

DPVIm: Differentially Private Variational Inference Improved

Anonymous authors

Paper under double-blind review

Abstract

Differentially private (DP) release of multidimensional statistics typically considers an aggregate sensitivity, e.g. the vector norm of a high-dimensional vector. However, different dimensions of that vector might have widely different magnitudes and therefore DP perturbation disproportionately affects the signal across dimensions. We observe this problem in the gradient release of the DP-SGD algorithm when using it for variational inference (VI), where it manifests in poor convergence as well as high variance in outputs for certain variational parameters, and make the following contributions: (i) We mathematically isolate the cause for the difference in magnitudes between gradient parts corresponding to different variational parameters. Using this as prior knowledge we establish a link between the gradients of the variational parameters, and propose an efficient while simple fix for the problem to obtain a less noisy gradient estimator, which we call *aligned* gradients. This approach allows us to obtain the updates for the covariance parameter of a Gaussian posterior approximation without a privacy cost. We compare this to alternative approaches for scaling the gradients using analytically derived preconditioning, e.g. natural gradients. (ii) We suggest using iterate averaging over the DP parameter traces recovered during the training, to reduce the DP-induced noise in parameter estimates at no additional cost in privacy. Finally, (iii) to accurately capture the additional uncertainty DP introduces to the model parameters, we infer the DP-induced noise from the parameter traces and include that in the learned posteriors to make them *noise aware*. We demonstrate the efficacy of our proposed improvements through various experiments on real data.

1 Introduction

Differential privacy (DP) (Dwork et al., 2006) protects privacy of data subjects by limiting how much about the input data can be learned from the output of an algorithm. Additive noise mechanisms achieve DP by adding noise calibrated to the maximum change in function output due to a single individual, known as sensitivity. When releasing high-dimensional data through such mechanisms, different variables may have widely different sensitivities, which is typically ignored or unknown. Instead, the sensitivity of the release is computed as an aggregate over all the dimensions, which we call *global sensitivity*, in contrast to variable-specific *local sensitivity*. As the DP noise is subsequently scaled with this global sensitivity, it affects dimensions with lower local sensitivities more. A prominent example where this occurs is the gradient release in DP stochastic gradient descent (DP-SGD) (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016). Furthermore, the final parameters released from DP-SGD are noisy estimators of the optimal parameters and the noisy deviation is usually treated as an unavoidable trade-off of providing privacy (Abadi et al., 2016). This means that large amounts of non-quantified noise for certain parameter dimensions can lead to large errors in the results of DP-SGD as well as making them sensitive to both, initialisation and luck of the draw in the last DP gradient estimate.

We discover these issues in the perturbed gradients used in DP variational inference (DPVI) (Jälkö et al., 2017) where they lead to severe errors in capturing posterior uncertainty. This results e.g. in poor predictive uncertainty estimation, making the predictions of the learned model less accurate. DPVI is a widely applicable state-of-the-art algorithm for privacy-preserving (approximate) Bayesian inference, based on DP-SGD.

We mathematically isolate the cause for these problems in DPVI and propose and evaluate two ways of alleviating the problem of gradient scales in DPVI: one scales gradients with a preconditioning matrix before applying the DP mechanism, the other is based on insights into the mathematical structure of the gradients, which reveals that their components are mathematically linked and can be derived from each other in a post-processing step.

Additionally, we theoretically and experimentally evaluate the method of iterate averaging as a way to further improve the parameter estimate as well as approximate the additional variance induced by DP perturbations to DPVI to make the posterior approximation noise aware at no additional cost in privacy.

1.1 Related work

In the context of DP-SGD, the following previous works acknowledge the different local sensitivities of different parts of the full gradient: McMahan et al. (2018) suggested clipping the gradients of a neural network separately for each layer to avoid the clipping-induced bias (Chen et al., 2020). Other lines of work (Andrew et al., 2021; Wang et al., 2022) suggest adaptive clipping, where the global sensitivity is re-evaluated throughout the optimisation process to avoid adding excessive amounts of noise to the gradients. However, since in all these the perturbation is still scaled with the global sensitivity aggregated over the dimensions, this approach does not improve the disparate effect that the Gaussian noise will have on the dimensions with smaller gradients, so we see these approaches as orthogonal to our work.

For noise aware DP Bayesian inference, the most related work is by Bernstein and Sheldon (Bernstein and Sheldon, 2018; 2019) and Kulkarni et al. (2021). These works include the DP perturbation mechanism into a probabilistic model using perturbed sufficient statistics as the inputs. This allows capturing the DP-induced additional uncertainty in the posterior distribution of model parameters.

2 Preliminaries

2.1 Differential privacy

Definition 2.1. For $\epsilon \geq 0$ and $\delta \in [0, 1]$, a randomised mechanism \mathcal{M} satisfies (ϵ, δ) -differential privacy (Dwork et al., 2006) if for any two data sets different in only one element, $D, D' \in \mathcal{D}$, and for all outputs $S \subseteq \text{im}(\mathcal{M})$, the following constraint holds:

$$\Pr(\mathcal{M}(D) \in S) \leq e^\epsilon \Pr(\mathcal{M}(D') \in S) + \delta. \quad (1)$$

2.2 Variational inference

Variational inference is a commonly applied technique in probabilistic inference, where the aim is to learn an approximation for a (typically intractable) posterior distribution of the parameters of a probabilistic model (Jordan et al., 1999). This is done by maximising a quantity called *evidence lower bound* (ELBO) over the parameters of the variational approximation. For a probabilistic model $p(D, \theta)$, where D denotes the observed variables and θ the model parameters, and for a variational approximation $q(\theta)$ of the posterior, the ELBO is given as

$$\mathcal{L}(q) = \log p(D) - \text{KL}(q(\theta) \parallel p(\theta \mid D)) = \mathbb{E}_{q(\theta)} [\log p(D, \theta)] + H(q), \quad (2)$$

where KL denotes the Kullback-Leibler divergence and H the (differential) entropy.

In the following we first restrict ourselves to the commonly used *mean-field* variational inference, i.e., using a Gaussian with diagonal covariance as the posterior approximation. We will later generalise this to a full-rank covariance approximation. For d -dimensional data the diagonal approximation is parametrised by the means $\mathbf{m}_q \in \mathbb{R}^d$ and the dimension-wise standard deviations $\sigma_q \in \mathbb{R}^d$. We further reparametrise the model with $\mathbf{s}_q = T^{-1}(\sigma_q)$, where $T : \mathbb{R} \rightarrow \mathbb{R}_+$ is monotonic, in order to facilitate optimisation in an unconstrained domain. Both T and T^{-1} are applied element-wise for each of the parameters. Common choices for T are the exponential function $T(s) = \exp(s)$ and the softplus function $T(s) = \log(\exp(s) + 1)$ (used e.g. in the Pyro

probabilistic programming package (Bingham et al., 2019)). We use $\xi = (\mathbf{m}_q, \mathbf{s}_q)$ to refer to the complete set of variational parameters.

A draw from this posterior distribution can then be written as (Kingma and Welling, 2014):

$$\boldsymbol{\theta} := \boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q) = \mathbf{m}_q + T(\mathbf{s}_q)\boldsymbol{\eta}, \quad (3)$$

where $\boldsymbol{\eta} \sim N(0, I_d)$, $\boldsymbol{\theta}, \boldsymbol{\eta} \in \mathbb{R}^d$ and I_d is a d -dimensional identity matrix. Kucukelbir et al. (2017) use this reparametrisation trick together with single-sample MC integration to give the ELBO a differentiable form with gradients:

$$\mathbf{g}_m := \nabla_{\mathbf{m}_q} \mathcal{L}(q) = \nabla_{\mathbf{m}_q} \log p(D, \boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) \quad (4)$$

$$\mathbf{g}_s := \nabla_{\mathbf{s}_q} \mathcal{L}(q) = \nabla_{\mathbf{s}_q} \log p(D, \boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) + \nabla_{\mathbf{s}_q} H(q), \quad (5)$$

where $\boldsymbol{\eta} \sim N(0, I)$. Throughout this work we assume that the likelihood factorises as: $p(D | \boldsymbol{\theta}) = \prod_{\mathbf{x} \in D} p(\mathbf{x} | \boldsymbol{\theta})$. Using N to denote the size of D , we can now further decompose the gradients in (4) and (5) as

$$\mathbf{g}_m = \sum_{\mathbf{x} \in D} \left(\nabla_{\mathbf{m}_q} \log p(\mathbf{x} | \boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) + \frac{1}{N} \nabla_{\mathbf{m}_q} \log p(\boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) \right) \quad (6)$$

$$\mathbf{g}_s = \sum_{\mathbf{x} \in D} \left(\nabla_{\mathbf{s}_q} \log p(\mathbf{x} | \boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) + \frac{1}{N} (\nabla_{\mathbf{s}_q} \log p(\boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)) + \nabla_{\mathbf{s}_q} H(q)) \right). \quad (7)$$

We denote the per-example gradient components (i.e., those for each individual \mathbf{x}) that appear in the above sums with $\mathbf{g}_{m,\mathbf{x}}$ and $\mathbf{g}_{s,\mathbf{x}}$ respectively.

A common approach to performing variational inference in practice is to initialise \mathbf{s}_q to small values, which allows the algorithm to move \mathbf{m}_q quickly close to their optimal values due to large error in the KL term of the ELBO induced by the narrow approximation.

Assumptions for the probabilistic model and variational posterior For the rest of the paper, we make the following assumptions:

- **Exchangeability:** We assume that there is no ordering in observing the elements of data set D : $p(D | \boldsymbol{\theta}) = \prod_{\mathbf{x} \in D} p(\mathbf{x} | \boldsymbol{\theta})$
- **Gaussian posterior approximation:** $q(\boldsymbol{\theta}) = \mathcal{N}_d(\mathbf{m}_q, \Sigma_q)$, where \mathcal{N}_d denotes the pdf of a d -dimensional Gaussian. When working with isotropic Gaussian posterior approximation (i.e. diagonal covariance), we denote the dimension-wise standard deviations with $\boldsymbol{\sigma}_q$.
- **Optimising the $\boldsymbol{\sigma}_q$:** For the isotropic Gaussian, we use a mapping function $T : \mathbb{R} \rightarrow \mathbb{R}^+$ to optimize the variational standard deviations. We apply this function element-wise to the parameter vector \mathbf{s}_q . Commonly used examples for this are the softplus function $T(s) = \log(1 + \exp(s))$ as well as the exponential function $T(s) = \exp(s)$.

