



Variational Bayes on manifolds

Minh-Ngoc Tran¹ · Dang H. Nguyen² · Duy Nguyen³

Received: 17 December 2020 / Accepted: 22 August 2021 / Published online: 14 September 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

Variational Bayes (VB) has become a widely-used tool for Bayesian inference in statistics and machine learning. Nonetheless, the development of the existing VB algorithms is so far generally restricted to the case where the variational parameter space is Euclidean, which hinders the potential broad application of VB methods. This paper extends the scope of VB to the case where the variational parameter space is a Riemannian manifold. We develop an efficient manifold-based VB algorithm that exploits both the geometric structure of the constraint parameter space and the information geometry of the manifold of VB approximating probability distributions. Our algorithm is provably convergent and achieves a convergence rate of order $\mathcal{O}(1/\sqrt{T})$ and $\mathcal{O}(1/T^{2-2\epsilon})$ for a non-convex evidence lower bound function and a strongly retraction-convex evidence lower bound function, respectively. We develop in particular two manifold VB algorithms, Manifold Gaussian VB and Manifold Wishart VB, and demonstrate through numerical experiments that the proposed algorithms are stable, less sensitive to initialization and compares favourably to existing VB methods.

Keywords Marginal likelihood · Variational Bayes · Natural gradient · Stochastic approximation · Riemannian manifold

1 Introduction

Increasingly complicated models in modern statistics and machine learning have called for more efficient Bayesian estimation methods. Of the Bayesian tools, Variational Bayes (VB) (Waterhouse et al. 1996; Jordan et al. 1999) stands out as one of the most versatile alternatives to conventional Monte Carlo methods for statistical inference in complicated models. VB approximates the posterior probability distribution by a member from a family of tractable distributions indexed by variational parameters λ belonging to a parameter space \mathcal{M} . The best member is found by minimizing the Kullback-Leibler divergence from the candidate member to

the posterior. VB methods have found their application in a wide range of problems including variational autoencoders (Kingma and Welling 2013), text analysis (Hoffman et al. 2013), Bayesian synthetic likelihood (Ong et al. 2018a), deep neural nets (Tran et al. 2019), to name but a few.

Most of the existing VB methods work with cases where the variational parameter space \mathcal{M} is (a subset of) the Euclidean space \mathbb{R}^d . This paper considers the VB problem where \mathcal{M} is a Riemannian manifold, which naturally arises in many modern applications. For example, in Gaussian VB where the VB approximating distribution is a multivariate Gaussian with mean μ and covariance Σ , $\lambda = (\mu, \Sigma)$ belongs to the product manifold $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ where \mathcal{M}_1 is an Euclidean manifold and \mathcal{M}_2 is the manifold of symmetric and positive definite matrices. We develop manifold-based VB algorithms that cast Euclidean-based constrained VB problems as manifold-based unconstrained optimization problems under which the solution can be efficiently found by exploiting the geometric structure of the constraints. Optimization algorithms that work on manifolds often enjoy better numerical properties. See the monograph of Absil et al. (2009) for recent advances.

Many Euclidean-based VB methods employ (Euclidean) stochastic gradient decent (SGD) for solving the required optimization problem, and it is well-known that the natural

✉ Minh-Ngoc Tran
minh-ngoc.tran@sydney.edu.au

Dang H. Nguyen
dangnh.maths@gmail.com

Duy Nguyen
nducduy@gmail.com

¹ Discipline of Business Analytics, University of Sydney, Business School, Sydney, NSW, Australia

² Department of Mathematics, University of Alabama, Tuscaloosa, AL 35487, USA

³ Department of Mathematics, Marist College, Poughkeepsie, NY 12601, USA

gradient (Amari 1998) is of major importance in SGD. The natural gradient, a geometric object itself, takes into account the information geometry of the family of approximating distributions to help stabilize and speed up the updating procedure. For a comprehensive review and recent development of the natural gradient descent in Euclidean spaces, the reader is referred to Martens (2020). Extending natural gradient descent for use in Riemannian stochastic gradient descent is a non-trivial task and of interest in many VB problems. This paper develops a mathematically formal framework for incorporating the natural gradient into manifold-based VB algorithms. The contributions of this paper are threefold:

- We develop a doubly geometry-informed VB algorithm that exploits both the geometric structure of the manifold constraints of the variational parameter space, and the information geometry of the manifold of the approximating family, which leads to a highly efficient VB algorithm for Bayesian inference in complicated models.
- The proposed manifold VB algorithm is provably convergent and achieves a convergence rate of order $\mathcal{O}(1/\sqrt{T})$ and $\mathcal{O}(1/T^{2-2\epsilon})$, with $\epsilon \in (0, 1)$ and T the number of iterations, for a non-convex lower bound function and a strongly retraction-convex lower bound function, respectively.
- We develop in detail a Manifold Gaussian VB algorithm and a Manifold Wishart VB algorithm, both can be used as a general estimation method for Bayesian inference. The numerical experiments demonstrate that these manifold VB algorithms work efficiently, are more stable and less sensitive to initialization as compared to some existing VB algorithms in the literature. We would like to emphasize that making VB more stable and less initialization-sensitive is of major importance in the current VB literature. We also apply our VB method to estimating a financial time series model and demonstrate its high accuracy in comparison with a “gold standard” Sequential Monte Carlo method.

The paper is organized as follows. Section 2 reviews the VB method on Euclidean spaces and sets up notations. Section 3 develops the manifold-based VB algorithm and Sect. 4 studies its convergence properties. Section 5 presents the Manifold Gaussian VB and the Manifold Wishart VB algorithms, and their applications. Section 6 concludes. The technical proofs are presented in the Appendix.

2 VB algorithms on Euclidean spaces

This section gives a brief overview of VB methods where the variational parameter λ lies in (a subset of) a Euclidean space. It also gives the definition of the natural gradient, and the

motivation for extending the Euclidean-based VB problem into manifolds.

Let y be the data and $p(y|\theta)$ the likelihood function based on a postulated model, with θ the set of model parameters to be estimated. Let $p(\theta)$ be the prior. Bayesian inference requires computing expectations with respect to the posterior distribution whose density (with respect to some reference measure such as the Lebesgue measure) is

$$p(\theta|y) = \frac{p(\theta)p(y|\theta)}{p(y)},$$

where $p(y) = \int p(\theta)p(y|\theta)d\theta$, called the marginal likelihood. It is often difficult to compute such expectations, partly because the density $p(\theta|y)$ itself is intractable as the normalizing constant $p(y)$ is often unknown. For simple models, Bayesian inference is often performed using Markov Chain Monte Carlo (MCMC), which estimates expectations w.r.t. $p(\theta|y)$ by sampling from it. For models where θ is high dimensional or has a complicated structure, MCMC methods in their current development are either not applicable or very time consuming. In the latter case, VB is often an attractive alternative to MCMC. VB approximates the posterior $p(\theta|y)$ by a probability distribution with density $q_\lambda(\theta)$, $\lambda \in \mathcal{M}$ - the variational parameter space, belonging to some tractable family of distributions such as Gaussian. The best λ is found by minimizing the Kullback-Leibler (KL) divergence from $q_\lambda(\theta)$ to $p(\theta|y)$

$$\lambda^* = \arg \min_{\lambda \in \mathcal{M}} \left\{ \text{KL}(q_\lambda \| p(\cdot|y)) = \int q_\lambda(\theta) \log \frac{q_\lambda(\theta)}{p(\theta|y)} d\theta \right\}.$$

It is easy to see that

$$\text{KL}(q_\lambda \| p(\cdot|y)) = - \int q_\lambda(\theta) \log \frac{p(\theta)p(y|\theta)}{q_\lambda(\theta)} d\theta + \log p(y),$$

thus minimizing KL is equivalent to maximizing the lower bound on $\log p(y)$

$$\mathcal{L}(\lambda) = \int q_\lambda(\theta) \log \frac{p(\theta)p(y|\theta)}{q_\lambda(\theta)} d\theta. \quad (2.1)$$

SGD techniques are often employed to solve this optimization problem. The VB approximating distribution $q_\lambda(\theta)$ with the optimized λ is then used for Bayesian inference. See Tran et al. (2021) for an accessible tutorial introduction to VB.

Let

$$\mathcal{N} = \{q_\lambda(\theta) : \lambda \in \mathcal{M}\}$$

be the set of VB approximating probability distributions parameterized by λ , and

$$I_F(\lambda) := \text{Cov}_{q_\lambda}(\nabla_\lambda \log q_\lambda(\theta)) = \mathbb{E}_{q_\lambda}[\nabla_\lambda \log q_\lambda(\theta)(\nabla_\lambda \log q_\lambda(\theta))^\top] \quad (2.2)$$

be the Fisher information matrix w.r.t. q_λ . By the Taylor expansion, we have that

$$\begin{aligned} \text{KL}(q_\lambda || q_{\lambda+\epsilon}) &\approx \text{KL}(q_\lambda || q_\lambda) + (\nabla_{\lambda'} \text{KL}(q_\lambda || q_{\lambda'})|_{\lambda'=\lambda})^\top \epsilon + \\ &\quad + \frac{1}{2} \epsilon^\top I_F(\lambda) \epsilon \\ &= -\mathbb{E}_{q_\lambda}(\nabla_\lambda \log q_\lambda(\theta))^\top \epsilon + \frac{1}{2} \epsilon^\top I_F(\lambda) \epsilon \\ &= \frac{1}{2} \epsilon^\top I_F(\lambda) \epsilon. \end{aligned} \quad (2.3)$$

This shows that the local KL divergence around the point $q_\lambda \in \mathcal{N}$ is characterized by the Fisher matrix $I_F(\lambda)$. Formally, \mathcal{N} can be made into a Riemannian manifold with the Riemannian metric induced by the Fisher information matrix (Rao 1945; Amari 1998).

Assume that the objective function \mathcal{L} is smooth enough, then

$$\mathcal{L}(\lambda + \epsilon) \approx \mathcal{L}(\lambda) + \nabla_\lambda \mathcal{L}(\lambda)^\top \epsilon.$$

The steepest ascent direction ϵ for maximizing $\mathcal{L}(\lambda + \epsilon)$ among all the directions with a fixed length $\|\epsilon\| := \epsilon^\top I_F(\lambda) \epsilon = l$ is

$$\arg \max_{\epsilon: \epsilon^\top I_F(\lambda) \epsilon = l} \left\{ \nabla_\lambda \mathcal{L}(\lambda)^\top \epsilon \right\}. \quad (2.4)$$

By the method of Lagrangian multipliers, this steepest ascent is

$$\epsilon = \nabla_\lambda^{\text{nat}} \mathcal{L}(\lambda) := I_F^{-1}(\lambda) \nabla_\lambda \mathcal{L}(\lambda). \quad (2.5)$$

Amari (1998) termed this the natural gradient and popularized it in machine learning. In the statistics literature, the steepest ascent in the form (2.5) has been used for a long time and is often known as Fisher's scoring in the context of maximum likelihood estimation (see, e.g., Longford (1987)). We adopt the term natural gradient in this paper. The efficiency of the natural gradient over the ordinary gradient has been well documented (Sato 2001; Hoffman et al. 2013; Tran et al. 2017). A remarkable property of the natural gradient is that it is invariant under parameterization (Martens 2020), i.e. it is coordinate-free and an intrinsic geometric object. This further motivates the use of natural gradient in optimization on manifolds.

Most of the VB methods and natural gradient descent are developed for cases where the variational parameter λ lies in an unconstrained Euclidean space. In many situations, however, λ belongs to a non-linear constrained space that forms a differential manifold. A popular example is Gaussian VB where the covariance matrix Σ of size $d \times d$ is subject to the symmetric and positive definite constraint. Ong et al. (2018b) avoid the difficulty of dealing with this constraint by using a factor decomposition $\Sigma = BB^\top + D^2$, where B a full-rank matrix of size $d \times p$ with $p \leq d$, and D a diagonal matrix. Such a decomposition is invariant under orthogonal transformations of B , i.e.

$$\Sigma = BB^\top + D^2 = B'B'^\top + D^2$$

for all $B' = BO$ with O an orthogonal matrix, i.e. $OO^\top = I_p$. That is, the variational parameter B lies in a quotient manifold where each point in this manifold is an equivalence class

$$[B] = \{BO : OO^\top = I_p\}. \quad (2.6)$$

This manifold structure is not considered in Ong et al. (2018b). Zhou et al. (2020) take into account this manifold structure and report some improvement over the plain VB methods. Another example is Wishart VB where the VB distribution $q_\lambda(\theta)$ is an inverse-Wishart distribution $IW(v, \Sigma)$. Here, the variational parameter Σ lives in the manifold of the symmetric and positive definite matrices.

Related work

As we employ the SGD method for optimizing the lower bound $\mathcal{L}(\lambda)$, our paper is related to the recent development of SGD algorithms on Riemannian manifolds. Bonnabel (2013) is one of the first to develop SGD where the cost function is defined on a Riemannian manifold. It is showed in his paper that under some suitable conditions the Reimannian SGD algorithm converges to a critical point of the cost function. In a recent paper, Kasai et al. (2019) propose an adaptive SGD on Riemannian manifolds, which uses different learning rates for different coordinates. Their method is proved to converge to a critical point of the cost function at a rate $\mathcal{O}(\log(T)/\sqrt{T})$. For a recent discussion of generalization of Euclidean adaptive SGD algorithms, such as Adam and Adagrad, to Riemannian manifolds, see Bécigneul and Ganea (2018). The monograph of Absil et al. (2009) provides an excellent account of recent development on optimization on matrix manifolds. Companion user-friendly software such as Manopt (Boumal et al. 2014) has been developed to assist fast growing research in Riemannian optimization.

