Amortized variance reduction for doubly stochastic objectives

Ayman Boustati [*]
University of Warwick
Coventry, UK

Sattar Vakili
Prowler.io
Cambridge, UK

James Hensman
Prowler.io
Cambridge, UK

ST John
Prowler.io
Cambridge, UK

Abstract

Approximate inference in complex probabilistic models such as deep Gaussian processes requires the optimisation of doubly stochastic objective functions. These objectives incorporate randomness both from mini-batch subsampling of the data and from Monte Carlo estimation of expectations. If the gradient variance is high, the stochastic optimisation problem becomes difficult with a slow rate of convergence. Control variates can be used to reduce the variance, but past approaches do not take into account how mini-batch stochasticity affects sampling stochasticity, resulting in sub-optimal variance reduction. We propose a new approach in which we use a recognition network to cheaply approximate the optimal control variate for each mini-batch, with no additional model gradient computations. We illustrate the properties of this proposal and test its performance on logistic regression and deep Gaussian processes.

1 INTRODUCTION

Many machine learning tasks such as regression and classification can be cast into a form in which we infer model parameters $\boldsymbol{\theta}$ by optimising an objective function $\mathcal{L} = \sum_{n=1}^{N} \ell_n(\boldsymbol{\theta})$ which is a sum over contributions from each data point n. We focus on objectives that contain an analytically intractable expectation, $\ell_n(\boldsymbol{\theta}) = \mathbb{E}_{p(\boldsymbol{\epsilon})}[f_n(\boldsymbol{\epsilon}, \boldsymbol{\theta})]$, such as in Black Box Variational Inference (Ranganath et al., 2014), Variational Auto-Encoders (Kingma and Welling, 2014), or Deep Gaussian Processes (Salimbeni and Deisenroth, 2017).

In practice, such objectives are treated using Monte Carlo (MC) sampling to obtain an unbiased stochastic estimate
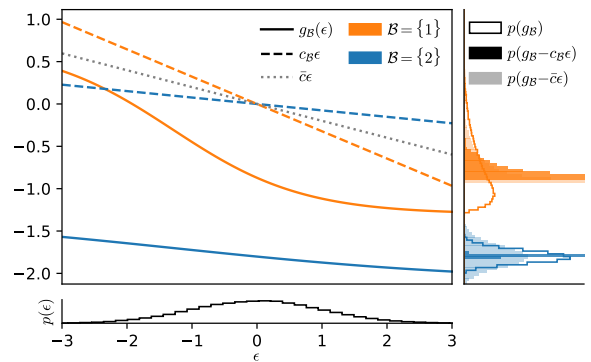
[*]Work done while at Prowler.io.

Figure 1: In reparameterized variational inference, the gradient value g is a function of the randomness sample $\epsilon \sim p(\epsilon)$. This functional relationship $g_{\mathcal{B}}(\epsilon)$ (solid lines) depends on the mini-batch $\mathcal{B}$ (orange vs blue). Here we show linear control variates (CVs) with batch-dependent coefficients $c_{\mathcal{B}}\epsilon$ (dashed lines) and the best batch-independent CV $\bar{c}\epsilon$ (dotted grey line). The right-hand plot shows the distribution of these expectation estimators for each mini-batch: without CV (outline), with batch-independent CV (shaded), and with batch-dependent CV (filled). The batch-dependent CV significantly reduces the variance, whereas here the batch-independent CV actually increases the variance for the blue mini-batch.

of the expectation, $\hat{\ell}_n = \frac{1}{S} \sum_{s=1}^{S} f_n(\boldsymbol{\epsilon}_n^{(s)}, \boldsymbol{\theta})$, where $\boldsymbol{\epsilon}_n^{(s)} \sim p(\boldsymbol{\epsilon})$. We can then optimise using Stochastic Gradient Descent (SGD) on the noisy gradients (Robbins and Monro, 1951). For large N, the evaluation of the full sum in $\mathcal{L}$ is often computationally intractable. This can be addressed by subsampling mini-batches $\mathcal{B} \subset \{1, \dots, N\}$ of size $|\mathcal{B}|$ from the full data set, introducing additional noise and leading to a doubly stochastic objective function:

$$\hat{\mathcal{L}} := \frac{N}{|\mathcal{B}|S} \sum_{b \in \mathcal{B}} \sum_{s=1}^{S} f_b(\boldsymbol{\epsilon}_b^{(s)}, \boldsymbol{\theta}), \qquad (1)$$

with $\mathbb{E}[\hat{\mathcal{L}}] = \mathcal{L}$.

The variance of the gradients of $\mathcal{L}$ affects both the rate of convergence of the optimisation and how close can the optimiser get to the optimum. This motivates various approaches for reducing either mini-batch variance (e.g., Johnson and Zhang (2013)) or the variance due to MC estimation of the expectation (Ranganath et al., 2014; Roeder et al., 2017). A common approach for variance reduction are control variates (see Section 2.1), which have recently been adopted in the literature (Paisley et al., 2012; Miller et al., 2017; Grathwohl et al., 2018; Geffner and Domke, 2018). The focus for the latter work is on deriving and applying control variate schemes to MC objectives, specifically in the context of Variational Inference (VI).

However, to the best of our knowledge, the schemes in the literature do not consider the mini-batching case and do not explicitly take into account how the context of the data point b affects the dependence of f_b on ϵ. This dependence is illustrated in Figure 1 at the example of Bayesian logistic regression. The gradient $g_{\mathcal{B}}(\epsilon)$ as a function of the randomness ϵ of a doubly stochastic objective is shown for two different mini-batches $\mathcal{B}$. In this simplified case, each batch consists of a single context point. The two different context points induce different relationships between randomness and gradient value as shown by the solid lines. This means that the two gradients correlate differently with the randomness, resulting in different control variates represented by the dashed lines. For comparison, we include a batch-independent control variate for the expectation estimator (dotted line) which has to average over all contexts. Adapting the control variate to the batch significantly reduces the variance, shown in the right-hand panel in Figure 1.

In this work, we propose a novel idea for computing control variates that adapt to the context (mini-batch) of the controlled estimators (the gradient). The new formulation takes into account the dependence of the MC estimate on the data by using a recognition network to learn an adaptive control variate coefficient. We derive a low-variance objective function to train the network to approximate the optimal control variate coefficient per batch. Additionally, we propose two computationally cheaper alternatives to the network objective with higher variance. All control variate objectives re-use the already computed model objective gradient, and hence do not require extra backpropagation steps. We empirically test the properties of our proposed method in Section 5.

2 METHODOLOGY

In Section 2.1, we start with a review of control variates and highlight the importance of computing the optimal

control variate coefficient in the general case to allow for maximal variance reduction. We introduce the dependence of the gradients and the control variates on the selected mini-batch and propose learning context-aware control variate coefficients in Section 2.2 . Finally, in Section 2.3 we derive the objectives for the control variate coefficients that allow amortisation through a recognition network.

2.1 CONTROL VARIATES

We want to reduce the variance of an unbiased stochastic estimator $\hat{g}_\theta(\epsilon)^*$ for an intractable expectation $\mathbb{E}[g(\epsilon)]$, where $\epsilon \sim p(\epsilon)$ is a random variable. We consider a different function $w(\epsilon)$ whose expectation is known analytically, $\mathbb{E}[w(\epsilon)] = W$. Then $C(w(\epsilon) - W)$ has zero expectation for any C, and its unbiased estimator, $C(\hat{w}(\epsilon) - W)$, can be subtracted from the original estimator,

$$\tilde{g}(\epsilon) = \hat{g}(\epsilon) - C(\hat{w}(\epsilon) - W). \qquad (2)$$

This new estimator has the same expectation as the original estimator and is also unbiased. Minimising its variance $\mathrm{Var}[\tilde{g}]$ gives the optimal $C^* = \mathrm{Cov}[\hat{g}, \hat{w}]/\mathrm{Var}[\hat{w}]$, and $\tilde{g}$ will have lower variance than $\hat{g}$ if $g(\epsilon)$ and $w(\epsilon)$ are correlated. In particular, choosing C optimally results in variance reduction of

$$\mathrm{Var}[\tilde{g}] = (1 - \rho_{g,w})\,\mathrm{Var}[\hat{g}], \qquad (3)$$

where $\rho_{g,w}$ is the Pearson correlation coefficient between g and w. In practice, however, computing C^* is not possible, as $\mathrm{Cov}[\hat{g}, \hat{w}]$ and $\mathrm{Var}[\hat{w}]$ cannot be evaluated exactly, and are usually estimated from the optimisation statistics, e.g. running averages (Paisley et al., 2012). Another option is to pre-specify C and keeping it fixed (Miller et al., 2017; Grathwohl et al., 2018).

