to rely on approximate inference schemes to train them. Typically this is done using stochastic variational inference (SVI) methods (Hoffman et al., 2013; Salimbeni & Deisenroth, 2017), which recast the inference problem into that of stochastic optimization of an evidence lower bound (ELBO) with respect to the parameters of an approximating distribution. The resulting SVI-based approaches to training DGPs lead to state-of-the-art performance on a wide range of challenging problems requiring uncertainty quantification (Salimbeni & Deisenroth, 2017).

Despite the successes of these approaches, Salimbeni et al. (2019) find that they do not perform well when applied to DGP models that have been extended to include latent variables that allow them to represent non-Gaussianity and multimodality in data. To address this, they propose an *importance-weighted variational inference* (IWVI; Burda et al. (2016); Mnih & Rezende (2016); Domke & Sheldon

(2018)) approach that provides both a tighter variational bound and lower-variance gradient updates by importance sampling the latent variables. They show that this in turn leads to learning improved predictive distributions.

In this paper, we highlight a potential shortcoming of such approaches: Tightening the bound can cause a deterioration in the *signal-to-noise ratio* (SNR) of the gradient estimates associated with the variational parameters of the latent variables (see Figure 1 for a demonstration). Focusing on the particular example of IWVI, we see that even though using more importance samples both tightens the bound and reduces the variance of the gradients, it can actually *increase* the *relative variance* of the gradients estimates compared to the true gradient, which is itself also tending to zero as we use more samples. This weaker gradient signal leads to difficulties optimizing the variational parameters, a phenomenon analogous to known issues in training the inference networks of variational autoencoders (Rainforth et al., 2018b). We demonstrate this SNR degradation both theoretically in the limit of using a large number of samples, and empirically for samples sizes often used in practice.

To address this degradation, we show how the *doubly-reparameterized gradient* (DREG) estimator of Tucker et al. (2019) can be adapted to the DGP IWVI bound. The theoretical and empirical SNR of the resulting gradient estimator increases for *all* parameters as the number of importance samples is increased, thereby providing more accurate gradient updates in the training of DGPs with latent variables. It can be used as a drop-in replacement for the standard estimator at no computational or other cost. Empirically, we find that this fix can lead to improvements in predictive performance. In particular, we show that when the quality of latent variable variational approximation is important for prediction, then our estimator typically improves performance.¹

To summarize, our core contributions are as follows:

- We highlight the presence of a signal-to-noise ratio (SNR) issue in the training of DGPs with IWVI as the number of importance samples increases;
- We nullify this SNR issue by introducing a doubly-reparameterized gradient (DREG) estimator;
- We quantify the detrimental effect of this SNR deterioration on predictive performance, and show that our fix leads to predictive distributions conferring statistically significant improvements.

¹Our code is available at https://github.com/timrudner/snr_issues_in_deep_gps.

2. Preliminaries & Related Work

2.1. Latent Deep Gaussian Process Models

Let $X = [x_1, \dots, x_N]^\top$ denote a collection of N D -dimensional input data points and let $Y = [y_1, \dots, y_N]^\top$ denote a collection of N P -dimensional noisy observations. A DGP model is defined as the composition

$$y_n = f^{(L)}(f^{(L-1)}(\dots f^{(1)}(x_n)\dots)) + \varepsilon_n,$$

where $\varepsilon_n \sim \mathcal{N}(0, \sigma^2 I_P)$, L is the number of layers, and each $f^{(\ell)}(f^{(\ell-1)})$ in the composition denotes a GP evaluated at the draws from the previous layer, $f^{(\ell-1)}$. As per (Salimbeni et al., 2019), we further augment this model with a \tilde{D} -dimensional latent variable z which extends the input space. The resulting *latent-variable* DGP model is given by

$$y_n = f^{(L)}(f^{(L-1)}(\dots f^{(1)}([x_n, z_n])\dots)) + \varepsilon_n,$$

where $\varepsilon_n \sim \mathcal{N}(0, \sigma^2 I_P)$ and $z_n \sim \mathcal{N}(0, I_{\tilde{D}})$. For simplicity of exposition, our notation here and throughout the paper assumes a DGP with two GP layers, $f^{(1)}$ and $f^{(2)}$, but emphasize that our results apply to DGPs of *any depth* (including latent shallow DGPs) and we include experiments that use more than two layers. Assuming a Gaussian likelihood with scalar noise σ^2 ,

$$\begin{aligned} p(y_n | f^{(1)}, f^{(2)}, z_n; x_n) \\ = \mathcal{N}(y | f^{(2)}(f^{(1)}([x_n, z_n])), \sigma^2 I_P), \end{aligned}$$

the model then has joint distribution

$$\begin{aligned} p(y_n, f^{(1)}, f^{(2)}, z_n; x_n) \\ = p(y_n | f^{(1)}, f^{(2)}, z_n; x_n) p(f^{(2)} | f^{(1)}) p(f^{(1)}) p(z_n), \end{aligned}$$

$$\begin{aligned} \text{where } f^{(1)} &\sim \mathcal{GP}(f^{(1)} | m_1, K_1), \\ f^{(2)} &\sim \mathcal{GP}(f^{(2)} | m_2, K_2), \\ z_n &\sim \mathcal{N}(0, I_{\tilde{D}}). \end{aligned}$$

Here, $m_1(\cdot)$ and $m_2(\cdot)$ are mean functions, and $K_1(\cdot, \cdot)$ and $K_2(\cdot, \cdot)$ covariance functions. For simplicity of notation, we will drop any subscripts from identity matrices in the remainder of the paper.