The natural gradient has been widely used in the machine learning literature; see, e.g., Sato (2001), Hoffman et al.

(2013), Khan and Lin (2017) and Martens (2020). However, most of the existing work only consider cases where the variational parameter λ belongs to an Euclidean space. Zhou et al. (2020) is the only paper that we are aware of develops a VB method on manifolds. However, their paper only considers the Factor Gaussian VB for the particular quotient manifold in (2.6), and does not consider natural gradient. Our paper develops a general VB method for Riemannian manifolds that incorporates the natural gradient, and provides a careful convergence analysis.

3 VB on manifolds with the natural gradient

This section presents our proposed VB algorithm on manifolds. Recall that we are interested in a VB problem where the variational parameter λ lies in a Riemannian manifold \mathcal{M} , i.e. we wish to solve the following optimization problem

$$\arg \max_{\lambda \in \mathcal{M}} \mathcal{L}(\lambda).$$

In order to incorporate the natural gradient into Riemannian SGD, we view the manifold \mathcal{M} as embedded in a Riemannian manifold $\overline{\mathcal{M}} \subset \mathbb{R}^d$, whose Riemannian metric is defined by the Fisher information matrix $I_F(\lambda)$. Let $T_\lambda \overline{\mathcal{M}}$ be the tangent space to $\overline{\mathcal{M}}$ at $\lambda \in \overline{\mathcal{M}}$. The inner product between two tangent vectors $\zeta_\lambda, \xi_\lambda \in T_\lambda \overline{\mathcal{M}}$ is defined as

$$\langle \zeta_\lambda, \xi_\lambda \rangle = \zeta_\lambda^\top I_F(\lambda) \xi_\lambda. \quad (3.1)$$

For VB on manifolds without using the natural gradient, this inner product is the usual Euclidean metric, i.e. $\langle \zeta_\lambda, \xi_\lambda \rangle = \zeta_\lambda^\top \xi_\lambda$. The metric in (3.1) is often referred to as the Fisher-Rao metric. Let $\overline{\mathcal{L}}$ be a differentiable function defined on $\overline{\mathcal{M}}$ such that its restriction on \mathcal{M} is the lower bound \mathcal{L} . Similar to (2.4)–(2.5), it can be shown that the steepest ascent direction at $\lambda \in \overline{\mathcal{M}}$ for optimizing the objective function $\overline{\mathcal{L}}(\lambda)$, i.e. the direction of

$$\arg \max_{\eta_\lambda \in T_\lambda \overline{\mathcal{M}}, \|\eta_\lambda\|=1} D\overline{\mathcal{L}}(\lambda)[\eta_\lambda],$$

is the natural gradient

$$\nabla_\lambda^{\text{nat}} \overline{\mathcal{L}}(\lambda) = I_F^{-1}(\lambda) \nabla_\lambda \overline{\mathcal{L}}(\lambda), \quad \lambda \in \overline{\mathcal{M}}. \quad (3.2)$$

Here, $D\overline{\mathcal{L}}(\lambda)[\eta_\lambda]$ denotes the directional derivative of $\overline{\mathcal{L}}$ at λ in the direction of η_λ , and $\nabla_\lambda \overline{\mathcal{L}}(\lambda)$ is the usual Euclidean gradient vector of $\overline{\mathcal{L}}(\lambda)$. We note that, for $\lambda \in \mathcal{M}$,

$$\nabla_\lambda^{\text{nat}} \overline{\mathcal{L}}(\lambda) = I_F^{-1}(\lambda) \nabla_\lambda \overline{\mathcal{L}}(\lambda) = I_F^{-1}(\lambda) \nabla_\lambda \mathcal{L}(\lambda) = \nabla_\lambda^{\text{nat}} \mathcal{L}(\lambda).$$

We recall that the Riemannian gradient of a smooth function $f(\lambda)$ on a Riemannian manifold \mathcal{M} , embedded in \mathbb{R}^d and equipped with the Riemannian metric $\langle \cdot, \cdot \rangle$, is the unique vector $\text{grad} f(\lambda)$ in the tangent space $T_\lambda \mathcal{M}$ at $\lambda \in \mathcal{M}$ such that

$$\langle \text{grad} f(\lambda), \xi_\lambda \rangle = Df(\lambda)[\xi_\lambda], \quad \forall \xi_\lambda \in T_\lambda \mathcal{M}.$$

The following lemma is important for the purpose of this paper. It shows that the natural gradient $\nabla_\lambda^{\text{nat}} \overline{\mathcal{L}}(\lambda)$ is a Riemannian gradient defined in the ambient manifold $\overline{\mathcal{M}}$, which leads to a formal framework for associating the natural gradient to the Riemannian gradient of the lower bound \mathcal{L} defined in the manifold \mathcal{M} .

Lemma 3.1 *The natural gradient of the function $\overline{\mathcal{L}}$ on the Riemannian manifold $\overline{\mathcal{M}}$ with the Fisher-Rao metric (3.1) is the Riemannian gradient of $\overline{\mathcal{L}}$. In particular, the natural gradient at λ belongs to the tangent space to $\overline{\mathcal{M}}$ at λ .*

We now need to associate the Riemannian gradient $\nabla_\lambda^{\text{nat}} \overline{\mathcal{L}}(\lambda)$ to the Riemannian gradient of the lower bound $\mathcal{L}(\lambda)$ defined in \mathcal{M} ; the latter is what we need for using Riemannian SGD to optimize $\mathcal{L}(\lambda)$. This is done in the two cases: \mathcal{M} is a submanifold (Sect. 3.1) and \mathcal{M} is a quotient manifold (Sect. 3.2).

3.1 Riemannian submanifolds

Suppose that \mathcal{M} is a submanifold of $\overline{\mathcal{M}}$. In order to define the Riemannian gradient of the lower bound \mathcal{L} defined on the manifold \mathcal{M} , we need to equip \mathcal{M} with a Riemannian metric. In most cases, this metric is inherited from that of $\overline{\mathcal{M}}$ in a natural way. Since $T_\lambda \mathcal{M}$ is a subspace of $T_\lambda \overline{\mathcal{M}}$, the Riemannian metric of $\zeta_\lambda, \xi_\lambda \in T_\lambda \mathcal{M}$ can be defined as

$$\langle \zeta_\lambda, \xi_\lambda \rangle = \zeta_\lambda^\top I_F(\lambda) \xi_\lambda,$$

with $\zeta_\lambda, \xi_\lambda$ viewed as vectors in $T_\lambda \overline{\mathcal{M}}$. With this metric, we can define the orthogonal complement $(T_\lambda \mathcal{M})^\perp$ of $T_\lambda \mathcal{M}$ in $T_\lambda \overline{\mathcal{M}}$, i.e. $T_\lambda \overline{\mathcal{M}} = T_\lambda \mathcal{M} \oplus (T_\lambda \mathcal{M})^\perp$. Write

$$\text{grad} \overline{\mathcal{L}}(\lambda) = \text{Proj}_\lambda \text{grad} \overline{\mathcal{L}}(\lambda) + \text{Proj}_\lambda^\perp \text{grad} \overline{\mathcal{L}}(\lambda)$$

where Proj_λ and $\text{Proj}_\lambda^\perp$ denote the projections on $T_\lambda \mathcal{M}$ and $(T_\lambda \mathcal{M})^\perp$, respectively. Recall that

$$\text{grad} \overline{\mathcal{L}}(\lambda) = \nabla_\lambda^{\text{nat}} \overline{\mathcal{L}}(\lambda) = I_F^{-1}(\lambda) \nabla_\lambda \mathcal{L}(\lambda), \quad \lambda \in \mathcal{M}.$$

Then, the Riemannian gradient of \mathcal{L} is the projection of $\text{grad} \overline{\mathcal{L}}$ on $T_\lambda \mathcal{M}$

$$\text{grad} \mathcal{L}(\lambda) = \text{Proj}_\lambda \text{grad} \overline{\mathcal{L}}(\lambda). \quad (3.3)$$

This is because

$$\begin{aligned} \langle \text{grad} \mathcal{L}(\lambda), \eta_\lambda \rangle &= \langle \text{grad} \bar{\mathcal{L}}(\lambda) - \text{Proj}_\lambda^\perp \text{grad} \bar{\mathcal{L}}(\lambda), \eta_\lambda \rangle \\ &= \langle \text{grad} \bar{\mathcal{L}}(\lambda), \eta_\lambda \rangle \\ &= D\bar{\mathcal{L}}(\lambda)[\eta_\lambda] = D\mathcal{L}(\lambda)[\eta_\lambda], \quad \forall \eta_\lambda \in T_\lambda \mathcal{M}. \end{aligned}$$

In some cases, such as Gaussian VB in Sect. 5.1, $T_\lambda \mathcal{M} \cong T_\lambda \bar{\mathcal{M}}$, then $\text{grad} \mathcal{L}(\lambda) = \text{grad} \bar{\mathcal{L}}(\lambda)$, i.e. the natural gradient is the Riemannian gradient of \mathcal{L} . In some other cases, however, using the inherited metric might lead to a projection Proj_λ that is cumbersome to compute. In such cases, one needs to use an alternative Riemannian metric on \mathcal{M} such that the projection Proj_λ is easy to compute. Below we give an example in the case of Stiefel manifold, which is a popular manifold in many applications.

Stiefel manifolds. Suppose that \mathcal{M} is a Stiefel manifold $\mathcal{M} = \mathcal{S}(p, n)$ defined as

$$\mathcal{S}(p, n) = \{W \in \text{Mat}(n, p) : W^\top W = I_p\}, \quad (3.4)$$

where $\text{Mat}(n, p)$ is the set of real matrices of size $n \times p$. We can think of \mathcal{M} as embedded in $\bar{\mathcal{M}} = \mathbb{R}^d$, $d = n \cdot p$, equipped with the Fisher-Rao metric

$$\langle \zeta_W, \xi_W \rangle = (\text{vec}(\zeta_W))^\top I_F(W) \text{vec}(\xi_W), \quad (3.5)$$

$\zeta_W, \xi_W \in T_W \bar{\mathcal{M}} \cong \text{Mat}(n, p)$; where $\text{vec}(\cdot)$ denotes the vectorization operator and

$$I_F(W) = \text{cov}_{q_W}(\nabla_{\text{vec}(W)} \log q_W(\theta)),$$

is the Fisher matrix defined in (2.2) with $q_W(\theta)$ the variational distribution. The natural gradient of function $\bar{\mathcal{L}}$ at $W \in \mathcal{M}$ is

$$\begin{aligned} \text{grad} \bar{\mathcal{L}}(W) &= \text{vec}^{-1} \left(I_F(W)^{-1} \nabla_{\text{vec}(W)} \mathcal{L}(W) \right) \\ &\in \text{Mat}(n, p), \end{aligned} \quad (3.6)$$

where vec^{-1} is the inverse of vec , sending a np -vector to the corresponding matrix in $\text{Mat}(n, p)$. It is easy to see that the tangent space of \mathcal{M} at W is

$$T_W \mathcal{M} = \{Z \in \text{Mat}(n, p) : Z^\top W + W^\top Z = 0_{p \times p}\}.$$

If we equip \mathcal{M} with the Riemannian metric defined in (3.5), the projection on $T_W \mathcal{M}$ is cumbersome to compute. We therefore opt to use the usual Euclidean metric

$$\langle \zeta_W, \xi_W \rangle_{\text{Euc}} = \text{trace}(\zeta_W^\top \xi_W) = (\text{vec}(\zeta_W))^\top \text{vec}(\xi_W),$$

with $\zeta_W, \xi_W \in T_W \mathcal{M}$. The following lemma gives an expression for the Riemannian gradient of \mathcal{L} defined on the Stiefel

manifold, and is useful for many applications involving Stiefel manifolds. Similar results to Lemma 3.2 can be found in the literature (see., e.g., Edelman et al. (1998)), however, here we state and prove the results specifically for the case G is the natural gradient.

Lemma 3.2 *Let \mathcal{L} be a function on the Stiefel manifold \mathcal{M} equipped with the usual Euclidean metric. The Riemannian gradient of \mathcal{L} at W is*

$$\text{grad} \mathcal{L}(W) = (I_n - WW^\top)G + \text{Wskew}(W^\top G) \quad (3.7)$$

with $G = \text{grad} \bar{\mathcal{L}}(W)$ given in (3.6), and $\text{skew}(A) := (A - A^\top)/2$.

