

Log-density gradient covariance and automatic metric tensors for Riemann manifold Monte Carlo methods

Tore Selland Kleppe 

Department of Mathematics and Physics,
University of Stavanger, Stavanger,
Norway

Correspondence

Tore Selland Kleppe, Department of
Mathematics and Physics, University of
Stavanger, Stavanger, Norway.
Email: tore.kleppe@uis.no

Funding information

Finansmarkedsfondet, Grant/Award
Number: 337601

Abstract

A metric tensor for Riemann manifold Monte Carlo particularly suited for nonlinear Bayesian hierarchical models is proposed. The metric tensor is built from symmetric positive semidefinite log-density gradient covariance (LGC) matrices, which are also proposed and further explored here. The LGCs generalize the Fisher information matrix by measuring the joint information content and dependence structure of both a random variable and the parameters of said variable. Consequently, positive definite Fisher/LGC-based metric tensors may be constructed not only from the observation likelihoods as is current practice, but also from arbitrarily complicated nonlinear prior/latent variable structures, provided the LGC may be derived for each conditional distribution used to construct said structures. The proposed methodology is highly automatic and allows for exploitation of any sparsity associated with the model in question. When implemented in conjunction with a Riemann manifold variant of the recently proposed numerical generalized randomized Hamiltonian Monte Carlo processes, the proposed methodology is highly competitive, in particular for the more challenging

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Scandinavian Journal of Statistics* published by John Wiley & Sons Ltd on behalf of The Board of the Foundation of the Scandinavian Journal of Statistics.

target distributions associated with Bayesian hierarchical models.

KEYWORDS

generalized randomized Hamiltonian Monte Carlo, MCMC, metric Tensor, Riemann manifold Monte Carlo

1 | INTRODUCTION

Efficient posterior sampling for Bayesian statistical models has attracted a substantial amount of research the last decades (see e.g., Martin et al., 2022). Riemann manifold Monte Carlo methods (Girolami & Calderhead, 2011) are in particular well suited for posterior distributions exhibiting complicated nonlinear dependence structures and/or substantial differences in scale across the target distribution. Posterior distributions with these properties arise (among other) for Bayesian hierarchical models which are widely used to model dependent data (see e.g., Kleppe, 2019). The successful application of the RMMC methods relies on the selection of a suitable metric tensor, a symmetric, positive definite matrix-valued function that should reflect the local scaling properties of the target distribution, which in this case is the posterior distribution in question.

This article makes several contributions towards the end of selecting metric tensors that are both of high quality and are easily applied by nonexperts in computational methods. The first contribution is the introduction of the, to the author's knowledge, new concept Log-density Gradient Covariance (LGC) and the development of some of its properties. Informally, the LGC associated with some probability density, say $\pi(\mathbf{x}|\boldsymbol{\theta})$ is defined to be the (necessarily symmetric, positive semidefinite [SPSD]) covariance matrix of the gradient of $\log \pi$ with respect to *both* \mathbf{x} and $\boldsymbol{\theta}$. Consequently, the LGC generalizes the Fisher information matrix (see e.g., Pawitan, 2017) (which is the covariance matrix of the log-density gradient with respect to $\boldsymbol{\theta}$ only). Subject to regularity conditions, the LGC is equal to expected negative Hessian of $\log \pi$ with respect to *both* \mathbf{x} and $\boldsymbol{\theta}$, and may informally speaking be used to measure the information content and dependence structure between- and among both \mathbf{x} and $\boldsymbol{\theta}$ in cases where both \mathbf{x} and $\boldsymbol{\theta}$ are sampled (e.g., when \mathbf{x} is a latent variable).

Secondly, a metric tensor is constructed from LGCs for a very broad class of possibly nonlinear models specified in terms of a sequence of conditional distribution statements. Very few restrictions are imposed, and in particular the class of models considered includes nonlinear hierarchical models, even with multiple- nonlinearly coupled layers of latent variables/priors. Consequently, guaranteed positive definite Fisher/LGC-based metric tensors may be constructed not only from the contributions to the posterior associated with observed data as is current practice (Girolami & Calderhead, 2011). Rather, a metric tensor may be constructed from arbitrarily complicated non-linear prior/latent variable structures, provided the LGC may be derived for each conditional distribution composing the posterior distribution. The proposed metric tensor may be derived directly from the model specification and does not involve any tuning parameters. Third, an efficient and highly automatic numerical implementation of the metric tensor based on Automatic Differentiation (AD) is proposed. The implementation may exploit any sparsity of the metric tensor, which for large scale hierarchical models is essential in a performance perspective.

The proposed metric tensor could in principle be used in conjunction with any Riemann manifold Monte Carlo method, for example, Riemann manifold Hamiltonian Monte Carlo or Riemann manifold Langevin dynamics (Girolami & Calderhead, 2011). However, in this article, the illustrations are done based on a Riemann manifold variant of the numerical generalized randomized Hamiltonian Monte Carlo (NGRHMC) method of (Kleppe, 2022). NGRHMC processes are continuous time piecewise deterministic processes (see e.g., Fearnhead et al., 2018) with Hamiltonian deterministic dynamics (Bou-Rabee & Sanz-Serna, 2017) which are implemented using adaptive numerical ordinary differential equations (ODEs) solvers. The usage of such ODE solvers introduces small biases, but at the same time avoids computationally intensive- and difficult to tune implicit symplectic integrators commonly used in Riemann manifold Hamiltonian Monte Carlo. The application of NGRHMC allows a clean comparison between samplers based on the proposed Riemann manifold and Euclidean versions of Hamiltonian dynamics, as the same numerical ODE solver may be used in both cases.

Finally, the paper contains several numerical illustrations, which benchmarks the proposed methodology against relevant alternatives. It is demonstrated that the proposed methodology may lead to substantial speed-ups in sampling efficiency (or expand the set of target distributions that may reliably sampled using HMC-like methods without introducing complicated rescaling methodology), in particular for challenging target distributions associated with large, high-dimensional Bayesian hierarchical models.

Below, Section 2 provides background material and relation to literature, and Section 3 introduces the LGC and discusses some of its properties. Section 4 derives a metric tensor based on LGC provides some illustrations of the properties of the metric tensor and discusses automatic implementation. Numerical examples and benchmarking are found in Sections 5 and 6, and Section 7 provides discussion. The article is accompanied by online supplementary material S1 which provides proofs and additional information in several regards.

2 | BACKGROUND

This section provides necessary background and fixes notation. For the purpose of readability, the notation and language is as far as possible avoiding differential-geometric nomenclature. Further, the paper assumes familiarity with Markov chain Monte Carlo (MCMC) methods and in particular Hamiltonian Monte Carlo (HMC) methods, for which, for example, Neal (2010), Girolami and Calderhead (2011), and Bou-Rabee and Sanz-Serna (2018) may serve as references.

The paper considers a continuous target density $\pi(\mathbf{q})$ with density kernel $\bar{\pi}(\mathbf{q}) \propto \pi(\mathbf{q})$, $\mathbf{q} \in \mathbb{R}^D$ (with respect to the Euclidean geometry) that allows evaluation. In the following, $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the density of a $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ random vector evaluated at \mathbf{x} . $\mathbf{0}_d \in \mathbb{R}^d$ and $\mathbf{0}_{d,n} \in \mathbb{R}^{d \times n}$ denote vectors and matrices of only zeros. For stacking of two vectors, say $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathbb{R}^n$ into $\mathbf{z} = [\mathbf{x}^T \ \mathbf{y}^T]^T \in \mathbb{R}^{d+n}$, the shorthand notation (\mathbf{x}, \mathbf{y}) is sometimes used. Further, $\nabla_{\mathbf{x}} f(\mathbf{x}) \in \mathbb{R}^d$ denotes the gradient of $f : \mathbb{R}^d \mapsto \mathbb{R}$, and $\nabla_{\mathbf{x}}^2 f(\mathbf{x}) \in \mathbb{R}^{d \times d}$ the Hessian of f . Finally, $\nabla_{\mathbf{x}} g(\mathbf{x}) \in \mathbb{R}^{p \times d}$ denotes the Jacobian of $g : \mathbb{R}^d \mapsto \mathbb{R}^p$.

2.1 | Metric tensors and Riemann manifold Hamiltonian dynamics

Broadly speaking, Riemann manifold MCMC methods (Girolami & Calderhead, 2011) rely on defining the proposal mechanism of the MCMC method on a (non-trivial) Riemann manifold

rather than the conventional Euclidean space \mathbb{R}^D . The Riemann manifold under consideration here may be characterized in terms of the *metric tensor* $\mathbf{G}(\mathbf{q}) \in \mathbb{R}^{D \times D}$, a smooth symmetric, positive definite (SPD) matrix-valued function for each $\mathbf{q} \in \mathbb{R}^D$. For purposes of this paper, it suffices to think of the metric tensor as giving the distance between two infinitesimally separated points $\mathbf{q} \in \mathbb{R}^D$ and $\mathbf{q} + \delta\mathbf{q} \in \mathbb{R}^D$ to be $\sqrt{(\delta\mathbf{q})^T \mathbf{G}(\mathbf{q}) \delta\mathbf{q}}$ (rather than the conventional Euclidean distance $\sqrt{(\delta\mathbf{q})^T \delta\mathbf{q}}$).

To leverage the flexibility afforded by introducing a non-trivial Riemann manifold for constructing HMC-like RMMC methods targeting π , Girolami and Calderhead (2011) suggested using the dynamics associated with the Hamiltonian

$$\mathcal{H}(\mathbf{q}, \mathbf{p}) = -\log \bar{\pi}(\mathbf{q}) + \frac{1}{2} \log |\mathbf{G}(\mathbf{q})| + \frac{1}{2} \mathbf{p}^T [\mathbf{G}(\mathbf{q})]^{-1} \mathbf{p}, \quad (1)$$

as the proposal mechanism. Here $\mathbf{p} \in \mathbb{R}^D$ is the fictitious momentum variable. The dynamics associated with (1) are governed by Hamilton's equations, which amounts to

$$\dot{\mathbf{q}}(t) = \nabla_{\mathbf{p}} \mathcal{H}(\mathbf{q}(t), \mathbf{p}(t)) = \mathbf{G}(\mathbf{q}(t))^{-1} \mathbf{p}(t), \quad (2)$$

$$\dot{\mathbf{p}}(t) = -\nabla_{\mathbf{q}} \mathcal{H}(\mathbf{q}(t), \mathbf{p}(t)). \quad (3)$$

The joint distribution for $\mathbf{z} = (\mathbf{q}^T, \mathbf{p}^T)^T$, associated with (1) is given by

$$\pi(\mathbf{z}) = \pi(\mathbf{q}, \mathbf{p}) = \pi(\mathbf{q}) \mathcal{N}(\mathbf{p} | \mathbf{0}_D, \mathbf{G}(\mathbf{q})). \quad (4)$$

The dynamics (2,3) preserve both the Hamiltonian (i.e., total energy), and are also volume preserving. Consequently, (2,3) preserve (4) in the sense that for any initial configuration $\mathbf{z}(0) \sim \pi(\mathbf{z})$, then $\mathbf{z}(T) \sim \pi(\mathbf{z})$ for any $T > 0$ (provided $\mathbf{z}(t)$ solves Hamilton's equations (2,3) for each $t \in [0, T]$). Clearly, the original target $\pi(\mathbf{q})$ is the \mathbf{q} -marginal of (4).