2.2. Importance-Weighted Variational Inference for Latent-Variable Deep GPs

Posterior inference in DGP models is generally intractable, which necessitates the use of approximate inference methods. For the DGP model presented above, Salimbeni et al. (2019) proposed an *importance-weighted variational inference* (IWVI) approach, inspired by Burda et al. (2016) (see Section 2.3), to obtain a tractable evidence lower bound (ELBO) that can be used to learn a variational approximation of the posterior. Specifically, they present a *partially-collapsed* ELBO that provides both lower-variance estimates

and a tighter bound as the number of importance samples K increases, namely,

$$\mathcal{L}_K \stackrel{\text{def}}{=} \mathbb{E} \left[\sum_n \log \frac{1}{K} \sum_{k=1}^K \frac{\mathcal{F}(x_n, y_n, f_k^{(1)}, z_{n,k}) p(z_{n,k})}{q_\phi(z_{n,k})} \right] - \sum_{\ell=1}^2 D_{\text{KL}} \left(q(f^{(\ell)}) \parallel p(f^{(\ell)}) \right), \quad (1)$$

where

$$\mathcal{F}(x_n, y_n, f_k^{(1)}, z_{n,k}) \stackrel{\text{def}}{=} \exp \left(\mathbb{E}_{q(f^{(2)})} \left[\log p(y_n | f_k^{(2)}, f_k^{(1)}, z_{n,k}) \right] \right),$$

which can be calculated analytically for a Gaussian likelihood; the expectation is over $(f_{1:K}^{(1)}, z_{n,1:K})$, with $z_{n,k} \sim q_\phi(z_n)$ and $p(z_n) = \mathcal{N}(0, I)$; and $f_k^{(1)}$ are samples from the variational distribution $q(f_k^{(1)})$. See [Appendix A](#) for further details.

2.3. Importance Weighted Autoencoders and Their Signal-to-Noise Ratio Issues

[Burda et al. \(2016\)](#) propose importance weighted autoencoders (IWAEs) as a means of providing a tighter variational lower bound for VAE training given by

$$\mathcal{L}_{\text{IWAE}} \stackrel{\text{def}}{=} \sum_{n=1}^N \mathbb{E} \left[\log \frac{1}{K} \sum_{k=1}^K w_{n,k} \right], \quad (2)$$

where $w_{n,k} \stackrel{\text{def}}{=} \frac{p_\theta(x_n, z_{n,k})}{q_\phi(z_{n,k} | x_n)}$,

K is the number of importance samples, N is the number of data points, $p_\theta(x_n, z_{n,k})$ represents the generative model, $q_\phi(z_{n,k} | x_n)$ the amortized inference network, and the expectation is over $\prod_{k=1}^K q_\phi(z_{n,k} | x_n)$. Training is done by optimizing $\mathcal{L}_{\text{IWAE}}$ with respect to both θ and ϕ using stochastic gradient ascent. Assuming that reparameterization of $z_{n,k}$ in $\mathcal{L}_{\text{IWAE}}$ is possible ([Kingma & Welling, 2013](#)), and that it is possible to take mini-batches of the data, the gradient estimates for data point x_n are given by

$$\Delta_{n,M,K}^{\text{IWAE}}(\theta, \phi) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M \nabla_{\theta, \phi} \log \frac{1}{K} \sum_{k=1}^K w_{n,m,k},$$

where $z_{n,m,k} \stackrel{\text{i.i.d.}}{\sim} q_\phi(z_n | x_n)$

and M is the number of samples used for estimation of the outer expectation. [Burda et al. \(2016\)](#) show that increasing K provably tightens the variational lower bound $\mathcal{L}_{\text{IWAE}}$, and [Rainforth et al. \(2018a\)](#) confirm that it also reduces the variance of the resulting estimates.

