
Optimising Distributions with Natural Gradient Surrogates

Jonathan So
University of Cambridge

Richard E. Turner
University of Cambridge

Abstract

Natural gradient methods have been used to optimise the parameters of probability distributions in a variety of settings, often resulting in fast-converging procedures. Unfortunately, for many distributions of interest, computing the natural gradient has a number of challenges. In this work we propose a novel technique for tackling such issues, which involves reframing the optimisation as one with respect to the parameters of a *surrogate* distribution, for which computing the natural gradient is easy. We give several examples of existing methods that can be interpreted as applying this technique, and propose a new method for applying it to a wide variety of problems. Our method expands the set of distributions that can be efficiently targeted with natural gradients. Furthermore, it is fast, easy to understand, simple to implement using standard autodiff software, and does not require lengthy model-specific derivations. We demonstrate our method on maximum likelihood estimation and variational inference tasks.

1 INTRODUCTION

Many problems in scientific research require the optimisation of probability distribution parameters. Such problems are of fundamental importance in the fields of machine learning and statistics. The goal of these optimisations is to minimise a function $f(\theta)$ where $\theta \in \Theta$ is the parameter vector of a probability distribution q . Notable examples include maximum likelihood estimation (MLE) and variational inference (VI).

When f is differentiable, the archetypal optimisation

method is gradient descent (GD). GD takes steps in the direction in parameter space that gives the greatest decrease in f for an infinitesimally small step size, which is equivalent to following the direction of the negative gradient. While GD has many desirable properties, convergence can often be slow, and much effort has been spent in designing methods that offer improved convergence properties.

Natural gradient descent (NGD) is one such method for optimising probability distribution parameters (Amari, 1998). In contrast with GD, NGD moves the parameters in the direction of steepest descent on a manifold of probability distributions, where steepness is defined with respect to a divergence measure between distributions. This notion of steepness is intrinsic to the manifold, and so it is not dependent on parameterisation, dispensing with a well-known pathology of GD.

There are some distributions for which the natural gradient can be computed efficiently. However, in its most general form, computing the natural gradient involves instantiating and inverting the Fisher information matrix, defined as

$$F(\theta) = \mathbb{E}_{q_\theta(x)} \left[\nabla_\theta [\log q_\theta(x)] \nabla_\theta [\log q_\theta(x)]^\top \right], \quad (1)$$

posing several challenges.¹ First, for some distributions, the Fisher can become singular, in which case the natural gradient is undefined. Second, it may not be available in closed form, in which case it must be estimated using sampling or other numerical methods. Finally, instantiating and inverting it explicitly costs $O(m^2)$ in memory and $O(m^3)$ in computation, for m the number of parameters. For even moderately large m , this can be prohibitively expensive.

In this paper we propose a simple technique for tackling optimisation problems in which computing the natural gradient is problematic. Namely, we reframe

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹We use e.g. q , to refer to both a distribution and its density function, with the intention made clear from context. We will use a subscript, e.g. q_θ , when we wish to make the dependence on parameters explicit. Our gradient notation is explained in Appendix A.

the problem as an optimisation with respect to a *surrogate* distribution \tilde{q} , for which computing natural gradients is easy, and perform the optimisation in that space. With a judiciously chosen surrogate, convergence can be rapid.

We find that a number of existing methods can be interpreted as applications of this technique: the natural gradient VI method of Lin et al. (2019) for exponential family mixtures, stochastic natural gradient expectation propagation (Hasenclever et al., 2017), a fixed-point iteration scheme for optimising elliptical copulas (Hernández et al., 2014), as well as the typical use of natural gradients in large-scale supervised learning settings. We also describe a new method for applying the technique to a wide variety of problems, using exponential family (EF) surrogate distributions. Our method is easy to understand, simple to implement using standard autodiff software, and does not require the use of lengthy model-specific derivations.

We present a number of experiments in which we find that our method significantly reduces the time to convergence compared to existing best-practice methods, both in number of iterations and wall-clock time. Our experiments consist of a variety of MLE and VI tasks, but we note that the method is applicable more generally to optimisations involving probability distributions, or indeed to any optimisation in which we think a manifold of probability distributions may serve as a useful surrogate for the solution space.

Our main contributions can be summarised as follows:

1. We propose a novel technique for optimising probability distribution parameters.
2. We prove the validity of this technique under stated conditions.
3. We find several examples of existing methods that can be viewed as applying this technique.
4. We describe a simple new method for applying it, using known properties of EF distributions.
5. We present a variety of MLE and VI experiments in which our method achieves significantly faster convergence than existing best-practice methods.

This paper is structured as follows. In Section 2 we provide a brief overview of NGD and EF distributions. Section 3 covers items 1 and 4 of the contributions above, as well as an overview of item 2 (details found in Appendix B). Sections 4 and 5 cover items 5 and 3 respectively. Finally, in Section 6, we discuss limitations of our method, and avenues for future research.

2 BACKGROUND

In this section we begin with an overview of NGD. We then briefly describe a class of distribution families for which we can efficiently compute natural gradients, namely exponential families. NGD and EF distributions will together serve as a foundation for the method we introduce in Section 3.

2.1 Natural Gradient Descent

NGD updates the parameters of a distribution q by taking a step proportional to the gradient of the objective f , preconditioned by the inverse of the Fisher matrix of q . That is, the update in parameters at step t of the optimisation is given by

$$\theta_{t+1} = \theta_t - \epsilon_t [F(\theta_t)]^{-1} \nabla f(\theta_t) \quad (2)$$

where $F(\theta)$ is given by (1). The step size ϵ_t may follow some pre-defined schedule, or be found by line search. It can be shown that as $\epsilon_t \rightarrow 0^+$, update (2) moves θ in the direction of steepest descent in f , on the manifold of probability distributions spanned by q , where steepness is defined with respect to a divergence measure between distributions (Ollivier et al., 2017). Because this notion of steepness is intrinsic to the manifold, NGD is locally invariant to parameterisation.

When NGD is applied to MLE objectives, it is Fisher-efficient (Amari, 1998), and can be seen as a robust approximation to Newton’s method (Martens and Grosse, 2015). These properties do not, in general, apply outside of this setting. Nevertheless, natural gradient methods have been applied in several other settings, often resulting in rapidly converging procedures (Kakade, 2001; Hoffman et al., 2013; Khan et al., 2017; Hasenclever et al., 2017).

2.2 Exponential Family Distributions

In this section we provide a brief introduction to EF distributions, whose properties support efficient natural gradient computation.

The EF of distributions defined by the vector-valued statistic function t and base measure ν , has density

$$q_\eta(x) = \nu(x) \exp(t(x)^\top \eta - A(\eta)), \quad (3)$$

where A is the log partition function, and η are the *natural parameters*. When the components of t are linearly independent, the family is said to be *minimal*.

There is an alternative parameterisation of q , given by

$$\mu(\eta) = \mathbb{E}_{q_\eta(x)}[t(x)], \quad (4)$$

where μ are known as the *mean parameters* of q . In minimal families the correspondence between η and μ

is one-to-one, and we denote the reverse map as $\eta(\cdot)$.² In this paper we will only consider minimal EFs.

Notable examples of EF distributions include the multivariate normal, gamma, categorical, and Dirichlet distributions to name but a few.

Let f_η and f_μ be functions related by $f_\eta = f_\mu \circ \mu$. It is a remarkable property of EFs that the natural gradients of f_η and f_μ with respect to η and μ , respectively, are given by

$$[F(\eta)]^{-1} \nabla f_\eta(\eta) = \nabla f_\mu(\mu(\eta)) \quad (5)$$

$$[F(\mu)]^{-1} \nabla f_\mu(\mu) = \nabla f_\eta(\eta(\mu)). \quad (6)$$

That is, the natural gradient in one parameterisation is simply given by the regular gradient with respect to the other (Hensman et al., 2012). Provided that we have an efficient way of converting between natural and mean parameters, this allows us to compute the natural gradient without explicitly instantiating or inverting $F(\theta)$.

3 METHOD

The exponential families introduced in Section 2 provide us with a relatively rich set of distributions for which we can efficiently compute natural gradients. However, there are many distributions outside of this set for which computing natural gradients remains difficult. In this section, we present details of our technique, which expands the set of distributions that we can efficiently target with natural gradients.

3.1 Surrogate Natural Gradient Descent

The main idea in this paper is simple. When faced with an optimisation objective $f(\theta)$, where $\theta \in \Theta$ are the parameters of a distribution q for which computing the natural gradient update is problematic, we solve the problem in two steps. First, we re-parameterise f as a function of some other parameters $\tilde{\theta} \in \tilde{\Theta}$, related by $\theta = g(\tilde{\theta})$, and define the reparameterised objective

$$\tilde{f}(\tilde{\theta}) = f(g(\tilde{\theta})). \quad (7)$$

Second, we *interpret* $\tilde{\theta}$ as the parameters of a surrogate distribution \tilde{q} , for which computing natural gradients is easy, and then perform NGD in \tilde{f} with respect to \tilde{q} and $\tilde{\theta}$. That is, we perform a sequence of updates

$$\tilde{\theta}_{t+1} = \tilde{\theta}_t - \epsilon_t [\tilde{F}(\tilde{\theta}_t)]^{-1} \nabla \tilde{f}(\tilde{\theta}_t), \quad (8)$$

²Due to our overloaded use of μ and η to denote both specific parameters and maps between parameterisations, we will use e.g. $\mu(\cdot)$ to disambiguate the latter if the intention is not clear from context.

where $\tilde{F}(\tilde{\theta})$ is the Fisher information matrix of \tilde{q} . We refer to q and \tilde{q} as the *target* and *surrogate* distributions respectively. Update (8) has the straightforward interpretation of performing preconditioned GD in a reparameterised objective. Upon converging to a local minimiser $\tilde{\theta}^*$ of \tilde{f} , a solution to the original problem is obtained by $\theta^* = g(\tilde{\theta}^*)$. We call this technique *surrogate natural gradient descent* (SNGD).

As a simple example, when q is a multivariate Student's t distribution with known degrees of freedom ν , we can choose \tilde{q} to be multivariate normal. If $\tilde{\theta}$ contains the mean and scale matrix parameters of \tilde{q} , and θ the mean and covariance matrix of q , then a natural choice for g is simply the identity map on $\tilde{\Theta} = \Theta$.

It is natural to question which properties SNGD shares with NGD under q . Given that SNGD is NGD (with respect to \tilde{q}), it remains locally invariant to parameterisation, and performs steepest descent in a statistical manifold (that of \tilde{q}). What is not guaranteed, is that the statistical manifold of \tilde{q} remains useful for the optimisation of \tilde{f} . For example, in MLE settings, SNGD will not in general retain the asymptotic efficiency of NGD under q . However, what it gains is tractability, and as we demonstrate in Section 4, the practical performance benefits can be significant. In some cases, remarkably, SNGD can actually outperform NGD under q , with respect to both the original parameters θ , and the reparameterisation defined by g .

3.2 Choice of Surrogate

The performance of SNGD relies crucially on appropriate choices for \tilde{q} and g . A natural assumption would be that $\tilde{q}_{\tilde{\theta}}$ should be an approximation to $q_{g(\tilde{\theta})}$, and in some cases this can be a useful guide. However, often an effective surrogate can be found that has support over an entirely different space than that of q ; in fact, this is true for most of the examples in this paper, as can be seen in Table 1.

A more general principle is that a surrogate should be chosen such that $\tilde{\theta}$ has similar *local* effects on Kullback-Leibler (KL) divergences in q than it does in \tilde{q} . More specifically, the effect of a small (infinitesimal) perturbation δ in $\tilde{\theta}$ should have a similar impact on the KL divergence from $\tilde{q}_{\tilde{\theta}+\delta}$ to $\tilde{q}_{\tilde{\theta}}$ as it does on that from $q_{g(\tilde{\theta}+\delta)}$ to $q_{g(\tilde{\theta})}$.³ That is, at $\delta = \mathbf{0}$, we would like:

$$\nabla_\delta \left(\text{KL}[\tilde{q}_{\tilde{\theta}+\delta} \parallel \tilde{q}_{\tilde{\theta}}] \right) = \nabla_\delta \left(\text{KL}[q_{g(\tilde{\theta}+\delta)} \parallel q_{g(\tilde{\theta})}] \right). \quad (9)$$

If equality (9) is satisfied then SNGD will move in the same direction in $\tilde{\theta}$ as NGD under (reparameterised) q (see Appendix G for details). Finding tractable surrogates for which equality holds exactly will not typically

³See Appendix G for a definition of the KL divergence.

be possible, but this motivates choosing \tilde{q} , g , for which it is approximately true.