3.2 Quotient manifolds

This section derives the Riemannian gradient of \mathcal{L} when \mathcal{M} is a quotient manifold induced from the ambient manifold $\bar{\mathcal{M}}$. Suppose that $\bar{\mathcal{M}} \subset \mathbb{R}^d$ is a Riemannian manifold with the Riemannian metric $\langle \cdot, \cdot \rangle_{\bar{\mathcal{M}}}$. Suppose that there is an equivalence relation on $\bar{\mathcal{M}}$ defined as

$$\lambda, \lambda' \in \bar{\mathcal{M}}, \quad \lambda \sim \lambda' \text{ if and only if } q_\lambda = q_{\lambda'},$$

and thus $\mathcal{L}(\lambda) = \mathcal{L}(\lambda')$. This is the case of Gaussian VB with the covariance matrix Σ having a factor decomposition. Define the equivalence class

$$[\lambda] = \{\lambda' \in \bar{\mathcal{M}} : q_{\lambda'} = q_\lambda\},$$

i.e., the class of all parameterizations λ that represent the same distribution. Let

$$\mathcal{M} := \bar{\mathcal{M}} / \sim := \{[\lambda] : \lambda \in \bar{\mathcal{M}}\}$$

and define the canonical projection

$$\pi : \bar{\mathcal{M}} \rightarrow \mathcal{M} = \bar{\mathcal{M}} / \sim, \quad \lambda \mapsto [\lambda]. \quad (3.8)$$

Then we can endow the quotient set \mathcal{M} with the topology induced from $\bar{\mathcal{M}}$ by the projection π . This makes \mathcal{M} become a smooth manifold called the quotient manifold, see Absil et al. (2009). If we define $L : \mathcal{M} \rightarrow \mathbb{R}$, $[\lambda] \mapsto L([\lambda]) = \mathcal{L}(\lambda)$, i.e. $\mathcal{L} = L \circ \pi$, then L is the function defined on \mathcal{M} that we want to optimise. For optimisation on \mathcal{M} , one needs to be able to represent numerically tangent vectors at each $[\lambda] \in \mathcal{M}$. Geometrical objects in $\bar{\mathcal{M}}$, such as points λ and tangent vectors, are vectors in the usual sense, so they can be numerically represented in computer for numerical computation. However, geometrical objects in the quotient manifold \mathcal{M} are abstract, much of research in quotient manifolds has been focused on how to represent these geometrical objects

numerically. The key tool is the concept of *horizontal lift*; see, e.g. Absil et al. (2009); Kobayashi and Nomizu (1969).

By the level set theorem (Tu 2011, Chapter 2), $\pi^{-1}([\lambda])$ is an embedded submanifold in $\overline{\mathcal{M}}$, hence, it admits a tangent space

$$\mathcal{V}_\lambda := T_\lambda(\pi^{-1}([\lambda])),$$

called the vertical space, which is a linear subspace of $T_\lambda\overline{\mathcal{M}}$. Let \mathcal{H}_λ be the orthogonal complement of \mathcal{V}_λ in $T_\lambda\overline{\mathcal{M}}$, called the horizontal space, i.e. $T_\lambda\overline{\mathcal{M}} = \mathcal{H}_\lambda \oplus \mathcal{V}_\lambda$. The orthogonality here is w.r.t. the metric defined on $\overline{\mathcal{M}}$. For each tangent vector $\xi_{[\lambda]}$ at $[\lambda] \in \mathcal{M}$, there exists a unique vector $\bar{\xi}_\lambda$ in the horizontal space \mathcal{H}_λ such that (Kobayashi and Nomizu 1969, Prop. 1.2)

$$D\pi(\lambda)[\bar{\xi}_\lambda] = \xi_{[\lambda]},$$

$\bar{\xi}_\lambda$ is called the horizontal lift of $\xi_{[\lambda]}$. Then, as $\mathcal{L} = L \circ \pi$,

$$\begin{aligned} D\mathcal{L}(\lambda)[\bar{\xi}_\lambda] &= DL(\pi(\lambda))(D\pi(\lambda)[\bar{\xi}_\lambda]) \\ &= DL([\lambda])(\xi_{[\lambda]}), \end{aligned} \quad (3.9)$$

which shows that the directional derivative of L in the direction of $\xi_{[\lambda]}$ is characterised by the directional derivative of \mathcal{L} in the direction of the horizontal lift $\bar{\xi}_\lambda$. Intuitively, for the optimization purposes, we can ignore the vertical space and just focus on the horizontal space, as the objective function \mathcal{L} doesn't change along the vertical space. It's worth noting that the property (3.9) does not depend on any particular choice λ in $[\lambda]$.

Let $\text{grad}\mathcal{L}(\lambda)$ be the Riemannian gradient of \mathcal{L} at $\lambda \in \overline{\mathcal{M}}$. We have that, for all $\eta_\lambda \in \mathcal{V}_\lambda$

$$\langle \overline{\text{grad}\mathcal{L}(\lambda)}, \eta_\lambda \rangle_{\overline{\mathcal{M}}} = D\mathcal{L}(\lambda)[\eta_\lambda] = 0,$$

as $\mathcal{L}(\lambda)$ doesn't change along the vertical space, which shows that $\text{grad}\mathcal{L}(\lambda) \in \mathcal{H}_\lambda$. Let $\text{grad}L([\lambda])$ be the tangent vector to \mathcal{M} at $[\lambda]$ that has $\text{grad}\mathcal{L}(\lambda)$ as its horizontal lift. Then, by equipping \mathcal{M} with the inner product inherited from $\overline{\mathcal{M}}$,

$$\begin{aligned} \langle \text{grad}L([\lambda]), \eta_{[\lambda]} \rangle_{\mathcal{M}} &:= \langle \overline{\text{grad}\mathcal{L}(\lambda)}, \bar{\eta}_\lambda \rangle_{\overline{\mathcal{M}}} \\ &= D\mathcal{L}(\lambda)[\bar{\eta}_\lambda] \\ &= DL([\lambda])(\eta_{[\lambda]}), \quad \forall \eta_{[\lambda]} \in T_{[\lambda]}\mathcal{M} \end{aligned} \quad (3.10)$$

we have that $\text{grad}L([\lambda])$ is the Riemannian gradient of L at $[\lambda] \in \mathcal{M}$. We note that (3.10) does not depend on the choice of $\lambda \in [\lambda]$. So, with the inherited inner product from $\overline{\mathcal{M}}$, the usual Riemannian gradient of \mathcal{L} on $\overline{\mathcal{M}}$ is the horizontal lift of the Riemannian gradient of L on \mathcal{M} . This remarkable property of quotient manifolds makes it convenient for numerical optimisation problems.

Remark 3.1 Technically, in order for the inherited Riemannian metric on \mathcal{M} to be well-defined, it is often required in the literature that $\langle \bar{\xi}_\lambda, \bar{\eta}_\lambda \rangle_{\overline{\mathcal{M}}}$ does not depend on $\lambda \in [\lambda]$. This condition is typically not satisfied when $\overline{\mathcal{M}}$ is equipped with the Fisher-Rao metric as considered in this paper. However, as we show above, the Riemannian gradient of L is still well-defined without this requirement, as (3.10) holds for any $\lambda \in [\lambda]$.

3.3 Retraction

After deriving the Riemannian gradient, which is the steepest ascent direction of the lower bound function \mathcal{L} at the current point on the manifold, we need to derive the *exponential map*¹, denoted by $\text{Exp}_\lambda(\xi_\lambda)$, that projects a point on the tangent space back to the manifold. Exponential map is a standard concept in differential geometry. Intuitively, exponential maps are mappings that, given a point λ on a manifold and a tangent vector ξ_λ at λ , generalize the concept " $\lambda + \xi_\lambda$ " in Euclidean spaces. $\text{Exp}_\lambda(\xi_\lambda)$ is a point on the manifold that can be reached by leaving from λ and moving in the direction ξ_λ while remaining on the manifold. We refer to Absil et al. (2009) for a precise definition and examples. One major drawback of exponential maps is that their calculation is often cumbersome in practice. *Retraction*, the first order approximation of the exponential map, is often used instead. A retraction $R_\lambda : T_\lambda\mathcal{M} \rightarrow \mathcal{M}$ at $\lambda \in \mathcal{M}$ has the important property that it preserves gradients, i.e. the curve $\gamma_{\xi_\lambda} : t \mapsto R_\lambda(t\xi_\lambda)$ satisfies $D\gamma_{\xi_\lambda}(0)[\xi_\lambda] = \xi_\lambda$ for every $\xi_\lambda \in T_\lambda\mathcal{M}$. See (Absil et al. 2009, Chapter 4) for a formal definition of retraction and Manton (2002) for an interpretation of retraction from a local optimization perspective. Also see Fig. 1 (Left) for a visualization.

Closed-form formulae for retractions on common manifolds are available in the literature, see, for example, Absil et al. (2009). For instance, a popular retraction on the Stiefel manifold is

$$R_W(\xi_W) = \text{qf}(W + \xi_W), \quad W \in \mathcal{S}(p, n), \quad \xi_W \in T_W\mathcal{S}(p, n). \quad (3.11)$$

Here $\text{qf}(A) = Q$, where $A = QR$ is the QR decomposition of $A \in \text{Mat}(n, p)$, $Q \in \mathcal{S}(p, n)$ and $R \in \text{Mat}(n, p)$ is upper triangular. See Sato and Aihara (2019) for an efficient computation of this retraction based on the Cholesky QR factorization. For quotient manifolds, a popular retraction is

$$R_{[\lambda]}(\xi_{[\lambda]}) = \pi(\lambda + \bar{\xi}_\lambda), \quad [\lambda] \in \mathcal{M}, \quad \xi_{[\lambda]} \in T_{[\lambda]}\mathcal{M},$$

¹ In general, an exponential map at λ is defined locally near λ . In order to define exponential map on the entire tangent space, the manifold needs to be complete, see, e.g., Tu (2011).




Fig. 1 Left: Tangent space at λ and the retraction map. Right: Vector transport

$\bar{\xi}_\lambda$ is the horizontal lift of $\xi_{[\lambda]}$, and π is the canonical projection in (3.8).

3.4 Momentum

The momentum method, which uses a moving average of the gradient vectors at the previous iterates to accelerate convergence and also help reduce noise in the estimated gradient, is widely used in Euclidean-based stochastic gradient optimization. Extending the momentum method to manifolds requires *parallel translation*, a tool in differential geometry for moving tangent vectors from one tangent space to another, while still preserving the length and angle (to some fixed direction) of the original tangent vectors. Similar to exponential map, a parallel translation is often approximated by a *vector transport* which is much easier to compute; see (Absil et al. 2009, Chapter 8) for a formal definition. See Fig. 1 (Right) for a visualization. Let $\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(\xi_{\lambda_t})$ denote the vector transport of tangent vector $\xi_{\lambda_t} \in T_{\lambda_t}\mathcal{M}$ to tangent space $T_{\lambda_{t+1}}\mathcal{M}$. A simple vector transport is the projection of ξ_{λ_t} on $T_{\lambda_{t+1}}\mathcal{M}$, i.e. $\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(\xi_{\lambda_t}) = \text{Proj}_{T_{\lambda_{t+1}}\mathcal{M}}(\xi_{\lambda_t})$. Roy and Harandi (2017) is the first to use the momentum method in Riemannian SGD, but they do not provide any convergence analysis.

3.5 Manifold VB algorithm

The pseudo-code in Algorithm 1 summarizes our VB algorithm on manifolds. We use the “hat” notation in $\widehat{\nabla_\lambda^{\text{nat}} \mathcal{L}(\lambda)}$ to emphasize that the natural gradient is obtained from a noisy and unbiased estimator $\widehat{\nabla_\lambda \mathcal{L}(\lambda)}$ of the Euclidean gradient as we often don’t have access to the exact $\nabla_\lambda \mathcal{L}(\lambda)$. That is, starting from an initial λ_0 , the initial momentum gradient Y_0 can be found by projecting the natural gradient estimate $\widehat{\nabla_\lambda^{\text{nat}} \mathcal{L}(\lambda_0)}$ on $T_{\lambda_0}\mathcal{M}$. The relevant projections were described in Sects. 3.1 and 3.2. At a step t , from λ_t , we move on the tangent space $T_{\lambda_t}\mathcal{M}$ along the direction Y_t to find the next iterate λ_{t+1} by retraction, $\lambda_{t+1} = R_{\lambda_t}(\varepsilon Y_t)$. Then, we calculate the natural gradient estimate $\widehat{\nabla_{\lambda_{t+1}}^{\text{nat}} \mathcal{L}(\lambda_{t+1})}$ at λ_{t+1} , whose projection on $T_{\lambda_{t+1}}\mathcal{M}$ gives the Riemannian gradient $\text{grad} \mathcal{L}(\lambda_{t+1})$. Finally, the new momentum gradient Y_{t+1} is calculated from the vector transport of Y_t (to $T_{\lambda_{t+1}}\mathcal{M}$) and $\text{grad} \mathcal{L}(\lambda_{t+1})$.