Neither option is convincing for the doubly stochastic case. The first option has very high variance due to the presence of mini-batch stochasticity in addition to sampling stochasticity. The second option is unreliable as pre-specifying an arbitrary value for C does not guarantee optimal variance reduction as can be seen in (3).

In Section 2.2 we will specify C as a context-dependent adaptive parameter that is learned through the optimisation. In Section 2.3, we discuss the corresponding training objectives for C.

2.2 CONTROLLING MINI-BATCH GRADIENTS

For gradient-based optimisation we need the derivatives of the objective (1) with respect to the model parameters

---

*We use the $\hat{}$ symbol (as well as $\tilde{}$) on top of functions of random variables to denote the estimate of this function obtained by evaluating the relevant estimator. In the following we drop the dependence on θ to lighten the notation.

$\{\theta_i\}_{i=1}^P$. The estimated gradient contains a sum over mini-batch elements b,

$$\frac{\partial \hat{\mathcal{L}}}{\partial \theta_i} \propto \sum_{b \in \mathcal{B}} \frac{\partial f_b}{\partial \theta_i}(\boldsymbol{\epsilon}_b, \boldsymbol{\theta}) = \sum_{b \in \mathcal{B}} \hat{g}_{bi}(\boldsymbol{\epsilon}_b) =: \hat{G}_i, \quad (4)$$

where we chose $S = 1$ to simplify the equations (the extension to multiple MC samples is straightforward). Note that $\mathcal{B}$ is a random subset of $\{1, \dots, N\}$, i.e., the b's are indices into the full dataset, and each term gets its own realisation $\boldsymbol{\epsilon}_b$ of the randomness. We want to improve the optimisation performance by reducing the variance of this gradient. As demonstrated in Figure 1, each partial gradient estimator $\hat{g}_{bi}(\boldsymbol{\epsilon}_b)$ may have a different dependence on the randomness. To account for this, we introduce separate control variates for each term (data point) in the sum in (4). For a single partial gradient, we define the controlled gradient estimator

$$\tilde{g}_{bi}(\boldsymbol{\epsilon}_b) := \hat{g}_{bi}(\boldsymbol{\epsilon}_b) - \mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b). \quad (5)$$

(Here and in the following section we subsume the analytic expectation into the definition of the control variate such that $\hat{\mathbf{w}}(\boldsymbol{\epsilon})$ already has zero mean. We use the same type of control variates for all parameters; in principle, we could have a different $\hat{\mathbf{w}}_i$ per parameter θ_i. Note that in any case there are per-parameter coefficients $\mathbf{c}_{bi}$.) In general, the mapping $\hat{\mathbf{w}}(\boldsymbol{\epsilon})$ may have a different number of components than the randomness $\boldsymbol{\epsilon}$ itself. For simplicity, in the following we assume both $\boldsymbol{\epsilon}$ and $\hat{\mathbf{w}}(\boldsymbol{\epsilon})$ are D-dimensional. Note that $\hat{\mathbf{w}}(\cdot)$ does not depend on the batch element b; the dependence is captured in the coefficients $\mathbf{c}_{bi}$, which is a vector of length D for each index pair b, i.

Specifying the problem this way allows us to explicitly model each control variate coefficient per data point. Under this setting, the new estimator for the gradient is

$$\tilde{G}_i = \sum_{b \in \mathcal{B}} \tilde{g}_{bi}(\boldsymbol{\epsilon}_b) = \sum_{b \in \mathcal{B}} (\hat{g}_{bi}(\boldsymbol{\epsilon}_b) - \mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)). \quad (6)$$

The control variate coefficients $\mathbf{c}_{bi}$ can be set to optimally reduce the variance of $\tilde{G}_i$ by solving

$$\min_C \text{Tr}(\text{Cov}[\tilde{\mathbf{G}}]), \quad (7)$$

where C is the collection of $\mathbf{c}_{bi}$ and has shape $N \times P \times D$, as we need separate coefficients for all N data points.

Computing and storing these can be computationally prohibitive for large data sets, hence we propose to amortise the cost of this computation by using a recognition network $r_\phi : \mathcal{Y} \to \mathbb{R}^{P \times D}$ that outputs the coefficients for each mini-batch throughout the optimisation, where

$$\mathbf{c}_{bi} = [r_\phi(y_b)]_i \quad (8)$$

is a vector of dimension D and $y_b \in \mathcal{Y}$ are context points (e.g., feature vector and target for the bth data point in a supervised learning problem) and ϕ are the recognition network parameters.

As the control variate only adds terms to the gradients of the model's optimisation objective that are zero in expectation, we do not change the minima of the objective. This means that the extra parameters of the recognition network will *not* lead to overfitting.

2.3 TRAINING THE RECOGNITION NETWORK

Intuitively, we require the recognition network $r_\phi(\cdot)$ to output coefficients that minimise the variance of the controlled gradient estimator (6). This gives the training objective for the parameters ϕ:

$$\min_\phi \text{Tr}\left(\text{Cov}[\tilde{\mathbf{G}}]\right) = \min_\phi \sum_{i=1}^P \text{Var}[\tilde{G}_i]. \quad (9)$$

The ith term in the sum in (9) is

$$\begin{aligned} \text{Var}[\tilde{G}_i] &= \text{Var}\left[\sum_{b \in \mathcal{B}} \left(\hat{g}_{bi}(\boldsymbol{\epsilon}_b) - \mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\right)\right] \\ &= \sum_{b \in \mathcal{B}} \text{Var}\left[\hat{g}_{bi}(\boldsymbol{\epsilon}_b) - \mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\right] \\ &= \sum_{b \in \mathcal{B}} \Big(\text{Var}\left[\hat{g}_{bi}(\boldsymbol{\epsilon}_b)\right] + \text{Var}\left[\mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\right] \\ &\quad - 2\,\text{Cov}\left[\hat{g}_{bi}(\boldsymbol{\epsilon}_b), \mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\right]\Big) \\ &= \text{const} + \sum_{b \in \mathcal{B}} \Big(\mathbb{E}\left[(\mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))^2\right] \\ &\quad - 2\mathbb{E}\left[(\hat{g}_{bi}(\boldsymbol{\epsilon}_b))(\mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))\right]\Big), \quad (10) \end{aligned}$$

and we discard the terms that do not contain $\mathbf{c}_{bi}$ and hence do not give gradients for ϕ. For most problems, the expectations are intractable; we estimate these with MC sampling and define

$$\tilde{V}_i = \sum_{b \in \mathcal{B}} \left((\mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))^2 - 2(\hat{g}_{bi}(\boldsymbol{\epsilon}_b))(\mathbf{c}_{bi}^\top \hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))\right). \quad (11)$$

We can now learn the optimal recognition network parameters ϕ using SGD (or variants) on

$$\min_\phi \sum_{i=1}^P \tilde{V}_i. \quad (12)$$

To train the parameters ϕ, we need to compute gradients of $\sum_i \tilde{V}_i$. Examining the chain rule around the outputs of the recognition network, $\frac{\partial \tilde{V}_i}{\partial c_{bid}} \frac{\partial c_{bid}}{\partial \phi}$, the second term is

computed by backpropagation through the network, and the cost of computing the first term depends on the form of the estimator $\tilde{V}_i$.

The recognition network objective using (11) requires the partial gradients *per data point* $\hat{g}_{bi}(\boldsymbol{\epsilon}_b)$ of the original objective function, and we call this the **partial gradients estimator**. In common reverse-mode automatic differentiation libraries such as TensorFlow and PyTorch, it requires $|\mathcal{B}|$ additional backward passes on the model objective, each at least $O(|\mathcal{B}|)$, so this becomes prohibitively expensive when the mini-batch size is large. To overcome this limitation in current implementations, we derive two further estimators for the recognition network objective that are computationally cheaper, albeit with higher variance.