3.3 Equivalence with Optimisation of f

In Appendix B we prove a number of results regarding the validity of SNGD, which we summarise here.

Let Θ and $\tilde{\Theta}$ be open subsets of \mathbb{R}^i and \mathbb{R}^j respectively, with $i \leq j$. Let $g : \tilde{\Theta} \rightarrow \Theta$ be twice differentiable, with Jacobian of rank i everywhere, and with $g(\tilde{\Theta}) = \Theta$. Finally, let $f : \Theta \rightarrow \mathbb{R}$ be twice continuously differentiable, and define $\tilde{f} = f \circ g$.

Under these conditions, optimising \tilde{f} is equivalent to optimising f in the following sense: finding a local minimiser of \tilde{f} also gives us a local minimiser for f , and all local minima of f are attainable through \tilde{f} . Furthermore, \tilde{f} does not have any non-strict saddle points that are not also present at the corresponding points in f , and so \tilde{f} does not introduce any additional spurious attractors (Lee et al., 2016).

3.4 Exponential Family Surrogates

In this paper we choose \tilde{q} from the set of EF distributions described in Section 2, allowing us to perform update (8) efficiently without explicitly instantiating or inverting $\tilde{F}(\tilde{\theta})$. Algorithm 1 provides pseudocode for an implementation of SNGD when $\tilde{\Theta}$ is either the mean or natural domain of an EF. It assumes the existence of an autodiff operator `grad`, and an overloaded function `dualparams`, which maps from natural to mean parameters, or vice-versa, depending on the type of its argument. A full explanation of Algorithm 1 can be found in Appendix C.1.

Algorithm 1 SNGD with EF surrogate

Require: objective $f : \Theta \rightarrow \mathbb{R}$
Require: parameter mapping $g : \tilde{\Theta} \rightarrow \Theta$
Require: initial surrogate parameters $\tilde{\theta}_0 \in \tilde{\Theta}$
Require: step size schedule $\{\epsilon_t \in \mathbb{R}_+ : t = 0, 1, \dots\}$
 $\tilde{f}_{\text{dual}}(\cdot) := f(g(\text{dualparams}(\cdot)))$
 $t \leftarrow 0$
while not converged **do**
 $\tilde{\nabla} \leftarrow \text{grad}[\tilde{f}_{\text{dual}}](\text{dualparams}(\tilde{\theta}_t))$
 $\tilde{\theta}_{t+1} \leftarrow \tilde{\theta}_t - \epsilon_t \tilde{\nabla}$
 $t \leftarrow t + 1$
end while
return $g(\tilde{\theta}_t)$

The domains of valid natural and mean parameters of an EF are open convex sets. It is possible that finite-length NGD steps in these parameters can move outside of the valid domain. These steps can be corrected by using a backtracking linesearch, an approach taken

in previous work on natural gradients (Khan and Lin, 2017). However, we find that when this is required, the backtracking typically occurs so infrequently as to have a relatively small impact on performance.

3.5 Auxiliary Parameters

In some cases, a chosen surrogate’s parameters, $\tilde{\theta}$, may not be sufficient to fully specify θ . To handle such cases, we can generalise SNGD by augmenting $\tilde{\theta}$ with an additional vector of parameters, $\lambda \in \Lambda$, that are *not* optimised with natural gradients. Our reparameterised objective is then given by

$$\tilde{f}(\tilde{\theta}, \lambda) = f(g(\tilde{\theta}, \lambda)). \quad (10)$$

We optimise (10) by updating $\tilde{\theta}$ using natural gradients as before, and λ using standard gradient-based techniques, either jointly or in alternation. For example, we can apply the sequence of updates

$$\tilde{\theta}_{t+1} = \tilde{\theta}_t - \epsilon_t [\tilde{F}(\tilde{\theta}_t)]^{-1} \nabla_{\tilde{\theta}} \tilde{f}(\tilde{\theta}_t, \lambda_t), \quad (11)$$

$$\lambda_{t+1} = \lambda_t - \varepsilon_t \nabla_{\lambda} \tilde{f}(\tilde{\theta}_t, \lambda_t), \quad (12)$$

where ϵ_t , ε_t are the step sizes for $\tilde{\theta}$, λ , respectively at time t . Update (12) corresponds to GD in λ , but we may equally use any other first-order optimiser.

Revisiting our example from Section 3.1, if q is a multivariate Student’s t distribution, but now with *unknown* degrees of freedom ν , we can again directly map from $\tilde{\theta}$, which contains the location and scale matrix parameters of \tilde{q} , to the mean and covariance of q . ν has no analogue in the multivariate normal distribution however, and so we can capture this with $\lambda = (\nu)$. g can then simply be the identity map on $\tilde{\Theta} \times \Lambda = \Theta$.⁴

This extension expands the set of distributions that can be targeted by SNGD. Pseudocode for this extension is given in Appendix C.2, with examples appearing in the experiments of Sections 4.2 and 4.3.

4 RESULTS

In this section we present a number of experiments demonstrating the utility of SNGD. Natural gradient methods have been used in a variety of settings, but perhaps most extensively for MLE (Amari, 1998; Bernacchia et al., 2018; Martens and Grosse, 2015; Ren and Goldfarb, 2019; Roux et al., 2007) and VI (Hoffman et al., 2013; Khan and Lin, 2017; Khan et al., 2017, 2018; Lin et al., 2019; Salimbeni et al., 2018). We therefore used a variety of MLE and VI tasks as test cases.

⁴In practice it would be preferable to choose $\lambda = (\log \nu)$, with g adjusted accordingly, so that optimisation of λ is unconstrained.

Table 1: Summary of the surrogate-target pairs appearing in this paper, with links to the relevant sections. For the sake of readability we have omitted the conventional multivariate prefix from distribution names, as this information is conveyed by the support. Parameter mappings for these examples are found in Appendix E.

TARGET		SURROGATE		
DISTRIBUTION	SUPPORT	DISTRIBUTION	SUPPORT	SECTION
Negative binomial	\mathbb{N}	Gamma	\mathbb{R}_+	4.1
Negative binomial mixture	\mathbb{N}	Gamma mixture model	$\mathbb{R}_+ \times \{1, \dots, k\}$	4.4
Skew-normal	\mathbb{R}^d	Normal	\mathbb{R}^d	4.2
Skew-normal mixture	\mathbb{R}^d	Normal mixture model	$\mathbb{R}^d \times \{1, \dots, k\}$	4.4
Skew- t	\mathbb{R}^d	Normal	\mathbb{R}^d	4.2
Elliptical copula	$[0, 1]^d$	Zero-mean normal	\mathbb{R}^d	4.3

In MLE the goal is to find parameters θ of distribution q that maximise the (log) likelihood of observed data $\{x_i\}_{i=1}^n$ under q . That is, the objective function for MLE tasks is of the form

$$f(\theta) = -\sum_{i=1}^n \log q_{\theta}(x_i). \quad (13)$$

In VI we are given a generative model $p(x)p(\mathcal{D} | x)$, and observed data \mathcal{D} , and the goal is to find parameters θ of distribution q that minimise the Kullback-Leibler (KL) divergence from q to the posterior $p(x|\mathcal{D})$. The objective function for VI tasks is of the form

$$f(\theta) = -\mathbb{E}_{q_{\theta}(x)} \left[\log \frac{p(\mathcal{D}, x)}{q_{\theta}(x)} \right]. \quad (14)$$

It can be shown that minimising (14) is equivalent to maximising a lower bound on the log marginal likelihood: $\log p(\mathcal{D})$. In the VI experiments of this section, gradients of (14) were estimated by applying the reparameterisation trick (Kingma and Welling, 2014) to the target distribution.

In each experiment, we compared SNGD with a number of baselines. For tasks that were small scale, and had objectives that could be computed deterministically, we compared SNGD with two baselines: GD and BFGS (Nocedal and Wright, 2006). In these cases, we used exact line search to determine step sizes for each method, in order to compare search direction independent of hyperparameter settings. For larger tasks, or where only stochastic estimates of f were available (minibatched or VI experiments), we used Adam as the standard baseline, and used grid search to determine the best hyperparameters for each method.

Further details of all experiments are given in Appendix D. A summary of surrogate-target pairs used for SNGD in these experiments can be found in Table 1, and details of parameter mappings can be found in Appendix E.

We display average training curves for each experiment, with error bars corresponding to 2 standard errors, computed across 10 random seeds. In Appendix F we plot Pareto frontiers, showing the best performance achieved by *any* hyperparameter setting, as a function of both iteration count and wall-clock time.

In this section we use n to denote the number of training data points, and d for the dimensionality of the random variable under q . For brevity, we omit the conventional multivariate prefix from distribution names in this section, as this is implicit for $d > 1$.

4.1 Negative Binomial Distribution

The negative binomial distribution is widely used for modelling discrete data (Fisher, 1941; Lloyd-Smith et al., 2005; Lloyd-Smith, 2007; Orooji et al., 2021; Kendall et al., 2023). Maximum likelihood estimates of the negative binomial parameters are not available in closed form, and must be found numerically.

Let $\theta = (r, s)$ be the parameters of the negative binomial distribution with probability mass function

$$q(x) = \binom{x+r-1}{x} (1-s)^x s^r. \quad (15)$$

The gamma distribution can be seen as a continuous analogue of the negative binomial, and as it is a EF, it is a natural choice of surrogate when targeting the negative binomial with SNGD.

Let \tilde{q} , therefore, be the gamma distribution with parameters α, β and probability density function

$$\tilde{q}(x) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} x^{\alpha-1} \exp^{-\beta x}. \quad (16)$$

Guenther (1972) used the gamma distribution to approximate the CDF of the negative binomial, using the mapping $(r, s) = (\alpha/(1-\beta), \beta)$. Let g be the equivalent mapping, but defined in terms of *mean* parameters

of \tilde{q} . To ensure that r and s correspond to valid negative binomial parameters, we restrict Θ to the subset of gamma mean parameters for which $\beta < 1$.⁵ Finally, let f be defined as (13), where $\{x_i\}_{i=1}^n$ is the dataset of Fisher (1941), consisting of counts of ticks observed on a population of sheep ($n=82$, $d=1$).

Using these definitions, we applied Algorithm 1, modified to use a line search as discussed at the start of this section. We compared SNGD with our standard baselines: GD and BFGS. Additionally, this task was sufficiently small that we were able to compare with NGD under q by estimating and inverting the Fisher directly.

Figure 1 shows that SNGD outperformed all of the baselines. Remarkably, this included NGD under q ; this surprising result is explored in detail in Appendix H. There, we show that this was true even when NGD was performed with respect to the reparameterisation defined by g , implying that this is not simply an artefact of SNGD using a more effective parameterisation; rather, it is the change of parameterisation *and* distribution that results in such rapid convergence.




Figure 1: Negative binomial MLE on the sheep dataset. (left) Training curves. (right) Histogram of the observed counts, overlaid with the PDF of the gamma surrogate of SNGD at convergence.

4.2 Skew-Elliptical Distributions

The skew-elliptical distributions are a flexible family of unimodal multivariate distributions, allowing for features such as asymmetry and heavy tails (Azzalini, 2013). In this section we consider in particular the skew-normal and skew- t distributions. The skew-normal can be viewed as an asymmetric generalisation of the normal distribution. Similarly, the skew- t can be viewed as an asymmetric generalisation of the Student’s t distribution.

The normal distribution is a natural surrogate for these distributions. We used it as such in a number of experiments, mapping its mean and covariance to

⁵It can be shown that Θ remains an open convex set.




Figure 2: Training curves for MLE on the miniboone dataset ($n=32,840$, $d=43$) using (left) skew-normal, and (right) skew- t distributions.

parameters playing similar roles in the target distributions, with any other parameters captured by λ , as outlined in Section 3.5.

We performed two tasks on real data in our experiments. In the first, the skew-elliptical distributions were used for density estimation on the UCI miniboone dataset ($n=32,840$, $d=43$) (Roe, 2010), fitting their parameters via MLE. In the second, they were used as approximate posteriors in VI for a Bayesian logistic regression model on the UCI covtype dataset (Blackard, 1998); in this task we used a small subsample of observations in order to retain some uncertainty in the posterior ($n=500$, $d=53$). We also performed a further MLE experiment using high-dimensional synthetic data ($n=10,000$, $d=1,000$).