3 Differentially private variational inference

The first algorithm for differentially private variational inference for non-conjugate models (Jälkö et al., 2017) optimises the ELBO using gradients (4) and (5) with differentially private stochastic gradient descent (Abadi et al., 2016) to provide privacy. This involves concatenating each of the per-example gradients to obtain $\mathbf{g}_x = (\mathbf{g}_{m,x}^T, \mathbf{g}_{s,x}^T)^T$, clipping \mathbf{g}_x so that it has ℓ_2 norm no larger than a threshold C to limit the sensitivity, and finally adding Gaussian noise to the sum of these clipped per-example gradients to obtain $\tilde{\mathbf{g}}$, which is used for the parameter update. We refer to this algorithm in the following as *vanilla DPVI*.

This formulation induces a problem which, while seemingly minor at first glance, severely affects accuracy of solutions. We next isolate this problem and then propose a solution through detailed analysis on the gradients of the variational parameters.

3.1 Disparate perturbation of variational parameter gradients

While the clipping of the gradients allows us to bound the global sensitivity of the gradient vector, it completely ignores any differences in gradient magnitudes across the dimensions. As DPVI (and more generally DP-SGD) proceeds to add Gaussian noise with standard deviation proportional to the clipping threshold to *all of the dimensions*, the signal-to-noise-ratio can vary greatly across the parameter dimensions. Parameter dimensions that experience low signal-to-noise ratio will converge much slower than others (cf. Domke, 2019 and references therein).¹ Next, we will show that such a magnitude difference arises between the gradients of variational parameters \mathbf{m}_q and \mathbf{s}_q .

Note that the gradient of Equation (3) w.r.t. \mathbf{m}_q is $\nabla_{\mathbf{m}_q} \theta(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q) = 1$, which leads to the following proposition (a more detailed derivation can be found in Appendix B):

Proposition 3.1. Assume q to be diagonal Gaussian, then the gradient \mathbf{g}_s in Equation (5) becomes

$$\mathbf{g}_s = \boldsymbol{\eta} T'(\mathbf{s}_q) \mathbf{g}_m + \nabla_{\mathbf{s}_q} H(q), \quad (8)$$

where T' denotes the derivative of T .

As the entropy term is independent of the data, our update equation for \mathbf{s}_q depends on the data only through \mathbf{g}_m . Thus, in order to show that this term gets affected by the noise more than \mathbf{g}_m itself, it suffices to inspect the magnitudes of $\boldsymbol{\eta}$ and $T'(\mathbf{s}_q)$. As $\boldsymbol{\eta} \sim N(0, I_d)$, we have $\eta_j = O(1)$. We have $T'(\mathbf{s}_q) \leq \sigma_q$ for common choices of T discussed above (a proof for softplus and exp can be found in Appendix C). For the data dependent part of \mathbf{g}_s it then follows that

$$|\boldsymbol{\eta} T'(\mathbf{s}_q) \mathbf{g}_m| = O(\sigma_q) |\mathbf{g}_m|. \quad (9)$$

Therefore, it is easy to see that as σ_q becomes small, the gradient \mathbf{g}_s becomes small compared to \mathbf{g}_m . Note that this is especially problematic combined with the practice of initialising \mathbf{s}_q to small values to speed up the convergence of \mathbf{m}_q , discussed in Sec. 2.2.

3.1.1 Addressing the differing gradient scales

Above we have identified a magnitude difference between the gradient components, which leads to variational standard deviation parameters being disproportionately affected by DP noise. Next we use this structural knowledge to propose a method for scaling the gradients to more closely matching magnitudes, after discussing two alternatives based on standard techniques.

Natural Gradients As a first solution, we consider *natural gradients* (Amari, 1998). This is a common approach for improving convergence for VI (cf. e.g. Honkela et al., 2010; Khan and Nielsen, 2018; Salimbeni et al., 2018), which relies on scaling the gradients using the information geometry of the optimisation problem.

The natural gradients \mathbf{g}^{nat} are computed using the inverse of the Fisher information matrix \mathcal{I} as

$$\mathcal{I} = \mathbb{E}_{\boldsymbol{\theta}|\mathbf{s}_q, \mathbf{m}_q} [(\nabla_{\boldsymbol{\theta}} \log q(\boldsymbol{\theta}))(\nabla_{\boldsymbol{\theta}} \log q(\boldsymbol{\theta}))^T] \quad (10)$$

$$\mathbf{g}^{nat} = \mathcal{I}^{-1} \mathbf{g}. \quad (11)$$

For our setting this leads to

$$\mathbf{g}_m^{nat} = T(\mathbf{s}_q)^2 \nabla_{\mathbf{m}_q} \mathcal{L}(q) \quad (12)$$

$$\mathbf{g}_s^{nat} = \frac{1}{2T'(\mathbf{s}_q)} \left(\boldsymbol{\eta} \mathbf{g}_m^{nat} + \frac{T(\mathbf{s}_q)^2}{T'(\mathbf{s}_q)} \nabla_{\mathbf{s}_q} H(q) \right). \quad (13)$$

We observe that in the natural gradients the scaling by $T'(\mathbf{s}_q)$ in the gradients of \mathbf{s}_q is now reversed, meaning that for small $T'(\mathbf{s}_q)$ the gradients of \mathbf{s}_q will tend to dominate over those of \mathbf{m}_q . Therefore we expect natural gradients to result in a different instance of the problem of disproportionate DP noise instead of resolving it.

¹We also provide a high-level argument why this is the case in Appendix A.

Preconditioning of Gradients The simplest way to fix the disproportionate DP noise is preconditioning of the gradients to undo the downscaling of the data-dependent part in Eq. (8), by multiplying with $(T'(\mathbf{s}_q))^{-1}$, to obtain

$$\mathbf{g}_s^{precon} = \frac{1}{T'(\mathbf{s}_q)} \mathbf{g}_s = \boldsymbol{\eta} \mathbf{g}_m + \frac{\nabla_{\mathbf{s}_q} H(q)}{T'(\mathbf{s}_q)}. \quad (14)$$

We can see that the data dependent part of \mathbf{g}_s^{precon} (the first term) is of the same magnitude as \mathbf{g}_m , and thus the noise affects the gradient components equally.²

Aligned Gradients The preconditioning approach fixed the issue of different magnitudes in the gradients. However, it comes at the cost of increased ℓ_2 -norm of the full gradient, which requires a higher clipping threshold and therefore increases DP noise. We will now discuss a new alternative method for resolving the disproportionate DP noise problem that avoids this issue.

Equation (8) shows that we can write \mathbf{g}_s in terms of \mathbf{g}_m and an additional entropy term. Since neither the scaling factor $\boldsymbol{\eta} T'(\mathbf{s}_q)$ nor the entropy gradient $\nabla_{\mathbf{s}_q} H(q)$ depend on the data D , it suffices to release the gradients \mathbf{g}_m under DP as $\tilde{\mathbf{g}}_m$, from which we obtain the $\tilde{\mathbf{g}}_s$ via Eq. (8). As this is simply post-processing, it does not incur additional DP cost. Because $\tilde{\mathbf{g}}_s$ is now computed directly as a transformation of $\tilde{\mathbf{g}}_m$, the noise term in both gradients is aligned in proportion to the gradient signals. We refer to this approach as *aligned DPVI* for the rest of the paper. The procedure for computing the aligned DPVI gradients is summarized in Algorithm 1.

Algorithm 1 The aligned gradient procedure (single step)

- | | |
|---|--|
| 1: $\boldsymbol{\theta} \leftarrow \mathbf{m}_q + \boldsymbol{\eta} T(\mathbf{s}_q)$ where $\boldsymbol{\eta} \sim \mathcal{N}(0, I)$ | ▷ Draw sample from the variational posterior |
| 2: $\mathbf{g}_{m,\mathbf{x}} \leftarrow \nabla_{\mathbf{m}_q} \mathcal{L}(q)$ for $\mathbf{x} \in D$ | ▷ Compute the per-example gradients for \mathbf{m}_q |
| 3: $\gamma_{\mathbf{x}} \leftarrow \min(1, C/ \mathbf{g}_{m,\mathbf{x}})$ for $\mathbf{x} \in D$ | ▷ Compute the clipping multiplier |
| 4: $\tilde{\mathbf{g}}_m \leftarrow \sum_{\mathbf{x} \in D} \gamma_{\mathbf{x}} \mathbf{g}_{m,\mathbf{x}} + \sigma_{DP} C \boldsymbol{\psi}$, where $\boldsymbol{\psi} \sim \mathcal{N}(0, I)$ | ▷ Get DP release for \mathbf{g}_m |
| 5: $\tilde{\mathbf{g}}_s^{aligned} \leftarrow \boldsymbol{\eta} T'(\mathbf{s}_q) \tilde{\mathbf{g}}_m + \nabla_{\mathbf{s}_q} H(q)$ | ▷ Get DP aligned \mathbf{g}_s via post-processing |
-

The following theorem (proved in Appendix D) guarantees that the variance in the gradients of \mathbf{s}_q is reduced in aligned DPVI:

Theorem 3.1. Assume C is chosen such that the bias induced by clipping is the same in vanilla and aligned DPVI. Then for any fixed batch,

$$\text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\tilde{\mathbf{g}}_s^{aligned}] \leq \text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\tilde{\mathbf{g}}_s^{vanilla}], \quad (15)$$

where $\boldsymbol{\eta}$ is the random variable of the MC approximation to the ELBO and $\boldsymbol{\psi}$ that of the DP perturbation.

Aligned Natural Gradients Finally we also consider a combination of natural gradients and aligning, to enable the benefits of natural gradients for convergence while simultaneously removing the need to consider the gradient of \mathbf{s}_q for DP clipping and perturbation. The full procedure is given in Algorithm 2. We use $\mathcal{I}_m, \mathcal{I}_s$ to refer to the blocks of \mathcal{I} corresponding to the gradient components.

Algorithm 2 The aligned natural gradient procedure (single step)

- | | |
|---|--|
| 1: $\boldsymbol{\theta} \leftarrow \mathbf{m}_q + \boldsymbol{\eta} T(\mathbf{s}_q)$ where $\boldsymbol{\eta} \sim \mathcal{N}(0, I)$ | ▷ Draw sample from the variational posterior |
| 2: $\mathbf{g}_{m,\mathbf{x}}^{nat} \leftarrow \mathcal{I}_m^{-1} \nabla_{\mathbf{m}_q} \mathcal{L}(q)$ for $\mathbf{x} \in D$ | ▷ Compute per-example natural gradients for \mathbf{m}_q |
| 3: $\gamma_{\mathbf{x}} \leftarrow \min(1, C/ \mathbf{g}_{m,\mathbf{x}}^{nat})$ for $\mathbf{x} \in D$ | ▷ Compute the clipping multiplier |
| 4: $\tilde{\mathbf{g}}_m^{nat} \leftarrow \sum_{\mathbf{x} \in D} \gamma_{\mathbf{x}} \mathbf{g}_{m,\mathbf{x}}^{nat} + \sigma_{DP} C \boldsymbol{\psi}$, where $\boldsymbol{\psi} \sim \mathcal{N}(0, I)$ | ▷ Get DP release for \mathbf{g}_m^{nat} |
| 5: $\tilde{\mathbf{g}}_s^{nat,aligned} = \mathcal{I}_s^{-1} (\boldsymbol{\eta} T'(\mathbf{s}_q) \mathcal{I}_m \tilde{\mathbf{g}}_m^{nat} + \nabla_{\mathbf{s}_q} H(q))$. | ▷ Get DP aligned \mathbf{g}_s^{nat} via post-processing |
-

²Note that this scaling also affects the data-independent entropy term in the gradient for \mathbf{s}_q . While the scaling term $(T'(\mathbf{s}_q))^{-1}$ does increase the entropy part for small \mathbf{s}_q , the data-dependent term is still typically much larger and will dominate the gradient.