2.3.1 The Gradient Sum Estimator

To avoid the partial gradients in (11), we return to the ith term of the sum in (9). Instead of taking the sum out of the variance, we separate the sum over partial gradients from the control variates:

$$\text{Var}[\tilde{G}_i] = \text{Var}\Big[\sum_{b\in\mathcal{B}}\hat{g}_{bi}(\boldsymbol{\epsilon}_b) - \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\Big]$$
$$= \text{Var}\Big[\hat{G}_i - \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\Big]. \qquad (13)$$

We can expand the variance of a sum of two terms as

$$\text{Var}[\tilde{G}_i] = \text{Var}[\hat{G}_i] + \text{Var}\Big[\sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\Big]$$
$$- 2\,\text{Cov}\Big[\hat{G}_i, \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\Big]$$
$$= \text{const} + \sum_{b\in\mathcal{B}}\Big(\mathbb{E}[(\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))^2]$$
$$- 2\mathbb{E}[(\hat{G}_i)(\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))]\Big), \qquad (14)$$

and by replacing the expectations with MC estimates, we arrive at a new estimator

$$\tilde{V}_i^{\text{GS}} = \sum_{b\in\mathcal{B}}\big((\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))^2 - 2(\hat{G}_i)(\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b))\big). \quad (15)$$

This estimator is similar in form to the partial gradients estimator, replacing the gradient per data point with the sum over the whole mini-batch; we call this the **gradient sum** estimator. As it does not require any additional backward passes, it is much cheaper to compute. One can intuitively see that this estimator has a higher variance than the partial gradients estimator as it additionally includes cross terms that would be zero in expectation.

2.3.2 The Squared Difference Estimator

Alternatively, we can continue from (13) by expanding the variance into moment expectations:

$$\text{Var}[\tilde{G}_i] = \mathbb{E}\Big[\big(\hat{G}_i - \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\big)^2\Big]$$
$$- \big(\mathbb{E}\big[\hat{G}_i - \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\big]\big)^2,$$

where the control variate term has no contribution inside the second expectation by definition, and $\mathbb{E}[\hat{G}_i]$ is a constant with respect to the recognition network parameters ϕ. Evaluating the remaining expectation using MC gives us the **squared difference** estimator:

$$\tilde{V}_i^{\text{SD}} = \big(\hat{G}_i - \sum_{b\in\mathcal{B}}\mathbf{c}_{bi}^\top\hat{\mathbf{w}}(\boldsymbol{\epsilon}_b)\big)^2, \qquad (16)$$

which is also cheap to compute. In contrast to $\tilde{V}_i^{\text{GS}}$, it includes the second moment of $\hat{G}_i$. This is similar to a regression problem that uses $\hat{w}_d(\boldsymbol{\epsilon}_b)$ as basis functions to learn the gradient $\hat{G}_i$.

3 ILLUSTRATIVE EXAMPLE: CONTROL VARIATES FOR GAUSSIAN BASE RANDOMNESS

So far our discussion has been general. To implement a control variate, we need to specify both the distribution of the base randomness $\boldsymbol{\epsilon}$ and the functional form of the control variate $\mathbf{w}(\boldsymbol{\epsilon}_n)$. In principle, any functional form for control variates from the literature can be used with this method, e.g. Paisley et al. (2012); Ranganath et al. (2014); Miller et al. (2017). For the sake of simplicity, we illustrate our proposal on a simpler control variate form for the special case of Gaussian base randomness, which is of direct interest to many applications in VI.

We assume $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathrm{I}_D)$ without loss of generality.[†] In this section we introduce explicit forms for $\mathbf{w}(\boldsymbol{\epsilon}_n)$ for this case, starting with linear control variates, and then extend the discussion to higher-order polynomials.

3.1 LINEAR GAUSSIAN CONTROL VARIATES

The simplest control variate is an element-wise linear function of $\boldsymbol{\epsilon}_n$,

$$\mathbf{w}(\boldsymbol{\epsilon}_n) = \boldsymbol{\alpha} + \boldsymbol{\beta} \circ \boldsymbol{\epsilon}_n, \boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathrm{I}_D), \qquad (17)$$

with $\circ$ representing the element-wise product. Its expectation is $\mathbf{W} = \mathbb{E}[\mathbf{w}(\boldsymbol{\epsilon}_n)] = \boldsymbol{\alpha}$, and the control variate

[†]In the general case where $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, we can simply apply the reparameterisation $\boldsymbol{\epsilon} = \boldsymbol{\mu} + \text{Cholesky}(\Sigma)\boldsymbol{\epsilon}_0$ with $\boldsymbol{\epsilon}_0 \sim \mathcal{N}(\mathbf{0}, \mathrm{I}_D)$.

simplifies to $\boldsymbol{\beta} \circ \boldsymbol{\epsilon}_n$. We can also absorb $\boldsymbol{\beta}$ into the control variate coefficient $\mathbf{c}_{ni}$, which results in the following controlled version of the gradient component i, for data point n:

$$\tilde{g}_{ni}(\boldsymbol{\epsilon}_n) = \hat{g}_{ni}(\boldsymbol{\epsilon}_n) - \mathbf{c}_{ni}^\top \boldsymbol{\epsilon}_n. \tag{18}$$

Intuitively, one can think of control variates of this form as injecting the estimator with information on the linear dependence of the gradient on the noise. To understand this further, we take a look at the first-order Taylor expansion of the gradient component $g_{ni}(\boldsymbol{\epsilon}_n)$ around $\boldsymbol{\epsilon}_n = \mathbf{0}$,

$$g_{ni}(\boldsymbol{\epsilon}_n) = g_{ni}(\mathbf{0}) + \nabla g_{ni}(\mathbf{0})^\top \boldsymbol{\epsilon}_n + O(\boldsymbol{\epsilon}_n^2). \tag{19}$$

If the gradient is sufficiently linear with respect to $\boldsymbol{\epsilon}_n$ (i.e., the $O(\boldsymbol{\epsilon}_n^2)$ terms are negligible), and when $\mathbf{c}_{ni}$ is a good approximation to the Jacobian at $\mathbf{0}$, the estimator in (18) will have low variance.

3.2 HIGHER-ORDER POLYNOMIALS

In general, the gradient is unlikely to be linear with respect to the noise, especially for complicated models and objectives. To overcome this, we can use higher-order polynomials to capture some of the non-linear dependence of the gradient on the noise. Consider the following form for $\mathbf{w}(\boldsymbol{\epsilon}_n)$:

$$\mathbf{w}(\boldsymbol{\epsilon}_n) = \sum_{k=1}^{K} \boldsymbol{\alpha}_k \circ \boldsymbol{\epsilon}_n^k, \tag{20}$$

where the kth power is evaluated element-wise. $\mathbf{W}$ can be easily computed and would correspond to the sum of diagonal parts of the first K moment tensors of the multivariate Gaussian distribution, scaled by $\boldsymbol{\alpha}_k$. For instance, for $K = 2$ the control variate is given by

$$\boldsymbol{\alpha}_1 \circ \boldsymbol{\epsilon}_n + \boldsymbol{\alpha}_2 \circ (\boldsymbol{\epsilon}_n^2 - \text{diag}(I_D)). \tag{21}$$

We can again simplify by absorbing the $\boldsymbol{\alpha}_k$ into the control variate coefficient, with slight adjustments to the controlled gradient estimator.

We make the following observation:

Remark 1: A linear combination of control variates is also a valid control variate, i.e., $\tilde{g}_{ni}(\boldsymbol{\epsilon}_n) = \hat{g}_{ni}(\boldsymbol{\epsilon}_n) - \sum_{k=1}^{K} (\mathbf{c}_{ni}^{(k)})^\top (\hat{\mathbf{w}}_k(\boldsymbol{\epsilon}) - \mathbf{W}_k)$ is unbiased. By considering each term in (21) as a separate control variate, we can write the ith component of the controlled gradient at n as

$$\tilde{g}_{ni}(\boldsymbol{\epsilon}_n) = \hat{g}_{ni}(\boldsymbol{\epsilon}_n) - (\mathbf{c}_{ni}^{(1)})^\top \boldsymbol{\epsilon}_n - (\mathbf{c}_{ni}^{(2)})^\top (\boldsymbol{\epsilon}_n^2 - \text{diag}(I_D)). \tag{22}$$

The same construction trivially extends to $K > 2$.