Note that the numbers of free parameters in these optimisations were $O(d^2)$ due to the covariance-like parameters of the target distributions. The skew-normal distribution had 1,032, 1,537, and 502,500 parameters in the miniboone, covtype, and synthetic experiments, respectively. For the skew- t distribution those numbers were 1,033, 1,537, and 502,501.

Figures 2, 3 and 4 show training curves from the miniboone MLE, covtype VI, and synthetic MLE experiments, respectively. In all cases SNGD significantly outperformed Adam. In the particular case of the skew-normal VI experiment, it was also possible to apply the natural gradient method of Lin et al. (2019) based on *minimal conditional exponential family* (MCEF) distributions, and so we included this as an additional baseline; its performance was virtually identical to that of SNGD, however, we note that this method can itself be viewed as applying SNGD, a point we discuss further in Section 5.

4.3 Elliptical Copulas

Copula models define a distribution on the unit hypercube $[0, 1]^d$, for which each of the marginals is uniform on $[0, 1]$. If $x \in [0, 1]^d$ is distributed according to cop-




Figure 3: Training curves for Bayesian logistic regression VI on the covtype dataset ($n=500$, $d=53$) using (left) skew-normal, and (right) skew- t approximations. MCEF corresponds to the natural gradient VI method of Lin et al. (2019).




Figure 4: Training curves for MLE on a synthetic dataset ($n=10,000$, $d=1,000$) using (left) skew-normal, and (right) skew- t distributions.

ula q , and $y_i = P_i^{-1}(x_i)$, where P_i^{-1} is the inverse CDF of a distribution q_i , then y_i will have marginal distribution q_i . The y_i will be dependent in general, with dependence structure determined by $q(x)$.

Copula models are widely used in finance for modelling high-dimensional variables, in part because they allow marginal distributions to be modelled separately from the dependence structure (Dewick and Liu, 2022).

The class of elliptical copulas are defined as those copulas which can be used to generate elliptical distributions (Frahm et al., 2003). Common examples are the Gaussian and t copulas. Elliptical copulas are parameterised by a correlation matrix R , as well as a (possibly empty) set of additional parameters.

A simple way to target elliptical copulas with SNGD is to use a zero-mean normal surrogate, mapping its correlation matrix to R . Any additional parameters may then be captured by λ , as outlined in Section 3.5.

We used this approach to perform MLE of the t -copula on 5 years of daily stock returns from the FTSE 100 universe ($n=1,515$, $d=93$), with marginals estimated as (univariate) Student’s t distributions. In this experiment q had 4,279 free parameters. Figure 5 shows that SNGD converged significantly faster than Adam.




Figure 5: t -copula MLE using 5 years of daily stock return data ($n=1,515$, $d=93$). (left) Training curves. (right) Contours of a 2D marginal density from the fitted copula, overlaid with the training data.

4.4 Mixture Distributions

A mixture distribution expresses a complicated density as a convex combination of k simpler densities:

$$q(x) = \sum_{i=1}^k \pi_i q_{\theta_i}(x), \quad (17)$$

where $\pi \in \Delta^{k-1}$, and q_{θ_i} is known as the i -th *component* distribution. For each mixture, we can also define a corresponding mixture *model*, a joint distribution, with density

$$q(z, x) = \pi_z q_{\theta_z}(x), \quad (18)$$

where the mixture component identity, $z \in \{1, \dots, k\}$, is treated as a random variable. Note then, that $q(x) = \sum_{i=1}^k q(z = i, x)$. That is, a mixture model has a mixture as its *marginal*. A mixture of EFs is not in general itself an EF. However, it can be shown that the corresponding mixture model (joint distribution) *is* an EF (see Appendix J).

When an EF is an appropriate surrogate for applying SNGD to a given target family, it is therefore straightforward to consider the extension to mixtures of that family by using an EF mixture *model* as a surrogate. We emphasise that in doing so, the surrogate distribution has support over more variables (z and x) than the target distribution (just x).

In Section 4.1 we demonstrated the use of a gamma surrogate for optimising negative binomial distribution parameters with SNGD. It is therefore straightforward to use a gamma mixture model to target a negative binomial mixture. As an experiment, we performed MLE of a 5 component negative binomial mixture, using a dataset consisting of the number of daily COVID-19 hospital admissions in the UK over a 3 year period ($n=1,120$, $d=1$). Figure 6 shows that SNGD significantly outperformed our baselines.

As a further example, similarly, we can use a normal




Figure 6: Negative binomial mixture MLE ($k=5$) using the number of daily COVID hospital admissions in the UK over a 3 year period ($n=1,120$, $d=1$). (left) Training curves. (right) Histogram of the observed counts, overlaid with the PDF of (a marginal of) the gamma mixture model surrogate of SNGD at convergence.

mixture model as the surrogate for a skew-normal mixture. Figure 7 shows the qualitative improvement a skew-normal mixture can offer over a normal mixture on the Bayesian logistic regression VI task of Murphy (2012) ($n=60$, $d=2$). Figure 8 shows training curves for this task, as well as for the UCI covertime VI task of Section 4.2. In the latter, q had 4,613 and 7,689 free parameters for $k=3$, 5, respectively.




Figure 7: Comparison of (top) normal and (bottom) skew-normal mixture posterior approximations in the Bayesian logistic regression VI task of Murphy (2012). Ground truth (GT) generated using MCMC samples.

5 RELATED WORK

Several existing methods can be interpreted as examples of the general technique presented in this paper.

Lin et al. (2019) introduced the class of MCEF distributions, and showed that natural gradients with respect to a particular parameterisation of MCEF distributions can be efficiently computed using an identity analogous to (5). They also showed that several standard distributions can be expressed as *marginals* of MCEF distributions, including the multivariate Student’s t and multivariate skew-normal distributions.




Figure 8: Training curves for Bayesian logistic regression VI with a skew-normal mixture approximation. (left) the toy dataset of Murphy (2012) ($n=60$, $d=2$), and (right) UCI covertime ($n=500$, $d=53$). k is the number of mixture components.

They then used NGD with respect to the MCEF joint as a surrogate for its marginal in a number of VI tasks. This mismatch between the surrogate and target distributions (our terminology) was not explicitly discussed by Lin et al. (2019), possibly because the two were so closely related (by marginalisation). SNGD is more general in that it makes no such assumption about the relationship between surrogate and target distributions. In particular, it is not restricted to targets that can be expressed as the marginal of a MCEF (or any other) distribution.

The stochastic natural gradient expectation propagation (SNEP) algorithm of Hasenclever et al. (2017) solves a saddle-point optimisation problem, with the inner optimisation being over a set of parameters $\{\eta_i\}$ known as the *site parameters*. The site parameters jointly parameterise a set of distributions for which computing natural gradients is not straightforward. SNEP instead treats each η_i as if it were the natural parameter vector of a EF distribution in its own right, and performs NGD in the dual (mean parameter) space. These (pseudo)distributions acted as surrogates for the more problematic set of distributions that were the ultimate targets of the optimisation.

Hernández et al. (2014) devised a fixed-point iteration scheme for optimising the correlation matrix parameter of an elliptical copula. Interestingly, although their procedure was motivated in an entirely different way, the resulting updates for R are identical to those performed by SNGD when applied in the manner described in Section 4.3.

In supervised learning, the goal is to model the conditional density $q(y|x)$ given training data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$. In this setting the Fisher matrix is usually defined with respect to the distribution $p_{\mathcal{D}}(x)q_{\theta}(y|x)$, where $p_{\mathcal{D}}(x)$ is the empirical distribution of x under \mathcal{D} . In practice, when the data set is large, the Fisher must be computed with respect

to some subset of the training inputs \mathcal{S} , such as the current minibatch. This is equivalent to using $p_{\mathcal{S}}(x)q_{\theta}(y|x)$ as a surrogate for $p_{\mathcal{D}}(x)q_{\theta}(y|x)$, where $p_{\mathcal{S}}(x)$ is the empirical distribution of x in \mathcal{S} . In the context of supervised deep learning, Ren and Goldfarb (2019) showed that exact natural gradients with respect to $p_{\mathcal{S}}(x)q_{\theta}(y|x)$ can be computed efficiently when $|\mathcal{S}|$ is small.

Several works have attempted to handle intractability of the natural gradient in other ways. Some methods are based on structured approximations to the Fisher, that allow inverse vector products to be computed efficiently (Heskes, 2000; Martens and Grosse, 2015; George et al., 2018). Garcia et al. (2023) instead approximated the inverse Fisher matrix directly by expressing it in terms of the solution to an optimisation problem. The widely used Adam optimiser employs a moving average of squared gradients in a diagonal approximation to the Fisher (Kingma and Ba, 2014), although it has been cautioned that this is more accurately viewed as an approximation to the *empirical* Fisher (Kunstner et al., 2019).

6 DISCUSSION

In this work we proposed a novel technique for optimising functions of probability distribution parameters: reframing the objective as an optimisation with respect to a *surrogate* distribution for which computing natural gradients is easy, and performing optimisation in that space. We found several existing methods that can be interpreted as applying this technique, and proposed a new method based on EF surrogates. We demonstrated that our method is able to converge rapidly on a variety of MLE and VI tasks. We believe our method can be readily applied to distributions outside of the set of examples given here. We also expect that new methods can be found by applying the more general technique that motivated our method.

The main limitation of our method is the need to find a suitable surrogate and reparameterisation for a given target distribution. In some cases there is an obvious candidate, such as when an EF distribution could serve as an approximation for the target, or when the target can be viewed as an EF that has been warped or transformed in some way. However, for most of the examples in this paper, the target and surrogate distributions do not even have support over the same space (see Table 1). Our choices of surrogates were largely guided by the intuition that $\tilde{\theta}$ should have similar local effects on KL divergences in q than it does in \tilde{q} . However, finding more prescriptive or systematic methods for choosing surrogates would be beneficial, and may be an interesting direction for future work.

Acknowledgements

We would like to thank Runa Eschenhagen, Jihao Andreas Lin, and James Townsend for providing valuable feedback. Jonathan So is supported by the University of Cambridge Harding Distinguished Postgraduate Scholars Programme. Richard E. Turner is supported by Google, Amazon, ARM, Improbable and EPSRC grant EP/T005386/1.

References

- S.-i. Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 1998.
- A. Azzalini. *The Skew-Normal and Related Families*. Cambridge University Press, 2013.
- A. Bernacchia, M. Lengyel, and G. Hennequin. Exact natural gradient in deep linear networks and its application to the nonlinear case. In *Advances in Neural Information Processing Systems*, 2018.
- D. P. Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 1997.
- J. Blackard. Covertypes. UCI Machine Learning Repository, 1998.
- M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-Lopez, F. Pedregosa, and J.-P. Vert. Efficient and modular implicit differentiation. In *Advances in Neural Information Processing Systems*, 2022.
- J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2011.
- B. Christianson. Reverse accumulation and attractive fixed points. *Optimization Methods & Software*, 1994.
- P. R. Dewick and S. Liu. Copula modelling to analyse financial data. *Journal of Risk and Financial Management*, 2022.
- R. A. Fisher. The negative binomial distribution. *Annals of Eugenics*, 1941.
- G. Frahm, M. Junker, and A. Szimayer. Elliptical copulas: applicability and limitations. *Statistics & Probability Letters*, 2003.