3.1.2 Extending to full-rank covariance matrices

So far we have only considered a diagonal Gaussian as the variational posterior. Due to the low dimensionality of the variational parameters, this approach is computationally effective and often applied in practice, but it has limitations: It cannot capture correlations among different model parameters and, more importantly, it will underestimate the marginal variances of the parameters when the true covariance structure is non-diagonal. For those reasons, a full-rank covariance approximation would be favored to correctly capture the uncertainty of the parameters.

However, learning the full-rank covariance approximation results in a quadratic (in the number of dimensions d) expansion of the number of learnable parameters. This not only increases computational costs but also implies less accurate learning of the parameters under DP, as the available privacy budget has to be spread over more parameters. Fortunately, the aligning procedure can be extended to full-rank Gaussian approximations as well, which allows us to alleviate the issue of increased sensitivity. The proof is very similar to the diagonal case. Instead of the parameters \mathbf{s}_q corresponding to marginal standard deviations, we now consider a parameter vector $\mathbf{a}_q \in \mathbb{R}^{\frac{d(d+1)}{2}}$ and a transformation function $T : \mathbb{R}^{\frac{d(d+1)}{2}} \rightarrow \mathbb{R}^{d \times d}$ such that $T(\mathbf{a}_q)$ corresponds to the Cholesky factor of the posterior covariance. That is, T must guarantee that $T(\mathbf{a}_q)$ is a lower triangular with positive entries along its diagonal, which will require similar transformations as in the purely diagonal covariance case discussed previously. Now, the reparametrisation step in (3) becomes

$$\theta := \mathbf{m}_q + T(\mathbf{a}_q)\boldsymbol{\eta}, \quad (16)$$

and the gradient w.r.t \mathbf{a}_q can be written as

$$\mathbf{g}_a = J_a(T(\mathbf{a}_q)\boldsymbol{\eta})\mathbf{g}_m + \nabla_{\mathbf{a}_q}H(q), \quad (17)$$

where J_a denotes the Jacobian of T w.r.t \mathbf{a}_q . Therefore, the gradient \mathbf{g}_a can again be written as a data-independent transformation of \mathbf{g}_m . Thus under the post-processing immunity of DP, we can get the DP gradients for \mathbf{a}_q from DP versions of \mathbf{g}_m without suffering the quadratic increase of the size of the input to underlying the Gaussian mechanism present in vanilla DPVI.

3.2 DPVI samples noisy parameter estimates

As other applications of DP-SGD, vanilla DPVI does not take into account the uncertainty that the DP mechanism introduces. Instead, after a finite number of iterations, the values found in the last iteration are usually treated as the true variational parameters. We argue that treating them this way can lead to severe errors in accuracy because these values are merely a noisy estimate of the optimal values. Annealing the learning rate does not help: Depending on when the annealing is started, the distance to an optimum can still be large due to the random walk prior to annealing. We therefore turn to the known technique of iterate averaging (Polyak and Juditsky, 1992) to reduce variance of the obtained parameter estimate. We briefly review some theory about the random walk behavior around the optimum, which additionally suggests that we can quantify the additional variance introduced by DP noise and make the learned model approximately noise-aware.

Mandt et al. (2017) investigated a random walk behaviour around the optimum for regular (non-DP) SGD induced by the noise arising from subsampling. They assume that near the optimum $\boldsymbol{\xi}^*$ the loss function is well approximated by a quadratic approximation $L(\boldsymbol{\xi}) \approx \frac{1}{2}(\boldsymbol{\xi} - \boldsymbol{\xi}^*)^T \mathbf{A}(\boldsymbol{\xi} - \boldsymbol{\xi}^*)$, and show that the stochastic process around the optimum can be characterised as an Ornstein-Uhlenbeck (OU) process

$$d\boldsymbol{\xi}(t) = -\alpha \mathbf{A}(\boldsymbol{\xi}(t) - \boldsymbol{\xi}^*)dt + \frac{1}{\sqrt{S}}\alpha \mathbf{B}dW(t), \quad (18)$$

where $W(t)$ is a Wiener process, α is the step size of the SGD (assumed constant), S is the size of the subsampled data and \mathbf{B} the Cholesky decomposition of the covariance matrix \mathbf{Z} of the noise due to the subsampling. Directly adapting this analysis, we suggest that under the same regularity assumptions, DP-SGD still is an OU process. The principle of the proof is straightforward: DP-SGD adds an additional Gaussian

noise component, allowing us to add the (diagonal) covariance matrix of the DP noise to \mathbf{Z} and obtain a $\hat{\mathbf{B}}$ such that

$$\mathbf{Z} + S\sigma_{DP}^2\mathbf{I} = \hat{\mathbf{B}}\hat{\mathbf{B}}^T. \quad (19)$$

The more detailed proof can be found in Appendix E. This insight, combined with the fact that the privacy guarantees of DPVI hold for all released intermediate parameter values, allows us to make the following two suggestions to improve the parameter estimates of DPVI.

Iterate averaging to reduce noise in parameter estimate In order to reduce noise in our learned variational parameters, we apply iterate averaging and average the parameter traces, i.e., the sequence of parameters during optimisation, over the last $T_{\text{burn-out}}$ iterates for which we assume the trace has converged. As the OU process in Equation (18) is symmetric around the optimum, the mean of the trace is an unbiased estimator of ξ^* . Compared to using the final iterate of the chain (also an unbiased estimator of ξ^*), the averaged trace reduces the variance of the estimator by up to a factor of $T_{\text{burn-out}}^{-1}$.

Estimating the increased variance due to DP Finally, since our posterior approximation is Gaussian and the stationary distribution of the OU is Gaussian as well, we can add the variance of the averaged traces to the variances of our posterior to absorb the remaining uncertainty due to the inference process, and recover a noise-aware posterior approximation.

Now the remaining problem is to determine $T_{\text{burn-out}}$, the length of the trace where the parameters have converged. For this we suggest a simple convergence check based on linear regression: For each of the traces, we fit linear regression models over different candidate $T_{\text{burn-out}}$. The regressor $\mathbf{X}_{\text{linreg}}$ is set to interval $[0, 1]$ split to $T_{\text{burn-out}}$ points in an ascending order. The responses y are set to the corresponding parameter values in the trace, e.g. $\mathbf{y} = \{\mathbf{m}_q^{(t)}\}_{t=T-T_{\text{burn-out}}}^T$. If the linear regression model has a sufficiently small slope coefficient, we consider the trace as converged and pick the longest $T_{\text{burn-out}}$ for which this is the case.

4 Experiments

We experimentally test our methods for two different tasks using mean-field approximation with real data: learning a probabilistic generative model for private data sharing and learning a logistic regression model. We also experimentally explore aligned DPVI with full-rank Gaussian approximation using simulated data.

4.1 Implementation details

We implemented the different variants for DPVI introduced in Sec. 3.1.1 using the d3p package (Prediger et al., 2022) for the NumPyro probabilistic programming framework (Phan et al., 2019; Bingham et al., 2019). To compute the privacy cost in our experiments, we use the Fourier accountant method (Koskela et al., 2021). The hyperparameters used in our experiments are discussed in the Appendix F. We use the softplus function as our transformation T in all experiments.

In order to assess learning over multiple parameters which converge to different values, and over repeated runs with different initial values, we define a *mean proportional absolute error (MPAE)*: Let $\xi^{(t)} \in \mathbb{R}^D$ be the parameter vector at iteration t and ξ^* be the parameter vector at the optimum.³ We measure the MPAE at iteration t as

$$\text{MPAE}(\xi^{(t)}) = \frac{1}{D} \sum_{d=1}^D \frac{|\xi_d^{(t)} - \xi_d^*|}{|\xi_d^{(0)} - \xi_d^*|}. \quad (20)$$

An MPAE value of 0 indicates perfect recovery of the optimum, a value of 1 suggests that the parameters on average did not move away from their initialisation.

³Since the optimal value ξ^* is typically unknown, we instead use the results of classical non-DP variational inference in its place in practice.

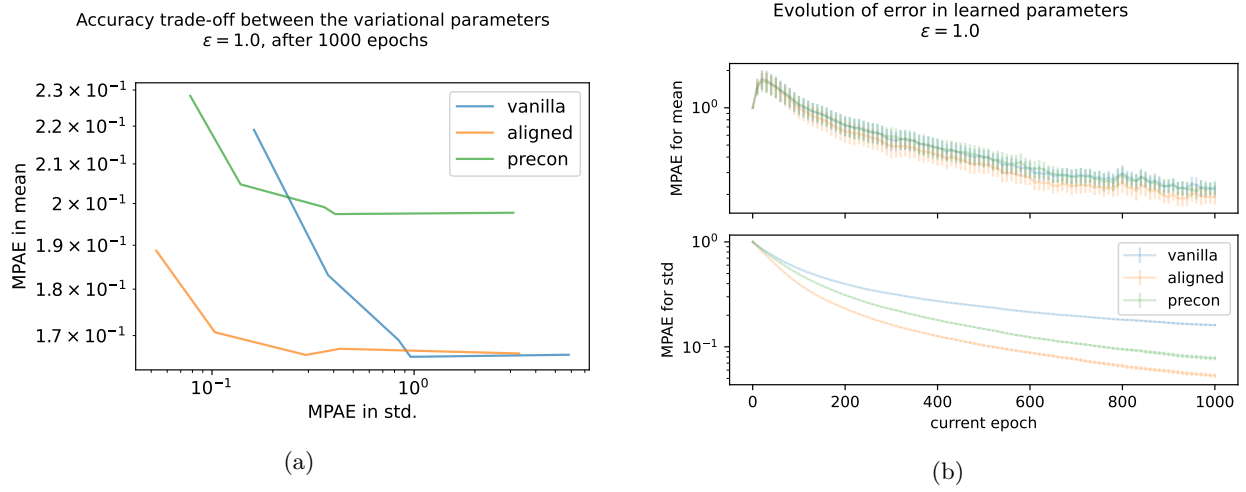


Figure 1: **UKB experiment:** (a) Aligned DPVI makes the smallest error in learning the variational parameters across all different initial values for σ_q , implying it is the most robust. (b) Aligned DPVI converges faster than the other variants while also having less deviation across the repeats (all initialised at $\sigma_q = 1$). Both subfigures show averaged MPAE for vanilla, aligned and preconditioned DPVI with error bars in (b) indicating standard error over repeats.