However, [Rainforth et al. \(2018b\)](#) show that the *relative* variance of the gradient estimates with respect to ϕ actually

increases with the number of importance samples K : As K increases, the expected gradient, $\mathbb{E}[\Delta_{n,M,K}^{\text{IWAE}}(\phi)]$, tends to zero faster than its standard deviation decreases. To formalize this, they introduce the notion of a signal-to-noise ratio (SNR)

$$\text{SNR}_{n,M,K}^{\text{IWAE}}(\psi) = \frac{\left| \mathbb{E} \left[\Delta_{n,M,K}^{\text{IWAE}}(\psi) \right] \right|}{\sqrt{\text{Var} \left[\Delta_{n,M,K}^{\text{IWAE}}(\psi) \right]}},$$

where $\psi \in \{\phi, \theta\}$, and show that

$$\begin{aligned} \text{SNR}_{n,M,K}^{\text{IWAE}}(\theta) &= \mathcal{O} \left(\sqrt{MK} \right) \\ \text{SNR}_{n,M,K}^{\text{IWAE}}(\phi) &= \mathcal{O} \left(\sqrt{M/K} \right), \end{aligned}$$

such that increasing K decreases the SNR of the inference network, thereby potentially removing its ability to train effectively. As well as being problematic in its own right, this pathology can further have a knock-on effect on the training of the generative network.

3. Signal-to-Noise Issues for Deep GPs

Given the similarity between the variational bounds in IWAEs and IWVI for DGPs, it is natural to ask whether the SNR issues in IWAEs also occur in IWVI for DGPs. We note here that for IWAEs, this pathology affected the encoder (i.e., ϕ) but not the decoder (i.e., θ) and so it is not immediately obvious how these results will translate; SNR issues are not a universal for all gradients of importance-weighted variational bounds.

Inspecting [Equation \(1\)](#), we see that the DGP setting shares a number of similarities, but also features some key differences to the IWAE setting. On one hand, both \mathcal{L}_K ([Equation \(1\)](#)) and $\mathcal{L}_{\text{IWAE}}$ ([Equation \(2\)](#)) contain expectations over a logarithm of an importance sampling estimate and both make use of a variational approximation. On the other hand, \mathcal{L}_K contains an additional KL term, preventing it from ever achieving a tight bound even when the variational approximation over the latent variable is optimal ([Salimbeni et al., 2019](#)). Furthermore, it contains the function $\mathcal{F}(x_n, y_n, f_k^{(1)}, z_{n,k})$ instead of a plain likelihood function, and the outer expectation contains additional stochasticity from the DGP layers. More specifically, the posterior predictive distribution of *each layer* needs to be inferred and the approximations produced are functions of the layer’s hyperparameters, its variational parameters, and the previous layer’s outputs. Establishing the implications of these differences will be key to unearthing whether DGPs suffer from SNR issues.

To assess if SNR deterioration of the gradient estimates occurs for \mathcal{L}_K as K increases, we must first identify which

gradients are potentially problematic. Considering Equation (1), it is relatively straightforward to see that, analogously to the decoder in a VAE, the gradients for the DGP hyperparameters will remain non-zero even as $K \rightarrow \infty$ and thus will not be susceptible to SNR issues. Similarly, it is reasonably straightforward to show that the parameters of $q(f^{(\ell)})$ will not be problematic. However, there is a potential for problems to occur for the parameters of the variational distribution over the latent variables, $q_\phi(z)$, with variational parameters ϕ .

It turns out that such a problem does indeed occur for the gradient estimates w.r.t. ϕ . To demonstrate this, the first key step is to reparameterize *all* of the stochastic quantities in the variational bound in Equation (1). We can then express the M -sample Monte Carlo estimate for the gradient associated with a single data point x_n as

$$\Delta_{n,M,K}^{\text{DGP}}(\phi) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M \nabla_\phi \log \frac{1}{K} \sum_{k=1}^K w_{n,m,k}, \quad (3)$$

$$\text{where } w_{n,m,k} \stackrel{\text{def}}{=} \frac{\mathcal{F}(x_n, y_n, f_{m,k}^{(1)}, z_{n,m,k}) p(z_{n,m,k})}{q_\phi(z_{n,m,k})},$$

$$z_{n,m,k} \sim q_\phi(z_n), \quad f_{m,k}^{(1)} \sim q(f^{(1)}).$$

We are now ready to present our main result which shows that IWVI for DGPs suffers from an SNR issue analogous to that seen in IWAES.

Theorem 1 (Asymptotic SNR in IWVI for DGPs). *Let $w_{n,m,k}$ be as defined in Equation (3). Assume that when $M = K = 1$, the expectation and variance of the gradients estimates in Equation (3) are non-zero, and that the first four moments of $w_{n,1,1}$ and $\nabla_\phi w_{n,1,1}$ are all finite and that their variances are also non-zero. Then the signal-to-noise ratio of each $\Delta_{n,M,K}^{\text{DGP}}(\phi)$ converges at the following rate*

$$\text{SNR}_{n,M,K}^{\text{DGP}}(\phi) = \sqrt{M} \left| \frac{\nabla_\phi \text{Var}[w_{n,1,1}] + \mathcal{O}\left(\frac{1}{K}\right)}{2Z_n \sqrt{K} \sqrt{\text{Var}[\nabla_\phi w_{n,1,1}] + \mathcal{O}\left(\frac{1}{\sqrt{K}}\right)}} \right|, \quad (4)$$

$$= \mathcal{O}\left(\sqrt{M/K}\right),$$

where $Z_n \stackrel{\text{def}}{=} \mathbb{E}[w_{n,1,1}]$ is a lower bound on the marginal likelihood of the n^{th} data point.

