Bayesian Estimation of Karhunen-Loève Expansions; A random subspace approach

Kenny Chowdhary, Habib N. Najm
Sandia National Laboratories, Livermore, CA

Abstract

One of the most widely-used procedures for dimensionality reduction of high dimensional data is Principal Component Analysis (PCA). More broadly, low-dimensional *stochastic* representation of random fields with finite variance is provided via the well known Karhunen-Loève expansion (KLE). The KLE is analogous to a Fourier series expansion for a random process, where the goal is to find an orthogonal transformation for the data such that the projection of the data onto this orthogonal subspace is optimal in the L^2 sense, i.e, which minimizes the mean square error. In practice, this orthogonal transformation is determined by performing an SVD (Singular Value Decomposition) on the sample covariance matrix or on the data matrix itself. Sampling error is typically ignored when quantifying the principal components, or, equivalently, basis functions of the KLE. Furthermore, it is exacerbated when the sample size is much smaller than the dimension of the random field. In this paper, we introduce a Bayesian KLE procedure, allowing one to obtain a probabilistic model on the principal components, which can account for inaccuracies due to limited sample size. The probabilistic model is built via Bayesian inference, from which the posterior becomes the matrix Bingham density over the space of orthonormal matrices. We use a modified Gibbs sampling procedure to sample on this space and then build probabilistic Karhunen-Loève expansions over random subspaces to obtain a set of low-dimensional surrogates of the stochastic process. We illustrate this probabilistic procedure with a finite dimensional stochastic process inspired by Brownian motion.

Keywords: Karhunen-Loève expansion, Principal Component Analysis, uncertainty quantification, Bayesian inference, matrix Bingham density, Gibbs sampling, Markov Chain Monte Carlo.

Email addresses: kchowdh@sandia.gov (Kenny Chowdhary), hnnajm@sandia.gov (Habib N. Najm)

Preprint submitted to Journal of Computational Physics

March 9, 2015

1. Introduction

The Karhunen-Loève theorem briefly states that a square integrable stochastic process can be represented by a linear combination of orthogonal functions, typically taken to be eigenfunctions of the covariance function of the stochastic process, with random coefficients [1]. To find the eigenfunctions of the covariance function, one can employ solvers for the Fredholm integral equation of the second kind [2], or, in the case of a discrete finite process, utilize Principal Component Analysis (PCA), which is a ubiquitous statistical procedure for model reduction of high-dimensional random data [3]. In this paper, we will consider the latter by only looking at discrete stochastic processes (or discretized versions of continuous stochastic processes).

The eigenfunctions in the Karhunen-Loève expansion (KLE) are known as the principal components or directions, and in many cases, since the covariance function is not known exactly, they are approximated from the sample covariance matrix. This can be done by performing an eigenvalue/ eigenvector decomposition of the sample covariance matrix or, more efficiently, by performing a singular value decomposition on the data itself. A consequence of working with sample data is that there is an underlying statistical uncertainty in computing these principal components. For different sets of realizations of the stochastic process, the sample covariance will change and thus the principal components will vary. Often this statistical variance, which is worse when the sample size is much smaller than the dimensionality of the stochastic process, is ignored, but a single set of principal components is still used. In the context of uncertainty quantification, it is important to understand and utilize the full probabilistic structure of the underlying quantity of interest - in this case being a stochastic process, which includes the principal components. Ignoring the full probabilistic structure while only using first or second order the summary statistics, e.g. means and/or variances, can lead to drastic under or over-estimates of quantities of interest [4]. More simply, it is incorrect to model the principal components deterministically. Other methods do exist, which take a slightly different route and place a distribution on the covariance matrix itself, incorporating the use of the Wishart and inverse Wishart distributions [5]. Often these methods require the sample covariance matrix to be fully ranked, and so one must explore algorithms that condition the covariance matrix, e.g., shrinkage methods, etc [6].

The method described in this paper does not require any type of conditioning or inversion of the covariance matrix directly.

In this paper, we derive the matrix Bingham density [7] for the principal components, which is a measure on the space of orthonormal matrices, i.e., the Stiefel manifold [8].
 35 This gives us a probabilistic characterization of the principal directions which best explain the data. In order to derive this density, we utilize the classical minimum reconstruction procedure [3] for determining the principal components. Once we obtain the appropriate density on the Stiefel manifold, we introduce a modified Gibbs sampling procedure, similar to the algorithm introduced by Hoff [7], to obtain samples of the principal components which
 40 can account for statistical uncertainty due to limited sample size. From this we obtain a collection of random subspaces onto which we can project our data and obtain random Karhunen-Loève expansions, i.e., low dimensional representations of our data.

Similar work can be found in topics related to probabilistic PCA algorithms and Factor Analysis, which formalize the problem in a more classical Bayesian framework, but utilize
 45 Expectation-Maximization algorithms to arrive at a single set of principal components [9]. Our work differs in that we can obtain multiple samples of the principal components on the manifold which describes the density of these principal components.

This paper is structured as follows. In Section 2 we briefly describe the mathematical setup and derivation of classical PCA. In Section 3 we derive the posterior density on the
 50 principal components in a Bayesian setting. Section 4 details the Gibbs sampling procedure used to sample from the density derived in the previous section. Section 5 describes how one can use the Bayesian framework to arrive at random KLE's. Finally, Section 6 illustrates these sampling methods on low and high-dimensional random processes.