4.2 Using DPVI to learn a generative model

Recently, Jälkö et al. (2021) suggested using DPVI to learn a probabilistic generative model for differentially private data sharing. Note that in this application it is especially crucial to learn the posterior variances well to faithfully reproduce the uncertainty in the original data in the synthetic data set.

A recent study by Niedzwiedz et al. (2020) on personal health data from the United Kingdom Biobank (UKB) (Sudlow et al., 2015) studied how socio-economic factors affect an individual’s risk of catching the SARS-CoV-2 virus. We aim to produce synthetic data, using DPVI to learn the generative model, from which we can draw similar discoveries.

Following Niedzwiedz et al. (2020), we consider a subset of UKB data which comprises of 58 261 individuals with $d = 7$ discrete (categorical) features. We split the features into a set of explanatory variables and a response variable indicating whether the individual was infected by SARS-CoV-2.

We place a mixture model for the explanatory variables \mathbf{X} , and a Poisson regression model mapping the explanatory variables to the responses \mathbf{y} , using $\theta_{\mathbf{X}}$, $\theta_{\mathbf{y}}$ and π to designate the model parameters:

$$p(\mathbf{X} \mid \theta_{\mathbf{X}}, \pi) = \sum_{k=1}^K \pi_k \prod_{j=1}^d \text{Categorical}(\mathbf{X}_j \mid \theta_{\mathbf{X}}^{(k)}) \quad (21)$$

$$p(\mathbf{y} \mid \mathbf{X}, \theta_{\mathbf{y}}) = \text{Poisson}(\mathbf{y} \mid \exp(\mathbf{X}\theta_{\mathbf{y}})). \quad (22)$$

In our experiments, we set the number of mixture components $K = 16$ which was chosen based on internal tests. Priors for the model parameters are specified in Appendix G.1.

Aligned DPVI is more robust to initialisation We first demonstrate that aligned DPVI improves robustness to initialisation over vanilla and preconditioned DPVI. To do so we fix a privacy budget of $\epsilon = 1$ and the number of passes over the entire data set, i.e., *epochs*, to 1000 and vary the initial value of \mathbf{s}_q such that σ_q is between 0.01 and 1. We perform 10 repetitions with different random seeds over which we keep the initialisation of \mathbf{s}_q fixed but initialise \mathbf{m}_q randomly. We compute the MPAE over the parameters of the Poisson regression part in the model, which corresponds directly to the downstream prediction task we are ultimately interested in.




Figure 2: **UKB experiment:** The aligned variant remains the most accurate method, even if we run the algorithm for longer. Initial $\sigma_q = 1$.




Figure 3: **UKB experiment:** RMSE of parameters found in downstream analysis when doing iterate averaging with different $T_{\text{burn-out}}$. Iterative averaging can reduce the error and significantly reduce variance of error compared to using only the last iterate. Initial $\sigma_q = 1$, 4000 epochs of aligned DPVI.

Figure 1a shows the trade-off different variants of DPVI make between the MPAE in variational means (\mathbf{m}_q) and stds (\mathbf{s}_q) averaged over the 10 repetitions for the different initial values of \mathbf{s}_q . We observe that the aligned variant is able to achieve small errors in \mathbf{m}_q and \mathbf{s}_q simultaneously while the alternatives cannot. To see how the MPAEs for \mathbf{m}_q and \mathbf{s}_q behave individually w.r.t. the initial value for \mathbf{s}_q refer to Appendix H. The natural gradient as well as the aligned natural gradient method performed slightly worse than the aligned method in this experiment and we report results for them in the appendix as well.

Longer runs do not help vanilla DPVI Figure 1b suggests that vanilla DPVI has not converged in terms of MPAE, in the allotted number of iterations for an initial $\sigma_q = 1$. An obvious solution then seems to be to run the inference for longer. We now fix the initialisation of σ_q to 1, which designates the least relative scaling of gradients at the beginning of training and thus a best-case scenario for vanilla DPVI. We vary the number of training epochs from 200 to 8000 while always keeping the privacy budget fixed at $\epsilon = 1$. Since longer runs require more accesses to the data, the DP perturbation scale increases with the number of iterations. As before, we repeat 10 inference runs for each parameter choice.

Figure 2 shows the final MPAE over all 10 repetitions and all parameters in the Poisson regression part of the model for the different numbers of total epochs.⁴ The upper panel shows that with an increasing number of epochs, the difference in MPAE of variational means between vanilla and aligned DPVI vanishes. However, the lower panel shows clearly that even in a long training regime, vanilla DPVI still does not converge in variational variances and is consistently beaten by our aligned variant.

Iterate averaging increases robustness of downstream task Next, we test the iterate averaging of noisy parameters traces for the generative model. We use the linear regression technique discussed in Sec. 3.2, individually for each parameter, to determine the length of the trace to average. We then use synthetic data from the generative model to learn the Poisson regression model used by Niedzwiedz et al. (2020) and compare the regression coefficients against those obtained from the original data. Further details on the downstream analysis setup are given in Appendix I. Figure 3 shows that the results from iterate averaged model are less noisy compared to just using the last iterate as the true parameters. However, the approach appears to be somewhat unstable: Changing the initial values of σ_q to 0.1 causes the variance of error for $\epsilon = 1$. to increase over the non-averaged case. This is likely due to the simple linear regression heuristic we used in this experiment not detecting convergence correctly in this case.

⁴Note that it is *not* showing the evolution of the error over a single training of 8000 epochs.

4.2.1 Experiments on the US Census data set

To ensure the results are not specific to only a single data set, we next applied DPVI on a large set of US Census 1990 data (from UCI Dua and Graff (2017)), in the same data sharing setting as the UK Biobank experiment. We focused on individuals having military background, and using the similar mixture model/Poisson regression combination as with the UK Biobank experiment, we tried to predict the poverty indicator of the individual given military service related and demographic features. After preprocessing the data comprised of 320 754 samples with 13 features. The hyperparameters for the different DPVI variants were identical to the ones used in UK Biobank experiment.

Figure 4 shows again that the aligned variant learns the variational scales better in terms of the mean proportional absolute error (MPAE) of DPVI parameters to the non-private optimum, as was the case with UK Biobank data.




Figure 4: **US Census experiment:** Aligned DPVI method learns the variational standard deviations faster than the alternatives. The plot shows the average MPAE and the standard deviation across 10 independent repeats of the variational inference with σ_q initialized to 1.0 applied on the US Census data sharing experiment.

4.3 Logistic regression with the Adult data set

As the UK Biobank data is access-restricted, we further demonstrate our methods on the publicly available Adult data set from the UCI machine learning repository (Dua and Graff, 2017), which contains 30 162 training records. We learn a logistic regression model, classifying whether the income feature of the data exceeds \$50k based on all other features.

Aligned natural gradient outperforms the other variants We compare our private logistic regression coefficients to the ones obtained using privacy-agnostic VI. We also test the aligned natural gradient and the natural gradient variants for this data. Figure 5a shows the ℓ_2 -norm between variational parameters learned




Figure 5: **Adult logistic regression experiment:** (a) The aligned natural gradient method is closest to the non-private variational parameters. Error is computed as a mean ℓ_2 -norm against non-private baseline over 20 repeats. Error bars show the standard error of the mean. (b) The standard deviation inferred from the converged traces (determined by the linear regression method) is close to that computed over last iterates across different repeats. Lines show the average MSE, error bars show the standard deviation across repeats. The baselines show the mean squared norm of the standard deviation estimated from repeated runs, corresponding to the MSE for not estimating DP-induced noise.

with and without DP. From this figure we see that the aligned natural gradient method clearly outperforms all the other variants in this setting. Additionally, we clearly see again that vanilla DPVI learns the stds poorly, and also that the natural gradient variant reverses the problem compared to vanilla and struggles in learning the variational means as suggested in Section 3.1.1.

DP noise can be inferred from the (converged) traces We test how well we can recover the DP noise effect from the parameter traces. We limit the test to the coefficients that have converged according to the linear regression test described in Section 3.2. Based on our internal tests, we chose a slope of 0.05 as the threshold for convergence. We compare the standard deviation of the converged parameter trace to an estimate of the DP-induced noise estimated by the standard deviation of the last iterates over 50 repeats in terms of mean squared error over different parameter sites.

Figure 5b shows that the noise std estimated from the converged traces is close to the noise std we have across the last iterates of multiple independent repeats.

4.4 Experiments with full-rank covariance

As a final experiment, we investigate aligned gradients for a full-rank Gaussian posterior approximation. We perform Bayesian linear regression over a simulated data set where we can control the number of feature dimensions as well as the strength of correlations. We control the latter by setting the rate of nonzero off-diagonal entries in the covariance matrix for simulated data points. Further details of the experimental setup can be found in Appendix J.

Figure 6 confirms that aligned gradients improve the average predictive log-likelihood of the posterior approximation over a held-out test set. This is true even when the data is not strongly correlated (panels in the right column), as the large increase in parameters over which vanilla DPVI has to split the privacy budget negatively impacts the learning.

5 Discussion

In this paper we introduced the aligned gradient solution for the specific task of learning a Gaussian variational posterior. The technique should be applicable also in other tasks where gradients with respect to different parameters depend on data through a common term. Detection of such cases could be even automated by inspecting how the data enters the computational graph of the task. This would be an interesting future direction.

A limitation of DP in general is that it guarantees indistinguishability among the individuals in the data set by aiming to preserve more common characteristics of the data. Therefore the utility of a DP algorithm might be worse for individuals from less common groups.

The poor performance of aligned natural gradients in the UKB example might be due to bad hyperparameter choices. While we performed some hyperparameter tuning without success, a more comprehensive search would be needed to fully assess the performance of these methods.

The literature on MCMC holds many existing diagnostics for the converge of chains, such as the (split) R-hat estimator (Gelman and Rubin, 1992; Vehtari et al., 2021), that could be used to test the convergence of the parameter traces as well. Tuning these methods to work well in diagnosing the converge of the parameter trace would require extensive testing which we leave for future work.




Figure 6: **Full-rank covariance experiment:** DPVI with aligned gradients achieves better predictive log-likelihood for a full-rank approximation than vanilla DPVI. Higher density means more nonzero entries in data covariance. $\epsilon = 1$. 50 repetitions.