Proof. We start our proof by noting that the average of the importance weights $w_{n,m,k}$ for a given n and m ,

$$\hat{Z}_{n,m,K} \stackrel{\text{def}}{=} \frac{1}{K} \sum_{k=1}^K w_{n,m,k}, \quad (5)$$

is an unbiased Monte Carlo estimator of Z_n as follows

$$\mathbb{E}[\hat{Z}_{n,m,K}] = \mathbb{E}[w_{n,m,1}] = \mathbb{E}[w_{n,1,1}] = Z_n. \quad (6)$$

Moreover, we also have $\lim_{K \rightarrow \infty} \hat{Z}_{n,m,K} = Z_n, \forall n, m$.

The key to establishing that the SNR pathology occurs in IWVI for DGPs, is now to show that Z_n is independent of ϕ , the parameters of the variational distribution over the latent variable z . For IWAES this was trivially true by construction: The expected importance weight is just the marginal likelihood of the model. In the DGP case, however, the dependence of Z_n on ϕ is not immediately obvious since the dependence of $\mathcal{F}(x_n, y_n, f^{(1)}, z)$ on z is more complicated than in the IWAES case. In particular, in IWVI for DGPs, z acts as an input to the approximate posterior predictive distribution of the DGP layer $f^{(1)}$, which itself needs to be inferred and with respect to which the expectation over the importance weight in Equation (6) is taken.

In short, we need to demonstrate that the following condition holds despite the dependence of $q(f^{(1)})$ on z :

Condition 1. $\nabla_\phi Z_n = 0 \forall n$.

To examine if this condition is indeed satisfied in IWVI for DGPs, the first hurdle that we need to overcome is that $f^{(1)}$ and $f^{(2)}$ represent stochastic functions, which means it is difficult to reason about their gradients or concretely establish the dependency relationships between variables. To get around this, the key step is to realize that we can reparameterize this stochasticity and then view the bound from the perspective of the pushforward distribution induced by passing an input-latent pair through the resulting realizations of the functions. Specifically, we have, for an arbitrary z ,

$$f^{(1)}(x_n, z) = \tilde{f}(\epsilon, x_n, z, \psi^{(1)}) \stackrel{\text{def}}{=} \mu^{f^{(1)}}(x_n, z, \psi^{(1)}) + \epsilon \odot \sqrt{\Sigma^{f^{(1)}}(x_n, z, \psi^{(1)})}, \quad (7)$$

where $\tilde{f}(\cdot)$ is a deterministic function; $\epsilon \sim \mathcal{N}(0, 1)$; $\mu^{f^{(1)}}$, $\Sigma^{f^{(1)}}$ denote the mean and covariance functions of the predictive distribution over $f^{(1)}$, respectively; and $\psi^{(1)}$ denotes the variational parameters associated with $f^{(1)}$. We thus have

$$Z_n = \mathbb{E}_{q_\phi(z)p(\epsilon)} \left[\frac{\mathcal{F}(x_n, y_n, \tilde{f}(\epsilon, x_n, z, \psi^{(1)})) p(z)}{q_\phi(z)} \right]$$

$$= \mathbb{E}_{p(z)p(\epsilon)} \left[\mathcal{F}(x_n, y_n, \tilde{f}(\epsilon, x_n, z, \psi^{(1)})) \right].$$

From here it is now straightforward to see that $\nabla_\phi Z_n = 0$ as the final form of Z_n above has no direct or indirect dependency on ϕ . We thus see that Condition 1 is satisfied.

This now allows us to invoke the proof of Rainforth et al. (2018b, Theorem 1) because a) our gradient estimator in Equation (3) is equivalent to that of IWAES except in the distribution of $w_{n,m,k}$, and b) the only VAE-context specific part of their proof is in showing that Condition 1 holds.