2.2 | Riemann manifold HMC

Arguably, the most promising general purpose RMMC method is Riemann manifold HMC (RMHMC) (Girolami & Calderhead, 2011). RMHMC is most easily explained as a discrete time MCMC algorithm targeting $\pi(\mathbf{z})$ (and samples targeting $\pi(\mathbf{q})$ may subsequently be obtained by discarding the \mathbf{p} -coordinates of samples targeting $\pi(\mathbf{z})$). Each transition of RMHMC, say from $\mathbf{z}_{(i)} = (\mathbf{q}_{(i)}, \mathbf{p}_{(i)})$ to $\mathbf{z}_{(i+1)}$ involves two steps, where the first step is updating the momentum $\mathbf{p}_{(i)}^* \sim \pi(\mathbf{p} | \mathbf{q}_{(i)}) = \mathcal{N}(\mathbf{p} | \mathbf{0}_D, \mathbf{G}(\mathbf{q}_{(i)}))$. In the second step, a fixed/random number of numerical time-integration steps applied to the ODE (2,3) with initial configuration $(\mathbf{q}_{(i)}, \mathbf{p}_{(i)}^*)$ are computed. The final state of the time-integration process is either accepted or rejected as $\mathbf{z}_{(i+1)}$ according to a Metropolis-Hastings (MH) mechanism in order to adjust for errors introduced by the numerical integration relative to the exact solution of (2,3).

Provided the time-integration is done using a symplectic/time-reversible method (see e.g., Leimkuhler & Reich, 2004; Sanz-Serna & Calvo, 1994), the accept probability of the MH step takes a particularly simple form. However, symplectic numerical methods for (2,3) are necessarily implicit, with each integration step requiring the iterative solution of a set of D non-linear equations involving the \mathbf{q} -gradient of the Hamiltonian (1). Further, to ensure stability and convergence of these iterative processes, it is typically necessary to use very short/many time-integration

steps in each transition. Consequently, unless special structures in the model may be exploited (see e.g., Kleppe, 2019; Zhang & Sutton, 2014), RMHMC may be very computationally demanding in practice for general model. It is also worth mentioning that MH adjustment mechanisms for RMHMC-like methods may be implemented with explicit time-reversible (but not volume preserving/symplectic) integrators, but this in general leads to highly nontrivial calculations for the accept probability (see e.g., Lan et al., 2015).

2.3 | Riemann manifold Numerical Generalized Randomized HMC processes

Rather than conventional discrete time RMHMC methods, the illustrations of the proposed metric tensor in this paper are done using Riemann manifold variants of numerical Generalized Randomized HMC (GRHMC) processes (Kleppe, 2022). The Riemann manifold variant of numerical GRHMC processes (see Kleppe, 2022, supplementary material), implemented with general purpose explicit adaptive first-order ODE solvers (see e.g., Hairer et al., 1993), bypasses the need for iterative nonlinear equation solving, while still exploiting the conservative nature (2,3) and the freedom to choose $\mathbf{G}(\mathbf{q})$. The savings in computing time comes at the cost of arbitrarily small errors introduced by the nonsymplectic and un-adjusted ODE solver.

For numerical GRHMC processes, the same adaptive time step integrator, with the same error tolerances may be used both for Riemann manifold-based methods endowed with the proposed metric tensor, and Euclidean metric (i.e., $\mathbf{G}(\mathbf{q}) = \mathbf{M}$ for some fixed SPD mass matrix \mathbf{M}). Consequently, an added benefit of considering numerical GRHMC processes, rather than RMHMC (based on implicit difficult to tune/computationally costly symplectic integrators), is that a clean comparison between Riemann manifold- and fixed metric numerical GRHMC methods may be carried out.

Riemann manifold GRHMC processes, say $\mathbf{Z}(t) = (\mathbf{Q}(t), \mathbf{P}(t))$, $t \geq 0$, are continuous time processes that may be specified so that $\mathbf{Q}(t)$ has an arbitrary continuous stationary distribution. The processes are special cases of piecewise deterministic Markov processes (Davis, 1993; Fearnhead et al., 2018; Vanetti et al., 2018). For simplicity, in this paper, only constant event rate processes, where events occur according to a time-homogenous Poisson process with intensity λ are considered. Between events, $\mathbf{Z}(t)$ solves (2,3). At events times t , the momentum coordinate $\mathbf{P}(t)$ is updated according to $\pi(\mathbf{p}|\mathbf{q} = \mathbf{Q}(t))$. Continuous time trajectories are simulated for a prespecified time interval $[0, T_{\max}]$, and position coordinate is subsequently sampled at discrete times, say $\{\mathbf{q}_i = \mathbf{Q}(i\Delta)\}_i$ for some suitable time increment Δ . The discrete time samples $\{\mathbf{q}_i\}_i$ may be used in the same manner as samples from conventional (discrete time) MCMC methods.

The (Riemann and Euclidean metric) numerical GRHMC processes used for illustration are implemented in the `pdmphmc` R-package (<https://github.com/torekleppe/pdmphmc>). For improved numerical performance, the simulations are done in standardized variables (see Appendix A for details on standardization and other aspects related to the numerical implementation). This standardization has the added benefit of making the unit of process time t comparable across many models/manifolds, and as a rule of thumb (obtained by trial and error), one should expect on the order of 0.2 effective samples per unit of process time.

2.4 | Metric tensors in the literature

So far, the metric tensor $\mathbf{G}(\mathbf{q})$ has been left unspecified. The overarching aim of working with a non-constant $\mathbf{G}(\mathbf{q})$ is to ensure an adequate representation of the geometry of the target distribution, thus resulting in that the dynamics (2,3) exhibit efficient exploration of the support of the target. In most applications of Riemann manifold HMC methods, $\mathbf{G}(\mathbf{q})$ is chosen to be some sort of positive definite approximation to the negative Hessian of the log-target density, that is, $-\nabla_{\mathbf{q}}^2 \log \pi(\mathbf{q})$. Such choices may be motivated by that in the flow of \mathbf{q} associated with (2,3), the log-target gradient gets scaled by $\mathbf{G}^{-1}(\mathbf{q})$ (see e.g., Kleppe, 2018, Eq. 9). It is well known from the numerical optimization literature (see e.g., Nocedal & Wright, 1999) that scaling the target function gradient using some form of positive definite approximation to the inverse target Hessian typically result in moves well adapted to the target distribution.

When the target distribution is the posterior distribution of a *nonhierarchical* statistical model, that is, $\pi(\mathbf{q}) \propto \pi(\mathbf{y}|\mathbf{q})p(\mathbf{q})$ (where $\pi(\mathbf{y}|\mathbf{q})$ is the likelihood function for observations \mathbf{y} and parameters \mathbf{q} , and $p(\mathbf{q})$ is the prior.), Girolami and Calderhead (2011) suggest using the metric tensor

$$\mathbf{G}(\mathbf{q}) = \mathcal{F}(\mathbf{q}) + \mathcal{R}. \quad (5)$$

Here $\mathcal{F}(\mathbf{q}) = -E_{\mathbf{y}|\mathbf{q}}[\nabla_{\mathbf{q}}^2 \log \pi(\mathbf{y}|\mathbf{q})]$ is the Fisher information matrix (see e.g., Pawitan, 2017, Chap. 8) associated with the likelihood function, and $\mathcal{R} = -\nabla_{\hat{\mathbf{q}}}^2 \log p(\hat{\mathbf{q}})$, $\hat{\mathbf{q}} = \arg \max_{\mathbf{q}} p(\mathbf{q})$, is the negative Hessian at the maximizer of the prior, a common approximation to the precision matrix of $p(\mathbf{q})$. Based on that $\mathcal{F}(\mathbf{q})$ is the natural metric tensor for the parameter space Riemann manifold associated with the statistical model $\pi(\mathbf{y}|\mathbf{q})$, Girolami and Calderhead (2011, Sec. 4) provide a discussion of why (5) constitutes a suitable metric tensor. See also Amari (1998) for further discussion of the application of Fisher information for the closely related natural gradient, but again restricted to non-hierarchical models.

Betancourt (2013) and Kleppe (2018) propose to use positive definite approximations to/modifications of $-\nabla_{\mathbf{q}}^2 \log \pi(\mathbf{q})$ as the metric tensor. Such procedures have the benefit of allowing for a high degree of automation, as $-\nabla_{\mathbf{q}}^2 \log \pi(\mathbf{q})$ may be computed from a program specifying the log-target density using automatic differentiation (AD) techniques (Griewank, 2000). Betancourt (2013) uses a full eigen-decomposition and modifies any small positive or negative eigenvalues of the negative Hessian. Kleppe (2018) on the other hand uses modified Cholesky factorization that exploits any sparsity of the Hessian, commonly present under hierarchical models (see e.g., Rue et al., 2009), to a similar end. Common for both techniques is that they require the non-trivial selection of a regularization parameter which chooses a tradeoff between the smoothness of the resulting $\mathbf{G}(\mathbf{q})$ against the difference between $-\nabla_{\mathbf{q}}^2 \log \pi(\mathbf{q})$ and $\mathbf{G}(\mathbf{q})$. Further, computing the required derivatives of the Hamiltonian (3) effectively amounts to third order AD, which may both be computationally demanding and require highly specialized techniques or additional input by the user if sparsity is to be exploited.

Recently, Hartmann et al. (2022) proposed the Monge metric, which in the present notation amounts to $\mathbf{G}(\mathbf{q}) = \mathbf{I}_D + \alpha^2 [\nabla_{\mathbf{q}} \log \pi(\mathbf{q})][\nabla_{\mathbf{q}} \log \pi(\mathbf{q})]^T$ for with α being a tuning parameter. The Monge metric also does not assume any particular structure on the model, and would allow implementation based on second order AD. The identity plus rank 1 update structure of the Monge metric affords substantial savings in the numerical linear algebra involved in each update, but it is not clear how to choose α for any given statistical model. Note also that the expectation of the

gradient outer product of the Monge metric is the Fisher information provided a similar model structure as for (5) and flat priors.

To reduce the cost of each RMHMC update for hierarchical models, certain structure can be imposed on the metric tensor. Zhang and Sutton (2014) propose semi-separable HMC, and the dynamic rescaling method of Kleppe (2019) may also be interpreted in terms a metric tensor with certain properties which would simplify RMHMC sampling. Both approaches are based on Fisher information matrices, but requires different, rather strict assumptions on the model which does not lend themselves easily to automatic implementation.

In what follows, a new metric tensor, along with an efficient and automatic method of computation of this metric tensor is proposed. The proposed metric tensor may be seen as a generalization of the Fisher-based metric (5) of Girolami and Calderhead (2011) to hierarchical/latent variable models that allows for a high degree of automation.

3 | LOG-DENSITY GRADIENT COVARIANCE

Before discussing metric tensors per se, the log-density gradient covariance (LGC) is introduced. The LGC generalizes the Fisher information matrix (see e.g., Pawitan, 2017) for sufficiently smooth probability densities, and will constitute an important building block for the proposed metric tensor.

3.1 | Log-density gradient covariance

Assumption 1. Probability density $\pi(\mathbf{x}|\theta)$ on \mathbb{R}^d has continuous first-order derivatives w.r.t. \mathbf{x} for each $\theta \in \Omega \subseteq \mathbb{R}^p$, where Ω is the set of allowed parameters.