- J. R. Garcia, F. Freddi, S. Fotiadis, M. Li, S. Vakili, A. Bernacchia, and G. Hennequin. Fisher-legendre (fishleg) optimization of deep neural networks. In *International Conference on Learning Representations*, 2023.
- T. George, C. Laurent, X. Bouthillier, N. Ballas, and P. Vincent. Fast approximate natural gradient descent in a kronecker-factored eigenbasis. In *Neural Information Processing Systems*, 2018.
- W. C. Guenther. A simple approximation to the negative binomial (and regular binomial). *Technometrics*, 1972.
- L. Hasenclever, S. Webb, T. Lienart, S. Vollmer, B. Lakshminarayanan, C. Blundell, and Y. W. Teh. Distributed bayesian learning with stochastic natural gradient expectation propagation and the posterior server. *Journal of Machine Learning Research*, 2017.
- J. Hensman, M. Rattray, and N. Lawrence. Fast variational inference in the conjugate exponential family. In *Advances in Neural Information Processing Systems*, 2012.
- J. Hensman, N. Fusi, and N. Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, 2013.
- L. Hernández, J. Tejero, and J. Vinuesa. Maximum likelihood estimation of the correlation parameters for elliptical copulas. *arXiv:1412.6316*, 2014.
- T. Heskes. On “Natural” Learning and Pruning in Multilayered Perceptrons. *Neural Computation*, 2000.
- M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 2013.
- S. M. Kakade. A natural policy gradient. In *Advances in Neural Information Processing Systems*, 2001.
- M. Kendall, D. Tsallis, C. Wymant, A. Di Francia, Y. Balogun, X. Didelot, L. Ferretti, and C. Fraser. Epidemiological impacts of the nhs covid-19 app in england and wales throughout its first year. *Nature Communications*, 2023.
- M. Khan and W. Lin. Conjugate-Computation Variational Inference : Converting Variational Inference in Non-Conjugate Models to Inferences in Conjugate Models. In *International Conference on Artificial Intelligence and Statistics*, 2017.
- M. Khan, D. Nielsen, V. Tangkaratt, W. Lin, Y. Gal, and A. Srivastava. Fast and scalable Bayesian deep learning by weight-perturbation in Adam. In *International Conference on Machine Learning*, 2018.
- M. E. Khan, W. Lin, V. Tangkaratt, Z. Liu, and D. Nielsen. Variational adaptive-newton method. In *NeurIPS Workshop on Advances in Approximate Bayesian Inference*, 2017.
- D. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2014.
- D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- F. Kunstner, P. Hennig, and L. Balles. Limitations of the empirical fisher approximation for natural gradient descent. In *Advances in Neural Information Processing Systems*, 2019.
- J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, 2016.
- W. Lin, M. E. Khan, and M. Schmidt. Fast and simple natural-gradient variational inference with mixture of exponential-family approximations. In *International Conference on Machine Learning*, 2019.
- J. Lloyd-Smith, S. Schreiber, P. Kopp, and W. Getz. Superspreading and the effect of individual variation on disease emergence. *Nature*, 2005.
- J. O. Lloyd-Smith. Maximum likelihood estimation of the negative binomial dispersion parameter for highly overdispersed data, with applications to infectious diseases. *PLOS ONE*, 2007.
- J. Martens and R. Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*, 2015.
- T. Minka. Estimating a dirichlet distribution. Technical report, Microsoft, September 2000.
- T. Minka. Estimating a gamma distribution. Technical report, Microsoft, April 2002.
- K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.
- Y. Ollivier, L. Arnold, A. Auger, and N. Hansen. Information-geometric optimization algorithms: A unifying picture via invariance principles. *Journal of Machine Learning Research*, 2017.

A. Orooji, E. Nazar, M. Sadeghi, A. Moradi, Z. Jafari, and H. a. Esmaily. Factors associated with length of stay in hospital among the elderly patients using count regression models. *Medical Journal of the Islamic Republic Of Iran*, 2021.

G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2017.

Y. Ren and D. Goldfarb. Efficient subsampled gauss-newton and natural gradient methods for training neural networks. *arXiv:1906.02353*, 2019.

B. Roe. MiniBooNE particle identification. UCI Machine Learning Repository, 2010.

N. Roux, P.-a. Manzagol, and Y. Bengio. Topmoumoute online natural gradient algorithm. In *Advances in Neural Information Processing Systems*, 2007.

H. Salimbeni, S. Eleftheriadis, and J. Hensman. Natural gradients in practice: Non-conjugate variational inference in gaussian process models. In *International Conference on Artificial Intelligence and Statistics*, 2018.

M.-A. Sato. Online model selection based on the variational bayes. *Neural Computation*, 2001.

L. Tatzel, P. Hennig, and F. Schneider. Late-phase second-order training. In *Has it Trained Yet? NeurIPS Workshop*, 2022.

M. Wainwright and M. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 2008.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes, see Section 3 and Appendix B.]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes, see Appendix C.1.]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes, see Appendix D.]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes, see Section 3.3 and Appendix B.]
 - (b) Complete proofs of all theoretical results. [Yes, see Appendix B.]
 - (c) Clear explanations of any assumptions. [Yes, see Section 3.3 and Appendix B.]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes, see Appendix D.]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes, see Appendix D.]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes, see Appendix D.]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes, see Appendix D.]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes, see Appendix D.]
 - (b) The license information of the assets, if applicable. [Yes, see Appendix D.]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable.]
 - (d) Information about consent from data providers/curators. [Not Applicable.]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable.]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable.]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable.]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable.]

A GRADIENT NOTATION

In this section we introduce our notation for gradients and related quantities. We largely follow the notation of Bertsekas (1997), with one addition. For function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, the gradient at x , assuming all partial derivatives exist, is given by

$$\nabla f(x) = \left[\frac{\partial f(x)}{\partial x_1}, \frac{\partial f(x)}{\partial x_2}, \dots, \frac{\partial f(x)}{\partial x_m} \right]^\top. \quad (19)$$

Note, in particular, that this implies $\nabla f(g(y))$ is the gradient of f evaluated at $x = g(y)$. The Hessian of f , denoted $\nabla^2 f(x)$, is the matrix with entries given by

$$[\nabla^2 f(x)]_{ij} = \frac{\partial^2 f(x)}{\partial x_i \partial x_j}. \quad (20)$$

If $f : \mathbb{R}^{m+n}$ is a function of $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^n$, then

$$\nabla_x f(x, y) = \left[\frac{\partial f(x, y)}{\partial x_1}, \frac{\partial f(x, y)}{\partial x_2}, \dots, \frac{\partial f(x, y)}{\partial x_m} \right]^\top \quad (21)$$

with $\nabla_y f(x, y)$ defined similarly.

When $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a vector valued function, the *gradient matrix* of f , denoted $\nabla f(x)$, is the transpose of the Jacobian of f . That is, the matrix with i -th column equal to the gradient of f_i , the i -th component of f :

$$\nabla f(x) = [\nabla f_1(x) \dots \nabla f_n(x)]. \quad (22)$$

Finally (our addition), where ∇ is immediately followed by a bracketed expression, we use this to denote the gradient of an anonymous function, with definition given by the bracketed expression, and gradient taken with respect to the subscript, e.g.

$$\nabla_x(2f(x)) = 2\nabla f(x), \quad (23)$$

and in these cases (only), evaluation of the gradient at, e.g. $x = g(y)$, is denoted

$$\nabla_x(2f(x))|_{g(y)}. \quad (24)$$

B EQUIVALENCE WITH OPTIMISATION OF f

In this appendix we show that under certain conditions, optimising \tilde{f} is equivalent to optimising f in the following sense: finding a local minimiser of \tilde{f} also gives us a local minimiser for f (Proposition 1), and all local minima of f are attainable through \tilde{f} (Proposition 2). Furthermore, we show that f does not have any non-strict saddle points that are not also present at the corresponding points in \tilde{f} (Proposition 5, with support from Propositions 3 and 4).

Note that the results derived here are more general than those that are summarised in Section 3.3. The conditions stated in that section are sufficient to cover all of the examples appearing in this paper. In this appendix we use notation mirroring the method of Section 3.1, but the results apply equally to the extension of Section 3.5 if we replace $\tilde{\Theta}$ with the product manifold $\tilde{\Theta} \times \Lambda$ below.

Let Θ and $\tilde{\Theta}$ be differentiable manifolds of dimension i and j respectively, where $i \leq j$. Let $g : \tilde{\Theta} \rightarrow \Theta$ be a twice differentiable submersion on $\tilde{\theta}$, with $g(\tilde{\Theta}) = \Theta$. Let $f : \Theta \rightarrow \mathbb{R}$ be twice continuously differentiable, and define $\tilde{f} = f \circ g$.

Proposition 1. \tilde{f} has a local minimum at $\tilde{\theta}^*$ if and only if f has a local minimum at $\theta^* = g(\tilde{\theta}^*)$

Proof. First we prove the statement: f has a local minimum at $\theta^* = g(\tilde{\theta}^*) \Rightarrow \tilde{f}$ has a local minimum at $\tilde{\theta}^*$.

From the definition of a local minimum, there exists a neighbourhood of θ^* , \mathcal{V} such that

$$f(\theta^*) \leq f(\theta) \quad \forall \theta \in \mathcal{V}. \quad (25)$$

Because g is a continuous map, the preimage of \mathcal{V} , $\mathcal{U} = g^{-1}(\mathcal{V})$, is an open set which by construction contains $\tilde{\theta}^*$, and is therefore a neighbourhood of $\tilde{\theta}^*$. The result then follows

$$\begin{aligned} f(\theta^*) &\leq f(\theta) \quad \forall \theta \in \mathcal{V} \\ \Rightarrow f(g(\tilde{\theta}^*)) &\leq f(g(\tilde{\theta})) \quad \forall \tilde{\theta} \in \mathcal{U} \\ \Rightarrow \tilde{f}(\tilde{\theta}^*) &\leq \tilde{f}(\tilde{\theta}) \quad \forall \tilde{\theta} \in \mathcal{U} \end{aligned} \tag{26}$$

Finally, we prove the statement: \tilde{f} has a local minimum at $\tilde{\theta}^* \Rightarrow f$ has a local minimum at $\theta^* = g(\tilde{\theta}^*)$.

From the definition of a local minimum, there exists a neighbourhood of $\tilde{\theta}^*$, \mathcal{U} such that

$$\tilde{f}(\tilde{\theta}^*) \leq \tilde{f}(\tilde{\theta}) \quad \forall \tilde{\theta} \in \mathcal{U} \tag{27}$$

By assumption g is a submersion and therefore a continuous open map. Because \mathcal{U} is an open set containing $\tilde{\theta}^*$, $\mathcal{V} = g(\mathcal{U})$ must also be an open set containing θ^* . Finally, then

$$\begin{aligned} \tilde{f}(\tilde{\theta}^*) &\leq \tilde{f}(\tilde{\theta}) \quad \forall \tilde{\theta} \in \mathcal{U} \\ \Rightarrow f(g(\tilde{\theta}^*)) &\leq f(g(\tilde{\theta})) \quad \forall \tilde{\theta} \in \mathcal{U} \\ \Rightarrow f(\theta^*) &\leq f(\theta) \quad \forall \theta \in \mathcal{V} \end{aligned} \tag{28}$$

□

Proposition 2. For any local minimiser θ^* of f , $\exists \tilde{\theta}^* \in \tilde{\Theta}$ that is a local minimiser of \tilde{f} s.t. $\theta^* = g(\tilde{\theta}^*)$

Proof. $g(\tilde{\Theta}) = \Theta \Rightarrow \exists \tilde{\theta}^* \in \tilde{\Theta}$ s.t. $\theta^* = g(\tilde{\theta}^*)$. The result then follows from Proposition 1. □

Proposition 3. \tilde{f} has a local maximum at $\tilde{\theta}$ if and only if f has a local maximum at $\theta = g(\tilde{\theta})$

Proof. This follows from Proposition 1 by symmetric arguments. □

Proposition 4. \tilde{f} has a saddle point at $\tilde{\theta}$ if and only if f has a saddle point at $\theta = g(\tilde{\theta})$

Proof. First, we prove the following statement: \tilde{f} has a critical point at $\tilde{\theta} \Leftrightarrow f$ has a critical point at $\theta = g(\tilde{\theta})$.

Let $\tilde{\theta}$ and $\theta = g(\tilde{\theta})$ be represented as co-ordinates for some charts at those points, with f , g , \tilde{f} defined similarly. Then the statement about critical points can be expressed as follows

$$\nabla \tilde{f}(\tilde{\theta}) = \mathbf{0} \Leftrightarrow \nabla f(g(\tilde{\theta})) = \mathbf{0} \tag{29}$$

where $\mathbf{0}$ is a vector of zeros. To show that $\nabla \tilde{f}(\tilde{\theta}) = \mathbf{0} \Rightarrow \nabla f(g(\tilde{\theta})) = \mathbf{0}$, we have

$$\nabla \tilde{f}(\tilde{\theta}) = \nabla g(\tilde{\theta}) \nabla f(g(\tilde{\theta})) \tag{30}$$

where $\nabla g(\tilde{\theta})$ is the transposed Jacobian of g (see Appendix A for an explanation of this notation) and has full column rank due to g being a submersion, and so:

$$\nabla g(\tilde{\theta}) \nabla f(g(\tilde{\theta})) = \mathbf{0} \Rightarrow \nabla f(g(\tilde{\theta})) = \mathbf{0}. \tag{31}$$

For the other direction, $\nabla f(g(\tilde{\theta})) = \mathbf{0} \Rightarrow \nabla \tilde{f}(\tilde{\theta}) = \mathbf{0}$, trivially,

$$\begin{aligned} \nabla \tilde{f}(\tilde{\theta}) &= \nabla g(\tilde{\theta}) \nabla f(g(\tilde{\theta})) \\ &= \nabla g(\tilde{\theta}) \mathbf{0} \\ &= \mathbf{0}. \end{aligned} \tag{32}$$

Having proven correspondence of critical points, we can proceed to prove the proposition statement in two parts.