4 Convergence analysis

To be consistent with the standard notation in the optimization literature, and with an abuse of notation, let us define the cost function as $\mathcal{L}(\cdot) := -\mathcal{L}(\cdot)$. That is, our optimization problem is

$$\arg \min_{\lambda \in \mathcal{M}} \mathcal{L}(\lambda). \quad (4.1)$$

In this section, for notational simplicity, we will denote by $\nabla \mathcal{L}(\lambda)$ the Riemannian gradient of the *cost function* \mathcal{L} , and by $\widehat{\nabla \mathcal{L}(\lambda)}$ its unbiased estimator. Let $\{\lambda_t, t \leq 0\}$ be the iterates from Algorithm 1, and \mathcal{F}_t be the σ -field generated by $\{\lambda_s, s \leq t\}$. Because of the unbiasedness, we can write

$$\widehat{\nabla_\lambda \mathcal{L}(\lambda_t)} = \nabla_\lambda \mathcal{L}(\lambda_t) + \Delta M_t,$$

with $\{\Delta M_t\}$ a martingale difference w.r.t. $\{\mathcal{F}_t\}$, i.e. $\mathbb{E}(\Delta M_{t+1} | \mathcal{F}_t) = 0$. For the purpose of convergence analysis, we write our manifold VB algorithm as follows

$$\begin{cases} \lambda_0 \in \mathcal{M}, Y_0 \in T_{\lambda_0}\mathcal{M} \\ \lambda_{t+1} = R_{\lambda_t}(-Y_t), \forall 0 \leq t \in \mathbb{N} \\ Y_{t+1} = \zeta \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) + \gamma (\nabla \mathcal{L}(\lambda_{t+1}) + \Delta M_{t+1}), \end{cases} \quad (4.2)$$

with $\zeta, \gamma \in (0, 1)$, $0 \leq t \in \mathbb{N}$. The minus sign in $R_{\lambda_t}(-Y_t)$ results from the change in the notation $\mathcal{L}(\cdot) := -\mathcal{L}(\cdot)$ above, and with the suitable choice of ζ and γ we can recover Algorithm 1.

We next need some definitions.

Definition 4.1 (Huang et al. 2015a, Section 3.3) A neighborhood $\mathcal{S} \subset \mathcal{M}$ of x^* is said to be totally retractive if there is $\delta > 0$ such that for any $y \in \mathcal{S}$, $R_y(B(0_y, \delta)) \supset \mathcal{S}$ and R_y is a diffeomorphism on $B(0_y, \delta)$, where $B(0_y, \delta)$ is the ball of radius δ in $T_y\mathcal{M}$ centered at the origin 0_y .

Definition 4.2 (Huang et al. 2015b, Definition 3.1) For a function $f : \mathcal{M} \mapsto \mathbb{R}$ on a Riemannian manifold \mathcal{M} with retraction R , define $m_{\lambda, \eta}(t) = f(R_\lambda(t\eta))$ for $\lambda \in \mathcal{M}$, $\eta \in T_\lambda\mathcal{M}$. The function f is retraction-convex with respect to the retraction R in a set \mathcal{S} if for all $\lambda \in \mathcal{M}$, $\eta \in T_\lambda\mathcal{M}$, $|\eta| = 1$, $m_{\lambda, \eta}(t)$ is convex for all t which satisfy $R_\lambda(s\eta) \in \mathcal{S}$ for all $s \in [0, t]$. Moreover, f is strongly retraction-convex in \mathcal{S} if $m_{\lambda, \eta}(t)$ is strongly convex, that is $\tilde{a}_0 \leq \frac{d^2 m_{\lambda, \eta}}{dt^2}(t) \leq \tilde{a}_1$ for some positive constants \tilde{a}_1, \tilde{a}_0 .

Remark 4.1 If $m_{\lambda, \eta}(t)$ is strongly convex with $\tilde{a}_0 \leq \frac{d^2 m_{\lambda, \eta}}{dt^2}(t) \leq \tilde{a}_1$ then

$$f(\lambda) - f(\lambda^*) \leq \frac{\tilde{a}_0}{2} \left| \frac{dm_{\lambda^*, \eta}}{dt}(t) \right|^2$$

Input: Learning rate ε , momentum weight ω , and a lower bound function $\mathcal{L}(\lambda)$ on a manifold \mathcal{M} .

Output: A local mode λ^* of $\mathcal{L}(\lambda)$.

Initialization: $\lambda_0 \in \mathcal{M}$, $Y_0 \in T_{\lambda_0}\mathcal{M}$;

$t = 0$, stop=false;

while not stop **do**

$\lambda_{t+1} = R_{\lambda_t}(\varepsilon Y_t)$; /* retraction */

if stopping rule is met **then**

 stop=true;

else

 Compute the natural gradient estimate $\widehat{\nabla_{\lambda}^{\text{nat}} \mathcal{L}(\lambda_{t+1})}$;

 Compute the Riemannian gradient $\text{grad} \mathcal{L}(\lambda_{t+1})$;

$Y_{t+1} := \omega \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) + (1 - \omega) \text{grad} \mathcal{L}(\lambda_{t+1})$; /* steepest ascent direction with momentum */

$t := t + 1$;

end

end

Algorithm 1: VB on Manifold Algorithm

where $\lambda = R_{\lambda^*}(t\eta)$. If we assume that f is strongly retraction-convex in \mathcal{S} and for any $\tilde{\lambda} \in \mathcal{S}$, there exists $\tilde{\eta}$ such that $R_{\lambda^*}(\tilde{\eta}) = \tilde{\lambda}$ and the derivative $DR_{\lambda^*}(\eta)$ is bounded, then the chain rule implies $|\frac{dm_{\lambda^*, \eta}}{dt}(t)| \leq c|\nabla f(\lambda)|$ with $\lambda = R_{\lambda^*}(t\eta)$. As a result,

$$f(\lambda) - f(\lambda^*) \leq \frac{\tilde{a}_0 c^2}{2} |\nabla f(\lambda)|^2, \lambda \in \mathcal{M}.$$

In Theorem 1, we show the convergence of (4.2) under suitable conditions imposed on the objective function \mathcal{L} . It is worth emphasizing that the convergence analysis in this section is done in a general setting for Riemannian SGD with momentum rather than restricting on the VB problem in the previous sections. It can therefore be applied to more general settings.

Theorem 1 Assume that

- (i) There exists a totally retractive neighborhood of λ^* , \mathcal{S} , such that $\lambda_t \in \mathcal{S}$ for any $t \geq 0$.
- (ii) $\nabla \mathcal{L}$ and ΔM_t are bounded such that $|\nabla \mathcal{L}| + |\Delta M| \leq b_{\mathcal{L}}$ almost surely for some constant $b_{\mathcal{L}} > 0$.
- (iii) $\nabla \mathcal{L}(\lambda)$ is \tilde{L} -Lipschitz with respect to retraction R , that is, $|\nabla \mathcal{L}(R_{\lambda}(\eta)) - \nabla \mathcal{L}(\lambda)| \leq \tilde{L}|\eta|$ for $\lambda \in \mathcal{S}$, $\eta \in T_{\lambda}\mathcal{M}$.

Consider the sequence $(\lambda_t)_{t \in \mathbb{N}}$ obtained from (4.2) using $\gamma = \frac{1}{\sqrt{T}}$. The following holds true:

$$\min_{t \in [1, T]} \mathbb{E} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \leq \frac{C}{\sqrt{T}}, \text{ for some } C > 0. \quad (4.3)$$

Moreover, when the objective function \mathcal{L} is strongly retraction-convex, for $\epsilon \in (0, 1)$, by choosing $\gamma = \frac{1}{T^\epsilon}$, there exists a constant C_ϵ such that

$$\mathbb{E}(\lambda_T - \lambda^*)^2 \leq C_\epsilon T^{2\epsilon-2}.$$

The proof can be found in the Appendix. We note that the first assumption in Theorem 1 is standard, see, for example Huang et al. (2015a). The condition $|\nabla \mathcal{L}| + |\Delta M| \leq b_{\mathcal{L}}$ is essential to make sure that Y_t stays bounded. That is, it does not diverge to infinity, which is the key property of the algorithm. While the Lipschitz property of $\nabla \mathcal{L}(\lambda)$ guarantee that we can expand \mathcal{L} using Taylor expansion. The assumption related to the martingale difference ΔM_t is justified by the fact that the estimator $\widehat{\nabla \mathcal{L}}$ is unbiased.

5 Applications

5.1 Manifold Gaussian VB

Gaussian VB (GVB) uses a multivariate Gaussian distribution $N(\mu, \Sigma)$ for the VB approximation q_λ , $\lambda = (\mu, \text{vec}(\Sigma))$. GVB has been extensively used in the literature, often with some simplifications imposed on Σ ; e.g., Σ is a diagonal matrix $\text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ or has a factor structure $\Sigma = BB^\top + D^2$. One of the reasons of imposing these simplifications is to deal with the symmetric and positive definiteness constraints on Σ . We do not impose any simplifications on Σ , and deal with these constraints by considering the VB optimization problem on the manifold \mathcal{M} of symmetric and positive definite matrices $\mathcal{M} = \{\Sigma \in \text{Mat}(d, d) : \Sigma = \Sigma^\top, \Sigma > 0\}$. We can think of \mathcal{M} as embedded in $\overline{\mathcal{M}} = \mathbb{R}^{d^2}$.

From (2.1), the gradient of the lower bound is Tran et al. (2021)

$$\nabla_{\lambda} \mathcal{L}(\lambda) = \mathbb{E}_{q_{\lambda}} [\nabla_{\lambda} \log q_{\lambda}(\theta) \times h_{\lambda}(\theta)],$$

with

$$h(\theta) := \log \frac{p(\theta)p(y|\theta)}{q_{\lambda}(\theta)}.$$

Here, we use the so-called *score-function* VB as in Tran et al. (2017), which does not require the gradient of the log-likelihood. We follow Tran et al. (2017) and use a control variate for the gradient of lower bound

$$\begin{aligned}\nabla_{\lambda} \mathcal{L}(\lambda) &= \mathbb{E}_{q_{\lambda}} [\nabla_{\lambda} \log q_{\lambda}(\theta) \times h(\theta)] \\ &= \mathbb{E}_{q_{\lambda}} [\nabla_{\lambda} \log q_{\lambda}(\theta) \times (h(\theta) - c)],\end{aligned}\quad (5.1)$$

where c is a vector selected to minimize the variance of the gradient estimate

$$c_i = \frac{\text{cov}(\nabla_{\lambda_i} \log q_{\lambda}(\theta), \nabla_{\lambda_i} \log q_{\lambda}(\theta) \times h(\theta))}{\mathbb{V}(\nabla_{\lambda_i} \log q_{\lambda}(\theta))},$$

for $i = 1, \dots, d_{\lambda}$, with d_{λ} the size of λ , which can be estimated by sampling from q_{λ} .

Mardia and Marshall (1984) show that the Fisher information matrix for the multivariate Gaussian distribution $N(\mu, \Sigma)$ is

$$I_F(\lambda) = \begin{pmatrix} \Sigma^{-1} & 0 \\ 0 & I_F(\Sigma) \end{pmatrix}$$

where $I_F(\Sigma)$ is an $d^2 \times d^2$ matrix with entries

$$(I_F(\Sigma))_{\sigma_{ij}, \sigma_{kl}} = \frac{1}{2} \text{tr} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \sigma_{ij}} \Sigma^{-1} \frac{\partial \Sigma}{\partial \sigma_{kl}} \right).$$

One can derive that $I_F(\Sigma) \approx \Sigma^{-1} \otimes \Sigma^{-1}$, with \otimes the Kronecker product. Therefore

$$I_F(\lambda)^{-1} \approx \begin{pmatrix} \Sigma & 0 \\ 0 & \Sigma \otimes \Sigma \end{pmatrix}, \quad (5.2)$$

which gives a convenient form for obtaining an approximate natural gradient. The natural gradient w.r.t. μ and Σ is approximated as

$$\nabla_{\mu}^{\text{nat}} \mathcal{L}(\lambda) = \Sigma \nabla_{\mu} \mathcal{L}(\lambda), \quad (5.3)$$

$$\begin{aligned}\nabla_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda) &= \text{vec}^{-1} \left((\Sigma \otimes \Sigma) \nabla_{\text{vec}(\Sigma)} \mathcal{L}(\lambda) \right) \\ &= \Sigma \nabla_{\Sigma} \mathcal{L}(\lambda) \Sigma.\end{aligned}\quad (5.4)$$

As $T_{\Sigma} \mathcal{M} = \text{Mat}(d, d) \cong T_{\Sigma} \overline{\mathcal{M}} \cong \mathbb{R}^{d^2}$, the projection in (3.3) is the identity, hence $\nabla_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda)$ is the Riemannian gradient of lower bound \mathcal{L} w.r.t. Σ . The Manifold GVB algorithm is outlined in Algorithm 2. One of the most popular retractions used for the manifold \mathcal{M} of symmetric and positive definite matrices is (see the Manopt toolbox of Boumal et al. (2014))

$$R_{\Sigma}(\xi) = \Sigma + \xi + \frac{1}{2} \xi \Sigma^{-1} \xi, \quad \xi \in T_{\Sigma} \mathcal{M} \quad (5.5)$$

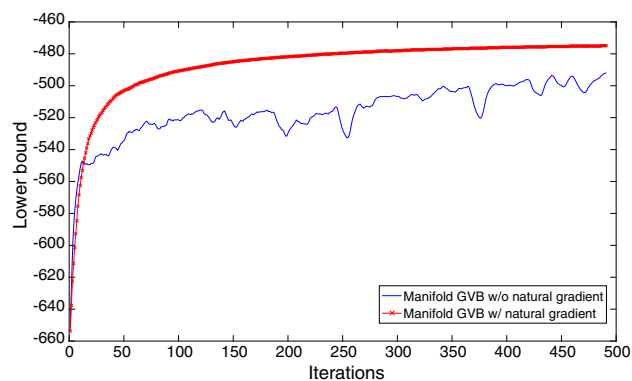


Fig. 2 Plots of lower bound over iteration for the manifold GVB algorithms with (cross red) and without (solid blue) using the natural gradient. Both algorithms were run for 500 iterations. The lower bound estimates have been smoothed by a moving average with a window of size 10

and vector transport

$$\Gamma_{\Sigma_1 \rightarrow \Sigma_2}(\xi) = E \xi E^{\top}, \quad E = (\Sigma_2 \Sigma_1^{-1})^{1/2}, \quad \xi \in T_{\Sigma_1} \mathcal{M}. \quad (5.6)$$

The Matlab code implementing the Manifold GVB algorithm is made available online at <https://github.com/VBayesLab>.