Under Assumption 1, the LGC $\mathbb{V}_\pi[\mathbf{x}|\theta]$ associated with probability density $\pi(\mathbf{x}|\theta)$ is defined as

$$\mathbb{V}_\pi[\mathbf{x}|\theta] = \text{Var}_{\pi(\mathbf{x}|\theta)} \left[\nabla_{(\mathbf{x},\theta)} \log \pi(\mathbf{x}|\theta) \right] = \begin{bmatrix} \mathcal{V}(\theta) & \mathcal{W}(\theta) \\ \mathcal{W}^T(\theta) & \mathcal{F}(\theta) \end{bmatrix},$$

where the blocks $\mathcal{V}(\theta) \in \mathbb{R}^{d \times d}$, $\mathcal{W}(\theta) \in \mathbb{R}^{d \times p}$ and $\mathcal{F}(\theta) \in \mathbb{R}^{p \times p}$ conform in sizes with the sizes of \mathbf{x} and θ . Clearly $\mathcal{F}(\theta)$ is the Fisher information matrix associated with $\pi(\mathbf{x}|\theta)$. Being proper covariance matrices, both $\mathbb{V}_\pi[\mathbf{x}|\theta]$ and $\mathcal{V}(\theta)$ (in addition to $\mathcal{F}(\theta)$ obviously) are symmetric and positive semi-definite. In general, the LGC is a SPSD matrix-valued function of θ , and sometimes the notation $\mathbb{V}_\pi[\mathbf{x}|\theta](\theta)$ is needed.

3.2 | Basic properties of the log-density gradient and the LGC

The log-density gradient, $\nabla_{(\mathbf{x},\theta)} \log \pi(\mathbf{x}|\theta)$, and the LGC have properties that mirror those of the score function and Fisher information:

Proposition 1. Under Assumption 1,

$$I : E_{\pi(\mathbf{x}|\theta)} \left[\nabla_{(\mathbf{x},\theta)} \log \pi(\mathbf{x}|\theta) \right] = \mathbf{0}_{d+p},$$

and

$$II : \mathbb{V}_\pi[\mathbf{x}|\boldsymbol{\theta}] = - \frac{E}{\pi(\mathbf{x}|\boldsymbol{\theta})} \left[\nabla^2_{(\mathbf{x},\boldsymbol{\theta})} \log \pi(\mathbf{x}|\boldsymbol{\theta}) \right].$$

The proof is provided in Appendix B.1.

The second part of the proposition indicates, via the established explicit relation to the Hessian matrix with respect to $(\mathbf{x}, \boldsymbol{\theta})$, that $\mathbb{V}_\pi[\mathbf{x}|\boldsymbol{\theta}]$ is a sensible “scale matrix” for statistical computing purposes in cases where variation in $\mathbf{x}, \boldsymbol{\theta}$ jointly is considered.

The above proposition relies critically on the smoothness of Assumption 1, which in turn implies that $\int \frac{\partial}{\partial x_i} \pi(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x} = \lim_{x_i \rightarrow \infty} \pi(\mathbf{x}|\boldsymbol{\theta}) - \lim_{x_i \rightarrow -\infty} \pi(\mathbf{x}|\boldsymbol{\theta}) = 0$ for all $i = 1, \dots, d$. Failures to be sufficiently smooth, for example, the exponential distribution (interpreted as a distribution on \mathbb{R} with density evaluating to 0 for negative arguments), may in certain cases be worked around by transformations of \mathbf{x} , see, for example, ExpGamma distribution below. The regular Gamma distribution is sufficiently smooth for shape parameter > 2 as then it will have continuous first order derivative with respect to \mathbf{x} everywhere.

3.3 | A transformation result

It is well known that the Fisher information matrix for some alternative parameter, say $\boldsymbol{\eta}$, may be expressed in terms the Fisher information associated with the original parameter, say $\boldsymbol{\theta} = \boldsymbol{\Psi}(\boldsymbol{\eta})$. A similar result can be derived for the LGC subject to transformations between $(\mathbf{x}, \boldsymbol{\theta})$ and $(\mathbf{z}, \boldsymbol{\eta})$ of the form

$$\mathbf{x} = \mathbf{a}(\boldsymbol{\eta}) + \mathbf{Bz}, \boldsymbol{\theta} = \boldsymbol{\Psi}(\boldsymbol{\eta}), \quad (6)$$

where it is assumed that matrix \mathbf{B} is invertible (and hence the dimensions of \mathbf{x} and \mathbf{z} are equal). Denote by $p(\mathbf{z}|\boldsymbol{\eta}) = \pi(\mathbf{a}(\boldsymbol{\eta}) + \mathbf{Bz}|\boldsymbol{\Psi}(\boldsymbol{\eta}))|\mathbf{B}|$ the density of $\mathbf{z}|\boldsymbol{\eta}$ implied by $\mathbf{x}|\boldsymbol{\theta}$ being distributed according to $\pi(\mathbf{x}|\boldsymbol{\theta})$ and (6). Then the LGC associated with $p(\mathbf{z}|\boldsymbol{\eta})$ may be expressed in terms of $\mathbb{V}_\pi[\mathbf{x}|\boldsymbol{\theta}]$, namely

$$\mathbb{V}_p[\mathbf{z}|\boldsymbol{\eta}] = \mathbf{U}(\boldsymbol{\eta})^T \{ \mathbb{V}_\pi[\mathbf{x}|\boldsymbol{\theta}](\boldsymbol{\Psi}(\boldsymbol{\eta})) \} \mathbf{U}(\boldsymbol{\eta}), \mathbf{U}(\boldsymbol{\eta}) = \nabla_{(\mathbf{z}, \boldsymbol{\eta})}(\mathbf{x}, \boldsymbol{\theta}) = \begin{bmatrix} \mathbf{B} & \nabla_{\boldsymbol{\eta}} \mathbf{a}(\boldsymbol{\eta}) \\ \mathbf{0} & \nabla_{\boldsymbol{\eta}} \boldsymbol{\Psi}(\boldsymbol{\eta}) \end{bmatrix}. \quad (7)$$

The rather elementary proof of (7) is detailed in Appendix B.2. Clearly, setting $\mathbf{a}(\boldsymbol{\eta}) = \mathbf{0}, \mathbf{B} = \mathbf{I}$ recovers the conventional re-parameterization formula for the Fisher information (Pawitan, 2017) (with the cross-information modified to be $\mathcal{W}[\mathbf{z}|\boldsymbol{\eta}] = \mathcal{W}[\mathbf{x}|\boldsymbol{\theta} = \boldsymbol{\Psi}(\boldsymbol{\eta})] \nabla_{\boldsymbol{\eta}} \boldsymbol{\Psi}(\boldsymbol{\eta})$). Further, the LGC exhibit intuitive behavior by being unchanged under constant (w.r.t. parameters) location shifts of the random variable ($\mathbf{B} = \mathbf{I}, \boldsymbol{\theta} = \boldsymbol{\eta}$ and $\nabla_{\boldsymbol{\eta}} \mathbf{a} = \mathbf{0}$ so that $\mathbf{U} = \mathbf{I}$). Even further, (7) entails that the LGC random variable block \mathcal{V} scales as conventional precision matrix under invertible linear transformations of the random variable.

3.4 | Examples of LGCs

This section gives some examples of LGCs for common probability distributions. The Gaussian distribution with density $\mathcal{N}(x|\mu, \sigma^2)$ has the LGC

$$\mathbb{V}_{\mathcal{N}}[x|(\boldsymbol{\mu}, \sigma)] = \sigma^{-2} \begin{bmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix}. \quad (8)$$

More generally, for a multivariate Gaussian distribution, say $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \mathbf{P}^{-1}(\boldsymbol{\omega}))$ where the precision matrix $\mathbf{P}(\boldsymbol{\omega})$ depends on a parameter vector $\boldsymbol{\omega}$, it is clear that

$$\mathbb{V}_{\mathcal{N}}[\mathbf{x}|(\boldsymbol{\mu}, \boldsymbol{\omega})] = \begin{bmatrix} \mathbf{P}(\boldsymbol{\omega}) & -\mathbf{P}(\boldsymbol{\omega}) & \mathbf{0} \\ -\mathbf{P}(\boldsymbol{\omega}) & \mathbf{P}(\boldsymbol{\omega}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{F}_{\boldsymbol{\omega}} \end{bmatrix},$$

where $\mathcal{F}_{\boldsymbol{\omega}}$ is the Fisher information of $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \mathbf{P}^{-1}(\boldsymbol{\omega}))$ with respect to $\boldsymbol{\omega}$.

For densities $\pi(\mathbf{y}|\boldsymbol{\theta})$ that do not have a everywhere continuous derivative with respect \mathbf{x} , LGCs may be derived after first transforming \mathbf{y} . Examples include the ExpGamma-distribution (named analogous with the LogNormal distribution), that is, if $X \sim \text{ExpGamma}(\alpha, \beta)$, then $Y = \exp(X) \sim \text{Gamma}(\alpha, \beta)$ where β is the scale parameter. The ExpGamma distribution has density $\pi(x|\alpha, \beta) \propto \exp(\alpha x - \beta^{-1} \exp(x))$, $x \in \mathbb{R}$, and yields the LGC

$$\mathbb{V}_{\pi}[x|(\alpha, \beta)] = \begin{bmatrix} \alpha & -1 & -\alpha\beta^{-1} \\ -1 & \Psi'(\alpha) & \beta^{-1} \\ -\alpha\beta^{-1} & \beta^{-1} & \alpha\beta^{-1} \end{bmatrix}, \quad (9)$$

where $\Psi'(a) = \frac{d^2}{da^2} \log(\Gamma(a))$.

Another such example would be the InverseLogitBeta, defined via $X \sim \text{InverseLogitBeta}(a, b) \Rightarrow Y = \text{logit}^{-1}(X) \sim \text{Beta}(a, b)$, where $\text{logit}^{-1}(x) = \frac{\exp(x)}{\exp(x)+1}$. The InverseLogitBeta distribution, which has density $\pi(x|a, b) \propto [\exp(x)/(1 + \exp(x))]^a [1/(1 + \exp(x))]^b$, $x \in \mathbb{R}$, has LGC given by

$$\mathbb{V}_{\pi}[x|(a, b)] = \begin{bmatrix} \frac{ab}{a+b+1} & -\frac{b}{a+b} & \frac{a}{a+1} \\ -\frac{b}{a+b} & \Psi'(a) - \Psi'(a+b) & -\Psi'(a+b) \\ \frac{a}{a+1} & -\Psi'(a+b) & \Psi'(b) - \Psi'(a+b) \end{bmatrix}.$$

Note that in the context of statistical computing using HMC-like methods, it is common practice to transform constrained variables into un-constrained ones (to obtain continuous first-order derivatives), as in the two latter examples, before sampling is performed. For example, Stan also uses internally the log- and logit-transforms to arrive at unconstrained variables from lower-bounded and compactly supported variables respectively (Carpenter et al., 2017). Hence, for application of the LGC within statistical computing, the requirement that the involved densities fulfill Assumption 1 is not too restrictive.

4 | METRIC TENSORS BASED ON LGC

Commonly, Bayesian hierarchical models are built from sequences of known conditional distributions, for which deriving LGCs (or Fisher information matrices in the case of discrete observation

likelihoods) is usually relatively easy. This section discusses how leverage such LGCs to arrive at the proposed metric tensor for models built from such sequences of conditional distributions with potentially non-linear interconnections which may generate complicated dependence structures.