3.3 BRIEF DISCUSSION

The simple examples of the linear and polynomial control variates presented above illustrate the importance of choosing a good control variate coefficient C. For instance, in the linear case in (18) the control variate function $\mathbf{w}(\boldsymbol{\epsilon}_n) = \boldsymbol{\epsilon}_n$ does not provide any extra information on the estimator on its own, since we are essentially just adding noise to the MC estimate. However, with the selection of a good control variate coefficient $\mathbf{c}_n$ for data point n, we introduce structure to the noise that contains information about the behaviour of the controlled quantity with respect to the Gaussian noise in the form of the Jacobian in (19). Indeed the optimal coefficient $\mathbf{c}_n^*$ for the linear control variate contains the Jacobian term.

4 RELATED WORK

Control variates are widely used to reduce the gradient variance of stochastic objectives, mainly motivated by VI. A comprehensive review can be found in Geffner and Domke (2018). Here, we highlight some relevant work and compare it to our contribution.

Paisley et al. (2012) first introduce the idea of using control variates to reduce the gradient variance in VI. They propose using a bound on the objective or an approximation of the model as control variates. Ranganath et al. (2014) build on this work, using the score function of the approximate posterior to control the gradient of Black Box Variational Inference objectives.

Inspiration for our work comes from Grathwohl et al. (2018), where they use a recognition network to approximate the model and its gradient as a control variate. Miller et al. (2017) derive an approximation to the reparameterisation gradient for Gaussian variational distributions by performing a first-order Taylor expansion of the gradient, using this approximation as a control variate. Our work is related to this construction where the recognition network can be viewed as a cheap approximation to the linear term in the Taylor expansion of the gradient (i.e., the Hessian of the model objective) in the case of the linear construction of Section 3.

The unifying work of Geffner and Domke (2018) categorises different control variate schemes for VI objectives. Additionally, they propose combining them to achieve greater variance reduction. They derive an optimal rule for this combination based on Bayesian risk minimisation.

These related works do not consider the effect of minibatching on the proposed control variates. Our work should be viewed as *complementary* to many of the methods mentioned above. Indeed, Geffner and Domke (2018) show that a combination of control variates is usually

more desirable that a single scheme. The method we proposed can be considered an extra addition to the control variate toolkit for doubly stochastic objectives, to take the effect of mini-batch stochasticity on the control variates into account. Our method can also be combined with other variance reduction methods such as extra sampling.

5 EXPERIMENTS

Our discussion thus far applied to the general class of doubly stochastic objectives. For our experiments we focus on objectives arising from VI problems. Amortising the computation of the control variate coefficients in this setting is advantageous since context arises naturally from the data in the underlying models.

In this section, we aim to answer three questions: a) To what extent can amortising with a recognition network reduce the variance compared to a fixed context-free control variate coefficient? b) How well can we train the recognition network in an online setting? c) What difference can an amortised control variate make in practice?

5.1 SETUP

We investigate (a), (b) and (c) on a classification task on the *titanic* dataset using a Bayesian logistic regression model and on a regression task on the *airfoil* dataset using a Deep Gaussian Process (DGP).

For the Bayesian logistic regression model, we use the reparameterisation gradient formulation of the VI problem. We choose a Gaussian approximate posterior, where we learn the mean vector and the full covariance matrix. We select a unit Gaussian prior on the weights.

For the DGP model, we use a 2-layer model with inner layer dimension of 5, and a Squared Exponential kernel for the GP priors. We use the doubly stochastic formulation of the VI problem (Salimbeni and Deisenroth, 2017). We learn the parameters of the approximate Gaussian posterior, keeping the hyperparameters fixed. The inducing locations are fixed and selected as the centroids of k-means clusters from the data.

Throughout, we use Adam (Kingma and Ba, 2015) for both the model objective function optimisation and the recognition network objective optimisation. We use a single-sample MC estimate of the gradients and control these when stated, applying the linear and quadratic control variates introduced in Section 3. We initialise the recognition network with Xavier initialisation (Glorot and Bengio, 2010) and use ReLU activations in the hidden layers.

We compare our proposal to a context-free control variate.

In this instance, this is implemented as an optimisable quantity that does not depend on data and uses the same optimisation objectives ((14) & (16)) as the recognition network, i.e. $\mathbf{c}$ is independent of the mini-batch $\mathcal{B}$ in these objectives. This is equivalent to approximating the coefficient with an exponentially weighted moving of the empirical covariance of the gradient and the control variate estimates.

5.2 VERIFICATION OF VARIANCE REDUCTION

The first question considered is whether the recognition network has the capacity to amortise the control variate coefficients and how well it can learn these versus a context-free coefficient? To test this, we freeze the model parameters at three points in the optimisation – early (10 steps), mid (200) steps, and late (1000 steps) – then optimise the control variate coeficient only. For each period, we iteratively sample a gradient value then perform an optimisation step on the recognition network. We repeat this procedure for 1000 steps and record the variance reduction at different steps. The variance reduction is measured by the ratio $\mathrm{Var}[\|\tilde{G}_N\|]/\mathrm{Var}[\|\hat{G}_N\|]$, where $\tilde{G}_N$ and $\hat{G}_N$ are the controlled and uncontrolled gradients, respectively, over the mini-batch (size 10), and $\|\cdot\|$ is the gradient norm. We compare different network sizes to see the effect this has on variance reduction.

Figure 2 shows that amortising the control variate coefficient computation induces greater variance reduction than optimising a context-free coefficient (labelled as *None* in the figure). The variance reduction does not occur immediately, as the control variate coefficients need to be optimised in all cases to reduce the variance. Also notable is that the amount of variance reduction depends on the optimisation stage of the model; at later stages of the model optimisation, the variance reduction is more pronounced. This is likely a property of both the model and the control variate where the gradients in the beginning of the optimisation have more pronounced non-linearities with respect to the noise. This can also be seen in the amount of variance reduction in logistic regression compared to the DGP. The gradients in the logistic regression models are approximately linear with respect to the noise, while in the DGP gradients have a more complex dependency on the noise. Finally, we can see that the variance reduction potential depends on the capacity of the network, where wider and deeper networks learn better control variate coefficients. Deeper networks reduce the variance more strongly than wider networks, which correspond to a highly non-linear mapping from the context points to the control variate coefficient.
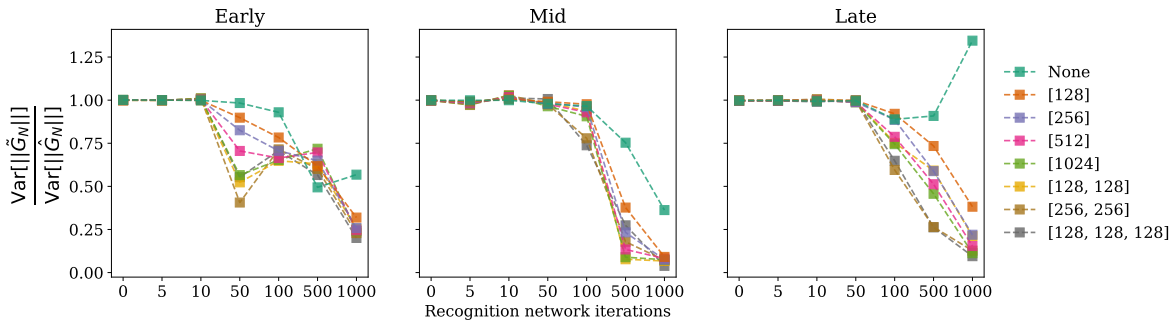
(a) Logistic regression on *titanic*.
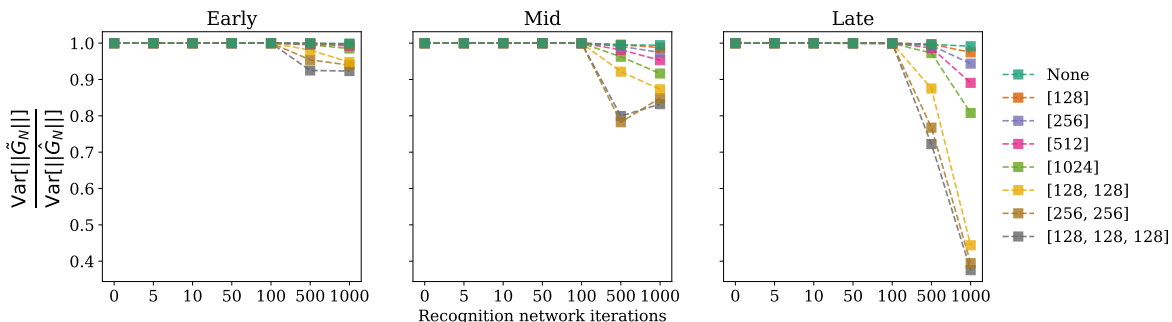
(b) DGP on *airfoil*.