Numerical experiments

We apply the Manifold GVB algorithm to fitting a logistic regression model using the German Credit dataset. This dataset, available on the UCI Machine Learning Repository <https://archive.ics.uci.edu/ml/index.php>, consists of observations on 1000 customers, each was already rated as being “good credit” (700 cases) or “bad credit” (300 cases). The covariates include credit history, education, employment status, etc. and lead to totally 25 predictors after using dummy variables to represent the categorical covariates.

A naive GVB implementation is to only update Σ when its updated value satisfies the symmetric and positive definiteness constraint. This naive implementation didn’t work at all in this example. To see the usefulness of incorporating the natural gradient into the Manifold GVB, we compare Algorithm 2 with a version without using the natural gradient. As shown in Fig. 2, using the natural gradient leads to a much faster and more stable convergence. Also, the Manifold GVB without the natural gradient requires a large number of samples used in estimating the gradient (5.1) (we used $S = 10,000$), compared to $S = 100$ for the Manifold GVB with the natural gradient. The CPU running time for the Manifold GVB algorithms with and without the natural gradient is 43 and 280 seconds, respectively.

Tran et al. (2019) develop a Gaussian VB algorithm where Σ is factorized as $\Sigma = B B^{\top} + D^2$ with B a vector, and term

Algorithm: Manifold Gaussian VB**Input:** Learning rate ε , momentum weight ω , prior $p(\theta)$ and likelihood $p(y|\theta)$ **Output:** An estimate μ and Σ Initialization: $\mu = \mu_0$ and $\Sigma = \Sigma_0$;Compute gradient estimates $\widehat{\nabla}_{\mu} \mathcal{L}(\lambda)$, $\widehat{\nabla}_{\Sigma} \mathcal{L}(\lambda)$;Compute natural gradients $\widehat{\nabla}_{\mu}^{\text{nat}} \mathcal{L}(\lambda) = \Sigma \widehat{\nabla}_{\mu} \mathcal{L}(\lambda)$ and $\widehat{\nabla}_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda) = \Sigma \widehat{\nabla}_{\Sigma} \mathcal{L}(\lambda) \Sigma$;Initialize the momentum: $m_{\mu} = \widehat{\nabla}_{\mu}^{\text{nat}} \mathcal{L}(\lambda)$ and $m_{\Sigma} = \widehat{\nabla}_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda)$;

stop=false;

while not stop **do** $\mu = \mu + \varepsilon m_{\mu}$; $\Sigma_{\text{old}} = \Sigma$, $\Sigma = \mathbf{R}_{\Sigma}(\varepsilon m_{\Sigma})$;**if** stopping rule is met **then**

| stop=true;

else| Compute gradient estimates $\widehat{\nabla}_{\mu} \mathcal{L}(\lambda)$, $\widehat{\nabla}_{\Sigma} \mathcal{L}(\lambda)$;| Compute natural gradient estimates $\widehat{\nabla}_{\mu}^{\text{nat}} \mathcal{L}(\lambda) = \Sigma \widehat{\nabla}_{\mu} \mathcal{L}(\lambda)$ and $\widehat{\nabla}_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda) = \Sigma \widehat{\nabla}_{\Sigma} \mathcal{L}(\lambda) \Sigma$;| Compute the momentum: $m_{\mu} = \omega m_{\mu} + (1 - \omega) \widehat{\nabla}_{\mu}^{\text{nat}} \mathcal{L}(\lambda)$, $m_{\Sigma} = \omega \Gamma_{\Sigma_{\text{old}} \rightarrow \Sigma}(m_{\Sigma}) + (1 - \omega) \widehat{\nabla}_{\Sigma}^{\text{nat}} \mathcal{L}(\lambda)$;**end****end**/* update μ *//* update Σ */**Algorithm 2:** Manifold Gaussian VB.

their algorithm NAGVAC. Figure 3 plots the lower bound estimates of the Manifold GVB and NAGVAC. Manifold GVB stopped after 921 iterations and NAGVAC stopped after 1280 iterations, and their CPU running times are 19 seconds and 12 seconds, respectively. As shown, the Manifold GVB algorithm converges quicker than NAGVAC and obtains a larger lower bound. We note, however, that NAGVAC is less computational demanding than Manifold GVB in high-dimensional settings such as deep neural networks.

To assess the training stability of the Manifold GVB and NAGVAC algorithms, we use the same initialization μ_0 and Σ_0 for both algorithms and run each for 20 different replications. The standard deviations of the estimates of μ (across the different runs, then averaged over the 25 coordinates) for NAGVAC and Manifold GVB are 0.03 and 0.01, respectively. This demonstrates that the Manifold GVB algorithm is more stable than NAGVAC. To assess their sensitivity to the initialization, in each algorithm, we now use a random initialization but fix the random seed in the updating stage. The standard deviations of the estimates of μ (across the different runs, then averaged over the 25 coordinates) for NAGVAC and Manifold GVB are 0.0074 and 0.0009, respectively. This demonstrates that the Manifold GVB algorithm is less sensitive to the initialization than NAGVAC.

Application to financial time series data

This section applies the MGVB method to analyze a financial stock return data set, and compares MGVB to Sequential Monte Carlo (SMC). We consider the GARCH model of Bollerslev (1986) for modelling the underlying volatility dynamics in the Standard & Poor 500 stock indices observed




Fig. 3 Lower bound plots for the manifold GVB algorithm (cross red) and the NAGVAC algorithm (solid blue). The lower bounds have been smoothed by a moving average with a window of size 10

from 4 Jan 1988 to 26 Feb 2007 (1000 observations). Let $\{y_t, t = 1, \dots, n\}$ be the stock returns. We consider the following GARCH model

$$y_t = \sigma_t \varepsilon_t, \quad \varepsilon_t \sim N(0, 1)$$

$$\sigma_t^2 = w + \alpha \sigma_{t-1}^2 + \beta y_{t-1}^2, \quad t = 1, 2, \dots, n.$$

The parameters θ are $w > 0$, $\alpha > 0$ and $\beta > 0$, with the constraint $\alpha + \beta < 1$ to ensure the stationarity. To impose this constraint, we parameterize α and β as $\alpha = \psi_1(1 - \psi_2)$ and $\beta = \psi_1\psi_2$ with $0 < \psi_1, \psi_2 < 1$. We use an inverse Gamma prior $\text{IG}(1, 1)$ for w and an uniform prior $U(0, 1)$ for ψ_1 and ψ_2 . Finally, we use the following transformation

$$\theta_w = \log(w), \quad \theta_{\psi_1} = \log \frac{\psi_1}{1 - \psi_1}, \quad \theta_{\psi_2} = \log \frac{\psi_2}{1 - \psi_2}$$




Fig. 4 The posterior estimates for the GARCH model parameters by SMC and MGVB




Fig. 5 Left: lower bound. Right: true posterior mean of v_{ij} v.s. their Manifold Wishart VB estimates

and work with the unconstrained parameters $\tilde{\theta} = (\theta_w, \theta_{\psi_1}, \theta_{\psi_2})$, but we will report the results in terms of the original parameters $\theta = (w, \alpha, \beta)$.

We compare the MGVB method to SMC. In this application with only three unknown parameters, SMC is applicable and can be considered as the “gold standard” as it produces an asymptotically exact approximation of the posterior $p(\tilde{\theta}|y)$. We implement the likelihood annealing SMC method of Tran et al. (2014), which is a robust SMC sampling technique that first draws samples from an easily-generated distribution and then moves these samples via annealing distributions towards the posterior distribution through weighting, resampling and Markov moving. We run SMC with 10,000 particles and the annealing distributions are adaptively designed such that the effective sample size is always at least 80%. Figure 4 plots the posterior estimates for the model parameters θ by SMC and MGVB, which shows that the MGVB estimates are almost identical to that of SMC. The CPU running time for MGVB and SMC is 10.4 and 380.5 seconds, respectively.

5.2 Manifold Wishart variational Bayes

Suppose that we are interested in approximating the posterior distribution of a covariance matrix V of size d by an inverse-Wishart distribution

$$q_\lambda(V) = \text{inverse-Wishart}(v_q, \Sigma_q).$$

The density q_λ is

$$q_\lambda(V) = \frac{|\Sigma_q|^{v_q/2}}{2^{dv_q/2} \Gamma_d(v_q/2)} |V|^{-\frac{1}{2}(v_q+d+1)} \exp\left(-\frac{1}{2}\text{trace}(\Sigma_q V^{-1})\right),$$

with

$$\Gamma_d(v) = \pi^{\frac{1}{4}d(d-1)} \prod_{j=1}^d \Gamma\left(v - \frac{j-1}{2}\right)$$

the multivariate gamma function. The gradient of $\log q_\lambda(V)$ w.r.t. v_q and Σ_q are

$$\nabla_{v_q} \log q_\lambda(V) = \frac{1}{2} \log |\Sigma_q| - \frac{1}{2} d \log 2 - \frac{1}{2} \psi_d\left(\frac{v_q}{2}\right) - \frac{1}{2} \log |V|,$$

and

$$\nabla_{\Sigma_q} \log q_\lambda(V) = \frac{1}{2} v_q \Sigma_q^{-1} - \frac{1}{2} V^{-1},$$

where $\psi_d(v) = \partial \log \Gamma_d(v) / \partial v$. From this, it is straightforward to estimate the gradients of the lower bound $\nabla_{v_q} \mathcal{L}(\lambda)$ and $\nabla_{\Sigma_q} \mathcal{L}(\lambda)$ using the score-function method. It is often much more efficient to use

$$\tilde{\nabla}_{\Sigma_q} \mathcal{L}(\lambda) = \Sigma_q \nabla_{\Sigma_q} \mathcal{L}(\lambda) \Sigma_q \quad (5.7)$$

instead of $\nabla_{\Sigma_q} \mathcal{L}(\lambda)$, which is analogous to the use of the natural gradient in (5.4). The natural-gradient-like expression in (5.7) is often found very useful in practice; without it, it is difficult for the manifold Wishart VB algorithm that only uses the Euclidean $\nabla_{\Sigma_q} \mathcal{L}(\lambda)$ to converge. As Σ_q lies on the manifold of symmetric and positive definite matrices, the retraction and vector transport in (5.5) and (5.6) can be used. To update v_q , it is recommended to use some adaptive learning method such as ADAM or AdaDelta, or an approximate natural gradient. If we set the correlation between $\nabla_{v_q} \log q_\lambda(V)$ and $\nabla_{\Sigma_q} \log q_\lambda(V)$ to be zero, then the natural gradient of the lower bound with respect to v_q is approximated by

$$\tilde{\nabla}_{v_q} \mathcal{L}(\lambda) = \left(\frac{1}{4} \psi'_d\left(\frac{v_q}{2}\right)\right)^{-1} \nabla_{v_q} \mathcal{L}(\lambda) \quad (5.8)$$

Table 1 The numbers in brackets show the posterior variances

v_{ij}	True	VB	v_{ij}	True	VB	v_{ij}	True	VB
v_{11}	0.89 (0.03)	0.91 (0.04)	v_{21}	-0.37(0.02)	-0.37(0.02)	v_{31}	0.39 (0.02)	0.39 (0.02)
v_{41}	-0.32(0.02)	-0.33(0.03)	v_{51}	0.25 (0.02)	0.26 (0.02)	v_{22}	0.83 (0.03)	0.84 (0.03)
v_{32}	-0.39(0.02)	-0.39(0.02)	v_{42}	0.29 (0.02)	0.30 (0.02)	v_{52}	-0.32(0.02)	-0.32(0.02)
v_{33}	0.93 (0.04)	0.94 (0.04)	v_{43}	-0.50(0.03)	-0.50(0.03)	v_{53}	0.35 (0.02)	0.35 (0.02)
v_{44}	1.18 (0.06)	1.19 (0.07)	v_{54}	-0.73(0.04)	-0.74(0.04)	v_{55}	0.96 (0.04)	0.97 (0.04)

The manifold Wishart Variational Bayes stopped after 75 iterations, with $S = 1000$ Monte Carlo samples used for estimating the score-function gradient of the lower bound. The gradients in (5.7) and (5.8) were used

where $\psi'_d(v) := \partial \psi_d(v) / \partial v$.

