Geodesic Lagrangian Monte Carlo over the space of positive definite matrices:
with application to Bayesian spectral density estimation

Andrew Holbrook[*1], Shiwei Lan[2], Alexander Vandenberg-Rodes[1], and Babak Shahbaba[1]

[1]Department of Statistics, University of California, Irvine
[2]Department of CMS, California Institute of Technology

We extend the application of Hamiltonian Monte Carlo to allow for sampling from probability distributions defined over symmetric or Hermitian positive definite matrices. To do so, we exploit the Riemannian structure induced by Cartan's century-old canonical metric. The geodesics that correspond to this metric are available in closed-form and—within the context of Lagrangian Monte Carlo—provide a principled way to travel around the space of positive definite matrices. Our method improves Bayesian inference on such matrices by allowing for a broad range of priors, so we are not limited to conjugate priors only. In the context of spectral density estimation, we use the (non-conjugate) complex reference prior as an example modeling option made available by the algorithm. Results based on simulated and real-world multivariate time series are presented in this context, and future directions are outlined.

Keywords: HMC; Riemannian geometry; spectral analysis

1. Introduction

In statistics, positive definite (PD) matrices primarily appear as covariance matrices parameterizing the multivariate Gaussian models. This type of models is the workhorse of modern statistics and machine learning: linear regression, probabilistic principal components analysis, Gaussian Markov random fields, spectral density estimation, and Gaussian process models all rely on the multivariate Gaussian distribution. The d-dimensional Gaussian distribution is completely specified by a mean vector μ and a covariance matrix Σ in $\mathcal{S}_d^+$, the space of d-by-d PD matrices. By imposing different structures on the covariance matrix, one can create different models. In some cases, it is possible to parameterize the covariance matrices in terms of a small number of parameters. However, learning of the *unstructured* covariance matrices, usually involved in inference on a large number of parameters, has remained as an issue. The conjugate Gaussian inverse-Wishart model has known deficiencies [1]. Outside of non-linear parameterizations of the Cholesky decomposition or matrix logarithm, there has not yet been a way to

---

˚holbroo@uci.edu

perform Bayesian inference directly on the space of PD matrices with flexible prior specifications using unstructured covariance matrices.

In this most general context the difficulty is in sampling from a posterior distribution on an abstract, high-dimensional manifold with boundary. It has not been clear how to propose moves from point to point within (and without leaving) this space. Our method takes advantage of the intrinsic, Riemannian geometry on the space of PD matrices. This space is incomplete under the Euclidean metric: following a straight trajectory will often result in matrices that are no longer positive definite. The space is, however, geodesically complete under the canonical metric: no matter how far the sampler travels along any geodesic, it never leaves the space of PD matrices. Intuitively, we redefine 'straight line' in a way that precludes leaving the set of interest. Moreover, the metric-induced geodesics provide a natural way to traverse the space of PD matrices, and these geodesics fit nicely with recent advances in Hamiltonian Monte Carlo on manifolds [2, 3, 4].

To this end, we use geodesic Lagrangian Monte Carlo (gLMC), which belongs to a growing class of Hamiltonian Monte Carlo algorithms. Hamiltonian Monte Carlo (HMC) provides an intelligent, partially deterministic method for moving around the parameter space while leaving the target distribution invariant. New Markov states are generated by numerically integrating a Hamiltonian system while Metropolis-Hastings steps account for the numerical error [5]. Riemannian manifold Hamiltonian Monte Carlo (RMHMC) adapts the proposal path by incorporating second-order information in the form of a Riemannian metric tensor [2]. Lagrangian Monte Carlo (LMC) builds on RMHMC by using a random velocity in place of RMHMC's random momentum. LMC's explicit integrator is no longer volume preserving; it therefore requires Jacobian corrections for each accept-reject step [6]. The embedding geodesic Monte Carlo (egMC) [3] is able to take the geometry of the parameter space into account while avoiding implicit integration by splitting the Hamiltonian [7] into a Euclidean and a geodesic component. Unfortunately, egMC is not applicable when a manifold's Riemannian embedding is unknown. gLMC, on the other hand, efficiently uses the same split Hamiltonian formulation as egMC but does not require an explicit Riemannian embedding [See 4, for example]. This last fact makes gLMC an ideal candidate for Bayesian inference on the space of PD matrices. We refer to the algorithm as positive definite Hamiltonian Monte Carlo (PDHMC).

PDHMC allows us to treat the entire covariance matrix as one would treat any other model parameters. We are no longer restricted to use a conjugate prior distribution or to specify a low-rank structure. We illustrate applications of PDHMC using both simulated and real-world data. First, we show that PDHMC provides the same empirical results as the closed-form solution for the conjugate Gaussian inverse-Wishart model. After this, we focus on applying Hermitian PDHMC to multivariate spectral density estimation and compare the results obtained from two different prior specifications: the inverse-Wishart prior and the complex reference prior. Then, we get credible intervals for the squared coherences (see Section 2) of simulated vector autoregressive time series for which the spectral density matrix is known. Finally, we apply PDHMC to learn the spectral density matrix associated with multivariate local field potentials (LFPs).

2. Motivation: learning the spectral density matrix

Given a stationary multivariate time series $y(t) = (y_1(t), \ldots, y_d(t))^\mathsf{T} \in \mathbb{R}^d$, $t = 1, \ldots, T$, one often wants to characterize the dependencies between vector elements through time. There are multiple ways to define such dependencies, and these definitions feature either the time series directly or the Fourier-transformed series in the frequency domain. In the time domain, one

characterization of dependence is provided by the lagged variance-covariance matrix Γ_ℓ. In terms of lag ℓ, this is written

$$\Gamma_\ell = \text{Cov}\big(y(t),\, y(t-\ell)\big) = \text{E}\Big(\big(y(t)-\mu\big)\big(y(t-\ell)-\mu\big)^T\Big) . \tag{1}$$

Note that Γ_ℓ and μ are invariant over time by stationarity. If the scientist has a reason to suspect that a certain lag ℓ is important *a priori*, then Γ_ℓ can be a useful measure. On the other hand, it is often more scientifically tractable to think in terms of frequencies rather than lags. In neuroscience, for example, one might hypothesize that two brain regions have 'correlated' activity during the performance of a specific task, but this co-activity may be too complex to describe in terms of a simple lagged relationship. The spectral density approach lends itself naturally to this kind of question. For a full discussion, see [8].

The power spectral density matrix is the Fourier transform of Γ_ℓ:

$$\Sigma(\omega) = \sum_{\ell=-\infty}^{\infty} \Gamma_\ell \, \exp(-i2\pi\omega\ell) . \tag{2}$$

$\Sigma(\omega)$ is a Hermitian PD matrix. A Hermitian matrix M is a complex valued matrix satisfying $M^H = \overline{M}^{\mathsf{T}} = M$, where $\overline{(\cdot)}$ denotes taking the complex conjugate. A Hermitian matrix M is defined to be PD if $z^H M z > 0$, $\forall z \in \mathbb{C}^d \setminus \{0\}$.