Figure 2: Variance reduction at different points in the objective optimisation (lower is better); early: 10 steps, mid: 200 steps, late: 1000 steps. The results are shown for the *linear* control variate from Section 3. Recognition network training uses the *squared difference* objective optimised with Adam with learning rate of 10^{-2} for the logistic regression and 10^{-3} for the DGP. For a small number of iterations on the recognition network, it struggles to learn a good control variate coefficient. Continuing the network optimisation, it is able to learn good control variate coefficients that significantly reduce the variance in comparison to the context-free coefficient. Also notable is that the variance reduction is more pronounced at the later stages of the model optimisation.

5.3 SIMULTANEOUS OPTIMISATION OF OBJECTIVE FUNCTION AND CONTROL VARIATE COEFFICIENT

In practice, the recognition network needs to be able to learn the control variate coefficients while the model is being optimised, giving a moving target. In this section, we investigate the viability of chasing this target by simultaneously optimising the model objective and recognition network. We use a recognition network with three layers of size 128 each, as this architecture showed the largest variance reduction in Section 5.2. In each step in the optimisation procedure, we compute one gradient estimate of the model objective for a mini-batch of size 10. We take one Adam step on the recognition network, then we apply the control variate correction to the sampled gradient and take an Adam step on the model parameters. We measure the variance of the gradient at different periods in the optimisation by sampling 100 gradient values at each period and taking the empirical variance of their norm.

The recognition network is able to learn good control variate coefficients in this dynamic regime, see Figure 3. The variance reduction improves later on in the optimisation as observed in Section 5.2. We again observe that the amortised control variate results in greater variance reduction than the context-free one.

5.4 APPLICATION

To show how our approach works in practice, we use it for training the logistic regression and DGP models. We apply the alternating optimisation procedure described in Section 5.3 on each for 2000 iterations with mini-batches size of 10. We record the mean value of the Negative Evidence Lower Bound (NELBO) from 100 MC samples for the logistic regression and 10 MC samples for the DGP at every iteration computed on the entire datasets.

The resulting traces are shown in Figure 4; in both cases we see that the optimisation with controlled gradients
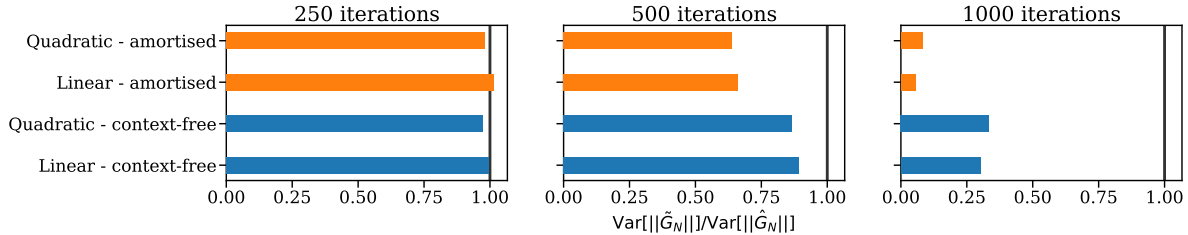
(a) Logistic regression on *titanic*.
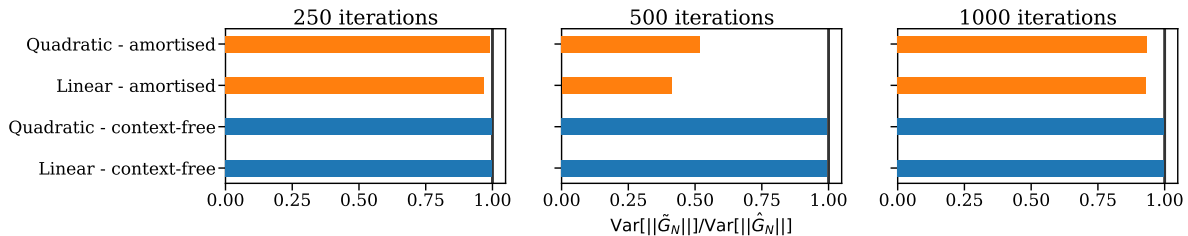
(b) DGP on *airfoil*.

Figure 3: Gradient variance ratio at three different points in joint optimisation of the model and control variate parameters (lower is better); the vertical line corresponds to a ratio of 1 (i.e. no reduction). Both the model and the control variate objectives are optimised with Adam with learning rate of 10^{-2} for the model and 10^{-2} and 10^{-3} for the logistic regression and DGP control variate coefficients respectively. The recognition network is able to learn a good control variate coefficient and continues to improve throughout the optimisation. The recognition network outperforms the context-free control variate.

starts off in a worse regime than the uncontrolled gradients (curves on or above the dashed line); however, it improves as better control variate coefficients are learned. The gap between the one-sample MC estimator and the controlled estimators widens later for logistic regression, and fluctuates for the DGP. This is because the linear control variate sufficiently approximates the dependence of the gradient on the randomness, whereas in the case of the DGP this dependence is more complex.

For both models, amortising the control variate coefficients result in lower NELBO values on average in comparison to the uncontrolled and the context-free controlled cases. We also see that the optimisation of the control variate coefficients is robust to the choice of objective function, with similar behaviour for the gradient sum and squared difference objectives for both the amortised and context-free cases.

Table 1 shows the average cost for the controlled optimisation steps for the two problems. Amortising the control variate coefficients with a recognition network of size [128, 128, 128] has an additional overhead of around 25% on the context-free coefficient on the CPU. The overhead depends on many factors such as the recognition network size, control variate formulation, mini-batch size and number of gradient components. These should all be taken into account when implementing this scheme.

Table 1: Average overall optimisation step time in *milliseconds* (on the CPU) for logistic regression and DGP for different *linear* control variate objective functions. The statistics are computed based on 100 repetitions of 10 runs. The implementation uses TensorFlow 2.0.

| Method	Logistic	DGP
Squared diff. - amortised	1.20(0.64)	3.77(0.22)
Grad. sum - amortised	1.25(0.11)	3.78(0.19)
Squared diff. - context-free	0.87(0.91)	3.17(0.13)
Grad. sum - context-free	0.84(0.77)	3.07(0.82)

6 CONCLUSIONS

We introduced a control variate formulation that exploits the structure of doubly stochastic objectives to remove Monte Carlo sampling variance from mini-batch gradient estimators. We proposed three objectives for an amortising recognition network that can learn context aware control variate coefficients. Training the network re-uses the gradients of the model objective and does not require additional passes through the model.