A numerical example

Data of size n is generated from $\mathcal{N}_d(0, V)$, where the elements of the true covariance matrix V are $v_{ij} = (-0.5)^{|i-j|}$. An inverse-Wishart(v_0, S_0) prior is used for V with $v_0 = d$ and $S_0 = 0.01I_d$. It is easy to see that the posterior distribution of V is inverse-Wishart with the degree of freedom $v = n + d$ and scale matrix $S = S_0 + \sum_i y_i y_i^\top$. We run the manifold Wishart VB algorithm with the initial value $v_q = n$ and $\Sigma_q = n \times S_n$, where S_n is the sample covariance matrix of the y_i .

In the first simulation, we consider $n = 50$ and $d = 5$. The estimation result is summarized in Table 1. We consider the second simulation with $n = 500$ and $d = 50$. The lower bound and a plot of the true posterior means of the v_{ij} v.s. their VB estimates are shown in Fig. 5. These results show that the manifold Wishart VB algorithm appears to work effectively and efficiently in this example.

6 Conclusion

We proposed a manifold-based Variational Bayes algorithm that takes into account both information geometry and geometric structure of the constraint parameter space. The algorithm is provably convergent for the case when the objective function either non-convex or strongly retraction-convex. Our numerical experiments demonstrate that the proposed algorithm converges quickly, is stable and compares favourably to the existing VB methods. An interesting research direction is to develop a Manifold VB method for the Sylvester normalizing flows of Berg et al. (2018). We leave it as an interesting project for future research.

Acknowledgements Tran would like to thank Prof. Pierre-Antoine Absil for useful discussions, and Nghia Nguyen for his help with the SMC implementation. The research of Tran was partially supported by ARC grant DP200103015.

Appendix

Proof of Lemma 3.1 By definition, the Riemannian gradient of $\bar{\mathcal{L}}$ at $\lambda \in \bar{\mathcal{M}}$, denoted by $\text{grad}\bar{\mathcal{L}}(\lambda)$, is the unique tangent vector in $T_\lambda \bar{\mathcal{M}}$ such that

$$\langle \text{grad}\bar{\mathcal{L}}(\lambda), \xi_\lambda \rangle = D\bar{\mathcal{L}}(\lambda)[\xi_\lambda], \quad \forall \xi_\lambda \in T_\lambda \bar{\mathcal{M}}.$$

That is

$$\text{grad}\bar{\mathcal{L}}(\lambda)^\top I_F(\lambda) \xi_\lambda = \nabla_\lambda \bar{\mathcal{L}}(\lambda)^\top \xi_\lambda \quad \forall \xi_\lambda \in T_\lambda \bar{\mathcal{M}}. \quad (6.1)$$

The natural gradient $\nabla_\lambda^{\text{nat}} \bar{\mathcal{L}}(\lambda)$ in (3.2) satisfies (6.1). Indeed, due to the symmetry of $I_F(\lambda)$, we have

$$\begin{aligned} \text{grad}\bar{\mathcal{L}}(\lambda)^\top I_F(\lambda) \xi_\lambda &= (I_F^{-1}(\lambda) \nabla_\lambda \bar{\mathcal{L}}(\lambda))^\top I_F(\lambda) \xi_\lambda \\ &= \nabla_\lambda \bar{\mathcal{L}}(\lambda)^\top (I_F^{-1}(\lambda))^\top (I_F(\lambda) \xi_\lambda) \\ &= \nabla_\lambda \bar{\mathcal{L}}(\lambda)^\top \xi_\lambda. \end{aligned}$$

If $\zeta_\lambda \in T_\lambda \bar{\mathcal{M}}$ also satisfies (6.1), then

$$(\zeta_\lambda - \text{grad}\bar{\mathcal{L}}(\lambda))^\top I_F(\lambda) \xi_\lambda = 0 \quad \forall \xi_\lambda \in T_\lambda \bar{\mathcal{M}}$$

which implies that $\zeta_\lambda = \text{grad}\bar{\mathcal{L}}(\lambda)$. \square

Proof of Lemma 3.2 Much of our proof is taken from Tagare (2011). The idea of the proof is that we want to find a vector in $T_W \mathcal{M}$ that represents the action of the differential $D\mathcal{L}(W)$ via the natural gradient G . Let $W_\perp \in \text{Mat}(n, n-p)$ such that its columns together with the columns of W form an orthonormal basis for \mathbb{R}^n . As $[W, W_\perp]$ is an orthogonal matrix, for any $U \in \text{Mat}(n, p)$, there exists a $C \in \text{Mat}(n, p)$ such that $U = [W, W_\perp]C$. Write $C = \begin{pmatrix} U_W \\ U_{W_\perp} \end{pmatrix}$ with U_W the $p \times p$ -matrix formed by the first p rows of C , U_{W_\perp} the $(n-p) \times p$ matrix formed by the last $n-p$ rows of C . That is, any matrix $U \in \text{Mat}(n, p)$ can be written as

$$\begin{aligned} U &= WU_W + W_\perp U_{W_\perp}, \quad U_W \in \text{Mat}(p, p), \\ U_{W_\perp} &\in \text{Mat}(n-p, p). \end{aligned}$$

If $Z = WZ_W + W_\perp Z_{W_\perp} \in T_W \mathcal{M}$, then $Z^\top W + W^\top Z = 0_{p \times p}$ implies that $Z_W + Z_W^\top = 0$. So $T_W \mathcal{M}$ is a subset of the set

$$\{Z = WZ_W + W_\perp Z_{W_\perp} : Z_W = -Z_W^\top, Z_{W_\perp} \in \text{Mat}(n-p, p)\}.$$

It's easy to check that this set is also a subset of $T_W \mathcal{M}$. We arrive at an alternative representation of the tangent space $T_W \mathcal{M}$

$$T_W \mathcal{M} = \{WZ_W + W_\perp Z_{W_\perp} : Z_W \in \text{Mat}(p, p), Z_{W_\perp} \in \text{Mat}(n-p, p), Z_W = -Z_W^\top\}. \quad (6.2)$$

We want to find a vector in $T_W \mathcal{M}$ that represents the action of the differential $D\mathcal{L}(W)$ on $T_W \mathcal{M}$, with $D\mathcal{L}(W)$ characterized by G ; i.e. find $U = WU_W + W_\perp U_{W_\perp} \in T_W \mathcal{M}$, $U_W = -U_W^\top$ such that

$$\langle U, Z \rangle_{\text{Euc}} = D\mathcal{L}(W)[Z], \quad \forall Z \in T_W \mathcal{M}. \quad (6.3)$$

As the gradient $G = \text{grad} \bar{\mathcal{L}}(W) \in \text{Mat}(n, p)$, it can be written as $G = WG_W + W_\perp G_{W_\perp}$. Based on the natural gradient G ,

$$\begin{aligned} D\mathcal{L}(W)[Z] &= \text{tr}(G^\top Z) \\ &= \text{tr}(G_W^\top Z_W) + \text{tr}(G_{W_\perp}^\top Z_{W_\perp}) \\ &= \text{tr}(\text{skew}(G_W)^\top Z_W) + \text{tr}(G_{W_\perp}^\top Z_{W_\perp}), \end{aligned} \quad (6.4)$$

where we have used the fact that $G_W = \text{skew}(G_W) + \text{sym}(G_W)$ with $\text{sym}(G_W) = (G_W + G_W^\top)/2$, and that $\text{tr}(\text{sym}(G_W)^\top Z_W) = 0$. We have that

$$\langle U, Z \rangle_{\text{Euc}} = \text{tr}(U_W^\top Z_W) + \text{tr}(U_{W_\perp}^\top Z_{W_\perp}). \quad (6.5)$$

Comparing (6.4) and (6.5) gives

$$U = W \text{skew}(G_W) + W_\perp G_{W_\perp}.$$

As $G_W = W^\top G$ and $W_\perp G_{W_\perp} = G - WG_W = (I_n - WW^\top)G$,

$$U = W \text{skew}(W^\top G) + (I_n - WW^\top)G.$$

From (6.3), U is the Riemannian gradient of \mathcal{L} . This completes the proof. \square

Proof of Theorem 1 (i) First, we consider the case \mathcal{L} is not assumed to be convex. Consider the update

$$\begin{aligned} \lambda_{t+1} &= R_{\lambda_t}(-Y_t), \\ Y_{t+1} &= \zeta \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) + \gamma \nabla \mathcal{L}(\lambda_{t+1}) + \gamma \Delta M_{t+1}, \end{aligned} \quad (6.6)$$

where $\zeta, \gamma \in (0, 1)$, ΔM_{t+1} is a martingale difference. Since $|\nabla \mathcal{L}| + |\Delta M| \leq b_{\mathcal{L}}$ a.s., we have $|Y_t| \leq \frac{2\gamma b_{\mathcal{L}}}{1-\zeta}$. This can be proved as follows. First for the rest of this proof, we use $O(a)$ to denote vector/scalar with $|O(a)| \leq a$. Let $P_{x \rightarrow y}$ be a parallel translation. From Lemma 6 Huang et al. (2015a), there exists a constant $a_1 > 0$ such that

$$|\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) - P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)| = O(a_1)|Y_t|^2, \quad (6.7)$$

which, together with $|P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)| = |Y_t|$ implies

$$\begin{aligned} |\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)| &\leq |P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)| + O(a_1)|Y_t|^2 \\ &= (1 + O(a_1)|Y_t|)|Y_t|. \end{aligned} \quad (6.8)$$

If we take the norm of the second equation in (6.6), and using (6.8) we have

$$\begin{aligned} |Y_{t+1}| &\leq \zeta |\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)| + \gamma (|\nabla \mathcal{L}(\lambda_{t+1})| + |\Delta M_{t+1}|) \\ &\leq \zeta (1 + O(a_1)|Y_t|)|Y_t| + \gamma (|\nabla \mathcal{L}(\lambda_{t+1})| + |\Delta M_{t+1}|) \\ &\leq (\zeta + a_1|Y_t|)|Y_t| + \gamma b_{\mathcal{L}}. \end{aligned} \quad (6.9)$$

With $\gamma < \frac{(1-\zeta)^2}{4a_1 b_{\mathcal{L}}}$ and $|Y_0| \leq \frac{2\gamma b_{\mathcal{L}}}{1-\zeta}$ if ζ, γ are sufficiently small, using induction, we can show that $|Y_t| \leq \frac{2\gamma b_{\mathcal{L}}}{1-\zeta}$. Indeed assuming $|Y_t| \leq \frac{2\gamma b_{\mathcal{L}}}{1-\zeta}$, using (6.9), we have

$$\begin{aligned} |Y_{t+1}| &\leq \left(\zeta + a_1 \frac{2\gamma b_{\mathcal{L}}}{1-\zeta} \right) \frac{2\gamma b_{\mathcal{L}}}{1-\zeta} + \gamma b_{\mathcal{L}} \\ &\leq \left(\zeta + a_1 \frac{2b_{\mathcal{L}}}{1-\zeta} \frac{(1-\zeta)^2}{4a_1 b_{\mathcal{L}}} \right) \frac{2\gamma b_{\mathcal{L}}}{1-\zeta} + \gamma b_{\mathcal{L}} \\ &\leq \frac{1+\zeta}{1-\zeta} \gamma b_{\mathcal{L}} + \gamma b_{\mathcal{L}} \\ &\leq \frac{2\gamma b_{\mathcal{L}}}{1-\zeta}. \end{aligned}$$

Next, applying the Taylor expansion to the function $t \mapsto \mathcal{L}(R_\lambda(t\eta))$ and using the Lipschitz continuity of $\nabla \mathcal{L}$ as well as the smoothness of $R_\lambda(\cdot)$ in \mathcal{S} , there exists a constant L such that

$$\mathcal{L}(\lambda_{t+1}) = \mathcal{L}(\lambda_t) - \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + O\left(\frac{L}{2}\right)|Y_t|^2. \quad (6.10)$$

From (6.6), squaring both sides of the second equation we obtain

$$\begin{aligned} |Y_{t+1}|^2 &= \zeta^2 |\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)|^2 + 2\gamma \zeta \langle \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &\quad + \gamma^2 |\nabla \mathcal{L}(\lambda_{t+1})|^2 + \Delta N_t + \gamma^2 |\Delta M_{t+1}|^2, \end{aligned} \quad (6.11)$$

where

$$\Delta N_t = 2\langle \zeta \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) + \gamma \nabla \mathcal{L}(\lambda_{t+1}), \Delta M_{t+1} \rangle.$$

From (6.8), with γ small enough so that $\frac{a_1\gamma b_{\mathcal{L}}}{1-\zeta} \leq 1$, we have

$$\begin{aligned} |\Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t)|^2 &= |Y_t|^2 + O(2a_1 + a_1^2|Y_t|)|Y_t|^3 \\ &= |Y_t|^2 + O(3a_1)|Y_t|^3 \\ &= |Y_t|^2 + O\left(\frac{6a_1b_{\mathcal{L}}\gamma}{1-\zeta}\right)|Y_t|^2. \end{aligned}$$

Consider the second term without the constant in (6.11), we have

$$\begin{aligned} &\langle \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &= \underbrace{\langle P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle}_{(I)} \\ &\quad + \underbrace{\langle \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t) - P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle}_{(II)}. \end{aligned}$$