A diagonal element $\Sigma_{ii}(\omega)$ is called the auto-spectrum of $y_i(t)$ at frequency ω, and an off-diagonal element $\Sigma_{ij}(\omega)$, $i \neq j$ is the cross-spectrum of $y_i(t)$ and $y_j(t)$ at frequency ω. The squared coherence is given by

$$\rho_{ij}^2(\omega) = \frac{|\Sigma_{ij}(\omega)|^2}{\Sigma_{ii}(\omega)\,\Sigma_{jj}(\omega)} , \tag{3}$$

where $|\cdot|$ denotes the complex modulus. There are a number of ways to estimate the spectral density matrix and, hence, the matrix of squared coherences. In this paper, we model the discrete Fourier transformed time series $Y(\omega_k) \in \mathbb{C}^d$ as following a complex multivariate Gaussian distribution:

$$Y(\omega_k) \overset{ind}{\sim} \text{CN}_d(0, \Sigma(\omega_k)) , \tag{4}$$

where, for $\omega_k = \frac{k}{T}$ and $k = -(\frac{T}{2}-1), \ldots, \frac{T}{2}$,

$$Y(\omega_k) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} y(t) \, \exp(-i2\pi\omega_k t) . \tag{5}$$

Three assumptions are made here. First, we assume that the $Y(\omega_k)$s are *exactly* Gaussian: this is true when the $y(t)$ follow any Gaussian process. Moreover, if $y(t)$ follow a linear process, then the $Y(\omega_k)$ are asymptotically Gaussian as T goes to infinity [9]. Second, we assume that for $\omega_k \neq \omega_{k'}$, $Y(\omega_k)$ and $Y(\omega_{k'})$ are independent, whereas [9] show that they are asymptotically uncorrelated. Third, we assume that $\Sigma(\cdot)$ varies smoothly and slowly as a function over the frequency domain, and take all $Y(\omega_k)$ to be approximately i.i.d. within a small enough frequency band. For example, if we are interested in the alpha band—the interval of frequencies ranging from 7.5 to 12.5 Hz—then we model

$$Y(\omega_k) \overset{iid}{\sim} \text{CN}_d(0, \Sigma_\alpha) , \tag{6}$$

where Σ_α denotes the spectral density matrix shared by the entire band.

Thus, having obtained samples $Y(\omega)$ from a fixed frequency band, we will use Hermitian PDHMC to perform inference on Σ. The posterior samples of Σ automatically provide samples for the distributions of the squared coherences, which can in turn elucidate dependencies between the univariate time series. Before discussing PDHMC, we establish necessary facts regarding the space of positive definite matrices.

3. The space of positive definite matrices[†]

Let $\mathcal{S}_d(\mathbb{C})$ denote the space of $d \times d$ Hermitian matrices, and $\mathcal{S}_d^+(\mathbb{C})$ denote its subspace of PD matrices. The space of Hermitian PD matrices, $\mathcal{S}_d^+(\mathbb{C})$, may be written as a quotient space $GL(d, \mathbb{C})/U(d)$ of the complex general linear group $GL(d, \mathbb{C})$ and the unitary group $U(d)$. The general linear group is the smooth manifold for which every point is a matrix with non-zero determinant. The unitary group is the space of all complex matrices U satisfying $U^H U = UU^H = I$. This quotient space representation is rooted in the fact that every PD matrix may be written as the product $\Sigma = GG^H = GUU^H G^H$ for a unique $G \in GL(d, \mathbb{C})$ and any arbitrary unitary matrix $U \in U(d)$. For the convenience of exposition, we drop the dependence on $\mathbb{C}$ of symbols in the following of this section. Related references are [10, 11, 12].

$\mathcal{S}_d^+$ is a homogeneous space with respect to the general linear group: this means that the group acts transitively on the $\mathcal{S}_d^+$. Here the group action is given by conjugation:

$$G^* \Sigma = G\Sigma G^H . \tag{7}$$

For any $\Sigma_1, \Sigma_2 \in \mathcal{S}_d^+$, it simply takes the composition $\Sigma_2^{1/2*} \circ \Sigma_1^{-1/2*}$ to move from Σ_1 to Σ_2. The space of Hermitian matrices, $\mathcal{S}_d$, happens to be the tangent space to the space of Hermitian PD matrices at the identity, denoted as $T_{Id}\mathcal{S}_d^+$, that is, $T_{Id}\mathcal{S}_d^+ = \mathcal{S}_d$. The action

$$\Sigma^{1/2*} : V \mapsto \Sigma^{1/2} V \Sigma^{1/2} \tag{8}$$

translates vector $V \in T_{Id}\mathcal{S}_d^+$ to its corresponding vector in $T_\Sigma \mathcal{S}_d^+$, the tangent space to the space of PD matrices at point Σ.

Élie Cartan constructed a natural Riemannian metric $g(\cdot, \cdot)$ on the tangent bundle $T\mathcal{S}_d^+$ that is invariant under group action (7). For two vectors $V_1, V_2 \in T_{Id}\mathcal{S}_d^+$, the metric is given by

$$g_I(V_1, V_2) = \mathrm{tr}(V_1 V_2) . \tag{9}$$

In this way the space of PD matrices is isometric to Euclidean space (equipped with the Frobenius norm) at the identity. Next define the metric at any arbitrary point Σ to be

$$g_\Sigma(V_1, V_2) = \mathrm{tr}(\Sigma^{-1}V_1\Sigma^{-1}V_2) \tag{10}$$

It is easy to check that $g_I(V_1, V_2) = g_\Sigma(\Sigma^{1/2*}V_1, \Sigma^{1/2*}V_2)$ and so $\Sigma^{1/2*}$ is a Riemannian isometry on $\mathcal{S}_d^+$.