Empirical assessment showed that an approximation to
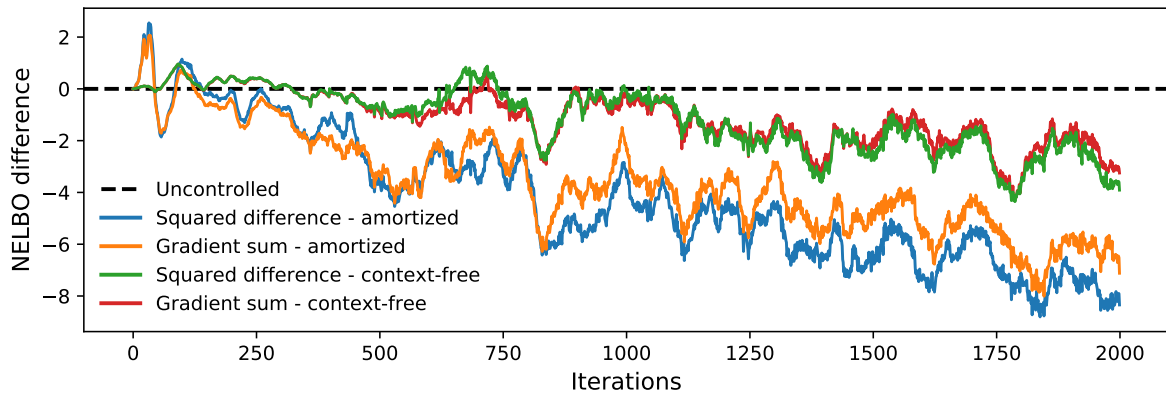
(a) Logistic regression on *titanic*.
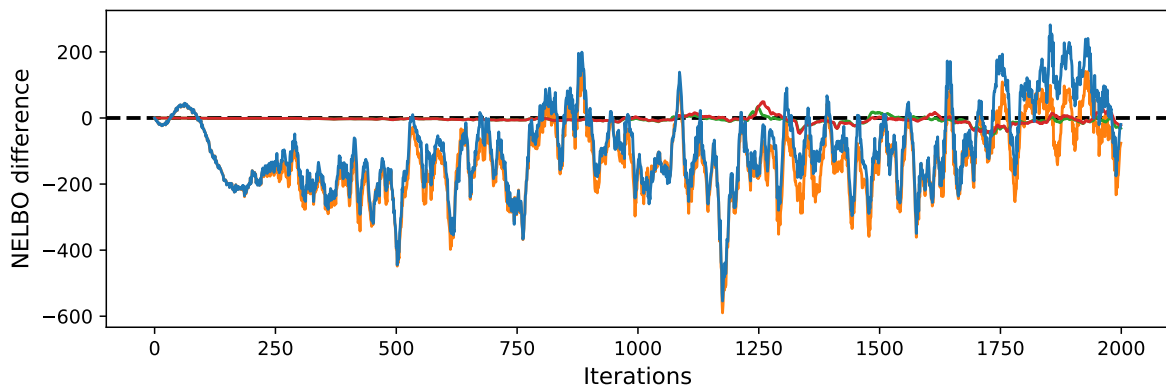
(b) DGP on *airfoil*.

Figure 4: Difference between optimisation traces for different control variate objectives, using the uncontrolled one-sample MC estimate of the gradient as a baseline (lower is better). *Linear* control variates are used in this experiment. The gap between the baseline and controlled models widens through the optimisation. Amortised control variate coefficients result in wider gaps indicating better optimisation performance.

the optimal control variate per mini-batch can be performed during optimisation and reduces the gradient variance in practice compared to a context-free global approach. In our experiments we used linear and quadratic control variates for Gaussian base randomness, but our approach is general and can be applied to other control variate formulae and randomness schemes.

References

T. Geffner and J. Domke. Using large ensembles of control variates for variational inference. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9960–9970. Curran Associates, Inc., 2018.

X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.

W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In *International Conference on Learning Representations*, 2018.

R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 315–323. Curran Associates, Inc., 2013.

D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.

D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

A. Miller, N. Foti, A. D' Amour, and R. P. Adams. Reducing reparameterization gradient variance. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3708–3718. Curran Associates, Inc., 2017.

J. Paisley, D. Blei, and M. Jordan. Variational bayesian inference with stochastic search. In J. Langford and J. Pineau, editors, *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, ICML '12, pages 1367–1374, New York, NY, USA, July 2012. Omnipress. ISBN 978-1-4503-1285-1.

R. Ranganath, S. Gerrish, and D. Blei. Black Box Variational Inference. In S. Kaski and J. Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 814–822, Reykjavik, Iceland, 22–25 Apr 2014. PMLR.

H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

G. Roeder, Y. Wu, and D. K. Duvenaud. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6925–6934. Curran Associates, Inc., 2017.

H. Salimbeni and M. Deisenroth. Doubly stochastic variational inference for deep gaussian processes. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4588–4599. Curran Associates, Inc., 2017.

Supplementary Material

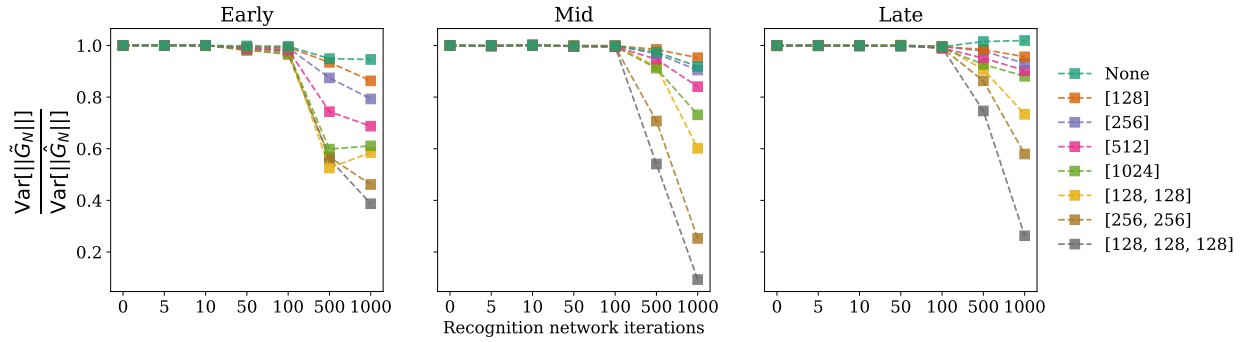## A VERIFICATION OF VARIANCE REDUCTION

Figure 5: Logistic regression. Squared difference objective. Recognition network learning rate $= 10^{-3}$.
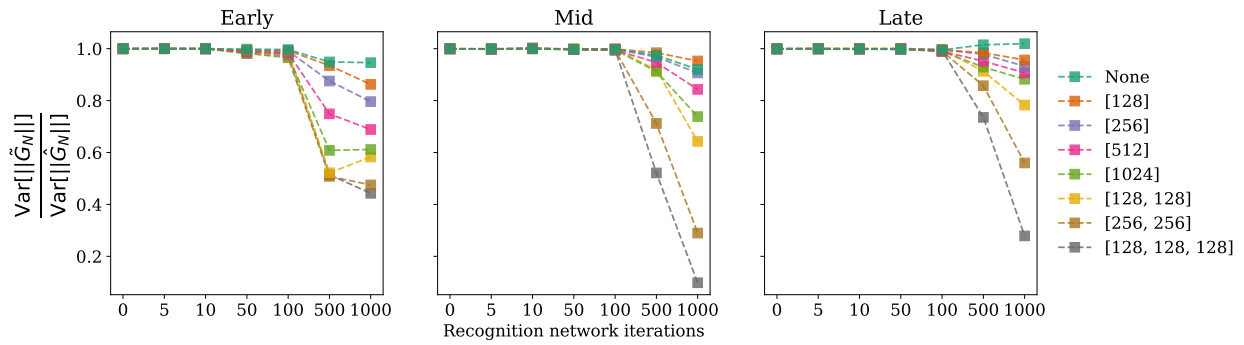

Figure 6: Logistic regression. Gradient sum objective. Recognition network learning rate $= 10^{-3}$.
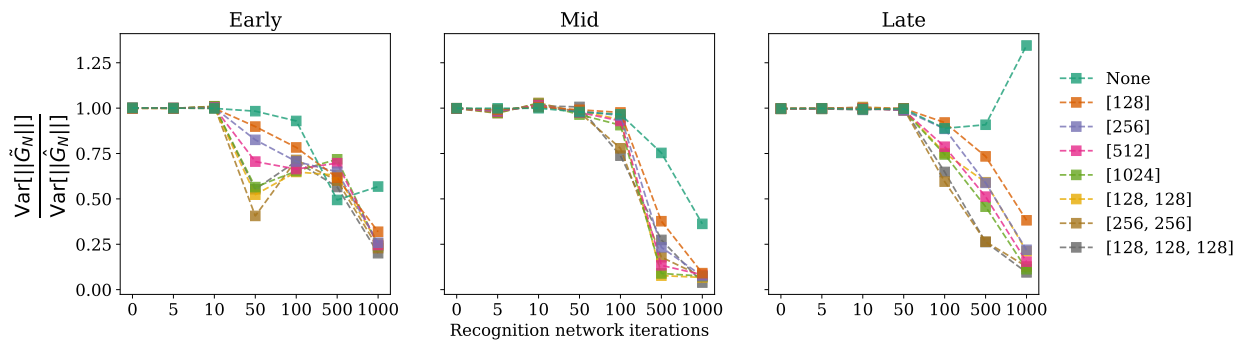