Recall that $P_{x \rightarrow y}$ is an isometry, see, for example Absil et al. (2009). As a result we have

$$\langle P_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle = \langle Y_t, P_{\lambda_t \rightarrow \lambda_{t+1}}^{-1}(\nabla \mathcal{L}(\lambda_{t+1})) \rangle.$$

In view of the fundamental theorem of calculus (Huang et al. 2015a, Lemma 8), we have $P_{\lambda_t \rightarrow \lambda_{t+1}}^{-1} \nabla \mathcal{L}(\lambda_{t+1}) = \nabla \mathcal{L}(\lambda_t) + O(a_2)|Y_t|$ for some constant $a_2 > 0$. Then, for $\gamma \leq \frac{1-\zeta}{a_2}$, we have

$$\begin{aligned} |\nabla \mathcal{L}(\lambda_{t+1})|^2 &= |P_{\lambda_t \rightarrow \lambda_{t+1}}^{-1} \nabla \mathcal{L}(\lambda_{t+1})|^2 \\ &= |\nabla \mathcal{L}(\lambda_t)|^2 + 2O(a_2)|Y_t||\nabla \mathcal{L}(\lambda_t)| + O(a_2)^2|Y_t|^2 \\ &= |\nabla \mathcal{L}(\lambda_t)|^2 + O(4a_2b_{\mathcal{L}})|Y_t| \quad (\text{since } |Y_t| \leq \frac{\gamma b_{\mathcal{L}}}{1-\zeta} \leq 2b_{\mathcal{L}}) \\ &= |\nabla \mathcal{L}(\lambda_t)|^2 + O\left(\frac{8a_2b_{\mathcal{L}}^2}{1-\zeta}\right)\gamma, \end{aligned} \quad (6.12)$$

and

$$\begin{aligned} (I) &= \langle Y_t, P_{\lambda_t \rightarrow \lambda_{t+1}}^{-1}(\nabla \mathcal{L}(\lambda_{t+1})) \rangle \\ &= \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + \langle Y_t, P_{\lambda_t \rightarrow \lambda_{t+1}}^{-1}(\nabla \mathcal{L}(\lambda_{t+1})) - \nabla \mathcal{L}(\lambda_t) \rangle \\ &= \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + O(a_2)|Y_t|^2. \end{aligned} \quad (6.13)$$

For (II), since $\nabla \mathcal{L}$ is a.s. bounded by $b_{\mathcal{L}}$, we derive from (6.8) that

$$(II) = O(a_1b_{\mathcal{L}})|Y_t|^2. \quad (6.14)$$

As a result of (I) and (II), we have

$$\begin{aligned} \langle \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle &= \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ &\quad + O(a_1b_{\mathcal{L}} + a_2)|Y_t|^2. \end{aligned} \quad (6.15)$$

Plugging (6.15) into (6.11),

$$\begin{aligned} |Y_{t+1}|^2 &= \zeta^2|Y_t|^2 + 2\gamma\zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ &\quad + \gamma^2|\nabla \mathcal{L}(\lambda_{t+1})|^2 + O(a_3)\gamma|Y_t|^2 \\ &\quad + \Delta N_t + \gamma^2 \Delta M_{t+1}^2. \end{aligned} \quad (6.16)$$

where $a_3 = \frac{6a_1b_{\mathcal{L}}}{1-\zeta} + 2(a_1b_{\mathcal{L}} + a_2)\zeta$.

Next, using the second equation in (6.6) and (6.15), we have

$$\begin{aligned} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle &= \\ &= \zeta \langle \Gamma_{\lambda_t \rightarrow \lambda_{t+1}}(Y_t), \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &\quad + \gamma|\nabla \mathcal{L}(\lambda_{t+1})|^2 + \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &= \zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + O(\zeta(a_1b_{\mathcal{L}} + a_2))|Y_t|^2 \\ &\quad + \gamma|\nabla \mathcal{L}(\lambda_{t+1})|^2 + \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &= \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle - (1-\zeta)\langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &\quad + O(\zeta(a_1b_{\mathcal{L}} + a_2))|Y_t|^2 + \gamma|\nabla \mathcal{L}(\lambda_{t+1})|^2. \end{aligned} \quad (6.17)$$

Multiplying (6.16), (6.17) with A (to be chosen later) and $-\frac{1}{1-\zeta}$ respectively and then adding to (6.10) we have

$$\begin{aligned} &\mathcal{L}(\lambda_{t+1}) + A|Y_{t+1}|^2 - \frac{1}{1-\zeta} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ &\leq \mathcal{L}(\lambda_t) + A|Y_t|^2 - \frac{1}{1-\zeta} \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ &\quad - \left(A(1-\zeta^2) - AO(a_3)\gamma - O\left(\frac{L}{2} + \zeta \frac{a_1b_{\mathcal{L}} + a_2}{1-\zeta}\right) \right) |Y_t|^2 \\ &\quad + 2A\gamma\zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ &\quad - \frac{\gamma}{1-\zeta} (1-(1-\zeta)\gamma A) |\nabla \mathcal{L}(\lambda_{t+1})|^2 \\ &\quad + A\Delta N + \gamma^2 A \Delta M_{t+1}^2 - \frac{1}{1-\zeta} \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle. \end{aligned} \quad (6.18)$$

Note that,

$$\begin{aligned} |2A\gamma\zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle| &\leq |2A\gamma \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle| \\ &\leq 2A\gamma(1-\zeta)|Y_t|^2 + A \frac{\gamma|\nabla \mathcal{L}(\lambda_t)|^2}{2(1-\zeta)} \\ &\leq 2A\gamma(1-\zeta)|Y_t|^2 + A \frac{\gamma|\nabla \mathcal{L}(\lambda_{t+1})|^2}{2(1-\zeta)} \\ &\quad + AO\left(\zeta \frac{8a_2b_{\mathcal{L}}^2}{(1-\zeta)^2}\right) \gamma^2 \quad \text{by (6.12)}. \end{aligned}$$

(In the second inequality above, we have used the fact that $|< p, q >| \leq \frac{p^2}{2\epsilon} + \frac{\epsilon q^2}{2}$ with $\epsilon = (1 - \zeta)$). Plugging back to (6.18), we obtain

$$\begin{aligned} & \mathcal{L}(\lambda_{t+1}) + A|Y_{t+1}|^2 - \frac{1}{1-\zeta} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ & \leq \mathcal{L}(\lambda_t) + A|Y_t|^2 - \frac{1}{1-\zeta} \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle - |Y_t|^2 \left(A(1-\zeta^2) \right. \\ & \quad \left. - 2(1-\zeta)\gamma A - A O(a_3)\gamma - O\left(\frac{L}{2} + \zeta \frac{a_1 b_{\mathcal{L}} + a_2}{1-\zeta}\right) \right) \\ & \quad - \frac{\gamma}{1-\zeta} \left(\frac{1}{2} - (1-\zeta)A\gamma \right) |\nabla \mathcal{L}(\lambda_{t+1})|^2 + \gamma^2 A \Delta M_{t+1}^2 \\ & \quad + A \Delta N + A \zeta \frac{8a_2 b_{\mathcal{L}}^2}{(1-\zeta)^2} \gamma^2 - \frac{1}{1-\zeta} \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle. \end{aligned} \quad (6.19)$$

Now select $A = \frac{4}{(1-\zeta^2)} \left(\frac{L}{2} - \zeta \frac{a_1 b_{\mathcal{L}} + a_2}{1-\zeta} \right)$ and let γ satisfy

$$\begin{aligned} & 2(1-\zeta)\gamma - a_3\gamma \\ & \leq \frac{(1-\zeta^2)}{4}; \quad (1-\zeta)\gamma A < \frac{1}{4} \quad \text{and} \quad \gamma \leq \frac{1-\zeta}{a_2}, \end{aligned}$$

taking the expectation of both sides of (6.19), we have

$$\begin{aligned} & \mathbb{E} \left(\mathcal{L}(\lambda_{t+1}) + A|Y_{t+1}|^2 - \frac{1}{1-\zeta} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \right) \\ & \leq \mathbb{E} \left(\mathcal{L}(\lambda_t) + A|Y_t|^2 - \frac{1}{1-\zeta} \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \right) \\ & \quad - \frac{A(1-\zeta^2)}{4} \mathbb{E}|Y_t|^2 - \frac{\gamma}{(1-\zeta)} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \\ & \quad + \gamma^2 \left(A b_{\mathcal{L}}^2 + A \zeta \frac{8a_2 b_{\mathcal{L}}^2}{(1-\zeta)^2} \right). \end{aligned} \quad (6.20)$$

Take the sum, we have

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E} \left(\frac{A(1-\zeta^2)}{4} |Y_t|^2 + \frac{\gamma}{(1-\zeta)} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \right) \\ & \leq (\mathcal{L}(\lambda_0) - \mathcal{L}(\lambda_*)) \\ & \quad - \frac{1}{1-\zeta} \langle Y_{T+1}, \nabla \mathcal{L}(\lambda_{T+1}) \rangle \\ & \quad + \gamma^2 \left(A b_{\mathcal{L}}^2 + A \zeta \frac{8a_2 b_{\mathcal{L}}^2}{(1-\zeta)^2} \right) T \\ & \leq C(\gamma^2 T + 1). \end{aligned}$$

Then

$$\sum_{t=1}^T \mathbb{E} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \leq C(\gamma T + \frac{1}{\gamma}).$$

Hence

$$\begin{aligned} T \min_{t \in [1, T]} \mathbb{E} |\nabla \mathcal{L}(\lambda_{t+1})|^2 & = \sum_{t=1}^T \min_{t \in [1, T]} \mathbb{E} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \\ & \leq C(\gamma T + \frac{1}{\gamma}). \end{aligned}$$

If we choose $\gamma^2 = \frac{1}{T}$, then

$$\min_{t \in [1, T]} \mathbb{E} |\nabla \mathcal{L}(\lambda_{t+1})|^2 \leq \frac{C}{\sqrt{T}}.$$

(ii) Now we assume that \mathcal{L} is strongly retraction convex with

$$\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*) \leq \tilde{\mu} |\nabla \mathcal{L}(\lambda_t)|^2, \quad (6.21)$$

for some $\tilde{\mu} > 0$ (see Remark 4.1). Let $\rho := 1 - \frac{\gamma}{4\tilde{\mu}(1-\zeta)} > \frac{1+\zeta}{2}$ when γ is small. Rewrite (6.17)

$$\begin{aligned} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle & = \rho \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle - (\rho - \zeta) \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ & \quad + \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ & \quad + O(a_1 b_{\mathcal{L}} + a_2) |Y_t|^2 + \gamma |\nabla \mathcal{L}(\lambda_{t+1})|^2. \end{aligned} \quad (6.22)$$

Multiplying (6.16), (6.22) with A (to be chosen later) and $-\frac{1}{\rho-\zeta}$ respectively and then adding to (6.10) we have

$$\begin{aligned} & \mathcal{L}(\lambda_{t+1}) - \mathcal{L}(\lambda_*) + A|Y_{t+1}|^2 - \frac{1}{\rho-\zeta} \langle Y_{t+1}, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ & \leq \mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*) + A \zeta^2 |Y_t|^2 - \frac{\rho}{\rho-\zeta} \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \\ & \quad + 2A\gamma \zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle + A\gamma O(a_3) Y_t^2 + A\gamma^2 M_t^2 \\ & \quad + O\left(\frac{L}{2}\right) |Y_t|^2 + O\left(\frac{a_1 b_{\mathcal{L}} + a_2}{\rho-\zeta}\right) |Y_t|^2 + A \Delta N_t \\ & \quad - \frac{1}{\rho-\zeta} \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle \\ & \quad - \frac{\gamma}{\rho-\zeta} (1 - (\rho - \zeta)\gamma A) |\nabla \mathcal{L}(\lambda_{t+1})|^2 \\ & \leq \rho \left(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*) + A|Y_t|^2 - \frac{1}{\rho-\zeta} \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle \right) \\ & \quad + \left((1-\rho)(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*)) \right. \\ & \quad \left. - \frac{1}{\rho-\zeta} (1 - (\rho - \zeta)\gamma) |\nabla \mathcal{L}(\lambda_t)|^2 \right) + A \Delta N_t + A b_{\mathcal{L}}^2 \gamma^2 \\ & \quad - \left(A(\rho - \zeta^2) - A\gamma a_3 - \frac{L}{2} - \frac{a_1 b_{\mathcal{L}} + a_2}{\rho-\zeta} \right) |Y_t|^2 \\ & \quad + 2A\gamma \zeta \langle Y_t, \nabla \mathcal{L}(\lambda_t) \rangle - \frac{1}{\rho-\zeta} \gamma \langle \Delta M_t, \nabla \mathcal{L}(\lambda_{t+1}) \rangle. \end{aligned} \quad (6.23)$$

We have

$$\begin{aligned} |2A\gamma\zeta\langle Y_t, \nabla\mathcal{L}(\lambda_t) \rangle| &\leq 2A\gamma(\rho - \zeta)|Y_t|^2 + A\frac{\gamma|\nabla\mathcal{L}(\lambda_t)|^2}{2(\rho - \zeta)} \\ &\leq 2A\gamma(\rho - \zeta)|Y_t|^2 + A\frac{\gamma|\nabla\mathcal{L}(\lambda_{t+1})|^2}{2(\rho - \zeta)} \\ &\quad + O\left(\frac{Aa_2b_{\mathcal{L}}^2}{(1 - \zeta)(\rho - \zeta)}\right)\gamma^2. \end{aligned}$$