4.1 | Model formulation and notation

The proposed methodology assumes that:

Assumption 2. The joint posterior distribution of statistical model under consideration may be written as

$$\log \pi(\mathbf{q}) \propto \sum_{\ell} \log \pi_{\ell}(\boldsymbol{\phi}_{\ell}(\mathbf{q}) | \boldsymbol{\psi}_{\ell}(\mathbf{q})), \quad (10)$$

for suitably chosen probability densities/mass functions $\{\pi_{\ell}\}_{\ell}$, “argument functions” $\{\boldsymbol{\phi}_{\ell}(\mathbf{q})\}_{\ell}$ and “parameter functions” $\{\boldsymbol{\psi}_{\ell}(\mathbf{q})\}_{\ell}$.

Assumption 3. Provided $\nabla_{\mathbf{q}} \boldsymbol{\phi}_{\ell}(\mathbf{q}) \neq \mathbf{0}$, then $\pi_{\ell}(\mathbf{x}_{\ell} | \boldsymbol{\theta}_{\ell})$ admit a LGC $\mathbb{V}_{\pi_{\ell}}[\mathbf{x}_{\ell} | \boldsymbol{\theta}_{\ell}]$. If $\nabla_{\mathbf{q}} \boldsymbol{\phi}_{\ell}(\mathbf{q}) = \mathbf{0}$, then $\pi_{\ell}(\mathbf{x}_{\ell} | \boldsymbol{\theta}_{\ell})$ admits a Fisher information matrix.

Note that both $\boldsymbol{\psi}_{\ell}(\mathbf{q})$ and $\boldsymbol{\phi}_{\ell}(\mathbf{q})$ may be constant with respect to the sampled quantity \mathbf{q} . For example, $\boldsymbol{\psi}_{\ell}$ being some fixed hyper-parameters if π_{ℓ} is a prior, or $\boldsymbol{\phi}_{\ell}(\mathbf{q})$ being equal to a set of observations/data. In cases where $\nabla_{\mathbf{q}} \boldsymbol{\phi}_{\ell}(\mathbf{q}) = \mathbf{0}$ and π_{ℓ} does not admit a LGC (e.g., a discrete distribution), \mathcal{V}_{ℓ} and \mathcal{W}_{ℓ} are taken to be the zero-matrices in the subsequent derivations.

4.2 | Proposed metric tensor

Based on the above model formulation and Assumptions 2 and 3, this Section proposes a metric tensor suitable for statistical computing applications. Define the Jacobian matrix

$$\mathbf{J}_{\ell}(\mathbf{q}) = \begin{bmatrix} \nabla_{\mathbf{q}} \boldsymbol{\phi}_{\ell}(\mathbf{q}) \\ \nabla_{\mathbf{q}} \boldsymbol{\psi}_{\ell}(\mathbf{q}) \end{bmatrix}.$$

In the current paper, it is proposed to use

$$\mathbf{G}(\mathbf{q}) = \sum_{\ell} \mathbf{G}_{\ell}(\mathbf{q}), \text{ where } \mathbf{G}_{\ell}(\mathbf{q}) = \mathbf{J}_{\ell}^T(\mathbf{q}) \{ \mathbb{V}_{\pi_{\ell}}[\mathbf{x}_{\ell} | \boldsymbol{\theta}_{\ell}](\boldsymbol{\psi}(\mathbf{q})) \} \mathbf{J}_{\ell}(\mathbf{q}), \quad (11)$$

as the metric tensor. Note that $\mathbf{G}(\mathbf{q})$ is the sum of induced (pull-back) pseudo-metric tensors from (the domain of) $(\mathbf{x}_{\ell} | \boldsymbol{\theta}_{\ell})$ to the sampling space characterized by \mathbf{q} .

Before proceeding, some remarks are in order.

- Under certain additional assumptions, (certain rows/columns of) \mathbf{G}_{ℓ} may itself be interpreted as a LGC. More precisely, momentarily assuming that there exist subsets I and J of $\{1, \dots, d\}$ where $I \cap J = \emptyset$, functions \mathbf{a} , \mathbf{c} and invertible matrix \mathbf{B} so that $\boldsymbol{\phi}_{\ell}(\mathbf{q}) = \mathbf{a}(\mathbf{q}_I) + \mathbf{B}\mathbf{q}_J$ and $\boldsymbol{\psi}_{\ell}(\mathbf{q}) = \mathbf{c}(\mathbf{q}_J)$. Then it follows from (7, with $\boldsymbol{\Psi} = \mathbf{c}$) that the $I \cup J$ -rows/columns of \mathbf{G}_{ℓ} is the LGC of the density of $\mathbf{q}_I | \mathbf{q}_J$ implied by π_{ℓ} , $\boldsymbol{\phi}_{\ell}$ and $\boldsymbol{\psi}_{\ell}$.

- Informally speaking, (11) correctly represents the precision matrix of any (possibly degenerate) linear Gaussian structure. More precisely, in the case characterized by $\pi_\ell = \mathcal{N}(\mathbf{a} + \mathbf{A}\mathbf{q}_I | \mathbf{b} + \mathbf{B}\mathbf{q}_J, \mathbf{P}^{-1}(\mathbf{q}_I))$, where again $I \cap J = \emptyset$, and \mathbf{A}, \mathbf{B} not necessarily invertible, the I -rows/columns of \mathbf{G}_ℓ are equal to $-\nabla_{\mathbf{q}_I}^2 \log \pi_\ell$.
- In general, (10) admit nonlinear ϕ_ℓ s, and \mathbf{G}_ℓ will be a SPSD matrix also in the case of such nonlinear ϕ_ℓ s. Inclusion of this possibility is mainly done as the automatic implementation of (11) (see Section 4.4) does not distinguish between linear and nonlinear ϕ_ℓ s. Still, the interpretation of nonlinear ϕ_ℓ s is not obvious as (10) does not involve the log-Jacobian determinants of the ϕ_ℓ s. Hence, the use of nonlinear ϕ_ℓ s is advised against. Further, if ϕ_ℓ is bijective from some subset of \mathbf{q} , the need for nonlinear ϕ_ℓ s may be alleviated by choosing an equivalent base distribution on the pre-image of ϕ_ℓ .
- For models not involving latent variables so that \mathbf{q} may be interpreted as a parameter vector, (11) reduces to (5) with the modification the matrix \mathcal{R} is now the sum of prior argument log-gradient covariances \mathcal{V}_ℓ for ℓ s corresponding to prior terms in (10). Further, for terms corresponding to the log-likelihood function, say $\ell \in \mathcal{L}$, $\sum_{\ell \in \mathcal{L}} \mathbf{G}_\ell(\mathbf{q})$ is exactly the Fisher of the observations with respect to the parameter \mathbf{q} . In this sense, the proposed methodology may be seen as a generalization of the suggestion of Girolami and Calderhead (2011) to a much more general class of models involving latent variables.

Further, for a given model with target distribution $\pi(\mathbf{q})$, the metric tensor (11) is in general *not invariant* to how the factorization (10) is carried out, as will be illustrated with following example. Consider the simple funnel-model $q_1 \sim N(0, 1), q_2 | q_1 \sim N(0, \exp(-3q_1))$. The factorization (10) may be done either as

1. $\pi(\mathbf{q}) = \pi_1(q_1)\pi_2(q_2|q_1)$, which leads to $\mathbf{G}(\mathbf{q}) = \text{diag}(11/2, \exp(-3q_1))$, or simply
2. $\pi(\mathbf{q}) = \pi_1(q_1, q_2)$, which leads to $\mathbf{G}(\mathbf{q}) = \mathbb{V}_{\pi_1}[q_1, q_2] = \text{diag}(11/2, \exp(-9/2))$.

Clearly, factorization 1 above leads to metric tensor provides useful (\mathbf{q} -dependent) scaling information for sampling from the funnel distribution. For factorization 2 above, on the other hand, the metric tensor is constant/leads to Euclidean sampling which is not desirable for such target distributions. From this simple example, it seems advisable if possible, to factorize highly non-Gaussian joint distributions rather than to derive and use the LGC of the joint distribution. In fact, working with a single factor with $\phi_1 = \mathbf{q}$ in (10) would, as exemplified by factorization 2 above, always result in a constant (Euclidean) metric tensor. The factorization issue is further explored in Section 6.

4.3 | Examples

To get a sense of the workings of the proposed methodology, some small examples are considered here.

4.3.1 | Nonlinear parameter transformations

First, a nonhierarchical model similar to that of Bornn and Cornebise (2011) is considered. The model involves two parameters $\mathbf{q} = (\theta_1, \theta_2)$ and may be summarized by

$$\ell = 1, \dots, n : y_\ell | \theta_1, \theta_2 \sim N(\theta_1 + \theta_2^2, 1) \text{ and } \ell = n + 1 : (\theta_1, \theta_2) \sim N(\mathbf{0}, 100\mathbf{I}_2). \quad (12)$$

Then (11) result in

$$\mathbf{G}_\ell = \begin{bmatrix} 1 & 2\theta_2 \\ 2\theta_2 & 4\theta_2^2 \end{bmatrix}, \ell = 1, \dots, n, \text{ and } \mathbf{G}_{n+1} = 100^{-1}\mathbf{I}_2. \quad (13)$$

The resulting metric tensor is the same as (5) obtained in Bornn and Cornebise (2011), which owes to the fact that for the Gaussian prior π_{n+1} , the \mathcal{V}_{n+1} is equal to the negative Hessian at the mode \mathcal{R} . In Figure 1a, it is seen that (13) appears to accurately represent the local scaling properties of the target distributions.

4.3.2 | A simple hierarchical model

Now consider a simple hierarchical model with $\mathbf{q} = (\lambda, z)$ where λ is the a priori log-precision of the latent variable z . The model is characterized by

$$\ell = 1 : \lambda \sim N(0, 3^2), \ell = 2 : z | \lambda \sim N(0, \exp(-\lambda)) \text{ and } \ell = 3 : y | z \sim N(z, 1). \quad (14)$$

In this case, (11) results in

$$\mathbf{G}_1 = \begin{bmatrix} 3^{-2} & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{G}_2 = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & \exp(\lambda) \end{bmatrix}, \mathbf{G}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}. \quad (15)$$

Figure 1b illustrates (15), where it is seen that the scaling properties of the target distribution appears well represented.