Two geometric quantities are required for our purposes: the Riemannian metric tensor and its corresponding geodesic flow, specified by a starting point and an initial velocity vector. The computational details involving the metric tensor are presented in Section 5. Here we present

---

[†]In this section we focus on the space of Hermitian positive definite matrices, since the class of symmetric matrices belongs to the broader class of Hermitian matrices. If the reader is primarily interested in the smaller class, then she is free to substitute $\mathbb{R}$ for $\mathbb{C}$, transpose $(\cdot)^\mathsf{T}$ for conjugate transpose $(\cdot)^H$, and the orthogonal group $O(d)$ for the unitary group $U(d)$.

the closed form solution for the geodesic flow as found in [10]. $\mathcal{S}_d^+$ is an affine symmetric space [13]. As such the geodesics under the invariant metric are generated by the one-parameter subgroups of the acting Lie group [10, 12]. These one-parameter subgroups are given by the group exponential map which, at the identity, is given by the matrix exponential $\exp tG$. In order to calculate the unique geodesic curve with starting position $\Sigma(0)$ and initial velocity $V(0)$, all one needs is to translate the velocity to the identity, compute the matrix exponential, and translate it back to the point of interest. In sum, the geodesic is given by

$$\Sigma(t) = \exp_\Sigma tV(0) = \Sigma(0)^{1/2} \exp\left(t\Sigma(0)^{-1/2}V(0)\Sigma(0)^{-1/2}\right)\Sigma(0)^{1/2} . \tag{11}$$

The corresponding flow on the tangent bundle will also be useful. This is obtained by taking the derivative with respect to t:

$$V(t) = \dot{\Sigma}(t) = \frac{d}{dt}\exp_\Sigma tV(0) = V(0)\Sigma(0)^{-1/2}\exp\left(t\Sigma(0)^{-1/2}V(0)\Sigma(0)^{-1/2}\right)\Sigma(0)^{1/2} . \tag{12}$$

For any lie group, the exponential map (on which the above formula is based) is a local diffeomorphism between the tangent space at a point on the manifold and the manifold itself. Given a tangent vector V at Σ, $\exp_\Sigma V$ is a point on the manifold. Incidentally, for the spaces of PD matrices, this diffeomorphism is global. The inverse of the exponential map is the logarithmic map. Whereas the exponential map on the manifold takes Hermitian matrices $(\mathcal{S}_d)$ to Hermitian PD matrices $(\mathcal{S}_d^+)$, the logarithmic map takes Hermitian PD matrices $(\mathcal{S}_d^+)$ to Hermitian matrices $(\mathcal{S}_d)$.

Together these are most of the geometric quantities required for PDHMC. The next section presents Hamiltonian Monte Carlo, its geometric extension RMHMC, and its Lagrangian manifestations.

4. Bayesian inference using the geodesic Lagrangian Monte Carlo

Given data $y_1, \ldots, y_N \in \mathbb{R}^n$, one may specify a generative model by a likelihood function, $p(y|q)$. In the following we allow $q \in \mathcal{M}^m$ to be an m-dimensional vector on a manifold that parameterizes the likelihood. Endowing q with a prior distribution $p(q)$ renders the posterior distribution

$$\pi(q) = p(q|y) = \frac{p(y|q)p(q)}{\int p(y|q)p(q)\,dq} . \tag{13}$$

The integral is often referred to as the evidence and may be interpreted as the probability of observing data y given the model. In most interesting models the evidence integral is intractable and high dimensional models do not lend themselves to easy numerical integration. Non-quadrature sampling techniques such as importance sampling or even random walk MCMC also suffer in high dimensions. HMC is an effective sampling tool for higher dimensional models over continuous parameter spaces [14, 5]. Here we discuss HMC and its geometric variants (see Section 1) in detail.

In HMC, a Hamiltonian system is constructed that consists of the parameter vector q and an auxiliary vector p of the same dimension. The negative-log transform turns the probability density functions into a potential energy function $U(q) = -\log\pi(q)$ and corresponding kinetic function $K(p)$. Thus q and p become the position and momentum of Hamiltonian function

$$H(q,p) = U(q) + K(p) . \tag{14}$$

By Euler's method or extensions, the system is numerically advanced according to Hamilton's equations:

$$\frac{dq}{dt} = \frac{\partial H}{\partial p}$$

$$\frac{dp}{dt} = -\frac{\partial H}{\partial q}. \tag{15}$$

Riemannian manifold HMC uses a slightly more complicated Hamiltonian to sample from posterior $\pi(q)$:

$$H(q, p) = -\log \pi(q) + \frac{1}{2} \log |G(q)| + \frac{1}{2} p^{\mathsf{T}} G(q)^{-1} p. \tag{16}$$

Here, $G(q)$ is the Fisher information matrix at point q (in Euclidean space) and may be interpreted as a Riemannian metric tensor induced by the curvature of the log-probability. Exponentiating and negating $H(q, p)$ reveals p to follow a Gaussian distribution centered at origin with metric tensor $G(q)$ for covariance. The corresponding system of first-order differential equations is given by

$$\frac{dq}{dt} = G(q)^{-1} p, \tag{17}$$

$$\frac{dp}{dt} = \nabla_q \left(\log \pi(q) - \frac{1}{2} \log |G(q)| - \frac{1}{2} p^{\mathsf{T}} G(q)^{-1} p \right).$$

The Hamiltonian is not separable in p and q. To get numerical solutions, one may split it into a potential term $H^{[1]}$, featuring q alone, and a kinetic term, $H^{[2]}$, featuring both variables [7, 3]. The two systems are then simulated in turn. The first term is given by

$$H^{[1]}(q, p) = -\log \pi(q) + \frac{1}{2} \log |G(q)| \tag{18}$$

and starting at $(q(0), p(0))$ the associated system has solutions

$$q(t) = q(0) \quad \text{and} \quad p(t) = p(0) + t \nabla_q \left(\log \pi(q) - \frac{1}{2} \log |G(q)| \right)|_{q=q(0)}. \tag{19}$$

The second component is the quadratic form

$$H^{[2]}(q, p) = \frac{1}{2} p^{\mathsf{T}} G(q)^{-1} p \tag{20}$$

The solutions to the system associated with $H^{[2]}$ are given by the geodesic flow under the Levi-Civita connection with respect to metric G and with momentum $p(t) = G(q(t))\dot{q}(t)$. There is, however, no *a priori* reason to restrict $G(q)$ to be the Fisher information as is done in the [2]. In fact, by allowing $G(q)$ to take on other forms, one may perform HMC on a number of manifold parameterized models.

4.1. Geodesic Lagrangian Monte Carlo

[3] show how to extend the RMHMC framework to manifolds that admit a known Riemannian isometric embedding into Euclidean space. The algorithm is especially efficient when there exists a closed form linear projection of vectors in the ambient space onto the tangent space at any

---
Algorithm 1 Geodesic Lagrangian Monte Carlo

1: $v \sim N(0, G^{-1}(q))$
2: $e \leftarrow -\log \pi(q) - \frac{1}{2} \log |G(q)| + \frac{1}{2} v^{\mathsf{T}} G(q) v$
3: $q^* \leftarrow q$
4: **for** $\tau = 1, \ldots, T$ **do**
5:      $v \leftarrow v + \frac{\epsilon}{2} G(q^*)^{-1} \nabla_q \left(\log \pi(q^*) + \frac{1}{2} \log |G(q^*)| \right)$
6:      Progress (q^*, v) along the geodesic flow for time ϵ.
7:      $v \leftarrow v + \frac{\epsilon}{2} G(q^*)^{-1} \nabla_q \left(\log \pi(q^*) + \frac{1}{2} \log |G(q^*)| \right)$
8: **end for**
9: $e^* \leftarrow -\log \pi(q^*) - \frac{1}{2} \log |G(q^*)| + \frac{1}{2} v^{\mathsf{T}} G(q^*) v$
10: $u \sim U(0, 1)$
11: **if** $u < \exp(e - e^*)$ **then**
12:      $q \leftarrow q^*$
13: **end if**
---

point. Although this embedding will always exist [15], it is rarely known. When equipped with the canonical metric, the space of PD matrices does not admit a known isometric embedding. Moreover, we are unaware of a closed-form projection onto the manifold's tangent space at a given point. We therefore opt for an intrinsic approach instead.