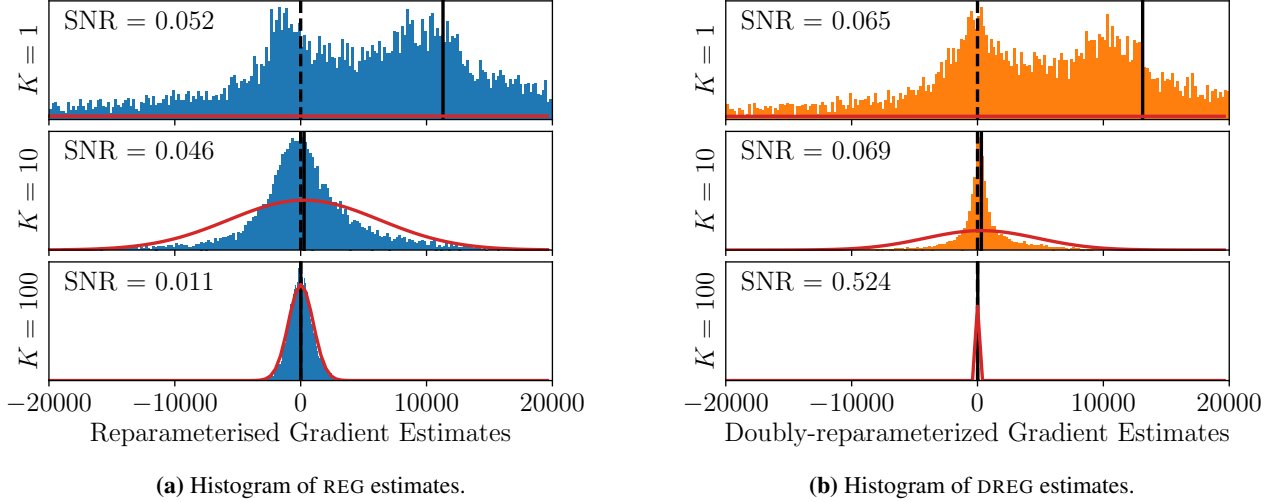


Figure 2. Histograms of REG (a) and DREG (b) estimates for varying numbers of importance samples, K , with $K = 1, 10, 100$. The solid vertical lines denote the mean of the empirical distributions of gradient estimates, and the dashed vertical lines mark zero. The red curves denote Gaussian probability densities with the same mean and standard deviation as the empirical distributions. Note that we have truncated the histograms for $K = 1$ in order to use a consistent x -axis for each row, see Figure 8 (in Appendix F) of a copy of this figure with a separate x -axis for each value of K .

Starting from equations (29) and (31) in their proof and using $\nabla_\phi Z = 0$, we can now derive the mean and variance of the gradient estimator of a single data point (i.e. for a particular n) as

$$\begin{aligned}\mathbb{E}[\Delta_{n,M,K}^{\text{DGP}}(\phi)] &= -\frac{1}{2KZ_n^2} \nabla_\phi (\text{Var}[w_{n,1,1}]) + \mathcal{O}(K^{-2}) \\ \text{Var}[\Delta_{n,M,K}^{\text{DGP}}(\phi)] &= \frac{1}{MKZ_n^2} \mathbb{E}[(\nabla_\phi w_{n,1,1})^2] + \frac{\mathcal{O}(K^{-2})}{M}.\end{aligned}$$

Because $w_{n,1,1}$ has been reparameterized, we also have $\mathbb{E}[(\nabla_\phi w_{n,1,1})^2] = \text{Var}[\nabla_\phi w_{n,1,1}]$. Substituting these into the definition of the SNR yields

$$\text{SNR}_{n,M,K}^{\text{DGP}}(\phi) = \left| \frac{-\frac{1}{2KZ_n^2} \nabla_\phi (\text{Var}[w_{n,1,1}]) + \mathcal{O}(K^{-2})}{\sqrt{\frac{1}{MKZ_n^2} \text{Var}[\nabla_\phi w_{n,1,1}] + \frac{\mathcal{O}(K^{-2})}{M}}} \right|$$

from which Equation (4) follows by straightforward manipulations. \square

Remark 1. Note that the additional results of Rainforth et al. (2018b, Section 3.2), which show that the SNR improves as more data points are used in the gradient calculation for IWAEs (namely as $\mathcal{O}(\sqrt{N})$), do not directly carry over to the DGP setting because $f^{(1)}$ induces correlations between the different $\Delta_{n,M,K}^{\text{DGP}}(\phi)$.

3.1. Empirical Confirmation

We now investigate the signal-to-noise ratio of the reparameterization gradient (REG) estimates in IWVI for DGPs

empirically, and show that the empirical results are consistent with Theorem 1.

We begin by examining the behavior of the gradient estimates for the parameters ϕ as we increase the number of importance samples, K , when training a two-layer latent-variable DGP using IWVI. For illustrative purposes, we start by investigating the convergence behavior of the gradient estimates on an easy-to-visualize synthetic dataset specifically designed to exhibit non-Gaussianity in the target values. We train the DGP until the variational parameters are near convergence and then take 10,000 gradient samples for each parameter in ϕ using the reparameterization gradient estimator. For further details about the experiment setup and a visualization of the dataset, see Appendix E. In Figure 2a, we present histograms of the empirical distribution of gradient estimates for a single (representative) variational parameter of q_ϕ .

As we would expect by the Central Limit Theorem and knowledge that Condition 1 is satisfied, the mean and the standard deviation of the distributions of gradient estimates approach zero as K increases. However, the means of the gradient estimates appear to be approaching zero more rapidly than their standard deviations, which would suggest a decrease in the SNR as K increases. To assess if such a deterioration in fact occurs, we compute the mean SNR across parameters ϕ for varying K and multiple DGP depths on a range of real-world datasets. We collect the resulting mean SNRs in the top row of Figure 3, and find that—across datasets and DGP depths—the SNR consistently approaches zero as K increases, confirming that IWVI for DGPs indeed suffer from SNR deterioration.