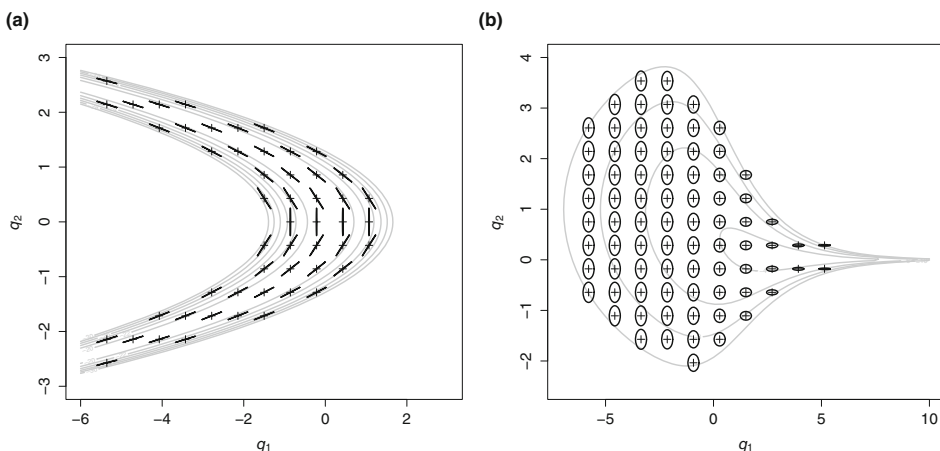


FIGURE 1 Illustrations of the proposed metric tensor for (a): model (12) and (b): model (14) with $y = 1$. In both plots, gray lines indicated target log-density contours, whereas the black closed curves are equi-probability ellipses associated with $N(\mathbf{q}, \mathbf{G}^{-1}(\mathbf{q}))$ for values of \mathbf{q} indicated by + signs. (a) Nonlinear parameters; (b) Simple hierarchical model.

4.3.3 | Intrinsic Gaussian

In the final small example, consider the intrinsic Gaussian model (see e.g., Rue & Held, 2005) for $\mathbf{q} \in \mathbb{R}^3$ characterized by

$$\ell = 1 : q_1 - q_2 \sim N(0, \kappa^{-1}), q_1 - q_3 \sim N(0, \kappa^{-1}) \text{ and } q_2 - q_3 \sim N(0, \kappa^{-1}), \quad (16)$$

for some fixed κ . This formulation still fits directly into (10), even with noninvertible argument functions, for example, $\phi_1(\mathbf{q}) = q_1 - q_2$ and with π_1 corresponding to $\mathcal{N}(x|0, \kappa^{-1})$. Further, in line with the second comment of Section 4.2, the metric tensor (11),

$$\mathbf{G} = \begin{bmatrix} 2\kappa & -\kappa & -\kappa \\ -\kappa & 2\kappa & -\kappa \\ -\kappa & -\kappa & 2\kappa \end{bmatrix},$$

is the (degenerate) precision matrix associated with (16).

4.4 | Automatic computation of $\mathbf{G}(\mathbf{q})$

Equation (11) may at first glance appear somewhat intimidating to compute for a general nonlinear model (10). However, it may be computed in a completely automatic fashion based on AD, provided the distribution families $\{\pi_\ell(\mathbf{x}|\theta)\}_\ell$ composing the posterior distribution all have known LGCs $\{\mathbb{V}_{\pi_\ell}\}_\ell$ (or Fisher information matrices in cases of discretely distributed observations). More specifically, the proposed methodology leading to (11) has been implemented in the C++ library `amt`, which is a part of the `pdmphmc` package (<https://github.com/torekleppe/pdmphmc>) which will be described shortly. With access to library `amt`, the user is only responsible for providing C++ code for specifying the model (10).

As an illustration, a working implementation of the model in Section 4.3.2 using library `amt` is given in the code column of Table 1. The computations arriving at (15) are done in an automatic manner as described in the comments column of Table 1. The library consist firstly of an AD-type `amtVar`, which is used to store both the value and also the gradient (w.r.t. \mathbf{q}) of every quantity in the model that depends on \mathbf{q} . The `amtVar` type is based on an implementation of first-order forward mode AD, which is sparse in the sense that it only stores the nonzero elements of each given gradient. This practice (as opposed to the celebrated backward mode AD) is informed by the fact that the rows of \mathbf{J}_ℓ are typically very sparse (each parameter/latent variable plays only a limited number of “roles”) for hierarchical models. In Table 1, lines 1, 2, and 4 illustrates the how the `amtVar` type is used to maintain the gradient of λ, z and $\sigma = \exp(-\lambda/2)$ with respect to \mathbf{q} .

Secondly, the `amt` library consist of a collection of probability distributions with known LGCs, illustrated here by the univariate Gaussian distribution `normal_1d()` in Table 1, lines 3, 5, and 6. Whenever such a function is called, the posterior log density kernel (10) is incremented by the appropriate $\log \pi_\ell$. Further, if the function is called with arguments and/or parameters of `amtVar` type, (11) is incremented by the appropriate \mathbf{G}_ℓ , computed from the LGC of the distribution and the gradients of the arguments and/or parameters.

TABLE 1 An implementation of the model in Section 4.3.2 using the `amt` library, where \mathbf{G}_1 , \mathbf{G}_2 and \mathbf{G}_3 are given in (15) and $\mathbf{q} = (\lambda, z)^T$.

Line number	Code	Comments
1	<code>PARAMETER_SCALAR(lambda);</code>	Stores λ as an <code>amtVar</code> variable and sets $\nabla_{\mathbf{q}}\lambda = [1, 0]$.
2	<code>PARAMETER_SCALAR(z);</code>	Stores z as an <code>amtVar</code> variable and sets $\nabla_{\mathbf{q}}z = [0, 1]$.
3	<code>model__ += normal_ld(lambda, 0.0, 3.0);</code>	Computes $\log \pi_1 = \log \mathcal{N}(\lambda 0, 3^2)$ and $\mathbf{G}_1 = [\nabla_{\mathbf{q}}\lambda]^T 3^{-2} [\nabla_{\mathbf{q}}\lambda]$.
4	<code>amtVar sigma = exp(-0.5*lambda);</code>	Computes $\sigma = \exp\left(-\frac{\lambda}{2}\right)$ and sets $\nabla_{\mathbf{q}}\sigma = -\frac{1}{2}\exp\left(-\frac{\lambda}{2}\right)\nabla_{\mathbf{q}}\lambda$.
5	<code>model__ += normal_ld(z, 0.0, sigma);</code>	Computes $\log \pi_2 = \log \mathcal{N}(z 0, \sigma^2)$, $\mathbf{J}_2 = \left[[\nabla_{\mathbf{q}}z]^T, [\nabla_{\mathbf{q}}\sigma]^T\right]^T$ and $\mathbf{G}_2 = \mathbf{J}_2^T \text{diag}(\sigma^{-2}, 2\sigma^{-2})\mathbf{J}_2$.
6	<code>model__ += normal_ld(1.0, z, 1.0);</code>	Computes $\log \pi_3 = \log \mathcal{N}(1 z, 1)$ and $\mathbf{G}_3 = [\nabla_{\mathbf{q}}z]^T [1] [\nabla_{\mathbf{q}}z]$.

The methodology for computing log-densities $\{\log \pi_\ell\}_\ell$, LGCs $\{\mathbb{V}_{\pi_\ell}\}_\ell$, Jacobians $\{\mathbf{J}_\ell\}_\ell$ and the Cholesky-factorization of the resulting $\mathbf{G}(\mathbf{q})$ required to compute (1) are in turn differentiated using the backward mode Stan AD (Carpenter et al., 2017) to obtain the gradient in (3). The library supports both dense- and sparse (Davis, 2006) storage of $\mathbf{G}(\mathbf{q})$, where the latter is in particularly relevant when the model exhibit some sort of Markov structure (see e.g., Rue & Held, 2005). Consequently, the relative merits (in terms of computational complexity and storage requirements) of dense- and sparse storage and associated Cholesky factorizations is highly model dependent, and should be kept in mind when specifying the model.

From a computational performance perspective, it is still advisable to keep the parameter functions $\{\psi_\ell(\mathbf{q})\}_\ell$ as simple as possible to avoid lengthy forward mode AD Jacobian calculations. This may be accomplished by defining new base distributions based on (6,7, with $\mathbf{a} = \mathbf{0}$, $\mathbf{B} = \mathbf{I}$) taking into account the reparameterization represented by $\Psi = \psi_\ell$. As an example, consider a model involving a normal linear model $\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$ with constant design matrix \mathbf{X} and $\boldsymbol{\beta} = \mathbf{q}_I$ for some index set I . Then it would be more effective to use the base distribution $\mathbf{y} | (\boldsymbol{\beta}, \sigma)$ with $\boldsymbol{\psi}(\mathbf{q}) = (\mathbf{q}_I, \sigma(\mathbf{q}))$, rather than base distribution $\mathbf{y} | (\boldsymbol{\mu}, \sigma)$ with $\boldsymbol{\psi}(\mathbf{q}) = (\mathbf{X}\mathbf{q}_I, \sigma(\mathbf{q}))$. This follows from that large savings may be realized by precomputing the factor $\mathbf{X}^T\mathbf{X}$ of the LGC of $\mathbf{y} | (\boldsymbol{\beta}, \sigma)$. In the $\mathbf{y} | (\boldsymbol{\mu}, \sigma)$, $\boldsymbol{\psi}(\mathbf{q}) = (\mathbf{X}\mathbf{q}_I, \sigma(\mathbf{q}))$ case, on the other hand, both calculating the nontrivial Jacobian of ψ using the forward mode AD routines, and also calculating the matrix product (11) would have to be repeated for each evaluation of \mathbf{G} . Deriving and implementing LGCs for the most common non-trivial “submodels” is an ongoing effort.

5 | EXAMPLES

This section considers real data example problems chosen in order to illustrate several aspects of the proposed methodology. In addition to the examples below, a further example, a mixed effects model for the Salamander mating data of McCullagh and Nelder (1989, Chapter 14.5)

may be found in Appendix E. For this model, due to rather high CPU time usage, the proposed methodology does not outperform the benchmark Euclidean metric sampler.

5.1 | Implementation details

The examples are all implemented using the `pdmphmc` C++ library (development version available at <https://github.com/torekleppe/pdmphmc>) which provides both Riemann manifold (RM)- and Euclidean metric (EM) numerical GRHMC (NGRHMC) processes, along with the library for automatic computation of (11) and an interface to R to facilitate building and running of models.

A snapshot of the version of `pdmphmc` used in this paper, along with R code, data sets etc used in this paper is also available at <https://github.com/torekleppe/AMTPaperCode>.

If not otherwise mentioned, in all the cases below, RM NGRHMC processes, along with EM NGRHMC processes as reference, are run with $T_{\max} = 10,000$ with the former half discarded as burn in. The “sampling”-part of the trajectories are sampled at 1000 equidistant times. In all cases, eight independent trajectories were used, and the reported figures are calculated across these eight trajectories. The eight trajectories were run in parallel in two batches of four trajectories on a 2020 macbook pro. The reported CPU times are the sum across trajectories of the CPU times required for generating the “sampling”-parts of the trajectories. The effective sample sizes (ESS) and (modified) Gelman-Rubin \hat{R} -statistics (Gelman et al., 2014) are calculated using the `rstan::monitor()`-function (Stan Development Team, 2017). Note that the reported results involve more trajectories than one would use in a typical application of the methodology, in order to reliably compare ESSes and time-weighted ESSes across sampling methods. For example, for a typical application one would rather use say four trajectories computed in parallel (resulting in wall-clock time being 1/8 of reported CPUtimes) to obtain roughly half the reported ESSes.

As a further benchmark, the examples were also implemented and sampled using Stan through the R interface `rstan` (version 2.26.23 with StanHeaders version 2.26.28) on the same computer. In all cases, eight chains of 1000 transitions (post warmup) and otherwise default settings were used for Stan.