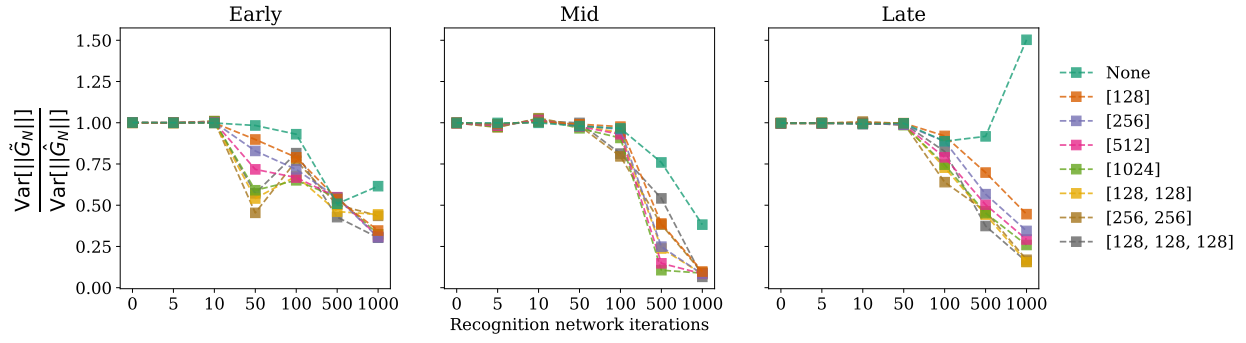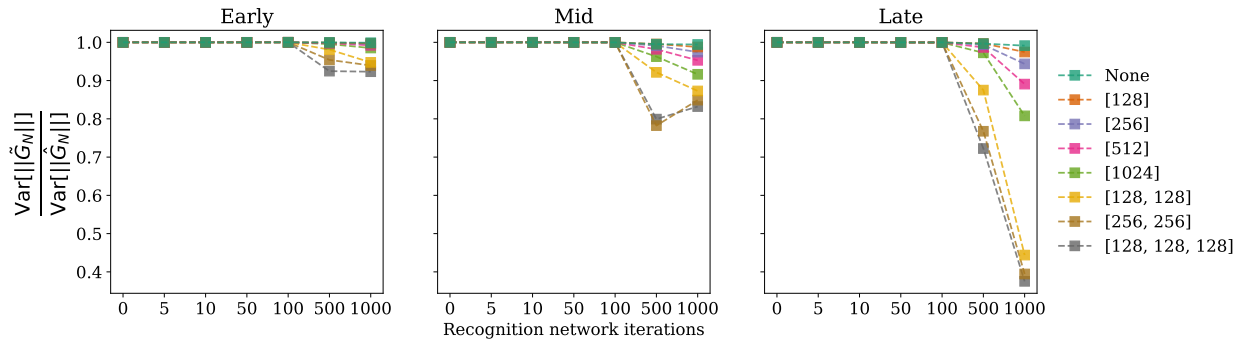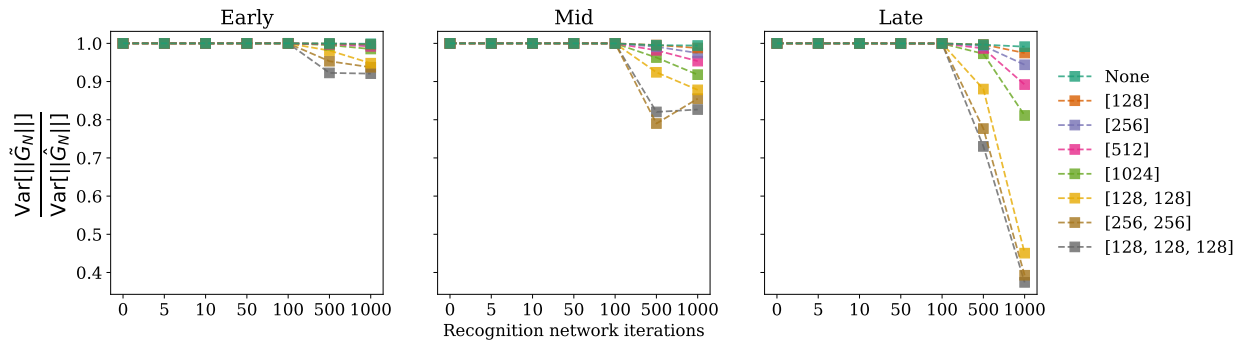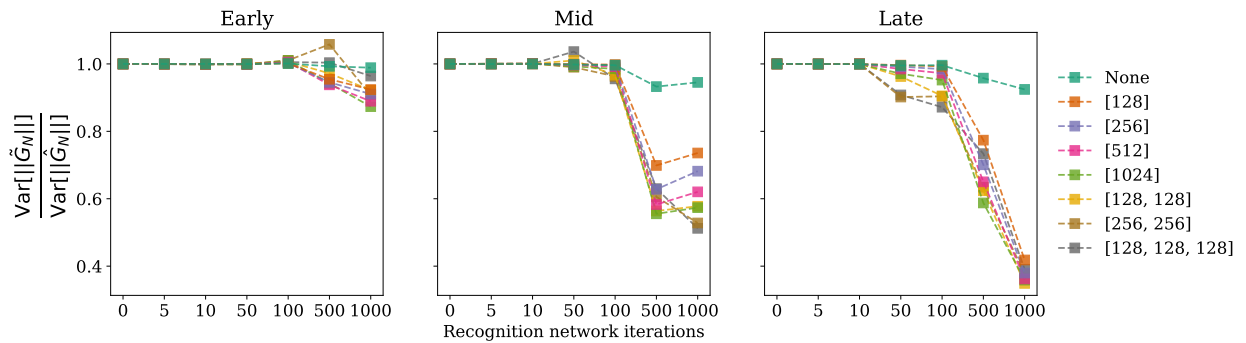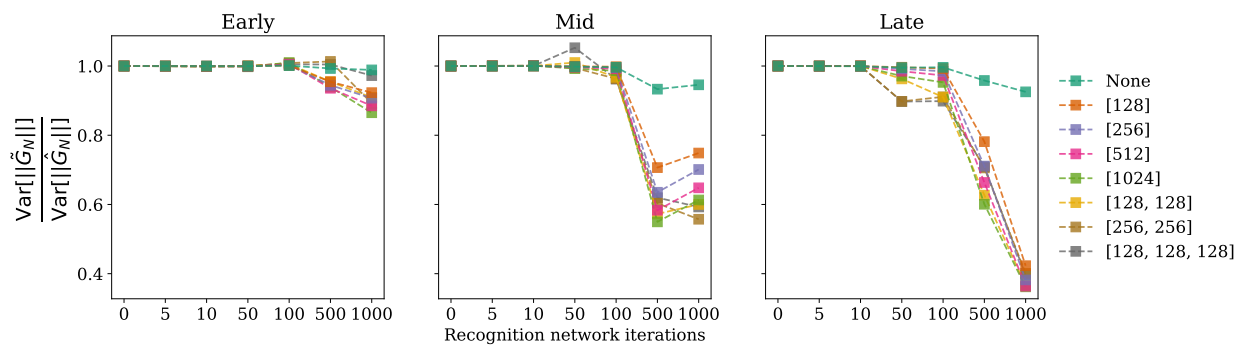Figure 7: Logistic regression. Squared difference objective. Recognition network learning rate $= 10^{-2}$.

Figure 8: Logistic regression. Squared difference objective. Recognition network learning rate $= 10^{-2}$.

Figure 9: DGP. Squared difference objective. Recognition network learning rate $= 10^{-3}$.

Figure 10: DGP. Gradient sum objective. Recognition network learning rate $= 10^{-3}$.

Figure 11: DGP. Squared difference objective. Recognition network learning rate $= 10^{-2}$.

Figure 12: DGP. Gradient sum objective. Recognition network learning rate $= 10^{-2}$.