Plugging this into (6.23), we have

$$\begin{aligned} \mathcal{L}(\lambda_{t+1}) - \mathcal{L}(\lambda_*) + A|Y_{t+1}|^2 - \frac{1}{\rho - \zeta}\langle Y_{t+1}, \nabla\mathcal{L}(\lambda_{t+1}) \rangle \\ \leq \rho\left(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*) + A|Y_t|^2 - \frac{1}{\rho - \zeta}\langle Y_t, \nabla\mathcal{L}(\lambda_t) \rangle\right) \\ + \left((1 - \rho)(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*))\right. \\ \left. - \frac{1}{\rho - \zeta}\left(\frac{1}{2} - (\rho - \zeta)\gamma\right)|\nabla\mathcal{L}(\lambda_t)|^2\right) - \left(A(\rho - \zeta^2)\right. \\ \left. - A\gamma a_3 - \frac{L}{2} - \frac{a_1b_{\mathcal{L}} + a_2}{\rho - \zeta} - 2A\gamma(\rho - \zeta)\right)|Y_t|^2 \\ + A\Delta N_t - \frac{1}{\rho - \zeta}\gamma\langle \Delta M_t, \nabla\mathcal{L}(\lambda_{t+1}) \rangle \\ + \left(\frac{Aa_2b_{\mathcal{L}}^2}{(1 - \zeta)(\rho - \zeta)} + Ab_{\mathcal{L}}^2\right)\gamma^2. \end{aligned} \quad (6.24)$$

We deduce from (6.21) that

$$\begin{aligned} (1 - \rho)(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*)) &= \frac{\gamma}{4\tilde{\mu}(1 - \zeta)}(\mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*)) \\ &\leq \frac{\gamma}{4(1 - \zeta)}|\nabla\mathcal{L}(\lambda_t)|^2 \\ &\leq \frac{1}{4(\rho - \zeta)}|\nabla\mathcal{L}(\lambda_t)|^2. \end{aligned} \quad (6.25)$$

On the other hand due to (6.12) we have

$$\begin{aligned} \frac{\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_{t+1})|^2 \\ = \frac{\rho\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_{t+1})|^2 + \frac{(1 - \rho)\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_{t+1})|^2 \\ = \frac{\rho\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_t)|^2 + \frac{(1 - \rho)\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_{t+1})|^2 \\ + O\left(\frac{8a_2b_{\mathcal{L}}^2}{1 - \zeta}\right)\frac{\gamma\rho}{\rho - \zeta} \\ = \frac{\rho\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_t)|^2 + \frac{(1 - \rho)\gamma}{\rho - \zeta}|\nabla\mathcal{L}(\lambda_{t+1})|^2 \\ + O\left(\frac{8a_2b_{\mathcal{L}}^2\rho}{(1 - \zeta)^2(\rho - \zeta)}\right)\gamma^2. \end{aligned} \quad (6.26)$$

Adding (6.24) and (6.26) and using (6.25) we have for

$$\begin{aligned} V_t &:= \mathcal{L}(\lambda_t) - \mathcal{L}(\lambda_*) + A|Y_t|^2 \\ &\quad - \frac{1}{\rho - \zeta}\langle Y_t, \nabla\mathcal{L}(\lambda_t) \rangle + \frac{\gamma}{(\rho - \zeta)}|\nabla\mathcal{L}(\lambda_t)|^2, \end{aligned}$$

that

$$\begin{aligned} V_{t+1} &\leq \rho V_t - \left(\frac{1}{\rho - \zeta}\left(\frac{1}{4} - (\rho - \zeta)\gamma\right)\right)|\nabla\mathcal{L}(\lambda_t)|^2 \\ &\quad - |Y_t|^2\left(A(\rho - \zeta^2) - A\gamma a_3 - \frac{L}{2} - \frac{a_1b_{\mathcal{L}} + a_2}{\rho - \zeta} - 2A\gamma(\rho - \zeta)\right) \\ &\quad + \left(\frac{Aa_2b_{\mathcal{L}}^2}{(1 - \zeta)(\rho - \zeta)} + Ab_{\mathcal{L}}^2 + \frac{8a_2b_{\mathcal{L}}^2\rho}{(1 - \zeta)^2(\rho - \zeta)}\right)\gamma^2 \\ &\quad + A\Delta N_t - \frac{1}{\rho - \zeta}\gamma\langle \Delta M_t, \nabla\mathcal{L}(\lambda_{t+1}) \rangle. \end{aligned} \quad (6.27)$$

Since \mathcal{L} is strongly retraction-convex, there exists $\bar{\mu} > 0$ such that $|\nabla\mathcal{L}(\lambda)|^2 \leq \tilde{L}^2|R_{\lambda^*}^{-1}(\lambda)|^2 \leq \tilde{L}^2\bar{\mu}(\mathcal{L}(\lambda_T) - \mathcal{L}(\lambda_*))$. As a result, if $A > \frac{2\tilde{L}^2\bar{\mu}}{(1 - \zeta)^2} > \frac{\tilde{L}^2\bar{\mu}}{2(\rho - \zeta)^2}$ then we have from Cauchy's inequality that

$$\begin{aligned} V_t &= A|Y_t|^2 - \frac{1}{\rho - \zeta}\langle Y_t, \nabla\mathcal{L}(\lambda_t) \rangle + (\mathcal{L}(\lambda_T) - \mathcal{L}(\lambda_*)) \\ &\geq \frac{1}{2}(\mathcal{L}(\lambda_T) - \mathcal{L}(\lambda_*)). \end{aligned}$$

Let A and γ be such that

$$A(\rho - \zeta^2) - A\gamma a_3 - 2A\gamma(\rho - \zeta) - \frac{L}{2} - \frac{a_1b_{\mathcal{L}} + a_2}{\rho - \zeta} > 0,$$

and

$$A > \rho, A > \frac{2\tilde{L}^2\bar{\mu}}{(1 - \zeta)^2}, \frac{1}{\rho - \zeta}\left(\frac{1}{4} - (\rho - \zeta)\gamma\right) > 0,$$

we have

$$\mathbb{E}V_{t+1} \leq \rho\mathbb{E}V_t + C\gamma^2.$$

As a result, recall the definition of ρ , we have

$$\mathbb{E}V_t \leq \rho^t C_0 + \frac{C}{1 - \rho}\gamma^2 \leq \left(1 - \frac{\tilde{\mu}\gamma}{2(1 - \zeta)}\right)^t C_0 + C\gamma^2,$$

where $C_0 = \mathbb{E}V_0$.

With $\gamma = \frac{1}{T^\varepsilon}$, we have

$$\mathbb{E}V_T \leq \left((1 - \tilde{\mu}T^{1-\varepsilon})^{T^{1-\varepsilon}}\right)^{T^\varepsilon} + CT^{2\varepsilon-2}.$$

When T is large then

$$\left((1 - \tilde{\mu} T^{1-\varepsilon}) T^{1-\varepsilon} \right)^{T^\varepsilon} \approx e^{-\frac{T^\varepsilon}{\tilde{\mu}}} < T^{2\varepsilon-2}.$$

Thus for $C_\varepsilon = \max\{C_0, C\}$,

$$\mathbb{E}V_T \leq C_\varepsilon T^{2\varepsilon-2}.$$

we have

$$\frac{1}{\tilde{L}^2 \tilde{\mu}} \mathbb{E}|\nabla \mathcal{L}(\lambda_t)|^2 \leq \frac{1}{2} \mathbb{E} \mathcal{L}(\lambda_T) - \mathcal{L}(\lambda_*) \leq \mathbb{E}V_T \leq C_\varepsilon T^{2\varepsilon-2}.$$

□

References

- Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, New York (2009)
- Amari, S.: Natural gradient works efficiently in learning. *Neural Comput.* **10**(2), 251–276 (1998)
- Béginneul, G., Ganea, O.-E.: Riemannian adaptive optimization methods. (2018) arXiv preprint [arXiv:1810.00760](https://arxiv.org/abs/1810.00760)
- Berg, R.v.d., Hasenclever, L., Tomczak, J. M., Welling, M.: Sylvester normalizing flows for variational inference. (2018) arXiv preprint [arXiv:1803.05649](https://arxiv.org/abs/1803.05649)
- Bollerslev, T.: Generalized autoregressive conditional heteroskedasticity. *J. Econom.* **31**(3), 307–327 (1986)
- Bonnabel, S.: Stochastic gradient descent on Riemannian manifolds. *IEEE Trans. Autom. Control* **58**(9), 2217–2229 (2013)
- Boumal, N., Mishra, B., Absil, P.-A., Sepulchre, R.: Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* **15**, 1455–1459 (2014)
- Edelman, A., Arias, T., Smith, A.: The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.* **20**(2), 303–353 (1998)
- Hoffman, M.D., Blei, D.M., Wang, C., Paisley, J.: Stochastic variational inference. *J. Mach. Learn. Res.* **14**, 1303–1347 (2013)
- Huang, W., Absil, P.-A., Gallivan, K.A.: A Riemannian symmetric rank-one trust-region method. *Math. Program.* **150**(2), 179–216 (2015a)
- Huang, W., Gallivan, K.A., Absil, P.-A.: A Broyden class of quasi-Newton methods for Riemannian optimization. *SIAM J. Optim.* **25**(3), 1660–1685 (2015b)
- Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. *Mach. Learn.* **37**(2), 183–233 (1999)
- Kasai, H., Jawanpuria, P., Mishra, B.: Adaptive stochastic gradient algorithms on riemannian manifolds. (2019) arXiv preprint [arXiv:1902.01144](https://arxiv.org/abs/1902.01144)
- Khan, M.E., Lin, W.: Conjugate-computation variational inference: Converting variational inference in non-conjugate models to inferences in conjugate models. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). (2017)
- Kingma, D. P., Welling, M.: Auto-encoding variational bayes. (2013) arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114)
- Kobayashi, S., Nomizu, K.: *Foundations of Differential Geometry, Vols. 1 and 2*. Wiley, New York (1969)
- Longford, N.T.: A fast scoring algorithm for maximum likelihood estimation in unbalanced mixed models with nested random effects. *Biometrika* **74**(4), 817–827 (1987)
- Manton, J.H.: Optimization algorithms exploiting unitary constraints. *IEEE Trans. Signal Process.* **50**(3), 635–650 (2002)
- Mardia, K.V., Marshall, R.J.: Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika* **71**(1), 135–146 (1984)
- Martens, J.: New insights and perspectives on the natural gradient method. *J. Mach. Learn. Res.* **21**(146), 1–76 (2020)
- Ong, V.M., Nott, D.J., Tran, M.-N., Sisson, S.A., Drovandi, C.C.: Variational Bayes with synthetic likelihood. *Stat. Comput.* **28**(4), 971–988 (2018a)
- Ong, V.M.-H., Nott, D.J., Smith, M.S.: Gaussian variational approximation with factor covariance structure. *J. Comput. Graph. Stat.* **27**(3), 465–478 (2018b)
- Rao, C.R.: Information and accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc.* **37**, 81–91 (1945)
- Roy, S.K., Harandi, M.: Constrained stochastic gradient descent: The good practice. In: *Proceedings of the International Conference on Digital Image Computing: Techniques and Applications (DICTA)*. (2017)
- Sato, H., Aihara, K.: Cholesky QR-based retraction on the generalized Stiefel manifold. *Comput. Optim. Appl.* **72**(2), 293–308 (2019)
- Sato, M.: Online model selection based on the variational bayes. *Neural Comput.* **13**(7), 1649–1681 (2001)
- Tagare, H.D.: Notes on optimization on Stiefel manifolds. http://noodle.med.yale.edu/~hdtg/notes/stiefel_notes.pdf (2011)
- Tran, M.-N., Nguyen, N., Nott, D., Kohn, R.: Bayesian deep net GLM and GLMM. *J. Comput. Graph. Stat.* (2019)
- Tran, M.-N., Nguyen, T.-N., Dao, V.-H.: A practical tutorial on Variational Bayes. (2021) [arXiv:2103.01327](https://arxiv.org/abs/2103.01327)
- Tran, M.-N., Nott, D.J., Kohn, R.: Variational Bayes with intractable likelihood. *J. Comput. Graph. Stat.* **26**(4), 873–882 (2017)
- Tran, M.-N., Pitt, M.K., Kohn, R.: Annealed important sampling for models with latent variables (2014) [arXiv:1402.6035](https://arxiv.org/abs/1402.6035)
- Tu, L.W.: *An Introduction to Manifolds*. Springer, New York (2011)
- Waterhouse, S., MacKay, D., Robinson, T.: Bayesian methods for mixtures of experts. In: Touretzky, M.C.M.D.S., Hasselmo, M.E. (eds.) *Advances in Neural Information Processing Systems*, pp. 351–357. MIT Press (1996)
- Zhou, B., Gao, J., Tran, M.-N., Gerlach, R.: Manifold optimisation assisted Gaussian variational approximation. *J. Comput. Gr. Stat.* (to appear) (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.