Figure 3. SNR of reparameterization (top row) and doubly reparameterized (bottom row) gradient estimates for shallow GPs and DGPs of 2-4 layers on a selection of real-world datasets. The labels on the x -axes correspond to the depths of the models. The bars for each depth show the SNR for increasing numbers of importance samples, $K = 1, 10, 100, 1000$, from left to right. In the top row, for (D)GPs of any depth, larger K tends to correspond to lower SNRs. In the bottom row, for (D)GPs of any depth, larger K tends to correspond to higher SNRs. Note the difference in y -axis scales across plots in the bottom row. See Appendix E for further experiment details.

4. Quantifying & Fixing the SNR Pathology

In the previous section, we demonstrated that IWVI for DGPs suffers from SNR deterioration as the number of importance samples is increased. This result leads to two questions: (i) Can we construct an alternative estimator that avoids these issues; and (ii) will fixing the SNR issue lead to improved predictive performance?

To address the former question, we adapt the double-reparameterized IWAE estimator of Tucker et al. (2019) to DGP models and show that the resulting DREG estimator completely remedies the SNR issue. We then use this estimator to assess the impact SNR issues have had on the quality of the DGP’s posterior predictive distribution. We consider a selection of datasets used in Salimbeni et al. (2019) and find that our fix often leads to a statistically significant *improvement in predictive performance* for IWVI in DGPs. Specifically, we find that on tasks where fitting the data well requires a predictive distribution with *non-Gaussian marginals*, such that the latent variables are necessary, fixing the SNR deterioration improves the predictive performance.

4.1. Avoiding SNR Deterioration in Deep GPs

To avoid SNR deterioration, we extend prior work to derive a doubly reparameterized gradient (DREG, (Tucker et al., 2019)) estimator for IWVI in DGPs. This gradient estimator is equal to the REG estimator (i.e. Equation (3)) in expectation, but does not suffer from asymptotic deterioration of the SNR as K increases. Assuming that reparameterization of $q(z_n)$ in \mathcal{L}_K is possible, the DREG estimator of \mathcal{L}_K at a

single data point x_n can be expressed as (see Appendix D for the derivation)

$$\begin{aligned} & \tilde{\Delta}_{n,M,K}^{\text{DGP}}(\phi) \\ \stackrel{\text{def}}{=} & \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K \left(\frac{w_{n,m,k}}{\sum_{j=1}^K w_{n,m,j}} \right)^2 \frac{\partial \log w_{n,m,k}}{\partial z_{n,m,k}} \frac{\partial z_{n,m,k}}{\partial \phi}, \end{aligned} \quad (8)$$

where $w_{n,m,k}$ is as before. As we explain in Appendix D, analogously to the results of Tucker et al. (2019), the SNR of this gradient estimator scales as $\mathcal{O}(\sqrt{K})$ instead of $\mathcal{O}(1/\sqrt{K})$, that is, the SNR *improves* as K increases. The DREG estimator can be used as a drop-in replacement of the REG estimator and is guaranteed to be at least as good or better without incurring additional computational cost (Tucker et al., 2019).²

To show that the DREG estimates of the latent-variable variational parameters do not suffer from SNR deterioration in practice, we revisit the empirical investigations carried out in Section 3.1 and compute the SNR of the gradient estimates for an increasing number of importance samples across datasets and DGP depths. Figure 3 shows the resulting SNRs. We find that the effect of increasing K for REG and DREG is markedly different: Unlike for the REG estimator, the DREG estimate SNR values *increase* with K .

The difference between the SNRs of the two gradient estimators is explained by the speed at which the mean and

²Our implementation of the DREG estimator follows <https://sites.google.com/view/dregs>.




Figure 4. Comparison of predictive performance of 2-layer DGPs with a learned variational distribution over the latent variable (left of each pair, blue) and a variational distribution over the latent variable fixed to the prior (right of each pair, green). The shaded area shows the range of test log-likelihoods over 10 train–test splits, with the width indicating the distribution over the range. The central horizontal lines in each plot show the mean.

standard deviation of the empirical distributions over the gradient estimates converge to zero as K increases. Figure 2b shows the empirical distribution of DREG estimates. As can be seen in the histograms, the means for $K = 1, 10, 100$ decrease as K increases, but the standard deviations of the empirical distributions of DREG estimates are significantly smaller than those of the empirical distributions of the REG estimates shown in Figure 2a. This difference is particularly striking for $K = 100$, where the DREG estimates are so peaked that they are difficult to visualize without changing the x -axis range. The upshot of the results in Figures 2 and 3 is that the DREG estimator completely remedies the SNR issue exhibited by the REG estimates.