5.2 | Zero-inflated Poisson mixed regression

The first model considered is a mixed effect regression model with zero-inflated Poisson count responses. Specifically, for a response y_i the response distribution is given by

$$P(y_i = 0 | \eta_i, g_i) = \frac{\exp(g_i) + \exp(-\exp(\eta_i))}{1 + \exp(g_i)}, \quad (17)$$

$$P(y_i = x | \eta_i, g_i) = \frac{\exp(x\eta_i - \exp(\eta_i))}{(1 + \exp(g_i))^x}, x = 1, 2, \dots \quad (18)$$

which may be interpreted as mixture of a point-mass in $y_i = 0$ and a Poisson distribution with mean $\exp(\eta_i)$, were the mixture weight of the $y_i = 0$ point mass is $\exp(g_i)(1 + \exp(g_i))^{-1}$. Consequently, $E(y_i | \eta_i, g_i) = \exp(\eta_i)(1 + \exp(g_i))^{-1}$. The Fisher information of (17,18) with respect to (η_i, g_i) has closed (but complicated) form and is given in Appendix C.1.

The data set considered is the Salamander data set originally discussed by Price et al. (2016) which is included in the R-package `g1mmTMB` (see Brooks et al., 2017) and consist of $n = 644$ observations. The “mean” linear predictor $\boldsymbol{\eta}$ involves a total of seven fixed effects (including an intercept term, with corresponding parameter $\boldsymbol{\beta}_\eta \in \mathbb{R}^7$) and a total of 23 random effects $\mathbf{b} \in \mathbb{R}^{23}$ with common variance parameter σ^2 . The “zero inflation” linear predictor \mathbf{g} consist of the same seven fixed effects (with corresponding parameter $\boldsymbol{\beta}_g \in \mathbb{R}^7$). Due to a complicated sparsity structure in $\mathbf{G}(\mathbf{q})$, and the moderate dimension of $\mathbf{q} = (\log(\sigma^2), \mathbf{b}, \boldsymbol{\beta}_\eta, \boldsymbol{\beta}_g)$, dense storage of $\mathbf{G}(\mathbf{q})$ was used as it resulted in slightly better performance. Further details are provided in Appendix C.1.

Table 2 provides results for the RM, EM, and Stan samplers. The \hat{R} -statistics indicated that all samplers exhibits satisfactory mixing. It is seen that RM sampler is substantially slower than EM in terms of simulating the same (process time) amount of trajectory, owing to that each evaluation of Hamilton’s equations is substantially more costly than for the EM counterpart. Further, Stan uses an order of magnitude less time than EM. The ESS of σ along with the worst case ESSes across $\boldsymbol{\beta}_\eta$, $\boldsymbol{\beta}_g$, and \mathbf{b} are similar for RM and EM except for a much poorer ESS for $\boldsymbol{\beta}_g$ for the EM sampler. This shortfall is likely to be related to the non-linear interaction between $\boldsymbol{\beta}_g$ and the remaining sampled quantities. The RM sampler, on the other hand, exhibit no such inefficiencies indicating that the proposed metric tensor is able to reflect these interactions. Even if Stan has relatively moderate raw ESSs for $\boldsymbol{\beta}_\eta$, $\boldsymbol{\beta}_g$ and \mathbf{b} , the very small CPU time of Stan result in the smallest time-weighted ESS for all parameters/latent variables for this model.

5.3 | Random walk stochastic volatility with leverage effect

Next, a random walk stochastic volatility (SV) model with leverage effect (see e.g., Yu, 2005) is considered. The latent log-volatility $\mathbf{z} = (z_0, \dots, z_T)$ evolves according to a Gaussian random walk

$$z_t | z_{t-1}, \sigma \sim N(z_{t-1}, \sigma^2), t = 1, \dots, T. \quad (19)$$

TABLE 2 Results for the zero-inflated Poisson mixed regression for Salamander count data.

	CPU time (s)	$\max \hat{R}$	σ			$\boldsymbol{\beta}_\eta$	$\boldsymbol{\beta}_g$	\mathbf{b}
			post. mean	post. SD	ESS	min ESS	min ESS	min ESS
RM	7178	1.002	1.37	0.21	5811	5894	10,537	5557
					[0.8]	[0.8]	[1.5]	[0.8]
EM	1034	1.002	1.37	0.22	7740	4367	514	2545
					[7.5]	[4.2]	[0.5]	[2.5]
Stan	65	1.008	1.38	0.22	5522	1095	985	1166
					[85]	[17]	[15]	[17]

Notes: CPU time is the total computing time spent sampling (post warmup) by eight independent trajectories. $\max \hat{R}$ is the maximum \hat{R} -statistic of over all quantities that are sampled. The table provides the posterior mean and SD, and ESS of the random effects variance parameter σ , in addition to the worst-case ESS and CPU time-weighted ESS (measured ESS per second CPU time) in []-brackets.

Abbreviations: EM, Euclidean metric; ESS, effective sample sizes; RM, Riemann manifold.

Further, the log-return observations $\mathbf{y} = (y_1, \dots, y_T)$ are modeled as

$$y_t | z_t, z_{t-1}, \rho, \sigma \sim N\left(\rho \exp\left[\frac{z_t - z_{t-1}}{2}\right] \frac{z_t - z_{t-1}}{\sigma}, \exp(z_{t-1})(1 - \rho^2)\right), t = 1, \dots, T. \quad (20)$$

Finally, the priors $\rho \sim \text{Uniform}(-1, 1)$ and $\sigma^2 \sim 0.1/\chi_{10}^2$ completes the model. The dataset consisted of $T = 2515$ log-return \times 100 observations of the S&P500 index spanning October 1, 1999 to September 30, 2009 (previously used by Grothe et al., 2019). Note that scale/covariance of conditional posterior $\mathbf{z} | \mathbf{y}, \rho, \sigma$ depends nonlinearly on both ρ and σ , that is, the posterior will be “funnel shaped along two separate dimensions”. Consequently, posterior sampling may be troublesome for many MCMC methods, and the model may be considered a rather challenging one.

For the RM variant of the sampler, the LGCs of $\mathbf{z} | \sigma$ consistent with (19) and $y_t | z_t, z_{t-1}, \rho, \sigma, t = 1, \dots, T$ consistent with (20) were used. In addition, otherwise identical calculations were done based on the LGC $\mathbb{V}[x | (\mu, \sigma)]$ for the $x \sim N(\mu, \sigma^2)$ -distribution for the observation Equation (20) (and hence parameter functions $\boldsymbol{\psi} = (\mu, \sigma) = (\rho \exp\left[\frac{z_t - z_{t-1}}{2}\right] \frac{z_t - z_{t-1}}{\sigma}, \sqrt{\exp(z_{t-1})(1 - \rho^2)})$). The latter approach, where nontrivial calculations are done using the general-purpose forward mode AD system, leads to an increase in computing time by around 30% (but otherwise identical results, hence not reported). Sparse storage with variable ordering $\mathbf{q} = (\mathbf{z}, -1 + 2\logit(\rho), \log(\sigma))$ was used, so that the sparsity structure has an arrowhead shape, which lends itself well to the sparse Cholesky factorization. For the EM and Stan implementations, the sampled quantity corresponding to the latent variable was $\bar{\mathbf{z}}_{0:T} = (z_0, (z_1 - z_0)\sigma^{-1}, \dots, (z_T - z_{T-1})\sigma^{-1})$ (i.e., so that $\bar{\mathbf{z}}_{1:T} \sim N(\mathbf{0}, \mathbf{I}_T)$ a priori) rather than \mathbf{z} in order to reduce “funnel” effects determined by σ .

Table 3 provides effective sample sizes and other diagnostic information for the EM and RM samplers, whereas Stan failed to produce meaningful results, issued a large number of warning messages, and is hence not reported on. From Table 3, it is seen that computing times are roughly equal, whereas the EM sampler fails to properly explore the posterior distribution of ρ . This failure is likely to be related to that the $\bar{\mathbf{z}}$ -parameterization does not take into account how the scale of $\bar{\mathbf{z}} | \mathbf{y}, \sigma, \rho$ varies with ρ . No such deficiencies are seen for the RM sampler, as the dependence of the scale of $\mathbf{z} | \mathbf{y}, \sigma, \rho$ on the parameters (σ, ρ) is automatically accounted for in the metric tensor.

5.4 | CEV model with additive noise

This section considers a daily time-discretization of constant elasticity of volatility model (Chan et al., 1992) with additive Gaussian noise for interest rate data previously considered by

TABLE 3 Effective sample sizes and diagnostics for the stochastic volatility model with leverage effect (19,20) applied to a data set of S&P500 log-returns.

	CPU time		ρ	σ	z_0	z_T
	(s)	$\max \hat{R}$	ESS	ESS	ESS	ESS
RM	9397	1.006	1762 [0.19]	1864 [0.20]	11,240 [1.20]	13,306 [1.42]
EM	9238	1.101	60 [0.01]	1767 [0.19]	8809 [0.95]	4143 [0.45]

Abbreviations: EM, Euclidean metric; ESS, effective sample sizes; RM, Riemann manifold.

Kleppe (2018). The model is formulated in continuous time with unit of continuous time being 1 year, and time-discretized to (business day) daily observations with time steps $\Delta = 1/252$. The model may be summarized by the time-discretized nonlinear latent “true” short-term interest rate

$$x_t = x_{t-1} + \Delta(\alpha - \beta x_{t-1}) + \sigma_x \sqrt{\Delta} x_{t-1}^\gamma \varepsilon_t, \varepsilon_t \sim \text{iid } N(0, 1), t = 2, \dots, T, \quad (21)$$

$$x_1 \sim N(0.09569, 0.01^2), \quad (22)$$

and the daily observations contaminated with additive Gaussian noise:

$$y_t = x_t + \sigma_y \eta_t, \eta_t \sim \text{iid } N(0, 1), t = 1, \dots, T. \quad (23)$$

The dataset considered was $T = 3082$ observations of the 7-day Eurodollar deposit spot rates from January 2, 1983, to February 25, 1995 previously used by Ait-Sahalia (1996) and Kleppe (2018). Further details, including priors may be found in Appendix C.2, and were chosen to be identical to the setup of Kleppe (2018) to allow for comparison with modified Cholesky Riemann manifold HMC.

The proposed methodology was implemented using univariate Gaussian LGCs with the standard parameterization, that is, $\mathbb{V}[x|(\mu, \sigma)]$ for the $x \sim N(\mu, \sigma^2)$ -distribution, (and not a bespoke LGC for say $x_t|x_{t-1}, \alpha, \beta, \sigma_x, \gamma$ in the case of (21)) thus relying on the general purpose forward mode AD system to handle the nonlinear relations between the sampled quantities. Using the variable ordering $\mathbf{q} = (\mathbf{z}, \alpha, \beta, \log(\sigma_x^2), \gamma, \log(\sigma_y^2))$, the metric tensor $\mathbf{G}(\mathbf{q})$ again has an arrow head sparsity structure which lend itself well to the sparse Cholesky factorization used.

Only results for RM based sampler are presented in Table 4, as direct EM-based or Stan-based sampling methods for this model failed to be even remotely competitive/produce reliable results, and dynamic rescaling methods for EM/Stan are not directly applicable due to the nonlinear nature of (21). Table 4 indicate that the proposed methodology produces reliable output with ESSes being quite even across the reported dimensions.

As a benchmark for sampling efficiency, Kleppe (2018) reports sampling efficiencies about an order of magnitude slower than those reported in between Table 4 (0.062 and 0.035 ESS per second) for a Riemann manifold HMC method (based on reversible symplectic integrator). Disentangling the effect of the here proposed metric tensor versus the modified Cholesky applied

TABLE 4 Posterior distributions and diagnostics information CEV model with additive noise model (21–23).