First, if \tilde{f} has a saddle point at $\tilde{\theta}$, then f must have a critical point at $\theta = g(\tilde{\theta})$. However, from Propositions 1 and 3, we know that this cannot be a local minimum or local maximum, and hence must be a saddle point.

Second, in the other direction, if f has a saddle point at $\theta = g(\tilde{\theta})$, then \tilde{f} has a critical point at $\tilde{\theta}$, which by similar reasoning must also be a saddle point. □

Proposition 5. *If \tilde{f} has a non-strict saddle point at $\tilde{\theta}$ then f has a non-strict saddle point at $\theta = g(\tilde{\theta})$*

Proof. From Proposition 4, we know that if \tilde{f} has a non-strict saddle point at $\tilde{\theta}$, then f must have a saddle point at $\theta = g(\tilde{\theta})$. It remains to be proven that the saddle point at θ must be a non-strict saddle point. We do this by contradiction. Let us assume that the saddle point of f at $\theta = g(\tilde{\theta})$ is strict.

Let $\tilde{\theta}, \theta$ be represented as co-ordinates for some charts at those points, with f, g, \tilde{f} defined similarly. Furthermore, let $H_f = \nabla^2 \tilde{f}(\tilde{\theta})$, $H_{\tilde{f}} = \nabla^2 f(g(\tilde{\theta}))$ be the Hessians of f and \tilde{f} , at $\tilde{\theta}$ and $\theta = g(\tilde{\theta})$, respectively. Then,

$$\begin{aligned} H_{\tilde{f}} &= \nabla g(\tilde{\theta}) H_f \nabla g(\tilde{\theta})^\top + \sum_{k=1}^j [\nabla f(g(\tilde{\theta}))]_k \nabla^2 g_k(\tilde{\theta}) \\ &= \nabla g(\tilde{\theta}) H_f \nabla g(\tilde{\theta})^\top \end{aligned} \quad (33)$$

where $\nabla^2 g_k(\tilde{\theta})$ is the Hessian of the k -th component of g . The second equality follows from θ being a critical point of f .

A strict saddle point is a saddle point for which there is at least one direction of strictly negative curvature, and so given the assumption that $\theta = g(\tilde{\theta})$ is a strict saddle point of f , $\exists v \in \mathbb{R}^i$ such that $v^\top H_f v < 0$. Let $\tilde{v} = A^\top v$, where A is any left inverse of $\nabla g(\tilde{\theta})$, then

$$\tilde{v}^\top H_{\tilde{f}} \tilde{v} = (A^\top v)^\top \nabla g(\tilde{\theta}) H_f \nabla g(\tilde{\theta})^\top (A^\top v) \quad (34)$$

$$\begin{aligned} &= v^\top H_f v \\ &< 0, \end{aligned} \quad (35)$$

implying that $\tilde{f}(\tilde{\theta})$ is a strict saddle point, a contradiction, and so we conclude that the saddle point at $f(\theta)$ cannot be strict. \square

C ALGORITHMS

C.1 SNGD WITH EF SURROGATES

In Algorithm 1, restated with line numbers below, we provide pseudocode for an implementation of SNGD when \tilde{q} is an EF distribution, and $\tilde{\theta}$ are either natural or mean parameters of that family. We assume the existence of an autodiff operator `grad`, which takes as input a real-valued function, and returns another function for computing its gradient. Note that when f cannot be computed deterministically, such as in VI or minibatch settings, we assume `grad` returns a function that provides unbiased stochastic estimates of the gradient. In our VI experiments, where gradients had to be taken through samples, we used the reparameterisation trick (Kingma and Ba, 2014), applied to the target distribution.

We also assume the existence of an overloaded function `dualparams`, which converts from mean to natural parameters of the EF, or vice-versa, depending on the type of its argument.⁶ That is, in the notation of Section 2.2, `dualparams` resolves to either $\mu(\cdot)$ or $\eta(\cdot)$ as appropriate. We note that each `dualparams` pair only needs to be defined once for each EF, and is not dependent on e.g. the target distribution or loss function, meaning that if these are supplied as part of a software library, the end user is only required to supply f, g and $\tilde{\theta}_0$. Algorithm 1 also assumes a given step size schedule, but can easily be extended to incorporate line search or other methods for choosing ϵ_t .

On line 1 of Algorithm 1 we define a reparameterisation of \tilde{f} in terms of the dual parameters. That is, if $\tilde{\theta}$ are natural parameters, \tilde{f}_{dual} is a function of mean parameters, and vice-versa. Line 4 computes the natural gradient, given by equation (5) or (6), using automatic differentiation of the function \tilde{f}_{dual} . Note that *both* overloads of `dualparams` are called: one inside the auto-differentiated function \tilde{f}_{dual} , and the other outside (to compute its argument). It is often possible for the inner conversion to be elided; for example, the user can supply a function g_{dual} (instead of g) that can perform the map from dual parameters to θ directly, more efficiently than the composition $g \circ \text{dualparams}$. For example, this is often the case when f depends on covariance-like parameters, and the composition $g \circ \text{dualparams}$ would otherwise involve inverting a matrix twice (a no-op).

⁶This is purely for convenience, as it allow us to describe a single implementation handling both parameterisations.

Algorithm 1 SNGD with EF surrogate

Require: objective $f : \Theta \rightarrow \mathbb{R}$
Require: parameter mapping $g : \tilde{\Theta} \rightarrow \Theta$
Require: initial surrogate parameters $\tilde{\theta}_0 \in \tilde{\Theta}$
Require: step size schedule $\{\epsilon_t \in \mathbb{R}_+ : t = 0, 1, \dots\}$

- 1: $\tilde{f}_{\text{dual}}(\cdot) := f(g(\text{dualparams}(\cdot)))$
- 2: $t \leftarrow 0$
- 3: **while** not converged **do**
- 4: $\tilde{\nabla} \leftarrow \text{grad}[\tilde{f}_{\text{dual}}](\text{dualparams}(\tilde{\theta}_t))$
- 5: $\tilde{\theta}_{t+1} \leftarrow \tilde{\theta}_t - \epsilon_t \tilde{\nabla}$
- 6: $t \leftarrow t + 1$
- 7: **end while**
- 8: **return** $g(\tilde{\theta}_t)$

If $\tilde{\theta}$ are mean parameters of \tilde{q} , then the computational overhead of Algorithm 1 (relative to GD in f) is approximately $\text{cost}[\eta] + 3 \times \text{cost}[g \circ \mu]$, where cost returns the cost of its function argument. Similarly, if θ are natural parameters, then the overhead will be approximately $\text{cost}[\mu] + 3 \times \text{cost}[g \circ \eta]$. The factor of 3 in the second term in each case results from taking gradients through $g \circ \text{dualparams}$; however, as discussed above, it is often the case the composition is almost zero cost, in which case this term can be largely eliminated by implementing g_{dual} directly.

We have assumed here that we can take gradients through either $\mu(\cdot)$ or $\eta(\cdot)$ efficiently, depending on the choice of parameterisation. For $\mu(\cdot)$ this is true by assumption for tractable families. For some families the ‘reverse’ map $\eta(\cdot)$, is not available in closed form, but can be efficiently computed using an iterative optimisation procedure (Minka, 2000, 2002). In such cases, we can use implicit differentiation techniques to efficiently compute gradients (Christianson, 1994; Blondel et al., 2022).

C.2 SNGD WITH EF SURROGATES AND AUXILIARY PARAMETERS

In Algorithm 2 we provide pseudocode for the extension of Section 3.5, in which we augment $\tilde{\theta}$ with auxiliary parameters λ . $\tilde{\theta}$ are optimised using natural gradients, whereas λ are optimised with standard first-order methods. Algorithm 2 uses GD with a fixed learning rate schedule for λ , but the extension to any first-order optimiser is straightforward.

The structure of Algorithm 2 is largely the same as that of Algorithm 1. One notable difference is that \tilde{f}_{dual} now has 2 arguments, and so the call to **grad** on line 4 returns a function that returns a 2-tuple of gradients, one for each argument.

Algorithm 2 SNGD with EF surrogate and auxiliary parameters

Require: objective $f : \Theta \rightarrow \mathbb{R}$
Require: parameter mapping $g : \tilde{\Theta} \rightarrow \Theta$
Require: initial surrogate parameters $\tilde{\theta}_0 \in \tilde{\Theta}$
Require: initial auxiliary parameters $\lambda_0 \in \Lambda$
Require: θ step size schedule $\{\epsilon_t \in \mathbb{R}_+ : t = 0, 1, \dots\}$
Require: λ step size schedule $\{\varepsilon_t \in \mathbb{R}_+ : t = 0, 1, \dots\}$

- 1: $\tilde{f}_{\text{dual}}(-1, -2) := f(g(\text{dualparams}(-1), -2))$
- 2: $t \leftarrow 0$
- 3: **while** not converged **do**
- 4: $(\tilde{\nabla}_{\theta}, \nabla_{\lambda}) \leftarrow \text{grad}[\tilde{f}_{\text{dual}}](\text{dualparams}(\theta_t), \lambda_t)$
- 5: $\theta_{t+1} \leftarrow \theta_t - \epsilon_t \tilde{\nabla}_{\theta}$
- 6: $\lambda_{t+1} \leftarrow \lambda_t - \varepsilon_t \nabla_{\lambda}$
- 7: $t \leftarrow t + 1$
- 8: **end while**
- 9: **return** $g(\theta_t, \lambda_t)$

D EXPERIMENT DETAILS

In this appendix we provide additional details about the experiments presented in the main paper. The code for these experiments is available at <https://github.com/cambridge-mlg/sngd>.

We repeated all experiments with 10 different random seeds. In all cases this led to different random parameter initialisations. For VI experiments, this also seeded randomness in the Monte Carlo samples, and for experiments with minibatching, it also seeded randomness in the minibatch sampling. The mean and standard errors as displayed in the training curves and Pareto frontier plots were computed over the 10 runs. For VI experiments we used the reparameterisation trick to estimate gradients for each method (Kingma and Welling, 2014).

In experiments that were small scale and had objectives that could be computed deterministically, namely MLE experiments using the sheep and COVID datasets, we used exact line search to determine step sizes for each of the methods being tested. This allowed us to compare the search direction of each method without confounding results with choice of hyperparameter settings.

For all other experiments we chose hyperparameters using a grid search. The training curves in the main paper correspond to the ‘best’ hyperparameter settings for each method. The best setting was considered to be that which had the best average (across time steps) worst case (over random seeds) value of the evaluation metric (negative elbo or negative log-likelihood as appropriate). Although this choice is somewhat arbitrary, we found that it consistently chose settings with training curves that closely resembled those that we considered best for each method. In Appendix F we provide Pareto frontier plots which incorporate *all* of the hyperparameter settings tried, which qualitatively are very similar to the training curves in the main paper.

For SNGD we chose $\tilde{\theta}$ to be mean parameters of \tilde{q} for all MLE tasks, and natural parameters for all VI tasks. These choices were motivated by the results of Appendix I, and we found them to consistently perform better than alternatives.

All experiments were executed on a 76-core Dell PowerEdge C6520 server, with 256GiB RAM, and dual Intel Xeon Platinum 8368Q (Ice Lake) 2.60GHz processors. Each individual optimisation run was locked to a single dedicated core. Implementations were written in JAX (Bradbury et al., 2018).

Next we provide details specific to each task featured in the experiments. In the list below, n denotes the number of training observations, and d denotes the dimensionality of the distribution that is being optimised in the task.

Sheep ($n=82$, $d=1$) Taken from a seminal work on the negative binomial distribution by Fisher (1941), this dataset consists of the number of ticks observed on each member of a population of sheep. The task for this dataset was MLE of the negative binomial distribution.

UCI miniboone ($n=32,840$, $d=43$) Taken from the MiniBooNE experiment at Fermilab, this dataset consists of a number of readings that can be used to classify observations as either electron or muon neutrinos (Roe, 2010).⁷ We follow the pre-processing of Papamakarios et al. (2017). Using this dataset we performed MLE of skew-normal and skew- t distributions. We used 32,840 of the observations for training, and the remaining 3,648 observations for evaluation. We used a minibatch size of 256 for each method.