4.2. Does the SNR Issue Affect Training and Predictive Performance?

Armed with the DREG estimator, we are now able to investigate the impact of correcting the SNR deterioration on the training and test performance of the model. To do so, we consider a selection of datasets for which we either expect or do not expect the SNR issue to lead to a deterioration in predictive performance. In particular, we assess the impact of SNR deterioration on two datasets where fitting the data well requires a predictive distribution with *non-Gaussian marginals* (‘forest’ and ‘solar’), two datasets which can be fit well with just Gaussian marginals (‘winewhite’ and ‘winered’), and two datasets where a predictive distribution with non-Gaussian marginals could potentially help a modest amount (‘pol’ and ‘power’). By construction, latent-variable DGPs are able to learn highly non-Gaussian marginals and so learning a variational distribution over the latent variable should positively affect predictive performance whenever highly non-Gaussian marginals lead to a better fit of the data. Hence, we would expect the effect of a deterioration in the SNR of the gradient estimates to be highest whenever this is the case.

In Figure 4, we present the plots of test log-likelihoods for latent-variable DGPs with learned variational distributions over the latent variable (left of each pair, in blue) and with




Figure 5. Marginal predictive distributions of 2-layer DGPs with a learned variational distribution over the latent variable (top row, blue) and a variational distribution over the latent variable fixed to the prior (bottom row, green) for randomly selected test points from the ‘forest’ and ‘winewhite’ datasets. We note that learning the variational distribution helps fitting a better non-Gaussian marginal distribution for ‘forest’ (as evidenced by the test log-likelihoods shown in Figure 4), whereas for ‘winewhite’ the marginal is Gaussian and fit equally well in both cases.

latent variable distributions fixed to standard Gaussian priors (right of each pair, in green). As can be seen from the plots, the importance of learning a variational distribution for the latents sometimes has a significant effect on predictive performance, but sometimes has no noticeable effect at all. In Figure 5, we present sample marginal distributions of the best-performing models, which indicate that learning a distribution over the latents, $q_\phi(z)$, has a larger effect the more non-Gaussian the predictive distribution’s marginals need to be. We thus expect our SNR fix to be helpful for ‘forest’ and ‘solar’, but not for ‘winewhite’ and ‘winered.’

To test this, we train latent-variable DGPs with both REG and DREG estimators on the six datasets shown in Figure 4 and compare the resulting train ELBOs and test log-likelihoods. For these experiments, we consider a two-layer DGP with the hyperparameters that directly affect the SNR—the number of importance samples K and the minibatch size—set to 50 and 64, respectively. Table 1 shows a summary of the results. We see that our DREG estimator provides improvements for the datasets where the latent variational approximation is clearly important (‘forest’ and ‘solar’), while performing similarly to the REG estimator when it is not.

To assess whether the differences are statistically significant,

Table 1. Comparison of predictive performance of two-layer DGPs trained with REG and DREG estimators. We choose $K = 50$ to ensure the SNR deterioration occurs, which Figure 3 shows is the case for $K \geq 10$. For each dataset, we provide the mean ELBOs on the training dataset and log-likelihoods on the test dataset over 20 random train-test splits as well as the corresponding standard errors. Boldface indicates higher means. The rightmost column shows p -values for one-sided Wilcoxon signed-rank hypothesis tests on the log-likelihoods, as described in Section 4.2.

Dataset	Train ELBO ($K = 50$)				Test log-likelihood				Wilcoxon Test
	REG		DREG		REG		DREG		
	Mean	SE	Mean	SE	Mean	SE	Mean	SE	p -value
forest	-97.56	(11.04)	-92.53	(10.42)	0.59	(0.08)	0.63	(0.08)	0.1%
solar	1657.41	(27.56)	1707.75	(42.20)	2.33	(0.17)	2.57	(0.11)	2.8%
pol	34610.49	(66.18)	34665.08	(70.34)	2.99	(0.01)	2.99	(0.01)	24.7%
power	1510.50	(10.62)	1515.60	(10.16)	0.21	(0.01)	0.21	(0.01)	67.3%
winewhite	-4701.26	(4.92)	-4703.14	(4.98)	-1.11	(0.01)	-1.11	(0.01)	50.0%
winered	447.91	(249.81)	314.75	(216.32)	0.57	(0.27)	0.61	(0.20)	41.1%
Across Datasets:									1.2%

we perform Wilcoxon signed-rank tests for each dataset individually as well as across datasets (Wilcoxon, 1992). This is chosen in preference to a more conventional t -test because of the highly non-Gaussian nature of the variations across seeds, such that the criteria for a t -test to be representative are not met. Specifically, for each dataset and across datasets, we test the null hypothesis that the difference in test log-likelihoods under REG and DREG for each random seed is zero with the alternative hypothesis that the test log-likelihood under DREG is greater than the test log-likelihood under REG. We present the results from this one-sided hypothesis test in Table 1 (rightmost column). As can be seen from these results, the p -values for the datasets on which learning a variational distribution over the latent variable resulted in a better fit of the data (see Figure 4) are 0.1% and 2.8%, indicating that we can reject the null hypothesis for both at the standard 5% confidence level. While the statistical significance of the improvement in log-likelihood under DREG is obfuscated by large standard errors across train-test splits, Figure 6 shows that DREG leads to a *consistent improvement* in predictive performance across splits when the variational distribution is important. We further observe that, as expected, when the variational distribution is not important for obtaining a good fit, there is no statistically significant change.