In the prior section, we stated that the solution to Hamilton's equations associated with the kinetic term $H^{[2]}$ is given by the geodesic flow with respect to the Levi-Civita connection. This flow is easily written in terms of the exponential map with respect to a velocity vector (as opposed to the momentum covector). Given an arbitrary covector $p \in T_q^* \mathcal{M}$, one may obtain the corresponding vector $v \in T_q \mathcal{M}$ by the one-to-one transformation $v = G^{-1}(q)p$. Hence whereas RMHMC augments the system with $p \sim N(0, G(q))$, Lagrangian Monte Carlo makes use of $v = G^{-1}(q)p \sim N(0, G^{-1}(q))$. The energy function is then given by

$$E(q, v) \quad \propto \quad -\log \pi(q) - \frac{1}{2} \log |G(q)| + \frac{1}{2} v^{\mathsf{T}} G(q) v \,. \tag{21}$$

The probabilistic interpretation of the energy remains the same as in the case of RMHMC: the energy is the negative logarithm of the probability density functions of two independent random variables, one of which is the variable of interest, the other of which is the augmenting Gaussian variable. On the other hand, the physical interpretation is different. We use the term 'energy' in order to accommodate the two physical interpretations available for (21): $E(q, v)$ may be thought of either as a Hamiltonian or as a Lagrangian energy. In practice, which formulation is used is dictated by the geometric information available. The Lagrangian formulation provides efficient update equations when no closed-form geodesics are available. In this case, the Lagrangian (energy) is defined as the kinetic term T less the potential term V as follows

$$V(q, v) = \log \pi(q) + \frac{1}{2} \log |G(q)|, \quad \text{and} \quad T(q, v) = \frac{1}{2} v^{\mathsf{T}} G(q) v \,. \tag{22}$$

But when closed-form geodesics are available, it is useful to follow [3] and split the (now considered) Hamiltonian into two terms as in (18) and (20). Within this regime, $H^{[1]} = -V$ and $H^{[2]} = T$. In analogy with (19) and starting at $(q(0), v(0))$, the system defined by potential V has solution

$$q(t) = q(0), \qquad v(t) = v(0) + t \, G(q)^{-1} \nabla_q V(q, v)|_{q=q(0)} \,, \tag{23}$$

Algorithm 2 Symmetric PDHMC

1: $\mathrm{vech}(V) \sim N(0, G^{-1}(\Sigma_t))$
2: $e \leftarrow -\log \pi(\Sigma_t) - \frac{d+1}{2}\log|\Sigma_t| + \frac{1}{2}\mathrm{vech}(V)^\mathsf{T} G(\Sigma_t)\mathrm{vech}(V)$
3: $\Sigma^* \leftarrow \Sigma_t$
4: **for** $\tau = 1, \ldots, T$ **do**
5: $\quad \mathrm{vech}(V) \leftarrow \mathrm{vech}(V) + \frac{\epsilon}{2}G^{-1}(\Sigma^*)\mathrm{vech}\left(\nabla_\Sigma\left(\log\pi(\Sigma^*) + \frac{d+1}{2}\log|\Sigma^*|\right)\right)$
6: $\quad$ Progress (Σ^*, V) along the geodesic flow for time ϵ.
7: $\quad \mathrm{vech}(V) \leftarrow \mathrm{vech}(V) + \frac{\epsilon}{2}G^{-1}(\Sigma^*)\mathrm{vech}\left(\nabla_\Sigma\left(\log\pi(\Sigma^*) + \frac{d+1}{2}\log|\Sigma^*|\right)\right)$
8: **end for**
9: $e^* \leftarrow -\log\pi(\Sigma^*) - \frac{d+1}{2}\log|\Sigma^*| + \frac{1}{2}\mathrm{vech}(V)^\mathsf{T} G(\Sigma^*)\mathrm{vech}(V)$
10: $u \sim U(0,1)$
11: **if** $u < \exp(e - e^*)$ **then**
12: $\quad \Sigma_{t+1} \leftarrow \Sigma^*$
13: **end if**

---

and the system defined by kinetic term T has the unique geodesic path specified by starting position $q(0)$ and initial velocity $v(0)$ as a solution. The inverse metric tensor $G^{-1}(q)$ is used to 'raise the index', i.e. transform the covector $\nabla_q V(q, v)$ into a vector on the tangent space at q. Thus it plays a similar function to the orthogonal projection in [3]. We call this formulation geodesic Lagrangian Monte Carlo (gLMC) and detail its steps in Algorithm 1, where the term 'Lagrangian' is used to emphasize the fact that we use velocities in place of momenta. Note [4, 16] implemented the similar idea on the manifold of a d-dimensional sphere. To implement geodesic Lagrangian Monte Carlo, one must be able to compute the inverse metric tensor $G^{-1}(q)$ and the geodesic path given starting values. When the space of PD matrices is equipped with the canonical metric, $G^{-1}(q)$ is given in closed-form and the geodesic path is easily computable.

5. Positive definite Hamiltonian Monte Carlo

Positive definite Hamiltonian Monte Carlo (PDHMC) is geodesic Lagrangian Monte Carlo on the space of PD matrices, $\mathcal{S}_d^+$, equipped with the canonical metric. In order to signify that we are no longer dealing with gLMC in its full generality, we adopt the notation of Section 3. PD matrix Σ replaces q, and symmetric or Hermitian matrix V replaces v. All other notations remain the same. As stated in the previous section, we require the inverse metric tensor $G^{-1}(\Sigma)$. To compute this quantity, we need a couple more tools provided by [11]. Let $\mathrm{vech}(\cdot)$ take symmetric (Hermitian) $d \times d$ matrices to vectors of length $\frac{d}{2}(d+1)$ by stacking diagonal and subdiagonal matrix elements in the following way:

$$\mathrm{vech}(V) = (V_{11}, V_{21}, \ldots, V_{d1}, V_{22}, \ldots, V_{d2}, \ldots, V_{dd}). \tag{24}$$

Let $\mathrm{vec}(\cdot)$ take symmetric (Hermitian) $d \times d$ matrices to vectors of length d^2 by stacking all matrix elements:

$$\mathrm{vec}(V) = (V_{11}, V_{21}, \ldots, V_{d1}, V_{12}, \ldots, V_{d2}, \ldots, V_{1d}, \ldots, V_{dd}). \tag{25}$$