	Post. mean	Post. SD	ESS	
α	0.010	0.009	6163	[0.31]
β	0.171	0.174	6186	[0.31]
σ_x	0.404	0.061	7695	[0.39]
γ	1.180	0.060	6914	[0.35]
σ_y	0.00054	2.3×10^{-5}	3698	[0.19]
x_1	0.095	0.0005	6419	[0.32]
x_T	0.061	0.0005	6273	[0.31]

Note: Total CPU time for the eight trajectories was 19,933 s, and $\max \hat{R}$ for the stored states (parameters and x_1, x_T) was 1.001977. Abbreviation: ESS, effective sample sizes.

to Hessian approach of Kleppe (2018), from the effect of different ODE integration strategies is impossible based on this information. Still, the combination of Riemann manifold NGRHMC processes and the here proposed metric tensor is highly competitive while at the same time requiring minimal expertise and coding efforts from the user.

5.5 | The Stock and Watson (2007) model

The final smaller example model considered is the Stock and Watson (2007) quarterly inflation rate model. The model may be summarized by a pair of latent stochastic volatility processes with first-order Gaussian random walk structure

$$z_t | z_{t-1}, \sigma \sim N(z_{t-1}, \sigma^2), t = 2, \dots, T - 1, \quad (24)$$

$$x_t | x_{t-1}, \sigma \sim N(x_{t-1}, \sigma^2), t = 2, \dots, T. \quad (25)$$

Further, a latent stochastic trend process is modeled as a first order random walk with stochastic volatility

$$\tau_t | \tau_{t-1}, z_{t-1} \sim N(\tau_{t-1}, \exp(z_{t-1})), t = 2, \dots, T. \quad (26)$$

Finally, the observed time series of inflation rates \mathbf{y} is modeled as

$$y_t | \tau_t, x_t \sim N(\tau_t, \exp(x_t)), t = 1, \dots, T. \quad (27)$$

The model is completed by the prior $\sigma^{-2} \sim \text{Gamma}(5.0, 0.5)$, and is applied to the same dataset as in Kleppe (2019), namely quarterly log-returns $\times 100$ of the US CPI between 1955Q1 and 2018Q1. It is seen that the model involves two layers (\mathbf{z} and $(\mathbf{x}, \boldsymbol{\tau})$) of nonlinearly coupled latent variables, which poses substantial challenges for most MCMC methods.

The model was implemented for the RM sampler with $\mathbf{q} = (\mathbf{z}, \mathbf{x}, \boldsymbol{\tau}, \log(\sigma^{-2}))$ to obtain a tri-diagonal sparsity structure suitable for the sparse Cholesky factorization employed. As benchmarks, EM- and Stan samplers based on two modes of Dynamic Rescaling denoted DR0 and DR1 (see Kleppe, 2019, Sec. 6 for details) were considered. Direct EM or Stan sampling (i.e., with $\mathbf{q} = (\mathbf{z}, \mathbf{x}, \boldsymbol{\tau}, \log(\sigma^{-2}))$) was not competitive.

Diagnostic results are provided in Table 5. It is seen that the proposed methodology produces reliable results with minimal requirements of the user. Implemented both with EM and Stan, the DR1 method is more efficient than RM, but it is worth noticing that the implementation of the DR1 methodology in this case requires substantial user input- and expertise (essentially involving integrating out the complete $\boldsymbol{\tau} | \mathbf{z}, \mathbf{x}, \mathbf{y}, \sigma$ using bespoke tri-diagonal Cholesky algorithms). Further, when implemented in Stan, DR0 has performance roughly on par with RM, but again the DR0 is also here highly nontrivial to implement.

6 | A WISHART TRANSITION RANDOM WALK STOCHASTIC VOLATILITY MODEL

This section considers a restricted case of the multivariate stochastic volatility model of Philipov and Glickman (2006), where the precision matrix of the log-return vectors follows a random walk

TABLE 5 Effective sample sizes and diagnostics for the Stock and Watson (2007) model (24–27).

	CPU		σ			\mathbf{z}_t	\mathbf{x}_t	$\boldsymbol{\tau}_t$
	time (s)	$\max \hat{R}$	Post.	Post.	ESS	min	min	min
			mean	SD		ESS	ESS	ESS
EM DR0	876	1.020	0.31	0.05	1825	782	189	916
					[2.1]	[0.9]	[0.2]	[1.0]
EM DR1	214	1.004	0.31	0.05	3950	2885	2043	5317
					[18.5]	[13.5]	[9.6]	[24.9]
Stan DR0	86	1.061	0.31	0.05	842	186	143	387
					[9.8]	[2.2]	[1.7]	[4.5]
Stan DR1	30	1.038	0.32	0.05	1267	199	624	851
					[42.1]	[6.6]	[20.7]	[28.3]
RM	658	1.007	0.31	0.05	1917	977	2054	3273
					[2.9]	[1.5]	[3.1]	[5.0]

Note: Two variants of Dynamic Rescaling (DR0 and DR1) was applied for the EM sampler (see Kleppe, 2019, Sec. 6 for details). Abbreviation: ESS, effective sample sizes.

model with Wishart distributed transitions. Denote by $\mathcal{W}_p(\mathbf{V}, \nu)$ the Wishart distribution on $p \times p$ SPD matrices for $\nu > p - 1$ degrees of freedom and with SPD scale matrix \mathbf{V} (so that $E(\mathbf{P}) = \nu \mathbf{V}$ when $\mathbf{P} \sim \mathcal{W}_p(\mathbf{V}, \nu)$). Then the model considered here may be summarized by

$$\mathbf{P}_t | \mathbf{P}_{t-1}, \nu \sim \mathcal{W}_p(\nu^{-1} \mathbf{P}_{t-1}, \nu), t = 2, \dots, T, \quad (28)$$

$$\mathbf{y}_t | \mathbf{P}_t \sim N(\mathbf{0}, \mathbf{P}_t^{-1}), t = 1, \dots, T, \quad (29)$$

where $\mathbf{y}_t \in \mathbb{R}^p$, $t = 1, \dots, T$, are log-return vectors of p assets. The model is finalized with the prior $\nu \sim N(250.0, 20.0^2)$, with no special attention given to the constraint on ν as $\nu \approx p - 1$ is highly unlikely under the posterior distribution considered here. The dataset (with $p = 3$ and $T = 1095$) under consideration consist of daily observations of exchange rates of Australian Dollars (AUD), Canadian Dollars (CAD) and Swiss Francs (CHF) against the US Dollar between January 2, 2008 and April 4, 2012. The data are a subset of the `exrates` data set from the R package `stochvol` (Kastner, 2016).

6.1 | LGCs of SPD matrix-variate distributions

The $p \times p$ SPD matrices \mathbf{P}_t , $t = 1, \dots, T$, are represented in terms of unrestricted vectors $\mathbf{z}_t \in \mathbb{R}^{p(p+1)/2}$ via the transformation

$$\mathbf{P}_t = \mathcal{P}(\mathbf{z}_t) = \mathbf{L}(\mathbf{z}_t) \boldsymbol{\Lambda}(\mathbf{z}_t) \mathbf{L}^T(\mathbf{z}_t), \boldsymbol{\Lambda}(\mathbf{z}_t) = \text{diag}(\exp([\mathbf{z}_t]_1), \dots, \exp([\mathbf{z}_t]_p)),$$

and where $\mathbf{L}(\mathbf{z}_t)$ is unit lower triangular with the below diagonal columns filled with $[\mathbf{z}_t]_{\mathbf{L}} = [\mathbf{z}_t]_{p+1:p(p+1)/2}$. See Appendix D for details. Hence internally, the sampled quantities are $\mathbf{q} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_T, \nu)$, but the details of the representation of SPD matrices is hidden from the user in the model specification code.

Appendix D.2 provides the LGC $\mathbb{V}[\mathbf{y}_t|\mathbf{z}_t]$ consistent with $\mathbf{y}_t|\mathbf{z}_t \sim N(\mathbf{0}, [\mathcal{P}(\mathbf{z}_t)]^{-1})$ needed to implement (29). Appendix D.5 gives the distribution of $\mathbf{z}_t|\mathbf{z}_{t-1}, \nu$ so that $\mathcal{P}(\mathbf{z}_t) \sim \mathcal{W}(\nu^{-1}\mathcal{P}(\mathbf{z}_{t-1}), \nu)$, which is needed to implement the time dynamics of $\{\mathbf{z}_t\}_t$ consistent with (28). Two variants of the proposed methodology are considered for $\mathbf{z}_t|\mathbf{z}_{t-1}, \nu$ consistent with (28), corresponding to two different factorizations in target representation (10). In the former, denoted RM-J, $\pi(\mathbf{z}_t|\mathbf{z}_{t-1}, \nu)$ is considered a factor in (10), and the LGC $\mathbb{V}[\mathbf{z}_t|\mathbf{z}_{t-1}, \nu]$ may be found in Appendix D.5. In the second factorization, denoted RM-F, $\pi(\mathbf{z}_t|\mathbf{z}_{t-1}, \nu)$ is further factorized as $\pi([\mathbf{z}_t]_{1:p}|\mathbf{z}_{t-1}, \nu)\pi([\mathbf{z}_t]_L|[\mathbf{z}_t]_{1:p}, \mathbf{z}_{t-1}, \nu)$. Here $[\mathbf{z}_t]_i, i = 1, \dots, p$ under $\pi([\mathbf{z}_t]_{1:p}|\mathbf{z}_{t-1}, \nu)$ are independent ExpGamma-distributed (with LGC given in (9)), and $\pi([\mathbf{z}_t]_L|[\mathbf{z}_t]_{1:p}, \mathbf{z}_{t-1}, \nu)$ consist of a sequence of independent multivariate Gaussian distributions with covariance matrices having (different) \mathcal{P} -representations, whose LGCs are given in Appendix D.3. Note that $\{\mathbf{z}_t\}_t$ is Markovian, which leads to an arrow-head structure of $\mathbf{G}(\mathbf{q})$ which lends itself well to efficient sparse Cholesky factorization. Also for the EM and Stan benchmarks for this model were carried out using $\mathbf{q} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_T, \nu)$ as the sampled quantity.

6.2 | Results

Diagnostic results, and posterior moments of ν are provided in Table 6. It is seen that the EM sampler fails to properly explore the target distribution, whereas both RM samplers provide reliable results. The Stan sampler uses roughly double the amount of CPU time, and produces only low ESS for the parameter ν , which all in all results in that the RM-based methods have uniformly the best time weighted ESSs.

The difference in raw ESSs between the two RM samplers are rather small, whereas the RM-F sampling is somewhat slower, leading to slightly slower time-weighted ESSs. Figure 2 presents posterior quantiles of the marginal volatilities and correlations. It is seen that the model captures substantial time-variation in the correlations, which is missed in other multivariate SV specifications.

TABLE 6 Diagnostic results and posterior moments of ν under the Wishart transition RWSV model (28,29) applied to exchange rate data (AUS, CAD, CHF against USD).