UCI covertype ($n=500$, $d=53$) This dataset classifies the forest cover type of 581,024 pixels, based on 53 cartographic variables (Blackard, 1998).⁸ We used the ‘binary scale’ preprocessing of Chang and Lin (2011). The task was to perform VI in a Bayesian logistic regression model, with regularisation parameter 1.0. We found that using anything close to the full number of observations resulted in degenerate posteriors with virtually zero uncertainty, obviating the need for variational inference, and so we used a randomly chosen subset of 500 observations for our experiments. We used skew-normal, skew- t and skew-normal mixture distributions as approximate posteriors. All methods used 20 Monte Carlo samples for training and 1000 for evaluation.

Synthetic skew-normal ($n=10,000$, $d=1,000$) We generated synthetic data from a d -dimensional skew-normal distribution with parameters (ξ, Ω, η) , where ξ and η were drawn from $\mathcal{N}(\mathbf{0}, I_d)$. We chose $\Omega = d^{-1}W^\top W +$

⁷Licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.

⁸Licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.

$10^{-4}I_d$, where the components of $W \in \mathbb{R}^{d \times d}$ were drawn from independent standard normal distributions. The task was to perform MLE of a skew-normal distribution on this dataset.

Synthetic skew- t ($n=10,000$, $d=1,000$) We generated synthetic data from a d -dimensional skew- t distribution with parameters (ξ, Ω, η, ν) . ξ , Ω and η were generated in the same manner as the synthetic skew-normal task as detailed above, with ν chosen to be 10. The task was to perform MLE of a skew- t distribution on this dataset.

FTSE 100 stock returns ($n=1,515$, $d=93$) This dataset consists of daily stock price returns from 2017/01/01 to 2022/12/31 for the subset of FTSE 100 stocks that were members of the index during the entire period.⁹ We first fitted univariate Student’s t distributions to each dimension independently, and then transformed the data by converting each observed variable to its quantile value under the marginal distribution. The task was then MLE of a t copula on the quantile values.

COVID hospital admissions ($n=1,120$, $d=1$) This dataset consists of the number of daily COVID hospital admissions in the UK from 2020/4/1 to 2023/5/1.¹⁰ The task for this dataset was MLE of a 5-component negative binomial mixture.

Synthetic 2D logistic regression ($n=30$, $d=2$) This synthetic logistic regression dataset was generated using the same procedure as Murphy (2012). The task was VI in a Bayesian logistic regression model, with regularisation parameter 1. We used a skew-normal mixture approximate posterior, with 20 Monte Carlo samples for training and 1,000 for evaluation.

Now we provide details particular to the distributions being optimised in the tasks above. Note that the correspondence between tasks (above) and target distributions (below) is many to many: some target distributions were applied to more than one task, and some tasks were used for several target distributions. The parameter mappings used for SNGD are given in Table 2 in Appendix E; we do not repeat them here unless additional explanation is required. When the parameter mappings make use of the auxiliary parameter extension of Section 3.5, we used Adam to optimise λ .

With our baseline methods, when a target distribution required a positive definite covariance matrix parameter, we tried two different parameterisations; covariance square-root (e.g. $\Sigma = W^T W$) and precision square-root (e.g. $\Sigma^{-1} = W^T W$) (Salimbeni et al., 2018). This parameterisation choice was determined by a hyperparameter which we included in our grid search. Further parameterisation details are given below.

Negative binomial The PMF of the negative binomial with parameters $\theta = (r, s)$ is given by (15). We initialised negative binomial parameters for all methods by drawing $s \sim \text{Uniform}(0.05, 0.95)$, $r \sim \text{Gamma}(6.25, 1.25)$. With GD, BFGS and NGD we used a log parameterisation of r , and a logit parameterisation of s . For SNGD we chose $\tilde{\theta}$ to be mean parameters of \tilde{q} . In order to ensure that $g(\tilde{\theta}) \in \Theta \forall \tilde{\theta} \in \tilde{\Theta}$, we require $\beta < 1 \forall \tilde{\theta} \in \tilde{\Theta}$, therefore we chose $\tilde{\Theta}$ as the subset of \mathcal{M} (the mean domain of \tilde{q}) for which $\beta < 1$. It can be shown that this remains an open convex set.

Skew-normal The PDF of the multivariate skew-normal with parameters $\theta = (\xi, \Omega, \eta)$ is given by $q_\theta(x) = 2\mathcal{N}_d(x; \xi, \Omega)\Phi(\eta^T(x - \xi))$, where $\mathcal{N}_d(\cdot)$ is the PDF of the d -dimensional normal distribution, and $\Phi(\cdot)$ is the CDF of the standard normal distribution.¹¹ We used random initialisations of $\xi \sim \mathcal{N}(\mathbf{0}, 0.01^2 I_d)$, $\eta \sim \mathcal{N}(\mathbf{0}, 0.01^2 I_d)$, with Ω initialised to I_d .

Skew- t The PDF of the multivariate skew- t with parameters $\theta = (\xi, \Omega, \eta, \nu)$ is given by $q_\theta(x) = 2\mathcal{T}_d(x; \xi, \Omega, \nu)\Psi(\alpha^T x \sqrt{(\nu + d)/(\nu + x^T \bar{\Omega}^{-1} x)}; \nu + d)$, where \mathcal{T}_d is the PDF of the d -dimensional Student’s t , $\Psi(\cdot)$ is the CDF of the (univariate) Student’s t , $\omega = (\Omega \odot I_d)^{1/2}$, $\bar{\Omega} = \omega^{-1} \Omega \omega^{-1}$, and $\alpha = \omega \eta$.¹² We used random initialisations of $\xi \sim \mathcal{N}(\mathbf{0}, 0.01^2 I_d)$, $\eta \sim \mathcal{N}(\mathbf{0}, 0.01^2 I_d)$, with Ω and ν initialised to I_d and 50 respectively. For all methods we parameterised ν as $\log(\nu - 2)$.

⁹Downloaded from the Bloomberg Terminal.

¹⁰Downloaded from <https://coronavirus.data.gov.uk> and licensed under the Open Government License v3.0.

¹¹Equation (5.1) of Azzalini (2013).

¹²Equation (6.24) of Azzalini (2013).

t copula The d -dimensional t copula with parameters $\theta = (R, \nu)$ has PDF $q_\theta(x) = \mathcal{T}_d(z; 0, R) / \prod_i \mathcal{T}(z_i; 0, 1)$, where \mathcal{T} is the PDF of the univariate Student’s t . We initialised R to $\text{corr}(I + W^\top W)$, and ν to 50, where $W_{ij} \sim \mathcal{N}(0, .01^2)$, and $\text{corr}(\Sigma) = (\Sigma \odot I_d)^{-\frac{1}{2}} \Sigma (\Sigma \odot I_d)^{-\frac{1}{2}}$ projects a covariance matrix to its implied correlation matrix. For all methods we parameterised ν as $\log(\nu - 2)$.

Negative binomial mixture The negative binomial mixture with k components and parameters $\theta = (\pi_i, r_i, s_i)_{i=1}^k$ has PMF given by (17), where $q_{\theta_i}(x)$ is the negative binomial PMF given by (15) with parameters $\theta_i = (r_i, s_i)$. We initialised π_i to k^{-1} for $i = 1, \dots, k$, and the component negative binomials were initialised to have mean m_i and variance v_i , with $m_i \sim \text{Uniform}(0.1x_{\max}, 0.9x_{\max})$ and $v_i = m_i o_i$ where $o_i^{-1} \sim \text{Uniform}(0.001, 0.02)$ and x_{\max} is the maximum value observed in the training data. With GD and BFGS, we used a softmax parameterisation of the mixture probabilities, a log parameterisation of r_i , and a logit parameterisation of s_i .

Skew-normal mixture The skew-normal mixture with k components and parameters $\theta = (\pi_i, \xi_i, \Omega_i, \nu_i)_{i=1}^k$ has PDF given by (17), where $q_{\theta_i}(x)$ is the skew-normal PDF with parameters $\theta_i = (\xi_i, \Omega_i, \nu_i)$. We initialised π_i to k^{-1} for $i = 1, \dots, k$, with the remaining parameters initialised equivalently to the (non-mixture) skew-normal case. With Adam, we used a softmax parameterisation of the mixture probabilities.

E PARAMETER MAPPINGS

In Table 2 we provide parameter mappings for all of the examples appearing in this paper. For convenience, we express the mappings in terms of the standard parameterisation of \tilde{q} . For the method described in Section 3.1, $\tilde{\theta}$ would be the mean or natural parameters corresponding to the parameters stated here, with g similarly adjusted. Note that in Table 2 we follow convention by using μ to denote the mean of a normal distribution, whereas in the main paper it refers to the mean *parameters* (expected sufficient statistics) of an EF distribution. See Appendix D for target distribution definitions.

Table 2: Example surrogate-target parameter mappings

TARGET	SURROGATE	θ	$\tilde{\theta}$	λ	$g(\tilde{\theta}, \lambda)$
Neg. bin.	Gamma	r, s	α, β		$\alpha / (1 - \beta^{-1}), \beta^{-1}$
Neg. bin. mix.	Gamma mix. model	(π_i, r_i, s_i)	$(\pi_i, \alpha_i, \beta_i)$		$(\pi_i, \alpha_i / (1 - \beta_i^{-1}), \beta_i^{-1})$
Skew-normal	Normal	ξ, Ω, η	μ, Σ	η	μ, Σ, η
Skew-normal mix.	Normal mix. model	$(\pi_i, \xi_i, \Omega_i, \eta_i)$	(π_i, μ_i, Σ_i)	(η_i)	$(\pi_i, \mu_i, \Sigma_i, \eta_i)$
Skew- t	Normal	ξ, Ω, η, ν	μ, Σ	$\eta, \log(\nu - 2)$	μ, Σ, η, ν
Elliptical copula	Zero-mean normal	R, \dots	Σ	\dots	$\text{corr}(\Sigma), \dots$ ¹³

F ADDITIONAL RESULTS

In this appendix we provide additional results for the experiments presented in the main paper. In several of the experiments (those not using line-search), the competing methods were dependent on hyperparameters which were chosen by grid search. For the training curves in the main paper, we used the heuristic method described in Appendix D in order to choose the ‘best’ settings for each method.

In order to provide a more complete comparison, in Figures 9 through 16 we present Pareto frontier plots that show the best (average) value of the chosen metric (y axis) attained by *any* learning rate setting for each method, as a function of both iteration count and wall-clock time (x axis). The displayed error bars correspond to 2 standard errors, computed for the optimal setting at the corresponding point in time.

¹³ $\text{corr}(\cdot)$ projects a covariance matrix to its implied correlation matrix. See Appendix D for its definition.




Figure 9: Pareto frontiers for Bayesian logistic regression VI on UCI covtype ($n=500$, $d=53$), with a skew-normal approximation.




Figure 10: Pareto frontiers for Bayesian logistic regression VI on UCI covtype ($n=500$, $d=53$), with a skew- t approximation.




Figure 11: Pareto frontiers for skew-normal MLE on UCI miniboone ($n=32,840$, $d=43$).




Figure 12: Pareto frontiers for skew- t MLE on UCI miniboone ($n=32,840$, $d=43$).




Figure 13: Pareto frontiers for skew-normal MLE on synthetic dataset ($n=10,000$, $d=1,000$).




Figure 14: Pareto frontiers for skew- t MLE on synthetic dataset ($n=10,000$, $d=1,000$).




Figure 15: Pareto frontiers for t -copula MLE on 5 years of stock return data ($n=1,515$, $d=93$).




Figure 16: Pareto frontiers for Bayesian logistic regression VI on UCI covtype ($n=500$, $d=53$), with a skew-normal mixture approximation.




Figure 17: Illustration of SNGD. An infinitesimal KL ball (ellipse, in parameter space) around $\tilde{\theta}$ with respect to: \tilde{q} (shaded), and q (dashed). Contours are of $\tilde{f}(\tilde{\theta})$. δ^* points to the minimum point in the shaded region; this is the SNGD direction. In general this will differ from the direction of NGD in q , which would instead point in the direction of the minimum point in the dashed region (not marked).

G CHOOSING EFFECTIVE SURROGATES

In order to apply SNGD to a given target distribution q , we must choose a surrogate \tilde{q} , and reparameterisation $\theta = g(\tilde{\theta})$. Appropriate choices here determine the effectiveness of SNGD.