Let D_d be the unique $d^2 \times \frac{d}{2}(d+1)$ matrix satisfying

$$\mathrm{vec}(V) = D_d\mathrm{vech}(V). \tag{26}$$

Denote D_d^+ as the Moore-Penrose inverse of D_d satisfying

$$\text{vech}(V) = D_d^+ \text{vec}(V) , \tag{27}$$

with D_d^+ given by

$$D_d^+ = (D_d^\mathsf{T} D_d)^{-1} D_d^\mathsf{T} . \tag{28}$$

Then [11] show that the metric tensor and inverse metric tensor are given by the $\frac{d}{2}(d+1) \times \frac{d}{2}(d+1)$ dimensional matrices

$$G(\Sigma) = D_d^\mathsf{T} (\Sigma^{-1} \otimes \Sigma^{-1}) D_d \quad \text{and} \quad G^{-1}(\Sigma) = D_d^+ (\Sigma \otimes \Sigma) D_d^{+T} . \tag{29}$$

Finally, the determinant of $G(\Sigma)$ can be expressed in terms of Σ alone:

$$|G(\Sigma)| \propto |\Sigma|^{d+1} . \tag{30}$$

The metric tensor features in the energy function for both symmetric and Hermitian PDHMC. For symmetric PDHMC, the energy is given by

$$E(\Sigma, V) \propto -\log \pi(\Sigma) - \frac{1}{2} \log |G(\Sigma)| + \frac{1}{2} \text{vech}(V)^\mathsf{T} G(\Sigma) \text{vech}(V) \tag{31}$$
$$\propto -\log \pi(\Sigma) - \frac{d+1}{2} \log |\Sigma| + \frac{1}{2} \text{vech}(V)^\mathsf{T} G(\Sigma) \text{vech}(V) ,$$

but the energy associated with Hermitian PDHMC is slightly different. In this case, both Σ and V are complex valued, and $\text{vech}(V)$ follows a multivariate complex Gaussian distribution with covariance $G^{-1}(\Sigma)$. Therefore, the Hermitian PDHMC energy is given by

$$E(\Sigma, V) \propto -\log \pi(\Sigma) - \log |G(\Sigma)| + \text{vech}(V)^H G(\Sigma) \text{vech}(V) \tag{32}$$
$$\propto -\log \pi(\Sigma) - (d+1) \log |\Sigma| + \text{vech}(V)^H G(\Sigma) \text{vech}(V)$$

where $(\cdot)^H$ signifies the conjugate transpose. Notice that the log-determinant and quadradic terms are not multiplied by the factor $1/2$. This accords with the density function of a complex Gaussian random variable. See Appendix A for more details.

The metric tensor (29) and the geodesic equations (11) and (12) are the only geometric quantities required for PDHMC. The t-th iteration of the symmetric PDHMC algorithm is shown in Algorithm 2. The t-th iteration of the Hermitian PDHMC algorithm is shown in Algorithm 3. First, one generates a Gaussian initial velocity on $T_{\Sigma_t} \mathcal{S}_d^+$ (Step 1). Then, the energy function is evaluated and stored (Step 2). Next, the system is numerically advanced using the split Hamiltonian scheme. Following Equation (23), the velocity vector V is updated one half-step with the gradient of $H^{[1]}$ (Step 4). For Step 5, both Σ and V are updated with respect to $H^{[2]}$, i.e. they are transported along the geodesic flow given by Equations (11) and (12):

$$\left[\Sigma(0), V(0) \right] \mapsto \left[\Sigma(\epsilon), V(\epsilon) \right] . \tag{33}$$

Again, the velocity vector V is updated one half-step with the gradient of $H^{[1]}$ (Step 6). Finally, the energy is evaluated at the new Markov state (Step 9), and a Metropolis accept-reject step is implemented (Steps 10-12). It is important to note that, besides being over different algebraic fields, the symmetric and Hermitian PDHMC only differ in their respective energies. The general implementation is the same. See Appendix A for a short discussion on gradients.

Algorithm 3 Hermitian PDHMC

1: $\text{vech}(V) \sim \text{CN}(0, G^{-1}(\Sigma_t))$

2: $e \leftarrow -\log \pi(\Sigma_t) - (d+1)\log|\Sigma_t| + \text{vech}(V)^H G(\Sigma_t)\text{vech}(V)$

3: $\Sigma^* \leftarrow \Sigma_t$

4: **for** $\tau = 1, \ldots, T$ **do**

5:      $\text{vech}(V) \leftarrow \text{vech}(V) + \frac{\epsilon}{2}G^{-1}(\Sigma^*)\text{vech}\left(\nabla_\Sigma\left(\log\pi(\Sigma^*) + (d+1)\log|\Sigma^*|\right)\right)$

6:      Progress (Σ^*, V) along the geodesic flow for time ϵ.

7:      $\text{vech}(V) \leftarrow \text{vech}(V) + \frac{\epsilon}{2}G^{-1}(\Sigma^*)\text{vech}\left(\nabla_\Sigma\left(\log\pi(\Sigma^*) + (d+1)\log|\Sigma^*|\right)\right)$

8: **end for**

9: $e^* \leftarrow -\log\pi(\Sigma^*) - (d+1)\log|\Sigma^*| + \text{vech}(V)^H G(\Sigma^*)\text{vech}(V)$

10: $u \sim U(0,1)$

11: **if** $u < \exp(e - e^*)$ **then**

12:      $\Sigma_{t+1} \leftarrow \Sigma^*$

13: **end if**

---

6. Results

This section features empirical validation of the PDHMC algorithm as well as an application to learning the spectral density matrix for vector time series. For empirical validation, we present quantile-quantile and trace plots comparing the PDHMC sample to the closed-form solution made available by the conjugate prior. We then use Hermitian PDHMC to learn the spectral density matrices of both simulated and LFP time series. We use the posteriors thus obtained to get credible intervals on the squared coherences for the vector time series.

6.1. Empirical validation

Before applying Hermitian PDHMC to spectral density estimation, we demonstrate validity by comparing samples from empirical posterior distributions of the Gaussian inverse-Wishart model obtained by PDHMC and the closed-form solution. We compare element-wise distributions with quantile-quantile plots and whole-matrix distributions with two global matrix measures. The comparisons based on 10,000 samples by the different sampling methods over 3-by-3 PD matrices are illustrated in Figure 1. The first 200 samples are discarded, and, for better visualization, every tenth sample is kept.

On the left panel of Figure 1, a quantile-quantile plot is used to compare the Hermitian PDHMC posterior sample to the closed-form posterior. Points are jittered for easy visualization, and each color specifies a different matrix element. Note that some colors appear twice: these double appearances correspond to real and imaginary matrix elements. For example, pink appears at zero as well as the upper right of the plot: this color corresponds to a diagonal matrix element since, on account of the matrix being Hermitian, its imaginary part is fixed at zero. Most importantly, all matrix elements fit tightly around the line $y = x$, suggesting a perfect match in quantiles between empirical distributions.