	CPU		ν			$\mathbf{z}_t, t = 1, \dots, T$		
	time (s)	$\max \hat{R}$	Post.	Post.	ESS	Min	Median	Max
			mean	SD		ESS	ESS	ESS
RM-J	24,020	1.0036	256.8	16.78	3880	2602	12,738	31,225
					[0.162]	[0.108]	[0.530]	[1.300]
RM-F	27,451	1.0032	256.9	16.87	3874	2780	12,694	31225
					[0.141]	[0.101]	[0.462]	[1.137]
EM	21,635	1.1618	258.7	17.01	25	1727	5410	7093
					[0.001]	[0.080]	[0.250]	[0.328]
Stan	58,293	1.018	254.7	17.21	264	2706	5579	10,898
					[0.005]	[0.046]	[0.096]	[0.187]

Note: Figures in square brackets are time-weighted ESS (ESS/s computing time).

Abbreviations: EM, Euclidean metric; ESS, effective sample sizes; RM, Riemann manifold.

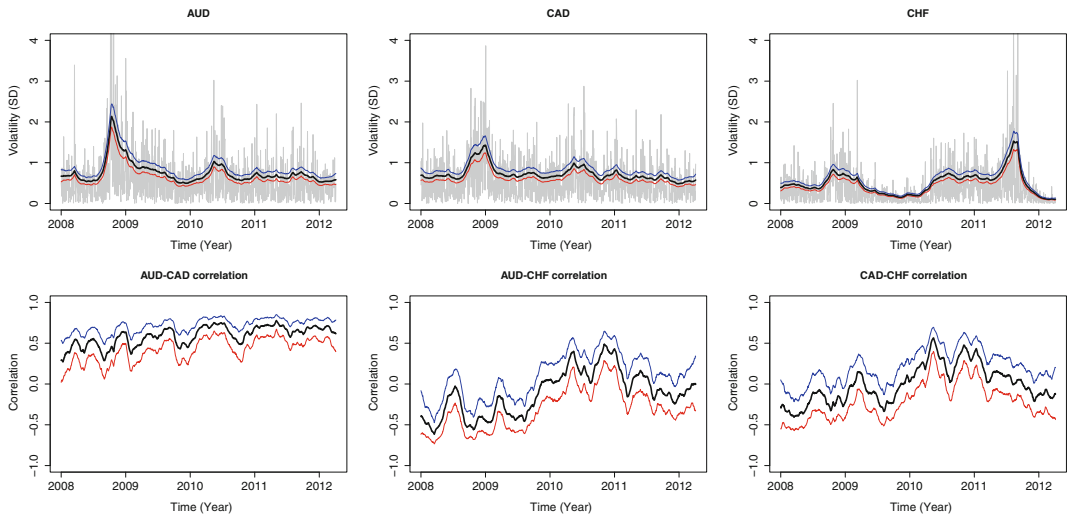


FIGURE 2 Marginal volatilities and correlations associated with the with the Wishart transition RWSV model model (28,29) for a single RM-GRHMC trajectory. In the upper panels, the absolute returns $|y_{i,t}|$ are given in grey, and the lines indicate 0.1,0.5,0.9-quantiles of the volatility $\sqrt{[\mathbf{P}_t^{-1}]_{i,i}}$ (i.e., measured in standard deviations). In the lower panel, the lines indicate the 0.1,0.5,0.9-quantiles of the correlations implied by the posterior distributions of each \mathbf{P}_t .

The MCMC method proposed by Philipov and Glickman (2006) was a Gibbs sampler involving each updating each \mathbf{P}_t using a random walk Metropolis steps, and which would be both time consuming and require substantial experience to develop. Though a direct comparison of the (probably highly autocorrelated) output from a Gibbs sampler against the output of the proposed methodology is not done here, it is at least clear that the proposed methodology can produce highly reliable results for large and complicated models with minimal user intervention. That is, the specification of the model requires only a handful of C++ lines corresponding to (28,29) and the prior, the rest is handled by software.

7 | DISCUSSION

Log-gradient covariances and a new metric tensor $\mathbf{G}(\mathbf{q})$ built from log-gradient covariances was proposed. Through numerical experiment and illustrations, it is shown that the metric tensor in conjunction with numerical generalized randomized HMC processes allows pushing the boundary for which hierarchical models can be fitted efficiently. The methodology is easy to use as the sole responsibility of the user is to specify the sequence of conditional distribution making up the model without much regard for imposing special structures in the model. The nuts and bolts of the proposed methodology, including derivative calculations and sparse matrix numerical linear algebra may be completely hidden from the user.

Deriving LGCs for further models and implementing these in the library holds scope for further work. Common structures such as linear regression models, logistic regression models, in addition to Gaussian spatial models are such examples. Further, adding functionality that hides the requirement to transform any variable to take values on the complete real line will also be developed. In addition, this paper only leverages a subset of what is possible within the

NGRHMC framework (Kleppe, 2022). Deriving more adaptive event rates, and corresponding methods for momentum updates, in the context Riemann manifold NGRHMC is an avenue that will be pursued.

ACKNOWLEDGMENTS

The author wishes to express warm thanks to the Editor, Professor Peltonen, an anonymous Associate Editor, two anonymous reviewers, Nawaf Bou-Rabee, Roman Liesenfeld and Hans J. Skaug for comments and suggestions for improvements on earlier versions of this paper. Kleppe acknowledges support from Finansmarkedsfondet, grant #337601.

ORCID

Tore Selland Kleppe  <https://orcid.org/0000-0001-8469-908X>

REFERENCES

- Ait-Sahalia, Y. (1996). Testing continuous-time models of the spot interest rate. *Review of Financial Studies*, 9(2), 385–426.
- Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10(2), 251–276.
- Betancourt, M. (2013). *A general metric for Riemannian manifold Hamiltonian Monte Carlo*. In F. Nielsen & F. Barbaresco (Eds.), *Geometric science of information Lecture Notes in Computer Science* (Vol. 8085, pp. 327–334). Springer.
- Bornn, L., & Cornebise, J. (2011). Comment on “Riemann manifold Langevin and Hamiltonian Monte Carlo methods”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2), 123–214.
- Bou-Rabee, N., & Sanz-Serna, J. M. (2017). Randomized Hamiltonian Monte Carlo. *Annals of Applied Probability*, 27(4), 2159–2194.
- Bou-Rabee, N., & Sanz-Serna, J. M. (2018). Geometric integrators and the Hamiltonian Monte Carlo method. *Acta Numerica*, 27, 113–206.
- Brooks, M. E., Kristensen, K., van Benthem, K. J., Magnusson, A., Berg, C. W., Nielsen, A., Skaug, H. J., Achler, M. M., & Bolker, B. M. (2017). glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *The R Journal*, 9(2), 378–400.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 1–32.
- Chan, K. C., Karolyi, G. A., Longstaff, F. A., & Sanders, A. B. (1992). An empirical comparison of alternative models of the short-term interest rate. *The Journal of Finance*, 47(3), 1209–1227.
- Davis, M. H. A. (1993). *Markov models and optimization*. Chapman & Hall.
- Davis, T. A. (2006). *Direct methods for sparse linear systems Fundamentals of algorithms* (Vol. 2). SIAM.
- Fearnhead, P., Bierkens, J., Pollock, M., & Roberts, G. O. (2018). Piecewise deterministic Markov processes for continuous-time Monte Carlo. *Statistical Science*, 33(3), 386–412.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. (2014). *Bayesian data analysis* (3rd ed.). CRC Press.
- Girolami, M., & Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2), 123–214.
- Griewank, A. (2000). *Evaluating derivatives: Principles and techniques of algorithmic differentiation*. SIAM.
- Grothe, O., Kleppe, T. S., & Liesenfeld, R. (2019). The gibbs sampler with particle efficient importance sampling for state-space models. *Econometric Reviews*, 38(10), 1152–1175.
- Hairer, E., Nørsett, S. P., & Wanner, G. (1993). *Solving ordinary differential equations I* (2nd Revised. Ed.) *Nonstiff problems*. Springer-Verlag.
- Hartmann, M., Girolami, M., & Klami, A. (2022). Lagrangian manifold Monte Carlo on Monge patches. *arXiv:2202.00755*.
- Kastner, G. (2016). Dealing with stochastic volatility in time series using the r package stochvol. *Journal of Statistical Software*, 69(5), 1–30.

- Kleppe, T. S. (2018). Modified Cholesky Riemann manifold Hamiltonian Monte Carlo: Exploiting sparsity for fast sampling of high-dimensional targets. *Statistics and Computing*, 28(4), 795–817.
- Kleppe, T. S. (2019). Dynamically rescaled Hamiltonian Monte Carlo for Bayesian hierarchical models. *Journal of Computational and Graphical Statistics*, 28(3), 493–507.
- Kleppe, T. S. (2022). Connecting the dots: Numerical randomized Hamiltonian Monte Carlo with state-dependent event rates. *Journal of Computational and Graphical Statistics*, 31(4), 1238–1253.
- Lan, S., Stathopoulos, V., Shahbaba, B., & Girolami, M. (2015). Markov chain Monte Carlo from Lagrangian dynamics. *Journal of Computational and Graphical Statistics*, 24(2), 357–378.
- Leimkuhler, B., & Reich, S. (2004). *Simulating Hamiltonian dynamics*. Cambridge University Press.
- Martin, G. M., Frazier, D. T., & Robert, C. P. (2022). Computing Bayes: From then ‘til now’. *arXiv:2208.00646*.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). Chapman & Hall.
- Neal, R. M. (2010). *MCMC using Hamiltonian dynamics*. In S. Brooks, A. Gelman, G. Jones, & X.-L. Meng (Eds.), *Handbook of Markov chain Monte Carlo* (pp. 113–162). CRC Press.
- Nocedal, J., & Wright, S. J. (1999). *Numerical optimization*. Springer.
- Pawitan, Y. (2017). *In all likelihood: Statistical modelling and inference using likelihood*. Oxford University Press.
- Philipov, A., & Glickman, M. E. (2006). Multivariate stochastic volatility via Wishart processes. *Journal of Business & Economic Statistics*, 24(3), 313–328.
- Price, S. J., Muncy, B. L., Bonner, S. J., Drayer, A. N., & Barton, C. D. (2016). Effects of mountaintop removal mining and valley filling on the occupancy and abundance of stream salamanders. *Journal of Applied Ecology*, 53(2), 459–468.
- Rue, H., & Held, L. (2005). *Gaussian Markov random fields: Theory and application*. Chapman and Hall-CRC Press.
- Rue, H., Martino, S., & Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2), 319–392.
- Sanz-Serna, J., & Calvo, M. (1994). *Numerical Hamiltonian problems*. Dover Publications Inc.
- Stan Development Team. (2017). *RStan: the R interface to Stan*. R package version 2.17.4.
- Stock, J. H., & Watson, M. W. (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- Vanetti, P., Bouchard-Côté, A., Deligiannidis, G., & Doucet, A. (2018). Piecewise-deterministic Markov chain Monte Carlo. *arXiv:1707.05296v2*.
- Yu, J. (2005). On leverage in a stochastic volatility model. *Journal of Econometrics*, 127, 165–178.
- Zhang, Y., & Sutton, C. (2014). *Semi-separable Hamiltonian Monte Carlo for inference in Bayesian hierarchical models*. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems 27* (pp. 10–18). Curran Associates, Inc.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Kleppe, T. S. (2024). Log-density gradient covariance and automatic metric tensors for Riemann manifold Monte Carlo methods. *Scandinavian Journal of Statistics*, 1–24. <https://doi.org/10.1111/sjost.12705>