References

- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, Proceedings*, 2006.
- Shuang Song, Kamalika Chaudhuri, and Anand D. Sarwate. Stochastic gradient descent with differentially private updates. In *IEEE Global Conference on Signal and Information Processing, GlobalSIP 2013, Austin, TX, USA, December 3-5, 2013*, pages 245–248. IEEE, 2013. doi: 10.1109/GlobalSIP.2013.6736861. URL <https://doi.org/10.1109/GlobalSIP.2013.6736861>.
- Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization, revisited. *CoRR*, abs/1405.7085, 2014. URL <http://arxiv.org/abs/1405.7085>.
- Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, page 308–318, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978318. URL <https://doi.org/10.1145/2976749.2978318>.
- Joonas Jälkö, Onur Dikmen, and Antti Honkela. Differentially private variational inference for non-conjugate models. In *Uncertainty in Artificial Intelligence 2017, Proceedings of the 33rd Conference (UAI),*, 2017.
- H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, Canada, Conference Track Proceedings*, 2018.
- Xiangyi Chen, Zhiwei Steven Wu, and Mingyi Hong. Understanding gradient clipping in private SGD: A geometric perspective. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/9ecff5455677b38d19f49ce658ef0608-Abstract.html>.
- Galen Andrew, Om Thakkar, Hugh Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://openreview.net/forum?id=RUQ1zwZR8_.
- Ning Wang, Yang Xiao, Yimin Chen, Ning Zhang, Wenjing Lou, and Y. Thomas Hou. Squeezing more utility via adaptive clipping on differentially private gradients in federated meta-learning. In *Proceedings of the 38th Annual Computer Security Applications Conference, ACSAC ’22*, page 647–657, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450397599. doi: 10.1145/3564625.3564652. URL <https://doi.org/10.1145/3564625.3564652>.
- Garrett Bernstein and Daniel R. Sheldon. Differentially private Bayesian inference for exponential families. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems, (NeurIPS) 2018,*, 2018.
- Garrett Bernstein and Daniel R. Sheldon. Differentially private Bayesian linear regression. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS) 2019*, 2019.
- Tejas Kulkarni, Joonas Jälkö, Antti Koskela, Samuel Kaski, and Antti Honkela. Differentially private Bayesian inference for generalized linear models. In *Proceedings of the 38th International Conference on Machine Learning, Proceedings of Machine Learning Research*. PMLR, 2021.
- Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Mach. Learn.*, 37(2):183–233, 1999. doi: 10.1023/A:1007665907178. URL <https://doi.org/10.1023/A:1007665907178>.

- Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.*, 20:28:1–28:6, 2019. URL <http://jmlr.org/papers/v20/18-403.html>.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M. Blei. Automatic differentiation variational inference. *J. Mach. Learn. Res.*, 18:14:1–14:45, 2017. URL <http://jmlr.org/papers/v18/16-107.html>.
- Justin Domke. Provable gradient variance guarantees for black-box variational inference. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Shun-ichi Amari. Natural gradient works efficiently in learning. 10(2):251–276, 02 1998. ISSN 0899-7667. doi: 10.1162/089976698300017746.
- Antti Honkela, Tapani Raiko, Mikael Kuusela, Matti Törnio, and Juha Karhunen. Approximate riemannian conjugate gradient learning for fixed-form variational bayes. *Journal of Machine Learning Research*, 11 (106):3235–3268, 2010. URL <http://jmlr.org/papers/v11/honkela10a.html>.
- Mohammad Emtiyaz Khan and Didrik Nielsen. Fast yet simple natural-gradient descent for variational inference in complex models. In *2018 International Symposium on Information Theory and Its Applications (ISITA)*, pages 31–35, 2018. doi: 10.23919/ISITA.2018.8664326.
- Hugh Salimbeni, Stefanos Eleftheriadis, and James Hensman. Natural gradients in practice: Non-conjugate variational inference in gaussian process models. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 689–697. PMLR, Apr 2018.
- B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992. doi: 10.1137/0330046.
- Stephan Mandt, Matthew D. Hoffman, and David M. Blei. Stochastic gradient descent as approximate Bayesian inference. 18(1):4873–4907, 2017. ISSN 1532-4435.
- Lukas Prediger, Niki Loppi, Samuel Kaski, and Antti Honkela. d3p - a python package for differentially-private probabilistic programming. *Proceedings on Privacy Enhancing Technologies*, 2022(2):407–425, 2022. doi: doi:10.2478/popets-2022-0052.
- Du Phan, Neeraj Pradhan, and Martin Jankowiak. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *arXiv preprint arXiv:1912.11554*, 2019.
- Antti Koskela, Joonas Jälkö, Lukas Prediger, and Antti Honkela. Tight differential privacy for discrete-valued mechanisms and for the subsampled gaussian mechanism using FFT. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 3358–3366. PMLR, 2021. URL <http://proceedings.mlr.press/v130/koskela21a.html>.
- Joonas Jälkö, Eemil Lagerspetz, Jari Haukka, Sasu Tarkoma, Antti Honkela, and Samuel Kaski. Privacy-preserving data sharing via probabilistic modeling. *Patterns*, 2(7):100271, 2021. ISSN 2666-3899. doi: <https://doi.org/10.1016/j.patter.2021.100271>. URL <https://www.sciencedirect.com/science/article/pii/S2666389921000970>.

- Claire L Niedzwiedz, Catherine A O'Donnell, Bhautesh Dinesh Jani, Evangelia Demou, Frederick K Ho, Carlos Celis-Morales, Barbara I Nicholl, Frances S Mair, Paul Welsh, Naveed Sattar, et al. Ethnic and socioeconomic differences in sars-cov-2 infection: prospective cohort study using uk biobank. *BMC medicine*, 18(1):1–14, 2020.
- Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins. Uk biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLOS Medicine*, 12(3):1–10, 03 2015. doi: 10.1371/journal.pmed.1001779.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Andrew Gelman and Donald B. Rubin. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457 – 472, 1992. doi: 10.1214/ss/1177011136.
- Aki Vehtari, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of mcmc (with discussion). *Bayesian analysis*, 16(2):667–718, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Shubhankar Mohapatra, Sajin Sasy, Xi He, Gautam Kamath, and Om Thakkar. The role of adaptive optimizers for honest private hyperparameter selection. 11 2021.
- Donald B Rubin. *Multiple imputation for nonresponse in surveys*, volume 81. John Wiley & Sons, 2004.
- Jerome P Reiter and Trivellore E Raghunathan. The multiple adaptations of multiple imputation. *Journal of the American Statistical Association*, 102(480):1462–1471, 2007.

Appendices

A Intuitive reasoning why larger noise would slow convergence

We consider only a single optimisation step in a single dimension for simplicity. Assume we are at $\theta^{(t)}$ and have noisy gradient

$$g^{(t)} = \nabla \mathcal{L}(\theta^{(t)}) + \eta, \eta \sim \mathcal{N}(0, \sigma^2) \quad (\text{A.1})$$

for some perturbation scale σ . We update the parameter as

$$g^{(t+1)} = \theta^{(t)} - \alpha g^{(t)} \quad (\text{A.2})$$

with learning rate α .

In order to get closer to the optimum, we want $\text{sign}(g^{(t)}) = \text{sign}(\nabla \mathcal{L}(\theta^{(t)}))$. Assume wlog that $\nabla \mathcal{L}(\theta^{(t)}) > 0$, then

$$\Pr [\text{sign}(g^{(t)}) = \text{sign}(\nabla \mathcal{L}(\theta^{(t)}))] = \Pr [\nabla \mathcal{L}(\theta^{(t)}) + \eta > 0] \quad (\text{A.3})$$

$$= \Pr [\eta > -\nabla \mathcal{L}(\theta^{(t)})] \quad (\text{A.4})$$

$$= 1 - \Pr [\eta \leq -\nabla \mathcal{L}(\theta^{(t)})] \quad (\text{A.5})$$

$$= 1 - \Phi(-\nabla \mathcal{L}(\theta^{(t)})) \quad (\text{A.6})$$

$$= \Phi(\nabla \mathcal{L}(\theta^{(t)})) \quad (\text{A.7})$$

$$= \frac{1}{2} \left(1 + \text{erf} \left(\frac{\nabla \mathcal{L}(\theta^{(t)})}{\sigma \sqrt{2}} \right) \right). \quad (\text{A.8})$$

$\text{erf}(\cdot)$ is a monotonically increasing function, so we see from the above that the probability of progressing towards the optimum decreases with decreasing $\frac{\nabla \mathcal{L}(\theta^{(t)})}{\sigma}$. I.e., for a fixed gradient, larger variance σ^2 will decrease the probability of progressing towards the optimum in each step.

B Proof of Proposition 3.1

We begin by restating Proposition 3.1.

Proposition B.1 (Proposition 3.1). Assume q to be diagonal Gaussian, then the gradient \mathbf{g}_s in Equation (5) becomes

$$\mathbf{g}_s = \boldsymbol{\eta} T'(\mathbf{s}_q) \mathbf{g}_m + \nabla_{\mathbf{s}_q} H(q),$$

where T' denotes the derivative of T .

Proof. We first recall the reparametrisation for the diagonal Gaussian approximation from Eq. (3) as $\boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q) = \mathbf{m}_q + T(\mathbf{s}_q) \boldsymbol{\eta}$ and observe that $\nabla_{\mathbf{m}_q} \boldsymbol{\theta} = \mathbf{1}$ and $\nabla_{\mathbf{s}_q} \boldsymbol{\theta} = \boldsymbol{\eta} T'(\mathbf{s}_q)$ (where we abbreviate $\boldsymbol{\theta}(\boldsymbol{\eta}; \mathbf{m}_q, \mathbf{s}_q)$ to simply $\boldsymbol{\theta}$). With this we obtain the gradient of the ELBO with respect to \mathbf{m}_q by applying the chain rule in Eq. (4) as:

$$\mathbf{g}_m = \nabla_{\mathbf{m}_q} \mathcal{L}(q) = \nabla_{\boldsymbol{\theta}} \log p(D, \boldsymbol{\theta}) \quad (\text{B.1})$$

Similarly applying the chain rule in Eq. 5 and inserting the above yields

$$\mathbf{g}_s = \nabla_{\mathbf{s}_q} \log p(D, \boldsymbol{\theta}) + \nabla_{\mathbf{s}_q} H(q) \quad (\text{B.2})$$

$$= \nabla_{\boldsymbol{\theta}} \log p(D, \boldsymbol{\theta}) \nabla_{\mathbf{s}_q} \boldsymbol{\theta} + \nabla_{\mathbf{s}_q} H(q) \quad (\text{B.3})$$

$$= \nabla_{\boldsymbol{\theta}} \log p(D, \boldsymbol{\theta}) \boldsymbol{\eta} T'(\mathbf{s}_q) + \nabla_{\mathbf{s}_q} H(q) \quad (\text{B.4})$$

$$= \boldsymbol{\eta} T'(\mathbf{s}_q) \mathbf{g}_m + \nabla_{\mathbf{s}_q} H(q). \quad (\text{B.5})$$