On the right panel of Figure 1, we present samples obtained from two whole-matrix statistics using symmetric PDHMC and the closed-form posterior. These measures are the effective variance (EV) and the effective dependence (ED):

$$\text{EV}(\Sigma) = |\Sigma|^{1/d}\,, \quad \text{and} \quad \text{ED}(\Sigma) = 1 - |\text{corr}(\Sigma)|^{1/d}\,. \tag{34}$$
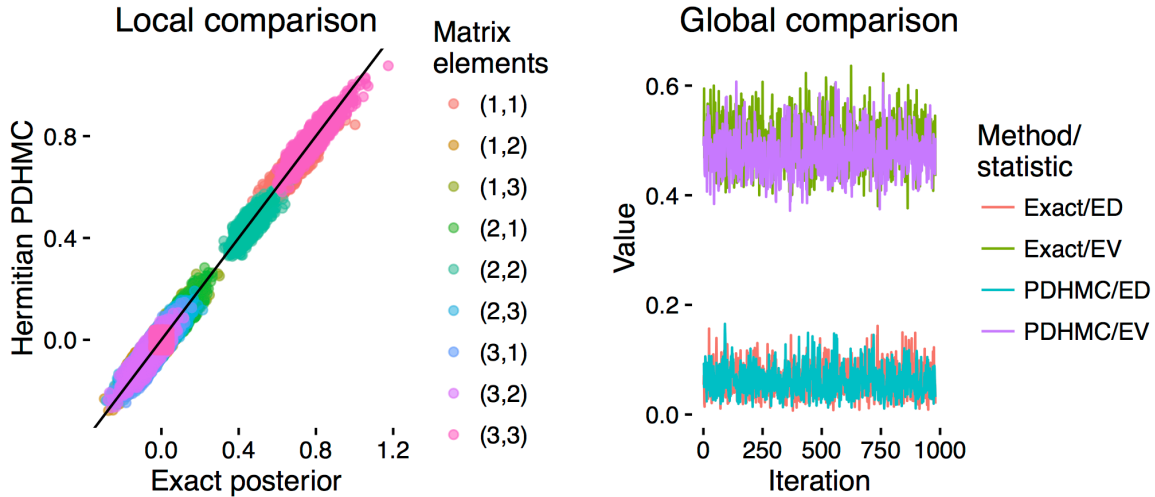
Figure 1: These figures provide empirical validation for the well-posedness of PDHMC. On the left is a quantile-quantile plot comparing the Hermitian PDHMC posterior sample with that of the closed-form posterior for the complex Gaussian inverse-Wishart model. Both real and imaginary elements are included, and points are jittered for visibility. On the right are posterior samples of 'global' matrix statistics pertaining both to (symmetric) PDHMC and the closed-form solution. These statistics are the effective variance and the effective dependence, built off the covariance matrix and the correlation matrix, respectively.

The EV is the geometric mean of the eigenvalues of the matrix Σ. It provides a dimension free measure of the total variance encoded in the matrix. The ED gets its name because the determinant of a correlation matrix is inversely related to the magnitude of the individual correlations that make up the off-diagonals. In addition to seeing that element-wise distributions match, one would also like to know that their joint distributions correspond. The EV and ED are good measures of global matrix features and here provide empirical evidence for the validity of PDHMC.

6.2. Learning the spectral density

An important benefit of PDHMC is that it enables practitioners to specify prior distributions other than the inverse-Wishart on PD matrices based on needs dictated by the problems at hand. PDHMC improves modeling flexibility. We use the problem of Bayesian spectral density estimation to demonstrate the possibility and advantage of using non-conjugate priors. The spectral density matrix $\Sigma(\omega)$ and its coherence matrix $R(\omega)$ are defined in Section 2. In the context of stationary, multivariate time series, the coherences that make up the off-diagonals of $R(\omega)$ provide a lag-free measure of dependence between univariate time series at a given frequency ω.

We compare posterior inference for these coherences between two models with different priors: the first model uses the complex inverse-Wishart prior; the second uses the complex reference prior [17]. The reference prior is an improper prior that has been proposed as an alternative to Jeffrey's prior for its superior eigenvalue shrinkage (which improves asymptotic efficiency of estimators). We use the reference prior to emphasize the flexibility allowed by PDHMC

11

Figure 2: Two 4-dimensional VAR1 time series and credible intervals for their 6 corresponding coherences measured at 20-40 Hz: the top row belongs to a block VAR1 process characterized by two independent 2-dimensional VAR1 time series; the bottom row belongs to a full VAR1 process. The left column shows the first 100 samples of both time series, each of which totals 5,000 samples in length. The right column shows credible intervals from posteriors obtained using the inverse-Wishart and reference priors.

but not as a modeling suggestion. [17] provide a Gibbs sampling routine based on the eigen-decomposition of the covariance matrix. The reference prior's form is provided in Appendix A.1.

We apply Hermitian PDHMC to learning the spectral density matrix for three distinct 4-dimensional time series. The first is a simulated first-order vector-autoregressive (VAR1) time series with block structure consisting of two independent, 2-dimensional VAR1 time series. The second time series is also VAR1 but with dependencies allowed between all four of the scalar time series of which it is composed. The third time series comes from local field potentials (LFP) recorded in the CA1 region of a rat hippocampus [18].

Figure 2 shows the first 100 samples from both VAR1 time series along with 95% posterior intervals for the complex moduli of the coherences. The time series are simulated with the form:

$$y(t) = \Phi\, y(t-1) + \epsilon_t, \qquad y(1) = \epsilon_1, \qquad \epsilon_t \sim N_4(0, I),\ t = 1, \ldots, 15000\,, \qquad (35)$$

where the eigenvalues of transition matrix Φ are bounded with absolute value less than 1 to induce stationarity. The first 10,000 data points are discarded to allow time for mixing. The

Figure 3: A 4-dimensional LFP signal with credible intervals for 6 coherences measured at 20-40 Hz (left) and 40-160 Hz (right). First 200 samples are shown for ease of visualization; the multi-dimensional time series totals 4,000 samples in length. Coherence profiles are remarkably similar between the two frequency bands considered.

first row of Figure 2 belongs to the block VAR1: Φ is a randomly constructed, block-diagonal matrix, so the first two scalar time series are independent from the second two. The second row of Figure 2 belongs the the second VAR1, all the scalar time series of which are dependent on all the others. Here Φ is also a randomly constructed matrix but is not block-diagonal. The intervals corresponding to the inverse-Wishart prior model are given in orange. The intervals corresponding to the reference prior model are given in blue. The true coherences are represented by black points and are obtained using the following closed-form formula for the spectral density of a VAR1 process [8]:

$$\Sigma(\omega) = \left(I - \Phi e^{-2\pi i \omega}\right)^{-1} Q \left(I - \Phi e^{-2\pi i \omega}\right)^{-H}. \tag{36}$$

Here Q is the covariance matrix of the additive noise ϵ_t, and $(\cdot)^{-H}$ denotes the inverse conjugate transpose. For the block VAR1 example, both models capture the true, non-null coherences (i.e. those given on the far left and the far right), but neither captures the null coherences. This is more than satisfactory, since coherences equal to zero imply the identity for a covariance matrix. By looking closely, one can see that the first and second time series (orange and green) are indeed strongly dependent on each other, as interval '1:2' suggests. For the full VAR1 example, both models capture five out of six true coherences, but the reference prior model gets closer to the truth than the inverse-Wishart model does.