One guiding principle that can be used is based on a view of NGD involving Kullback-Leibler (KL) divergences. The KL divergence from continuous¹⁴ distribution q to p is defined as

$$\text{KL}[q \parallel p] = \int q(x) \log \frac{q(x)}{p(x)} dx. \quad (36)$$

The NGD direction can be motivated by a result due to Ollivier et al. (2017), restated here in our notation.

Proposition 6. *Let f be a smooth function on the parameter space Θ . Let $\theta \in \Theta$ be a point where $\tilde{\nabla}f(\theta) = [F(\theta)]^{-1}\nabla f(\theta)$ does not vanish. Then, if*

$$\delta = -\frac{\tilde{\nabla}f(\theta)}{\|\tilde{\nabla}f(\theta)\|_F} \quad (37)$$

is the negative direction of the natural gradient of f (with $\|\cdot\|_F$ the Fisher norm), we have

$$\delta = \lim_{\tau \rightarrow 0^+} \frac{1}{\sqrt{2\tau}} \arg \min_{\delta' \in B(\tau)} f(\theta + \delta') \quad (38)$$

where $B(\tau) = \{\delta' : \text{KL}[q_{\theta+\delta'} \parallel q_\theta] < \tau\}$.

In words, this says that NGD moves in the direction in parameter space that gives the greatest decrease in the objective within an infinitesimally small KL-ball around the current point.

This suggests that we would ideally like to find \tilde{q} and g such that the KL divergence between the surrogate $\tilde{q}_{\tilde{\theta}}$ and a perturbed version $\tilde{q}_{\tilde{\theta}+\delta}$ is identical to the KL divergence from the target $q_{g(\tilde{\theta})}$ and its corresponding perturbation $q_{g(\tilde{\theta}+\delta)}$, for any small perturbation δ . However, because this is only required to hold in some infinitesimally small neighbourhood of $\tilde{\theta}$, it is sufficient to satisfy

$$\nabla_\delta (\text{KL}[q_{\tilde{\theta}+\delta} \parallel \tilde{q}_{\tilde{\theta}}])|_{\mathbf{0}} = \nabla_\delta (\text{KL}[q_{g(\tilde{\theta}+\delta)} \parallel q_{g(\tilde{\theta})}])|_{\mathbf{0}}. \quad (39)$$

If equality (39) holds then SNGD will move in the same direction in $\tilde{\theta}$ as ‘exact’ NGD under (reparameterised) q . While it will not typically be possible to find tractable surrogates for which equality holds exactly, this

¹⁴The discrete KL divergence is defined similarly, with summation replacing integration.

observation motivates choosing \tilde{q} , g for which it is approximately true. We illustrate this in Figure 17. This perspective also highlights the reason that q and \tilde{q} do not need to be distributions over the same space; all that matters is that the effect of (small) changes in $\tilde{\theta}$ on \tilde{q} and q is similar in a KL-divergence sense. If approximating NGD under q is our goal, then this perspective can help guide us toward a choice of \tilde{q} and g .¹⁵

As a concrete example, let q be a mixture distribution, with PDF given by

$$q_{\theta}(x) = \sum_{i=1}^k \pi_i q_{\theta_i}(x), \quad (40)$$

where q_{θ_i} are the component distributions, $\theta = (\pi_i, \theta_i)_{i=1}^k$, and $\pi \in \Delta^{k-1}$. Let us consider as a surrogate for q_{θ} the corresponding mixture *model* (joint distribution) with PDF

$$\tilde{q}_{\tilde{\theta}}(z, x) = \pi_z q_{\theta_z}(x), \quad (41)$$

where $\tilde{\theta} = (\pi_i, \theta_i)_{i=1}^k$, so that g is simply the identity map, and $\sum_{i=1}^k \tilde{q}_{\tilde{\theta}}(z = i, x) = q_{\tilde{\theta}}(x) = q_{g(\tilde{\theta})}(x)$. Using standard identities, we have

$$\begin{aligned} \nabla_{\delta}(\text{KL}[\tilde{q}_{\tilde{\theta}+\delta}(z, x) \parallel \tilde{q}_{\tilde{\theta}}(z, x)]) &= \nabla_{\delta} \left(\text{KL}[\tilde{q}_{\tilde{\theta}+\delta}(x) \parallel \tilde{q}_{\tilde{\theta}}(x)] + \mathbb{E}_{\tilde{q}_{\tilde{\theta}+\delta}(x)}[\text{KL}[\tilde{q}_{\tilde{\theta}+\delta}(z|x) \parallel \tilde{q}_{\tilde{\theta}}(z|x)]] \right) \\ &= \nabla_{\delta} \left(\text{KL}[q_{g(\tilde{\theta}+\delta)}(x) \parallel q_{g(\tilde{\theta})}(x)] + \mathbb{E}_{q_{g(\tilde{\theta}+\delta)}(x)}[\text{KL}[\tilde{q}_{\tilde{\theta}+\delta}(z|x) \parallel \tilde{q}_{\tilde{\theta}}(z|x)]] \right) \end{aligned} \quad (42)$$

$$\approx \nabla_{\delta} \left(\text{KL}[q_{g(\tilde{\theta}+\delta)}(x) \parallel q_{g(\tilde{\theta})}(x)] \right) \quad (43)$$

It follows that using a mixture model joint, as a surrogate for its marginal, can be viewed as using an approximate metric which includes an additional term. This term imposes an additional penalty on directions in parameter space that affect the expected *responsibility* distributions.¹⁶ In the experiments of Section 4.4 we take this approximation one step further, and use EF mixture models as surrogates for mixtures with components that are *not* EF distributions.

As an aside, we note that Lin et al. (2019) performed NGD with respect to a joint distribution in a VI objective for which the optimisation target was a *marginal* of the NGD distribution. This mismatch was not explicitly discussed by Lin et al. (2019). However, it is clear that this aspect of their method can be interpreted as an application of SNGD, and similar reasoning to that above shows that it implies making a specific approximation to the Fisher metric.

H COMPARISON WITH NGD UNDER q

In Appendix G we provided a motivation for SNGD that viewed it as approximating NGD under a reparameterisation of q . However, for the negative binomial MLE experiment of Section 4.1, we demonstrated that SNGD was actually able to *outperform* NGD under q . In Figure 18, we show that this was true when NGD was performed with respect to both a standard parameterisation (also shown in Figure 1), and the reparameterisation defined by $g(\tilde{\theta})$. This implies that the outperformance of SNGD is not simply an artefact of the reparameterisation. This surprising result deserves extra scrutiny.

We begin by highlighting that in fact, SNGD *is* NGD, but in the objective $\tilde{f}(\tilde{\theta}) = f(g(\tilde{\theta}))$ and with respect to distribution \tilde{q} . It is helpful therefore to consider the desirable properties of NGD:

1. It is locally invariant to parameterisation of the distribution being optimised.
2. It follows the direction of steepest descent in the objective on the statistical manifold of the distribution being optimised.¹⁷
3. For MLE objectives, it asymptotically approaches Newton’s method near the optimum.

¹⁵Approximating NGD under q should not *necessarily* be our ultimate goal. We elaborate on this in Appendix H.

¹⁶The responsibility of component i for x is the probability that x was generated by component i having observed x .

¹⁷Where steepness is defined with respect to a KL divergence.


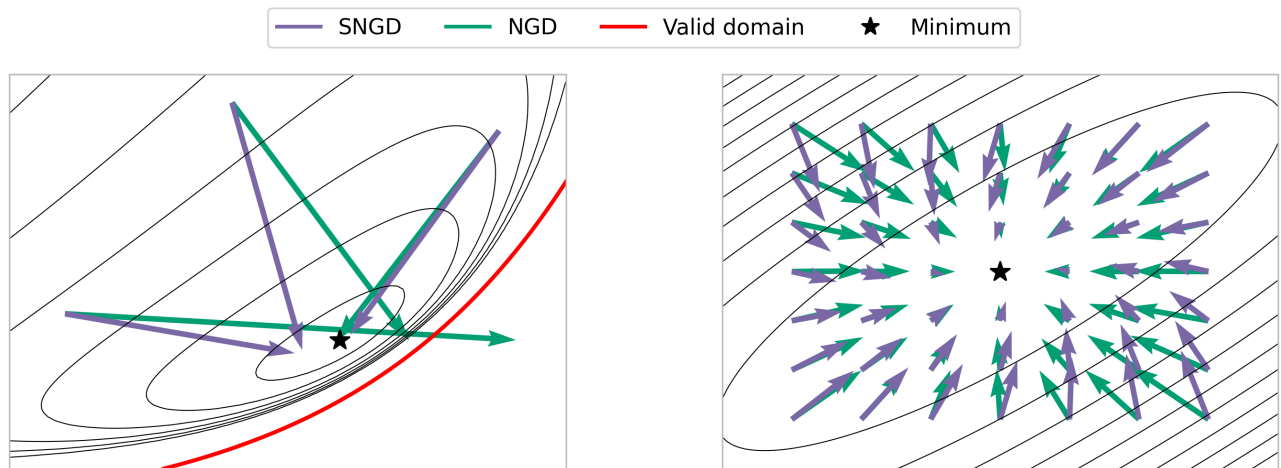


Figure 18: Training curves for negative binomial MLE on the sheep dataset ($n=82$, $d=1$), using surrogate natural gradient descent (SNGD) and natural gradient descent under q with respect to both a standard parameterisation (NGD), and the reparameterisation defined by g (NGD-g).

Properties 1 and 2 apply always, and so they are necessarily inherited by SNGD. However, when f is a MLE objective for q (as in the experiment of Section 4.1), \tilde{f} is *not*, in general, an MLE objective for \tilde{q} , and so SNGD loses property 3. Contrast this with NGD in \tilde{f} with respect to q : although the objective function is the same as that of SNGD, it is simply a reparameterised MLE objective for q , and so property 3 is retained. In other words, the asymptotic efficiency of NGD is retained under change of parameterisation, but not under change of distribution.



(a) Vectors represent a full (undamped) step under each method. When far from the optimum, a full step of SNGD more consistently moves to a low value of the objective, compared to that of NGD under q .

(b) The optimum magnified $\times 1000$. At this scale the loss is well approximated by a quadratic. The Fisher of q , $F(\tilde{\theta})$, approaches the Hessian. This is not true of $\tilde{F}(\tilde{\theta})$, because the loss is not a MLE objective for \tilde{q} . Here the gradient vector lengths have been scaled down by a factor of 3 for clarity.

Figure 19: Comparison of SNGD and NGD in the sheep experiment of Section 4.1. Contours are of $\tilde{f}(\tilde{\theta})$, an MLE objective for q .

On the face of it, then, when f is a MLE objective for q , SNGD appears less desirable than NGD under q , inheriting only two of the properties listed above. However, none of these properties say anything about performance over large (or indeed, any *non-infinitesimal*) steps in parameter space. But there is one special case in which we can say something about the performance of NGD over large steps in parameter space: when NGD is applied to a MLE objective of an EF distribution with respect to the mean parameters of that distribution,

a single undamped step will converge straight to the optimum (see Proposition 7 in Appendix I).

When we apply SNGD with EF \tilde{q} , and with respect to mean parameters $\tilde{\theta}$, we meet 2 of the criteria for single-step convergence. Although $\tilde{f}(\tilde{\theta})$ is not in general a MLE for \tilde{q} , if it behaves *approximately* like a MLE objective for \tilde{q} , then SNGD can make rapid progress over large distances in parameter space, and *approximately* converge in a single step. And so, while SNGD loses the asymptotic efficiency of NGD under q , it is possible to obtain improved performance in the early to mid stages. This is exactly what we observe in the experiment of Section 4.1, as is further illustrated in Figure 19.

The ability to make rapid progress early on is arguably more important for the overall performance of an optimiser in practice; however, if convergence to an exact (within machine precision range) optimum is required, it may be beneficial to switch to Newton-type methods during the late phase, an approach previously proposed by Tatzel et al. (2022).