B   SIMULTANEOUS OPTIMISATION OF MODEL AND RECOGNITION NETWORK
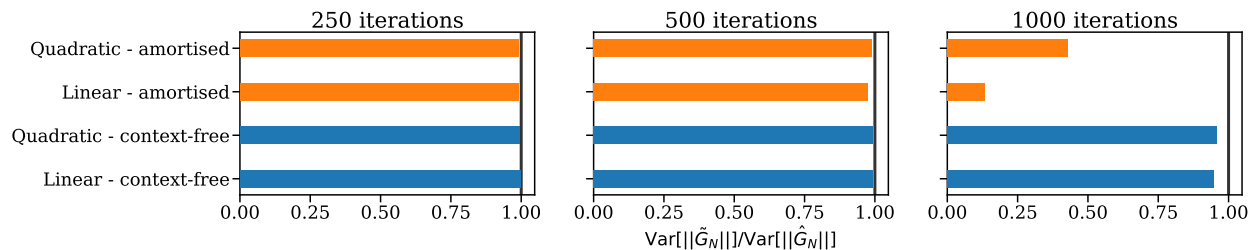
Figure 13: Logistic regression. Squared difference objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-3}.
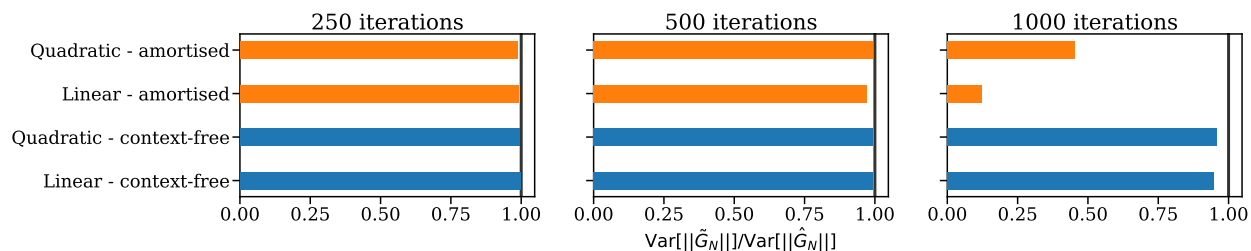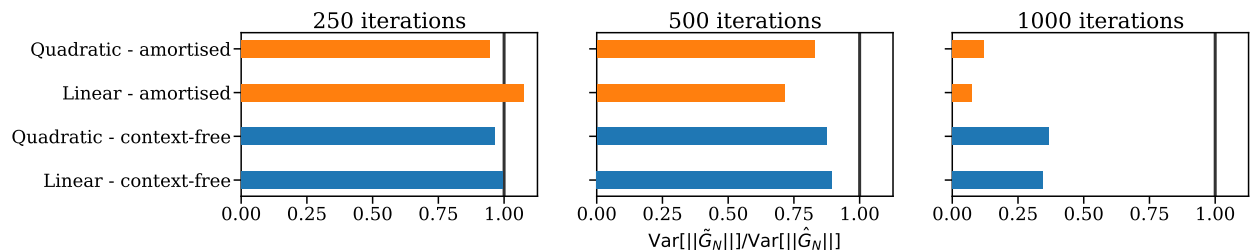

Figure 14: Logistic regression. Gradient sum objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-3}.
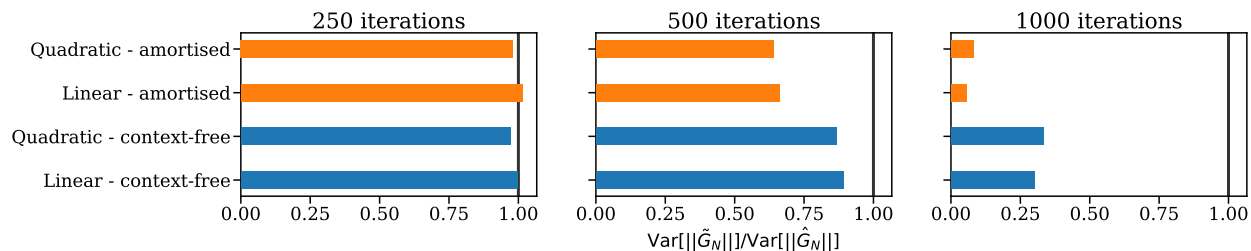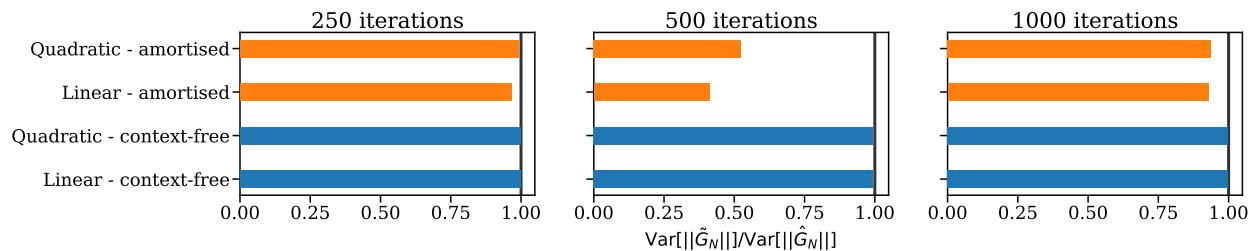

Figure 15: Logistic regression. Squared difference objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-2}.

Figure 16: Logistic regression. Gradient sum objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-2}.

Figure 17: DGP. Squared difference objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-3}.
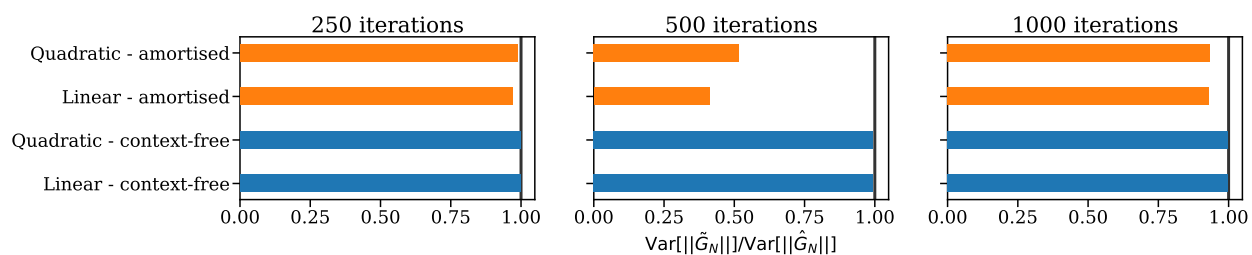

Figure 18: DGP. Gradient sum objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-3}.
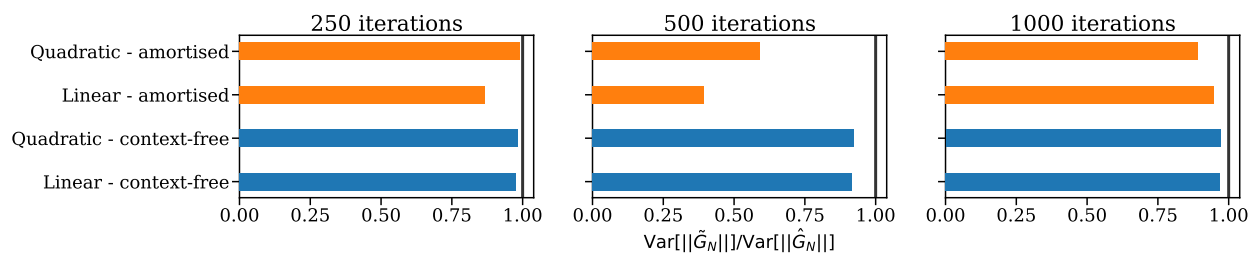

Figure 19: DGP. Squared difference objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-2}.
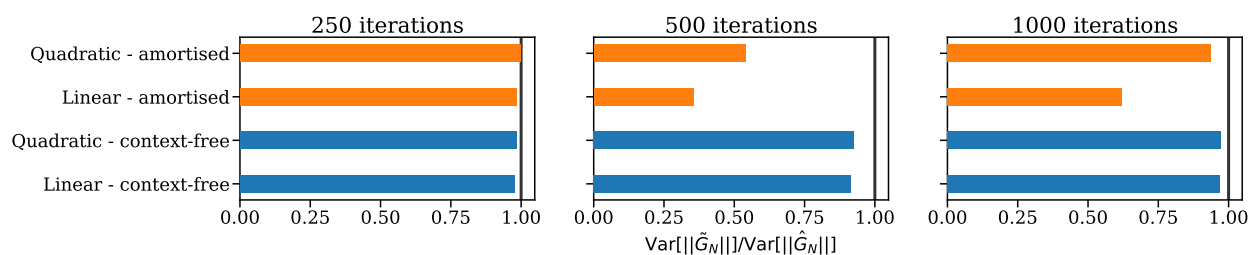

Figure 20: DGP. Gradient sum objective. Network of size [128, 128, 128]. Recognition network learning rate = 10^{-2}.