□

C Proof of $T'(\mathbf{s}_q) \leq T(\mathbf{s}_q)$ for softplus and exponential function

$T(\mathbf{s}_q) = \text{softplus}(\mathbf{s}_q)$ Consider we transform \mathbf{s}_q in positive real numbers using the softplus function:

$$T(\mathbf{s}_q) = \log(1 + \exp(\mathbf{s}_q)). \quad (\text{C.1})$$

First, we make the following observation which connects the softplus to the sigmoid function

$$\begin{aligned} T(\mathbf{s}_q) &= -\log\left(\frac{1}{1 + \exp(\mathbf{s}_q)}\right) \\ &= -\log\left(1 - \frac{1}{1 + \exp(-\mathbf{s}_q)}\right) \\ &= -\log(1 - \sigma(\mathbf{s}_q)), \end{aligned} \quad (\text{C.2})$$

where σ denotes the sigmoid function. We then get $T'(\mathbf{s}_q) = \sigma(\mathbf{s}_q)$. It is easy to see that $\log(x) \leq x - 1$ and hence

$$\begin{aligned} T(\mathbf{s}_q) &= -\log(1 - \sigma(\mathbf{s}_q)) \\ &\geq 1 - (1 - \sigma(\mathbf{s}_q)) = \sigma(\mathbf{s}_q) = T'(\mathbf{s}_q). \end{aligned} \quad (\text{C.3})$$

We have therefore shown that $T(\mathbf{s}_q) \geq T'(\mathbf{s}_q) \forall \mathbf{s}_q \in \mathbb{R}$.

$T(\mathbf{s}_q) = \exp(\mathbf{s}_q)$ For $T(\mathbf{s}_q) = \exp(\mathbf{s}_q)$ the proof follows immediately from the fact that $T'(\mathbf{s}_q) = \exp(\mathbf{s}_q) = T(\mathbf{s}_q)$.

D Proof of Theorem 3.1: Variance in aligned scale gradients is smaller

We begin by restating the theorem:

Theorem D.1 (Theorem 3.1). Assume C is chosen such that the bias induced by clipping is the same in vanilla and aligned DPVI. Then for any fixed batch,

$$\text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\hat{\mathbf{g}}_s^{\text{aligned}}] \leq \text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\hat{\mathbf{g}}_s^{\text{vanilla}}],$$

where $\boldsymbol{\eta} \in \mathcal{N}(0, 1)$ is the random variable of the MC approximation to the ELBO in the reparametrisation approach and $\boldsymbol{\psi} \in \mathcal{N}(0, 1)$ that of the DP perturbation.

Proof. We first point out that it is well known that clipping introduces bias to the solution and we aim to compare cases in which the solution is not changed. To that end we will first consider the case where the overall gradient magnitude is dominated by the gradient with respect to \mathbf{m}_q , i.e., $\|\mathbf{g}_m\|_2^2 \gg \|\mathbf{g}_s\|_2^2$. This implies, first, that $T'(\mathbf{s}_q) \ll 1$ and, second, that we can assume the clipping threshold C to be fixed when switching between vanilla and aligned DPVI variants. These two considerations allow us to start the analysis with relative simplicity. Later we will consider the case of larger $T'(\mathbf{s}_q)$ for which we consider a clipping threshold C that is adjusted to keep the bias induced by clipping the same when switching between vanilla and aligned DPVI.

Case $T'(\mathbf{s}_q) \ll 1$ and fixed C . In the vanilla approach, we directly clip and perturb the combined gradient $\mathbf{g} = (\mathbf{g}_m^T, \mathbf{g}_s^T)^T$. Focusing on \mathbf{g}_s only, we get

$$\tilde{\mathbf{g}}_s = \gamma \nabla_{\mathbf{s}_q} \mathcal{L}(q) + \boldsymbol{\psi} \sigma_{DPC} \quad (\text{D.1})$$

$$= \gamma (\boldsymbol{\eta} T'(\mathbf{s}_q) \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q)) + \boldsymbol{\psi} \sigma_{DPC} \quad (\text{D.2})$$

from inserting Eq. 8 to express \mathbf{g}_s in terms of \mathbf{g}_m . Here, $\gamma = \min(1, C/\|\mathbf{g}\|)$ is the clipping multiplier which ensures that $\|\gamma \mathbf{g}\| \leq C$.

Assuming a fixed batch and fixed parameters \mathbf{m}_q and \mathbf{s}_q , the variance of $\tilde{\mathbf{g}}_s$ with respect to $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$ is

$$\text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\gamma (\boldsymbol{\eta} T'(\mathbf{s}_q) \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q)) + \boldsymbol{\psi} \sigma_{DP} C] \quad (\text{D.3})$$

$$= \text{Var}_{\boldsymbol{\eta}} [\mathbb{E}_{\boldsymbol{\psi}|\boldsymbol{\eta}} [\gamma (\boldsymbol{\eta} T'(\mathbf{s}_q) \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q)) + \boldsymbol{\psi} \sigma_{DP} C]] \quad (\text{D.4})$$

$$+ \mathbb{E}_{\boldsymbol{\eta}} [\text{Var}_{\boldsymbol{\psi}|\boldsymbol{\eta}} [\gamma (\boldsymbol{\eta} T'(\mathbf{s}_q) \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q)) + \boldsymbol{\psi} \sigma_{DP} C]] \quad (\text{D.5})$$

$$= \text{Var}_{\boldsymbol{\eta}} [\gamma (\boldsymbol{\eta} T'(\mathbf{s}_q) \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q))] + \mathbb{E}_{\boldsymbol{\eta}} [\sigma_{DP}^2 C^2] \quad (\text{D.6})$$

$$= T'(\mathbf{s}_q)^2 \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + \sigma_{DP}^2 C^2. \quad (\text{D.7})$$

In contrast, the perturbed *aligned* gradient is obtained by inserting the clipped and perturbed $\tilde{\mathbf{g}}_m$ instead of \mathbf{g}_m in Eq. 8:

$$\tilde{\mathbf{g}}_s^{\text{aligned}} = \boldsymbol{\eta} T'(\mathbf{s}_q) \tilde{\mathbf{g}}_m + \nabla_{\mathbf{s}_q} H(q) \quad (\text{D.8})$$

$$= \boldsymbol{\eta} T'(\mathbf{s}_q) (\gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \boldsymbol{\psi} \sigma_{DP} C) + \nabla_{\mathbf{s}_q} H(q) \quad (\text{D.9})$$

and its variance is

$$\text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\boldsymbol{\eta} T'(\mathbf{s}_q) (\gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \boldsymbol{\psi} \sigma_{DP} C) + \nabla_{\mathbf{s}_q} H(q)] \quad (\text{D.10})$$

$$= \text{Var}_{\boldsymbol{\eta}} [\mathbb{E}_{\boldsymbol{\psi}|\boldsymbol{\eta}} [\boldsymbol{\eta} T'(\mathbf{s}_q) (\gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \boldsymbol{\psi} \sigma_{DP} C) + \nabla_{\mathbf{s}_q} H(q)]] \quad (\text{D.11})$$

$$+ \mathbb{E}_{\boldsymbol{\eta}} [\text{Var}_{\boldsymbol{\psi}|\boldsymbol{\eta}} [\boldsymbol{\eta} T'(\mathbf{s}_q) (\gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \boldsymbol{\psi} \sigma_{DP} C) + \nabla_{\mathbf{s}_q} H(q)]] \quad (\text{D.12})$$

$$= \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} T'(\mathbf{s}_q) \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q) + \nabla_{\mathbf{s}_q} H(q)] + \mathbb{E}_{\boldsymbol{\eta}} [\boldsymbol{\eta}^2 T'(\mathbf{s}_q)^2 \sigma_{DP}^2 C^2] \quad (\text{D.13})$$

$$= T'(\mathbf{s}_q)^2 \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + T'(\mathbf{s}_q)^2 \sigma_{DP}^2 C^2 \underbrace{\mathbb{E}_{\boldsymbol{\eta}} [\boldsymbol{\eta}^2]}_{=1} \quad (\text{D.14})$$

These variances differ only in the additional scaling of the DP perturbation term by a factor of $T'(\mathbf{s}_q)^2$. It follows directly that $\text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\tilde{\mathbf{g}}_s] \geq \text{Var}_{\boldsymbol{\eta}, \boldsymbol{\psi}} [\tilde{\mathbf{g}}_s^{\text{aligned}}]$ since $T'(\mathbf{s}_q)^2 \leq 1$ by our initial assumption.

Generalising to arbitrary $T'(\mathbf{s}_q)$. The above assumes that C stays constant between both variants, which is a reasonable assumption for small $T'(\mathbf{s}_q)$. We now argue that for larger values of $T'(\mathbf{s}_q)$ different clipping thresholds should be chosen for both variants and then the assertion about the variances holds for all $T'(\mathbf{s}_q)$.

We obtain the norm of the combined gradient as

$$\|\mathbf{g}\|_2 = \sqrt{\|\mathbf{g}_m\|_2^2 + \|\mathbf{g}_s\|_2^2} \geq \sqrt{1 + T'(\mathbf{s}_q)^2 \boldsymbol{\eta}^2} \|\mathbf{g}_m\|_2, \quad (\text{D.15})$$

where the last inequation simply omits the entropy term $\nabla_{\mathbf{s}_q} H(q) = T'(s)/T(s) > 0$ in \mathbf{g}_s .

Since for aligned DPVI we only need to clip and perturb the gradient \mathbf{g}_m , and because now $\|\mathbf{g}_m\|_2 \leq \|\mathbf{g}\|_2$, the effect of clipping changes compared to the vanilla case if C were kept constant. In order to keep the effect of clipping - and therefore the output of the algorithm - the same, we must adopt different clipping thresholds for the vanilla and the aligned DPVI variants, so that the clipping multiplier γ (and hence the gradient after clipping) stays the same.