We now attempt to characterise what it means for \tilde{f} to behave *approximately* like a MLE for \tilde{q} . Let \tilde{q} be an EF with mean parameters $\tilde{\theta}$. Furthermore, let $f(\tilde{\theta})$ be an MLE objective for q so that

$$\begin{aligned}\tilde{f}(\tilde{\theta}) &= f(g(\tilde{\theta})) \\ &= -\mathbb{E}_{p_{\mathcal{D}}(x)}[\log q_{g(\tilde{\theta})}(x)]\end{aligned}\quad (44)$$

where $p_{\mathcal{D}}(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - x_i)$ is the empirical density function. Using a trivial identity, we can rewrite this as

$$\begin{aligned}\tilde{f}(\tilde{\theta}) &= \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}] - \mathbb{E}_{p_{\mathcal{D}}(x)}[\log q_{g(\tilde{\theta})}(x)] - \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}] \\ &= \tilde{f}^*(\tilde{\theta}) + h(\tilde{\theta})\end{aligned}\quad (45)$$

where $\tilde{\theta}^*$ are optimal parameters of \tilde{f} , and we have defined


$$\tilde{f}^*(\tilde{\theta}) = \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}],\quad (46)$$

and

$$h(\tilde{\theta}) = -\mathbb{E}_{p_{\mathcal{D}}(x)}[\log q_{g(\tilde{\theta})}(x)] - \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}].\quad (47)$$

$\tilde{f}^*(\tilde{\theta})$ can be minimised by a single NGD step with respect to \tilde{q} and $\tilde{\theta}$ (see Proposition 7). The second term, $h(\tilde{\theta})$, can then be seen as distorting the natural gradient step, acting to move it away from the optimum $\tilde{\theta}^*$. If $h(\tilde{\theta})$ is (approximately) constant with respect to $\tilde{\theta}$, then SNGD will (approximately) converge in a single step.

In Figure 20 we show that in the sheep experiment of Section 4.1, the effect of the distortion term is indeed small, even when far from the optimum, illustrating why SNGD is able to make rapid progress on this problem.



(a) Contours of $\tilde{f}^*(\tilde{\theta})$. Vectors show $[\tilde{F}(\tilde{\theta})]^{-1} \nabla \tilde{f}^*(\tilde{\theta})$ at 3 distinct locations; these always move straight to the optimum of \tilde{f} .

(b) Contours of the distortion term $h(\tilde{\theta})$. Vectors show $[\tilde{F}(\tilde{\theta})]^{-1} \nabla h(\tilde{\theta})$ at a regular grid of points

(c) Contours of $\tilde{f}(\tilde{\theta})$. Vectors now show the resulting SNGD step, $[\tilde{F}(\tilde{\theta})]^{-1} \nabla \tilde{f}(\tilde{\theta})$. This is the sum of the previous two vector fields.

Figure 20: The objective, and SNGD steps for the sheep experiment of Section 4.1, each decomposed according to equation (45). The star marks the optimum of \tilde{f} .

To add some intuition behind $h(\tilde{\theta})$, let us assume that the data were generated from the model distribution q with parameters θ^* . Furthermore, take the infinite data limit,¹⁸ so that

$$\tilde{f}(\tilde{\theta}) = -\mathbb{E}_{q_{\theta^*}(x)}[\log q_{g(\tilde{\theta})}(x)] \quad (48)$$

and $\theta^* = g(\tilde{\theta}^*)$. The distortion term $h(\tilde{\theta})$ can then be written as

$$\begin{aligned} h(\tilde{\theta}) &= -\mathbb{E}_{q_{\theta^*}(x)}[\log q_{g(\tilde{\theta})}(x)] - \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}] \\ &= \text{KL}[q_{g(\tilde{\theta}^*)} \parallel q_{g(\tilde{\theta})}] - \text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}] + c \end{aligned} \quad (49)$$

where c is constant with respect to $\tilde{\theta}$. For $h(\tilde{\theta})$ to be roughly constant, we then require

$$\nabla_{\tilde{\theta}}(\text{KL}[\tilde{q}_{\tilde{\theta}^*} \parallel \tilde{q}_{\tilde{\theta}}]) \approx \nabla_{\tilde{\theta}}(\text{KL}[q_{g(\tilde{\theta}^*)} \parallel q_{g(\tilde{\theta})}]). \quad (50)$$

This is similar to the statement expressed by equation (39), concerning the effect of $\tilde{\theta}$ on *reverse* KL divergences between nearby points, whereas now we have a statement about the effect of $\tilde{\theta}$ on *forward* KL divergences from the (potentially distant) optimal point.

I EF NATURAL GRADIENTS

When NGD is applied to objectives involving a KL divergence (or related quantity), it has special properties when it is performed with respect to either the natural or mean parameters of an EF distribution, as the following propositions show.

Proposition 7. *For EF q with mean parameters $\mu \in \mathcal{M}$, and mean domain \mathcal{M} , let*

$$f(\mu) = -\mathbb{E}_{p^*(x)}[\log q_{\mu}(x)] + c \quad (51)$$

where p^* is any distribution with $\text{supp}[p^*] \subseteq \text{supp}[q]$, and c is constant with respect to μ . Then, $\forall \mu \in \mathcal{M}$,

$$\mu - [F(\mu)]^{-1} \nabla f(\mu) = \arg \min_{\mu' \in \mathcal{M}} f(\mu'). \quad (52)$$

Proof. First, let η be dually coupled with μ , so that $\mu = \mathbb{E}_{q_{\eta}(x)}[t(x)]$. From (6), we have that

$$\begin{aligned} \mu - [F(\mu)]^{-1} \nabla f(\mu) &= \mu + \nabla_{\eta}(\mathbb{E}_{p^*(x)}[\log q_{\eta}(x)]) \\ &= \mu + \nabla_{\eta}(\mathbb{E}_{p^*(x)}[t(x)^{\top} \eta - A(\eta) + \log \nu(x)]) \\ &= \mu - \nabla A(\eta) + \mathbb{E}_{p^*(x)}[t(x)] \\ &= \mathbb{E}_{p^*(x)}[t(x)] \end{aligned} \quad (53)$$

where we have used the standard EF identity $\nabla A(\eta) = \mu(\eta)$. It remains to show that $\mathbb{E}_{p^*(x)}[t(x)] = \arg \min_{\mu' \in \mathcal{M}} f(\mu')$. Let us reparameterise $f(\mu)$ as $f_{\eta}(\eta) = f(\mu(\eta))$, then

$$\begin{aligned} \nabla f_{\eta}(\eta) &= -\nabla_{\eta}(\mathbb{E}_{p^*(x)}[t(x)^{\top} \eta - A(\eta) + \log \nu(x)]) \\ &= \mu - \mathbb{E}_{p^*(x)}[t(x)], \end{aligned} \quad (54)$$

and equating to zero, we find the unique stationary point at $\mu = \mathbb{E}_{p^*(x)}[t(x)]$. To show that this is a minimum, note that

$$\nabla^2 f_{\eta}(\eta) = \nabla^2 A(\eta). \quad (55)$$

Given that A is strictly convex and twice-differentiable (see Wainwright and Jordan (2008)), then f_{η} is also strictly convex, and so must be minimised when $\mu = \mathbb{E}_{p^*(x)}[t(x)]$. \square

¹⁸We do this for convenience; for finite samples, the resulting expressions hold in expectation (over random data sets).

Proposition (7) says that (undamped) NGD with respect to the mean parameters of EF q will converge in a single step when the objective is a cross entropy from *any* distribution p^* to q_μ plus a constant that does not depend on μ . Two straightforward corollaries of this are that NGD also has single step convergence when $f(\mu) = \text{KL}[p^*(x) \parallel q_\mu(x)]$ (a forward KL divergence), or when $f(\mu) = -\sum_{i=1}^n \log q_\mu(x_i)$ (a MLE objective).

Typically, when performing MLE of an EF distribution, we simply think of it as ‘solving’ for the parameters, or ‘moment matching’; Proposition (7) shows that actually, this is equivalent to performing (undamped) NGD with respect to μ .

There is an *almost*-symmetry regarding NGD with respect to the natural parameters of q , as the following proposition states.

Proposition 8. *For EF q with natural parameters $\eta \in \mathcal{O}$, and natural domain \mathcal{O} , let*

$$f(\eta) = \text{KL}[q_\eta(x) \parallel q_{\eta^*}(x)] + c \quad (56)$$

where q_{η^*} is from the same EF as q with natural parameter η^* . Then, $\forall \eta \in \mathcal{O}$,

$$\eta - [F(\eta)]^{-1} \nabla f(\eta) = \arg \min_{\eta' \in \mathcal{O}} f(\eta'). \quad (57)$$

Proof. First, let μ be dually coupled with η , so that $\mu = \mathbb{E}_{q_\eta(x)}[t(x)]$. From (5), we have that

$$\begin{aligned} \eta - [F(\eta)]^{-1} \nabla f(\eta) &= \eta - \nabla_\mu (\text{KL}[q_\mu(x) \parallel q_{\eta^*}(x)]) \\ &= \eta - \nabla_\mu (A^*(\mu) + A(\eta^*) - \mu^\top \eta^*) \\ &= \eta - \nabla A^*(\mu) + \eta^* \\ &= \eta^* \end{aligned} \quad (58)$$

where: q_μ denotes the EF distribution with mean parameter μ ; A^* , the convex dual of A , is the negative entropy of q ; and we have used the identity $\nabla A^*(\mu) = \eta$. It is clear from the properties of the KL divergence that η^* must be the unique minimiser for $f(\eta)$. \square

Note that Proposition 8 requires q_{η^*} to be from the same family as q_η , which is not the case for the corresponding result with respect to mean parameters. This is a generalisation of a result found in Hensman et al. (2013); Salimbeni et al. (2018) for sparse Gaussian process VI. It is similar to to a result given by Sato (2001) for variational Bayes. Note that an analogous result to Proposition 8 can be derived for any Bregman divergence in the same manner.

J EF MIXTURE MODELS

An EF mixture with k components has density

$$q(x) = \sum_{i=1}^k \pi_i \nu(x) \exp(t(x)^\top \eta_i - A(\eta_i)) \quad (59)$$

where $\pi_i > 0$ are the mixture probabilities with $\pi \in \Delta^{k-1}$, and η_i are the natural parameters of mixture component i . We can also define a related *mixture model*, a joint distribution in which we consider the mixture identity $z \in \{1, \dots, k\}$ as a latent variable, with density

$$q(z, x) = \pi_z \nu(x) \exp(t(x)^\top \eta_z - A(\eta_z)). \quad (60)$$

A finite EF mixture is not, in general, an EF. However, the corresponding mixture model (joint distribution) *is* an EF. Consider its density,

$$\begin{aligned}
 q(z, x) &= \pi_z \nu(x) \exp(t(x)^\top \eta_z - A(\eta_z)) \\
 &= \sum_{i=1}^k \mathbb{I}_i(z) \pi_i \nu(x) \exp(t(x)^\top \eta_i - A(\eta_i)) \\
 &= \nu(x) \prod_{i=1}^k [\pi_i \exp(t(x)^\top \eta_i - A(\eta_i))]^{\mathbb{I}_i(z)} \\
 &= \nu(x) \exp\left(\sum_{i=1}^k \mathbb{I}_i(z) \log \pi_i + \mathbb{I}_i(z) t(x)^\top \eta_i - \mathbb{I}_i(z) A(\eta_i)\right) \\
 &= \nu(x) \exp\left[\left(\sum_{i=1}^{k-1} \mathbb{I}_i(z) \left(\log \frac{\pi_i}{\pi_k} - A(\eta_i) + A(\eta_k)\right)\right) + \left(\sum_{i=1}^k \mathbb{I}_i(z) t(x)^\top \eta_i\right) + \log \pi_k - A(\eta_k)\right],
 \end{aligned} \tag{61}$$

where $\mathbb{I}_i(z)$ is an indicator function for the singleton set $\{i\}$. We can then recognise (61) as having the form of an EF density with: base measure $\nu(\cdot)$; log-partition $A(\eta_k) - \log \pi_k$; and sufficient statistics and natural parameters as listed in Table 3. Note that the statistic functions are linearly independent, and so the EF defined by (60) is minimal.

Table 3: Sufficient statistics and natural parameters for an EF mixture model. Note that t is a vector-valued function, and each η_i is a vector of equal dimension.

Sufficient statistic	Natural parameter
$\mathbb{I}_1(z)$	$\log \frac{\pi_1}{\pi_k} - A(\eta_1) + A(\eta_k)$
...	...
$\mathbb{I}_{k-1}(z)$	$\log \frac{\pi_1}{\pi_k} - A(\eta_{k-1}) + A(\eta_k)$
$\mathbb{I}_1(z)t(x)$	η_1
...	...
$\mathbb{I}_k(z)t(x)$	η_k