13

We use the same tools to detect coherences between LFP signals simultaneously recorded from the CA1 region of a rat hippocampus prior to a memory experiment [18]. Two of the LFP signals are recorded on one end of the CA1 axis, and the other two LFP signals are recorded at the opposite end. Figure 3 shows the first 200 of 4,000 samples (recorded at 1,000 Hz) and 95% credible intervals for the coherences at two different frequency bands: 20-40 Hz and 40-160 Hz. *The spatial discrepancy is reflected in the posterior distributions of the individual coherences.* Both bands show similar coherence patterns, where spatial location appears to dictate strength of coherence: the leftmost and rightmost pairs are closer to each other in space, while the center pairs are farther from each other. This reflects what is apparent in the top of Figure (3), where the first and second time series (orange and green) are dependent, and the third and fourth time series (blue and purple) are dependent. These correspond to the intervals labeled '1:2' and '3:4', respectively. The credible intervals are smaller for the 20-40 Hz band because that band has only 1/6 the data of the 40-160 Hz band. Between prior models, the intervals differ more for the 40-160 Hz band. This is counter-intuitive since the influence of the prior distribution is often assumed to diminish with the size of the data set. One question is whether this surprising result is related to the reference prior's being the prior that is 'maximally dominated by the data' [19]. These differences—differences between posterior distributions for the two prior models—communicate that other prior distributions might provide tangible differences between results in spectral analysis and that it would be useful to understand which prior distributions are appropriate in which contexts.

7. Discussion

We presented PDHMC, a Lagrangian Monte Carlo method for Bayesian inference on the space of PD matrices. We outlined its relationship to other geometric extensions of HMC and showed how to apply PDHMC to both symmetric and Hermitian PD matrices. We demonstrated empirical validity using both element-wise and whole-matrix comparisons against the conjugate inverse-Wishart model. Finally, we applied Hermitian PDHMC to Bayesian spectral density estimation. The algorithm proved effective for detecting true coherences of simulated time series, as well as recovering spatial discrepancies between real-world LFP signals.

We see three branches of inquiry stemming from this work: the first is algorithmic; the second, theoretical; the third, methodological. First, what variations of HMC might help extend PDHMC into higher dimensions? There are multiple such extensions that are completely orthogonal to PDHMC. Examples are windowed HMC, geometric extensions to the NUTs algorithm, shortcut MCMC, and look-ahead HMC [5, 20, 21]. Auto-tuning will prove useful: even within the same dimension, different samples will dictate different numbers of leapfrog steps and step-sizes. From the theoretical standpoint, the canonical metric on the space of PD matrices is closely related to the Fisher information metric on covariance matrices: how should one characterize this intersection between information geometry and Riemannian symmetric spaces, and how might this relationship inform Bayes estimator properties or future variations on PDHMC? Finally, much work needs to be done in prior elicitation for Bayesian spectral density estimation. Which priors on Hermitian PD matrices should be used for which problems, what are the costs and benefits, and are there priors over symmetric PD matrices that need to be complexified (cf. [22, 23])? A clear delineation will be useful for practitioners in Bayesian time series research.

A. Real and complex matrix derivatives

The derivative of a univariate, real valued function with respect to a matrix is most cleanly calculated using the matrix differential. This is true whether $f : M_p(\mathbb{R}) \mapsto \mathbb{R}$ or $f : M_p(\mathbb{C}) \mapsto \mathbb{R}$, i.e. whether f is a function over real $p \times p$ matrices or complex $p \times p$ matrices. As an example, we consider the multivariate Gaussian distribution with mean 0 and covariance Σ. First, let f be the probability density function over real valued Gaussian random vectors $y_n \in \mathbb{R}^d, n = 1, \ldots, N$. Let Y be the $d \times N$ concatenation of these N i.i.d. random variables. Then the log density is given by

$$\log f(Y^N, \Sigma) \propto -\frac{N}{2} \log |\Sigma| - \frac{1}{2} \sum_{n=1}^{N} y_n^{\mathsf{T}} \Sigma^{-1} y_n \tag{37}$$

$$= -\frac{N}{2} \log |\Sigma| - \frac{1}{2} \mathrm{tr}\{\Sigma^{-1} Y Y^{\mathsf{T}}\} \ .$$

We apply the matrix differential to (37) using two general formulas:

$$d \log |\Sigma| = \mathrm{tr}\{\Sigma^{-1} \, d\Sigma\}, \quad \text{and} \quad d\Sigma^{-1} = -\Sigma^{-1} \, (d\Sigma) \, \Sigma^{-1} \ , \tag{38}$$

rendering

$$d \log f(Y, \Sigma) = -\frac{N}{2} \mathrm{tr}\{\Sigma^{-1} \, d\Sigma\} + \frac{1}{2} \mathrm{tr}\{\Sigma^{-1}(d\Sigma)\Sigma^{-1} Y Y^{\mathsf{T}}\} \tag{39}$$

$$= -\frac{N}{2} \mathrm{tr}\{(d\Sigma) \, \Sigma^{-1}\} + \frac{1}{2} \mathrm{tr}\{(d\Sigma) \, \Sigma^{-1} Y Y^{\mathsf{T}} \Sigma^{-1}\} \ .$$

Finally, we relate the matrix differential to the gradient with the fact that, for an arbitrary function g,

$$d \, g(\Sigma) = \mathrm{tr}\{(d\Sigma)A\} \quad \Longleftrightarrow \quad \nabla_{\Sigma} \, g(\Sigma) = A \ . \tag{40}$$

This gives the final form of the gradient of the log density function with respect to covariance Σ:

$$\nabla_{\Sigma} \log f(Y, \Sigma) = -\frac{N}{2} \Sigma^{-1} + \frac{1}{2} \Sigma^{-1} Y Y^{\mathsf{T}} \Sigma^{-1} \ . \tag{41}$$