Given some clipping threshold C' for the aligned variant, we get for the clipping threshold C of the vanilla variant that

$$C = C' \frac{\|\mathbf{g}\|}{\|\mathbf{g}_m\|} \geq C' \sqrt{1 + T'(\mathbf{s}_q)^2 \boldsymbol{\eta}^2}. \quad (\text{D.16})$$

Inserting into the variances calculated above in Eq. D.7 and D.14, we obtain

$$\text{Var}_{\boldsymbol{\eta}, \psi} [\tilde{\mathbf{g}}_s] = T'(\mathbf{s}_q)^2 \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + \sigma_{DP}^2 C^2 \quad (\text{D.17})$$

$$\geq T'(\mathbf{s}_q)^2 \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + \sigma_{DP}^2 C'^2 (1 + T'(s)^2) \quad (\text{D.18})$$

$$= T'(\mathbf{s}_q)^2 (\text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + \sigma_{DP}^2 C'^2) + \sigma_{DP}^2 C'^2, \text{ and} \quad (\text{D.19})$$

$$\text{Var}_{\boldsymbol{\eta}, \psi} [\tilde{\mathbf{g}}_s^{\text{aligned}}] = T'(\mathbf{s}_q)^2 \text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + T'(\mathbf{s}_q)^2 \sigma_{DP}^2 C'^2 \quad (\text{D.20})$$

$$= T'(\mathbf{s}_q)^2 (\text{Var}_{\boldsymbol{\eta}} [\boldsymbol{\eta} \gamma \nabla_{\mathbf{m}_q} \mathcal{L}(q)] + \sigma_{DP}^2 C'^2). \quad (\text{D.21})$$

Therefore $\text{Var}_{\boldsymbol{\eta}, \psi} [\mathbf{g}_s] \geq \text{Var}_{\boldsymbol{\eta}, \psi} [\mathbf{g}_s^{\text{aligned}}]$ holds for all $T'(\mathbf{s}_q)$ if we adapt C to keep the clipping-induced bias constant. \square

E Variance of DP from OU process on convergence

We follow closely the analysis performed by Mandt et al. (2017), Sec. 3.2., which makes the following assumptions for the loss function \mathcal{L} around its optimum $\boldsymbol{\xi}^*$:

1. Mini-batch gradients of the loss functions are well approximated by a zero-mean Gaussian distribution with covariance matrix $\frac{1}{S} \mathbf{Z}$, where S denotes the size of mini-batch,
2. \mathcal{L} is locally well approximated by a quadratic function.

We make the additional assumption that the clipping threshold C is chosen such that no clipping occurs for gradients close to the optimum in order to avoid clipping-induced bias.

We begin by stating the SGD parameter update equation and the resulting update step $\Delta \boldsymbol{\xi}(t)$ for update steps close to the optimum.

$$\begin{aligned} \boldsymbol{\xi}(t+1) &= \boldsymbol{\xi}(t) - \alpha \left(\nabla_{\boldsymbol{\xi}} \mathcal{L}(\boldsymbol{\xi}(t)) + \frac{1}{\sqrt{S}} \mathbf{B} \boldsymbol{\nu} + \sigma_{DP} \mathbf{I} \boldsymbol{\psi} \right), \quad \boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \boldsymbol{\psi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \\ &= \boldsymbol{\xi}(t) - \alpha \left(\nabla_{\boldsymbol{\xi}} \mathcal{L}(\boldsymbol{\xi}(t)) + \left(\frac{1}{\sqrt{S}} \mathbf{B} + \sigma_{DP} \mathbf{I} \right) \boldsymbol{\psi}' \right), \quad \boldsymbol{\psi}' \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \end{aligned} \quad (\text{E.1})$$

$$\Delta \boldsymbol{\xi}(t) := \boldsymbol{\xi}(t+1) - \boldsymbol{\xi}(t) = -\alpha \nabla_{\boldsymbol{\xi}} \mathcal{L}(\boldsymbol{\xi}(t)) - \alpha \left(\frac{1}{\sqrt{S}} \mathbf{B} + \sigma_{DP} \mathbf{I} \right) \boldsymbol{\psi}' \quad (\text{E.2})$$

We denote with \mathbf{B} the triangular matrix resulting from Cholesky decomposition of $\mathbf{Z} = \mathbf{B} \mathbf{B}^T$. Since we are performing DP-SGD, we have an additional independent noise term with scale σ_{DP} . However, as both sources of stochasticity are independent zero-mean Gaussians, we can easily reformulate using a single Gaussian source of noise with the total variance.

We now restate E.2 by the following stochastic differential equation:

$$d\boldsymbol{\xi}(t) = -\alpha \nabla_{\boldsymbol{\xi}} \mathcal{L}(\boldsymbol{\xi}(t)) - \alpha \left(\frac{1}{\sqrt{S}} \mathbf{B} + \sigma_{DP} \mathbf{I} \right) d\mathbf{W}(t) \quad (\text{E.3})$$

From our second assumption, we know that

$$\mathcal{L}(\boldsymbol{\xi}) \approx \frac{1}{2} (\boldsymbol{\xi} - \boldsymbol{\xi}^*)^T \mathbf{A} (\boldsymbol{\xi} - \boldsymbol{\xi}^*), \quad (\text{E.4})$$

where $\mathbf{A} = \frac{\partial^2}{\partial \boldsymbol{\xi}^2} \mathcal{L}(\boldsymbol{\xi}^*)$.

Inserting (E.3) in (E.4), we get

$$d\boldsymbol{\xi}(t) = -\alpha \mathbf{A} (\boldsymbol{\xi} - \boldsymbol{\xi}^*) dt + \alpha \left(\frac{1}{\sqrt{S}} \mathbf{B} + \sigma_{DP} \mathbf{I} \right) d\mathbf{W}(t) \quad (\text{E.5})$$

with describes an Ornstein-Uhlenbeck (OU) process with Gaussian stationary distribution

$$q(\boldsymbol{\xi}) \propto \exp \left\{ -\frac{1}{2}(\boldsymbol{\xi} - \boldsymbol{\xi}^*)^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{\xi} - \boldsymbol{\xi}^*) \right\} \quad (\text{E.6})$$

where $\boldsymbol{\Sigma}$ satisfies the Lyapunov equation

$$\boldsymbol{\Sigma} \mathbf{A} + \mathbf{A} \boldsymbol{\Sigma} = \alpha \left(\frac{1}{S} \mathbf{Z} + \sigma_{DP}^2 \mathbf{I} \right). \quad (\text{E.7})$$

Even without determining \mathbf{A} we can already make an important discovery from this: Since \mathbf{A} is fixed around the optimum, we see that the noise covariance of our OU process scales linearly with respect to σ_{DP}^2 .

Note that the above analysis holds for a constant learning rate. We have used the Adam optimization method (Kingma and Ba, 2015) in our experiments throughout the paper. While Adam does adapt the learning rate, recently Mohapatra et al. (2021) showed that the learning rate of Adam will converge to a static value, which means that the analysis above still holds as it is only concerned with the learning rate at convergence.

F Hyperparameters

We use Adam (Kingma and Ba, 2015) as the optimiser for all the experiments with starting learning rate of 10^{-3} . In all of our experiments, the δ privacy parameter was set to $1/N$ where N denotes the size of the training data.

For the UKB experiment In the experiments, we used various different training lengths (depicted e.g. in Figure 2). For all of our runs, we set the subsampling rate as 0.01. The clipping threshold C was set to $C = 2.0$ for the aligned and vanilla, 4.0 for preconditioned variant and to 0.1 for the natural gradient based variants.

For the Adult experiment The training was run for 4000 epochs with subsampling ratio of 0.01, corresponding to total of 400 000 gradient steps.

We chose the clipping thresholds for the gradient perturbation algorithm as the 97.5% upper quantile of the training data gradient norms at the non-private optima. This was done to avoid clipping-induced bias, thus making the models comparable to the non-private baseline. This lead to clipping thresholds C presented in Table F.1.

Table F.1: Clipping thresholds for the Adult data logistic regression model

Variant	C
Aligned	3.0
Aligned Natural Grad.	0.1
Natural Grad.	0.1
Vanilla	3.0
Preconditioned	4.0

G Model priors

G.1 For the UKB and the US Census experiments

Recall the probabilistic model used in the experiments:

$$p(\mathbf{X} \mid \boldsymbol{\theta}_X, \boldsymbol{\pi}) = \sum_{k=1}^K \pi_k \prod_{j=1}^d \text{Categorical}(\mathbf{X}_j \mid \boldsymbol{\theta}_X^{(k)}) \quad (\text{G.1})$$

$$p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}_y) = \text{Poisson}(\mathbf{y} \mid \exp(\mathbf{X} \boldsymbol{\theta}_y)). \quad (\text{G.2})$$




Figure H.1: The aligned variant has consistently low error across different initial values of q_σ . Figure shows the error for $\epsilon = 1$ for 1000 epochs of training.

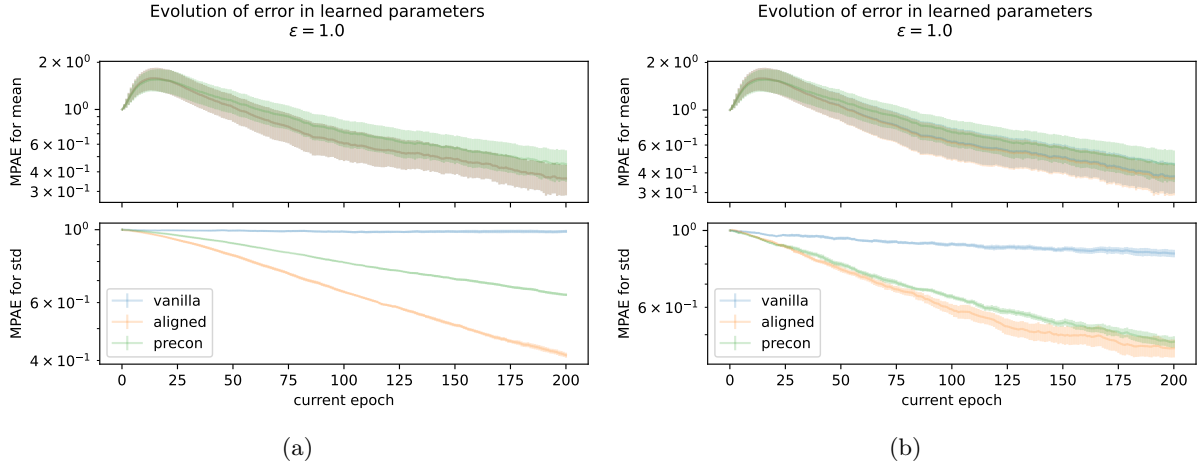


Figure H.2: Aligned DPVI consistently converges faster than the other methods for different initialisation of s_q . **On left**, the s_q is initialised such that $\sigma_q = 0.01$ and **on right** such that $\sigma_q = 0.1$.

The categorical probabilities $\theta_{\mathbf{X}_j}^{(k)}$ for each of the categorical features \mathbf{X}_j , were given a uniform Dirichlet(1) prior. Similarly the mixture weights π we assigned a uniform Dirichlet prior. The regression coefficients θ_y were given a std. normal $N(0, I)$ prior.

G.2 For the Adult experiment

We use the following model and prior

$$\mathbf{y} \sim \sigma(\mathbf{X}\mathbf{w}), \quad (\text{G.3})$$

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (\text{G.4})$$

where $\sigma(\cdot)$ denotes the logistic regression function, $\sigma(x) = 1/(1+\exp\{-x\})$.