Figure 6. Left: Test log-likelihoods across seeds on the ‘forest’ dataset, ordered from most negative to most positive. Right: Corresponding absolute difference in test log-likelihood between REG and DREG across seeds.

Finally, we find that the improvement in predictive performance under DREG *across datasets* for even a moderately small number of importance samples ($K = 50$) is statistically significant at the 5% confidence level (with a p -value of 1.2%), which is expected since, as Table 1 shows, fixing the SNR issue with the DREG estimator either improves predictive performance or does not affect it at all.

We finish by noting that that Salimbeni et al. (2019) already provided comparisons to other inference approaches for DGPs (Salimbeni & Deisenroth, 2017; Havasi et al., 2018)—along with other regressors—and found that using IWVI for DGPs with the REG estimator produced state-of-the-art predictive performance. It thus follows that our approach is able to further improve upon the state of the art for regression modeling and is especially well-suited for data that requires complex, highly-Gaussian predictive marginals.

5. Conclusions

We have shown that the gradient estimates used in training DGPs with IWVI are susceptible to *signal-to-noise ratio* issues: We demonstrated theoretically that the SNR for the latent-variable variational parameters increases as more importance samples are used, and confirmed this result empirically. We have further shown how this pathology can be remedied by adapting the doubly reparameterized approach of Tucker et al. (2019) to the DGP setting, resulting in a gradient estimator whose SNR *increases* with the number of importance samples for all variables. This estimator can be used as a drop in replacement without incurring any significant increase in computational cost, implementation challenges, or other negative effects. We find that it can provide improvements in predictive performance, even for small numbers of importance samples, when variational approximation of the latents is important, without damaging performance when it is not.

Acknowledgements

Tim G. J. Rudner is funded by the Rhodes Trust and the Engineering and Physical Sciences Research Council (EPSRC). Oscar Key is funded by the Engineering and Physical Sciences Research Council (EPSRC), grant number EP/S021566/1. We gratefully acknowledge donations of computing resources by the Alan Turing Institute.

References

- Bui, T., Hernandez-Lobato, D., Hernandez-Lobato, J., Li, Y., and Turner, R. Deep Gaussian processes for regression using approximate expectation propagation. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1472–1481, 20–22 Jun 2016.
- Burda, Y., Grosse, R. B., and Salakhutdinov, R. Importance weighted autoencoders. In *ICLR (Poster)*, 2016.
- Damianou, A. and Lawrence, N. Deep Gaussian processes. In Carvalho, C. M. and Ravikumar, P. (eds.), *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pp. 207–215, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR.
- Domke, J. and Sheldon, D. R. Importance weighting and variational inference. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256, 2010.
- Havasi, M., Hernández-Lobato, J. M., and Murillo-Fuentes, J. J. Inference in deep Gaussian processes using stochastic gradient Hamiltonian Monte Carlo. In *Advances in Neural Information Processing Systems 31*, pp. 7506–7516, 2018.
- Hensman, J., de G. Matthews, A. G., and Ghahramani, Z. Scalable variational Gaussian process classification. In *AISTATS*, 2014.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, May 2013. ISSN 1532-4435.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.
- Mnih, A. and Rezende, D. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, pp. 2188–2196. PMLR, 2016.
- Mockus, J. Application of bayesian approach to numerical methods of global and stochastic optimization. *J. Global Optimization*, 4(4):347–365, 1994.
- Rainforth, T., Cornish, R., Yang, H., Warrington, A., and Wood, F. On nesting Monte Carlo estimators. In *International Conference on Machine Learning*, pp. 4267–4276, 2018a.
- Rainforth, T., Kosiorek, A. R., Le, T. A., Maddison, C. J., Igl, M., Wood, F., and Teh, Y. W. Tighter variational bounds are not necessarily better. In *ICML*, 2018b.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, Massachusetts Institute of Technology, Cambridge, MA, USA, 2006. URL www.gaussianprocess.org/gpml.
- Salimbeni, H. and Deisenroth, M. Doubly stochastic variational inference for deep Gaussian processes. In *Advances in Neural Information Processing Systems 30*, pp. 4588–4599, 2017.
- Salimbeni, H., Dutordoir, V., Hensman, J., and Deisenroth, M. P. Deep Gaussian processes with importance-weighted variational inference. In *ICML*, 2019.
- Tucker, G., Lawson, D., Gu, S., and Maddison, C. J. Doubly reparameterized gradient estimators for Monte Carlo objectives. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- Wilcoxon, F. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pp. 196–202. Springer, 1992.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8(34):229256, May 1992. ISSN 0885-6125. doi: 10.1007/BF00992696.
- Wilson, A. and Nickisch, H. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1775–1784, Lille, France, 07–09 Jul 2015. PMLR.