For more on the matrix differential, see [24]. The complex matrix differential is treated in [25] and has a similar form real valued functions. The log density of the multivariate complex Gaussian with mean 0 is given by

$$\log f(Y, \Sigma) \propto -N \log |\Sigma| - \sum_{n=1}^{N} y_n^{H} \Sigma^{-1} y_n \tag{42}$$

$$= -N \log |\Sigma| - \mathrm{tr}\{\Sigma^{-1} Y Y^{H}\} \ ,$$

where $(\cdot)^{H}$ denotes the conjugate transpose. Note that the log density is scaled by a factor of two compared to the real case. The resulting gradient is

$$\nabla_{\Sigma} \log f(Y, \Sigma) = -N \Sigma^{-1} + \Sigma^{-1} Y Y^{H} \Sigma^{-1} \ . \tag{43}$$

A.1. The complex reference prior

Gradients of prior probabilities are calculated in a similar way. We demonstrate for the complex reference prior. Let $\lambda_i, i = 1, \ldots, d$ be the decreasing eigenvalues of Hermitian PD matrix Σ. Then the complex reference prior has the following form:

$$p(\Sigma) \propto \frac{d\Sigma}{|\Sigma| \prod_{k<j}(\lambda_k - \lambda_j)^2} \ . \tag{44}$$

To use the above approach for deriving the matrix derivatives, we need to be able to write the differential $d\lambda_i$ in terms of the matrix differential $d\Sigma$. [24] provides the formula when all eigenvalues are distinct:

$$d\lambda_i = \mathrm{tr}\Big(\sum_{j=1}^{d} V_{ij}^{-1} \Sigma^{j-1} d\Sigma \Big) \ , \tag{45}$$

where V is the Vandermonde matrix:

$$V^{\mathsf{T}} = \begin{vmatrix} 1 & \lambda_1 & \lambda_1^2 & \cdots & \lambda_1^{n-2} & \lambda_1^{n-1} \\ 1 & \lambda_2 & \lambda_2^2 & \cdots & \lambda_2^{n-2} & \lambda_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \lambda_n & \lambda_n^2 & \cdots & \lambda_n^{n-2} & \lambda_n^{n-1} \end{vmatrix} \ . \tag{46}$$

We now calculate the gradient of the log of the complex reference prior:

$$d \log p(\Sigma) = -d \log |\Sigma| - 2 \sum_{k<j} d \log(\lambda_k - \lambda_j) \tag{47}$$

$$= -\mathrm{tr}(\Sigma^{-1} d\Sigma) - 2 \sum_{k<j} \frac{d\lambda_k - d\lambda_j}{\lambda_k - \lambda_j}$$

$$= -\mathrm{tr}(\Sigma^{-1} d\Sigma) - 2 \sum_{k<j} \mathrm{tr}\Big(\sum_{i=1}^{d} \big(V_{ki}^{-1} - V_{ji}^{-1} \big) \Sigma^{i-1} d\Sigma \Big) / (\lambda_k - \lambda_j) \ .$$

Combining this with Equations (38) and (40) renders matrix gradient

$$\nabla_{\Sigma} \log p(\Sigma) \propto -\Sigma^{-1} - 2 \sum_{k<j} \Big(\sum_{i=1}^{d} \big(V_{ki}^{-1} - V_{ji}^{-1} \big) \Sigma^{i-1} \Big) / (\lambda_k - \lambda_j) \ . \tag{48}$$

References

[1] Tomoki Tokuda, Ben Goodrich, I Van Mechelen, Andrew Gelman, and F Tuerlinckx. Visualizing distributions of covariance matrices. *Dept. Statist., Columbia Univ., New York, NY, USA, Tech. Rep*, 2011.

[2] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

[3] Simon Byrne and Mark Girolami. Geodesic monte carlo on embedded manifolds. *Scandinavian Journal of Statistics*, 40(4):825–845, 2013.

[4] Shiwei Lan, Bo Zhou, and Babak Shahbaba. Spherical hamiltonian monte carlo for constrained target distributions. In *JMLR workshop and conference proceedings*, volume 32, page 629. NIH Public Access, 2014.

[5] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2:113–162, 2011.

[6] Shiwei Lan, Vasileios Stathopoulos, Babak Shahbaba, and Mark Girolami. Markov chain monte carlo from lagrangian dynamics. *Journal of Computational and Graphical Statistics*, 24(2):357–378, 2015.

[7] Babak Shahbaba, Shiwei Lan, Wesley O Johnson, and Radford M Neal. Split hamiltonian monte carlo. *Statistics and Computing*, 24(3):339–349, 2014.

[8] Hu L. Wang Y. and Ombao H. *Statistical Analysis of Electroencephalograms*, pages 523–565. CRC Press, 2016.

[9] Peter J Brockwell and Richard A Davis. *Time series: theory and methods*. Springer Science & Business Media, 2013.

[10] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006.

[11] Maher Moakher and Mourad Zéraï. The riemannian geometry of the space of positive-definite matrices and its application to the regularization of positive-definite matrix-valued data. *Journal of Mathematical Imaging and Vision*, 40(2):171–187, 2011.

[12] Sigurdur Helgason. *Differential geometry, Lie groups, and symmetric spaces*, volume 80. Academic press, 1979.

[13] Sebastian S Koh. On affine symmetric spaces. *Transactions of the American Mathematical Society*, 119(2):291–309, 1965.

[14] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.

[15] John Nash. The imbedding problem for riemannian manifolds. *Annals of mathematics*, pages 20–63, 1956.

[16] Shiwei Lan and Babak Shahbaba. *Sampling Constrained Probability Distributions Using Spherical Augmentation*, pages 25–71. Springer International Publishing, Cham, 2016.

[17] Lennart Svensson and Magnus Lundberg Nordenvaad. The reference prior for complex covariance matrices with efficient implementation strategies. *IEEE Transactions on Signal Processing*, 58(1):53–66, 2010.

[18] Timothy A Allen, Daniel M Salz, Sam McKenzie, and Norbert J Fortin. Nonspatial sequence coding in ca1 neurons. *The Journal of Neuroscience*, 36(5):1547–1563, 2016.

[19] James O Berger, José M Bernardo, and Dongchu Sun. The formal definition of reference priors. *The Annals of Statistics*, pages 905–938, 2009.

[20] Jascha Sohl-Dickstein, Mayur Mudigonda, and Michael R DeWeese. Hamiltonian monte carlo without detailed balance. *arXiv preprint arXiv:1409.5191*, 2014.

[21] MJ Betancourt. Generalizing the no-u-turn sampler to riemannian manifolds. *arXiv preprint arXiv:1304.1920*, 2013.

[22] Armin Schwartzman. Lognormal distributions and geometric averages of symmetric positive definite matrices. *International Statistical Review*, 2015.

[23] Farideh Fazayeli and Arindam Banerjee. The matrix generalized inverse gaussian distribution: Properties and applications. *arXiv preprint arXiv:1604.03463*, 2016.

[24] Jan R Magnus, Heinz Neudecker, et al. Matrix differential calculus with applications in statistics and econometrics. 1995.

[25] Are Hjorungnes and David Gesbert. Complex-valued matrix differentiation: Techniques and key results. *IEEE Transactions on Signal Processing*, 55(6):2740–2746, 2007.