Figure H.3: The natural gradient based variants struggle to converge in the UKB experiment. Figure shows the evolution of MPAAE for both of the variational parameters over 1000 epochs. Lines show the mean MPAAE over 10 independent repeats as well as std. of mean as error. In this experiment, the σ_q was initialised to 1 and clipping threshold for both the natural gradient variants was set to 0.1.

H More results for robustness

Figure H.1 shows how the MPAAE for the different DPVI variants behave for the different initial values for s_q for both, variational mean (upper panel) and standard deviation (lower panel) after 1000 epochs.

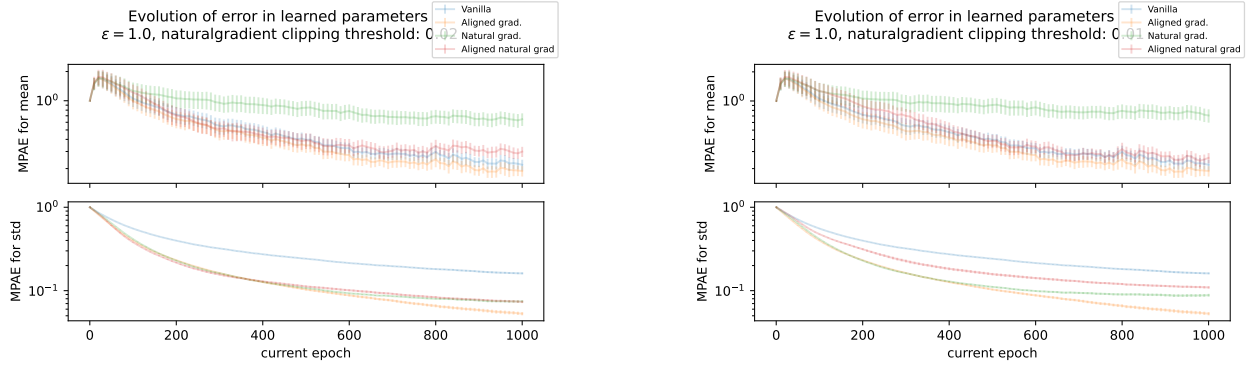
Figure H.2 shows the parameters traces for s_q initialised such that $\sigma_q = 0.01$ (left) and $\sigma_q = 0.1$ (right) with the same split in upper and lower panels, similar to Figure 1b for $\sigma_q = 1.0$ in the main body of the paper.

We observe that with decreasing initial values for s_q ($/\sigma_q$) it becomes increasingly difficult for vanilla DPVI to learn variational standard deviation but learning of means is slightly improved. Preconditioned DPVI performs better overall in terms of standard deviation but learns means worse. Aligned DPVI consistently outperforms both competing variants.

H.1 Natural gradients and the aligned natural gradients in the UKB experiment

Besides the vanilla and aligned variant, we also fitted the UKB model using the natural gradient and aligned natural gradient variants. From Figure H.3 we can see the trade-off natural gradient makes; the means are learned worse than standard deviations, which is what we expect based on the analysis of Section 3.1.1. Somewhat surprisingly, the aligned natural gradient variant performs worse than the aligned variant in this experiment. This might be due to poor choice of hyperparameter, for example the learning rate for the Adam optimiser used in the experiments was set to 10^{-3} for all the variants, while we know that the natural gradient variants tend to have smaller gradients than the others - although Adam should in theory be able to adapt to that.

To evaluate whether the difference between aligned and aligned natural gradient results is due to clipping threshold being too large for the aligned natural gradient approach, thus perturbing the gradients excessively, we repeat the experiment with lower clipping thresholds. Figure H.4a shows that reducing the clipping threshold helps the aligned natural gradient variant to converge in the variational std's. However, the method is still somewhat outperformed by the aligned variant in both measures. Figure H.4b shows that setting clipping threshold $C = 0.01$ natural gradient variants is too small, and the aligned natural gradient starts to suffer from clipping-induced bias.



(a) The aligned natural gradient variant performs slightly better when the clipping threshold is set to $C = 0.02$. The natural gradient variant appears to diverge from the true variational std. which might be caused by clipping-induced bias, while still struggling to learn the correct variational mean.

(b) Both natural gradient variants start to diverge if clipping is set too low $C = 0.01$.

Figure H.4: Tests with smaller clipping threshold for the natural gradient variants. The clipping threshold for the vanilla and aligned is still set to 2.0.

I Further details on downstream analysis for UKB data

In the UKB experiment, we use the learned variational posterior to sample a synthetic data set from the posterior predictive distribution (PPD) as suggested by Jälkö et al. (2021). We test the method by comparing the synthetic data in downstream analysis to the original data. As the downstream task, we fitted a Poisson regression model that aims to predict whether individual catches SARS-CoV-2 based on the predictors in the data. Note that this downstream perfectly overlaps with our generative model.

In order to properly reflect the uncertainty rising from the data generating process to the final results computed from the synthetic data, we will employ so called *Rubin's rules* (Rubin, 2004). In this procedure, we first sample multiple synthetic data sets from the PPD and compute the downstream analysis on each of the sampled synthetic data. Next, the results are aggregated according to a set of rules and we recover finally a more robust estimator for our downstream analysis. Further discussion about the Rubin's rules can be found for example in (Reiter and Raghunathan, 2007).

In our experiments, we sampled 100 data sets from the PPD learned using the aligned variant, and applied the Rubin's rules to compute a mean and std. estimate for the Poisson regression coefficients. Finally, the obtained means were compared to the Poisson regression coefficients learned using the original data.

J Experimental setup for full-rank Gaussian approximation

In this experiment we create simulated data where we control the amount of correlations between data dimensions as the ratio ρ of non-zero off-diagonal entries in the correlation matrix. To generate data with d dimensions and correlation density ρ , we

1. generate a correlation matrix \mathbf{C} using Algorithm J.1 with inputs $d, \rho, \alpha = 8, \beta = 10$,
2. sample a diagonal matrix \mathbf{D} of marginal variances, where $\{D\}_{ii} \sim \exp\{\mathcal{N}(0, 0.2^2)\}$,
3. obtain the covariance matrix $\mathbf{\Sigma} = \mathbf{DCD}$
4. sample $N = 10\,000$ data points $\mathbf{x}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$
5. sample random regression weight vector $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

6. sample $\mathbf{y} \sim \mathcal{N}(\mathbf{X}\mathbf{w}, \sigma_y^2)$, with $\sigma_y = 1$.

We perform the above for all combinations of $d = 100, 200$ and $\rho = 0.2, 0.8$. We then use vanilla DPVI and DPVI with aligned gradients to learn the full-rank Gaussian posterior approximation to the Bayesian linear regression model with priors

$$\mathbf{y} \sim \mathcal{N}(\mathbf{X}\mathbf{w}, \sigma_y^2), \quad (\text{J.1})$$

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (\text{J.2})$$

$$\sigma_y \sim \text{Gamma}(0.1, 0.1). \quad (\text{J.3})$$

We run the inference for 1 000 epochs, gradient clipping threshold 0.2 and subsampling ratio 0.01.

For the same d, ρ we then generate another 10 000 data points and compute the log-likelihood using the obtained posterior approximation. We repeat the inference and evaluation 50 times for each method and combination of d and ρ , keeping the generated training and testing set fixed.

Algorithm J.1 Routine to generate a d -dimensional correlation matrix with given density ρ and strength of correlations controlled by α, β .

Require: $d, \rho \in [0, 1], \alpha > 0, \beta > 0$

Ensure: correlation matrix \mathbf{C}

$K \leftarrow \rho \frac{d(d-1)}{2}$ ▷ number of non-zero off-diagonal entries

$U \leftarrow \emptyset$

$k \leftarrow 0$

$\mathbf{C} \leftarrow \mathbf{I}_d$

for $k = 1, \dots, K$ **do**

 sample $(i, j) \in T \setminus U$ at random

 sample $c \sim \text{Beta}(\alpha, \beta)$

▷ sample correlation strength, controlled by α and β

 sample $f \in \{-1, 1\}$ at random

▷ sample sign of correlation

$C_{ij} \leftarrow fc$

$C_{ji} \leftarrow C_{ij}$

$U \leftarrow U \cup \{(i, j), (j, i)\}$

end for

K Runtimes

K.1 UKB experiments

In this experiment, we ran all the variants separately for 10 seeds and 4 levels of privacy. Additionally, we experimented with different runtimes and initialisations. For a training of 1 000 epochs, a single repeat takes between 20 to 40 minutes. The runtime scales linearly with the number of epochs.

Further, we computed the downstream task for the aligned variants, which includes generating the 100 synthetic data sets and fitting the downstream Poisson regression model on those 100 synthetic data. This procedure takes between 30 to 60 minutes to complete.

A rough estimate of the runtimes for the UKB experiment is given in Table K.1. A single CPU core with 8gb of memory was used for all the runs.

K.2 Adult experiments

A single training repeat of learning the logistic regression model for all the different variants of DPVI, took between 10 and 30 minutes to finish on a single CPU core with 8gb of memory assigned. In total, the Adult experiment was repeated 50 times for four different levels of privacy. Therefore the total runtime of all the experiments is **between 2 000 and 6 000 minutes**.

Table K.1: Estimated runtimes for UKB experiment

#epochs	single repeat runtime	repeats	# epsilon values	# initial values
200	4-8min	10	4	5
400	8-16min	10	4	1
600	12-18min	10	4	1
800	16-32min	10	4	1
1000	20-40min	10	4	1
2000	40-80min	10	4	1
4000	80-160min	10	4	1
8000	160-320min	10	4	1

The variance in running times is likely due to differences in computation nodes in the clustered assigned by a automatic run scheduler.

K.3 Full-rank experiments

A single run for this experiment consisted of the inference using both vanilla and aligned DPVI with full-rank approximations. All runs were executed on a computing cluster utilising Nvidia K80, A100, P100, V100 GPU hardware, to which the runs were allocated automatically to balance overall load. As a result, runtimes varied slightly: Runs for and 100 dimensional data took 6-8 minutes to finish, runs for 200 dimensions took 8-10 minutes. With a total of 4 data set configurations and 50 repeats for each, the total runtime is 1 400 to 1 800 minutes.

L Gradient distributions for different variants




Figure L.1

Figure L.1 shows the distributions of gradient norms for variational means and scales for different variants of DPVI discussed in Section 3.1.1 when s_q is set to 0.1. Figure L.1a clearly shows the different magnitudes for variational standard deviation in vanilla DPVI. Figure L.1b demonstrates that natural gradients simply reverse the problem. Figure L.1c shows that the scaling approach achieves matching magnitudes quite well. However, it comes at the cost of increasing the norm of the full (combined) gradient and therefore increased sensitivity.