2. Setup and Derivation of the Principal Components

55 Consider the space of all $m \times R$ real, orthonormal matrices, referred to as the Stiefel manifold, denoted by $\mathcal{V}_{R,m}$. The Stiefel manifold consists of matrices whose columns live on $\mathcal{S}^{m-1}$, i.e., the surface of the m -dimensional unit sphere. Given a collection of n realizations of an m -dimensional zero-mean stochastic process, x , or random field, denoted by $\{x_1, \dots, x_n\}$, $x_i \in \mathbb{R}^m$, the principal components can be derived using the following min-
 60 imum reconstruction argument. Note that if x is not a zero-mean process, we can simply consider $x - \bar{x}$, where $\bar{x}$ is the exact or sample mean. Let $\Phi \in \mathbb{R}^{m \times R}$ be some element in

$\mathcal{V}_{R,m}$ so that its columns form an orthonormal basis for some R -dimensional subspace in an m -dimensional space. Consider the reconstruction or projection error of the data onto this R -dimensional subspace Φ , i.e., the mean square error:

$$\text{Err}(\Phi) = \sum_{i=1}^n \|x_i - \Phi \Phi^T x_i\|^2, \quad (1)$$

65 where $\Phi \Phi^T \in \mathbb{R}^{m \times m}$ is the projection matrix, and $\|\cdot\|^2$ is the usual Euclidean norm. PCA attempts to find an orthonormal matrix Φ which has the least projection error,

$$\Phi^* = \arg \min_{\Phi \in \mathcal{V}_{R,m}} \sum_{i=1}^n \|x_i - \Phi \Phi^T x_i\|^2, \quad (2)$$

where the columns of $\Phi^* \in \mathbb{R}^{m \times R}$ are referred to as the R principal components. Equivalently, in the derivation of the the Karhunen-Loève expansion, one tries to find the eigenfunctions which minimize the mean square error. In practice, in order to determine the
70 principal components, one can solve (2) analytically. Expanding the sum for $\text{Err}(\Phi)$ gives

$$\text{Err}(\Phi) = \sum_{i=1}^n x_i^T x_i - x_i^T \Phi \Phi^T x_i. \quad (3)$$

Since the minimization in (2) is only over $\Phi \in \mathcal{V}_{R,m}$, the first term on the right hand side can be ignored. This results in

$$\Phi^* = \arg \max_{\Phi \in \mathcal{V}_{R,m}} \sum_{i=1}^n x_i^T \Phi \Phi^T x_i. \quad (4)$$

Let $X \in \mathbb{R}^{m \times n}$ denote the data matrix where the i^{th} column is x_i and recall that the trace of a matrix is the sum of the diagonal elements. Then, (4) can be more compactly written

75 as

$$\Phi^* = \arg \max_{\Phi \in \mathcal{V}_{R,m}} \text{tr}(\Phi^T (nS) \Phi), \quad (5)$$

where tr denotes the matrix trace and $S \in \mathbb{R}^{m \times m}$ is the sample covariance matrix, i.e., $S = n^{-1} X X^T$, where we have assumed x is a zero-mean stochastic process. Now, using (5) it can be shown that Φ^* is exactly the set of eigenvectors of S which have the R largest eigenvalues [3]. It can also be shown that $\text{tr}(\Phi^T S \Phi)$ represents the sum of the variances
80 along each orthonormal column in Φ . Thus, (5) shows an equivalence between the minimum reconstruction derivation of PCA (2) and the maximum variance derivation.

One can equate the optimization problem (5) to an iterative optimization procedure where we solve for Φ^* one column at a time. We briefly detail this approach since it

will be illustrative in understanding how to sample matrices on the Stiefel manifold while
85 retaining orthogonality. In order to solve (5) one column at a time, we can first find the
 m -dimensional orthonormal vector, $\phi_1 \in \mathbb{R}^m$ with $\phi_1^T \phi_1 = 1$, that maximizes $\phi_1^T S \phi_1$. ϕ_1
represents the direction onto which the data exhibits maximum variance, which turns out
to be the eigenvector of S with the largest eigenvalue. Next, in order to find the second
principal component, we seek another orthonormal vector, $\phi_2 \in \mathbb{R}^m$ such that $\phi_2^T \phi_2 = 1$
90 and $\phi_1^T \phi_2 = 0$, which maximizes $\phi_2^T S \phi_2$. The method of Lagrange multipliers yields the
eigenvector with the second largest eigenvalue. This process can be repeated in order to
determine the remaining columns of Φ^* . In the next section, we introduce the Bayesian
procedure for obtaining samples on the principal components. We will utilize this iterative
Bayesian procedure to obtain orthogonal samples.

95 3. Bayesian approach to PCA

Let us assume that the projection error can be modeled by *i.i.d* Gaussian white noise.
That is,

$$x - \Phi \Phi^T x = \eta, \quad (6)$$

where $\eta \in \mathbb{R}^m$ and $\eta \sim \mathcal{N}(0, \sigma^2 I)$. This assumption is consistent with the minimum recon-
struction derivation discussed in the previous section in the following way. Consider the
100 conditional density $p(x|\Phi, \sigma)$, which we refer to as the likelihood of the data. Solving for x
in (6) gives

$$x = (I - \Phi \Phi^T)^{-1} \eta, \quad (7)$$

where I is the $m \times m$ identity matrix. Then, $x \sim \mathcal{N}(0, (I - \Phi \Phi^T)^2)$ and the log likelihood
is proportional to to (1):

$$\log p(x|\Phi, \sigma) \propto -\frac{1}{2\sigma^2} \eta^T \eta, \quad (8)$$

which follows from (7). In an Approximate Bayesian Computation framework [10], one may
105 think of the likelihood, $p(x|\Phi, \sigma)$, for a fixed x and σ , as purely a function of Φ , denoted by
 $g(\Phi) \doteq p(x|\Phi, \sigma)$. Then, $g(\Phi)$ can be interpreted as a penalization or cost function for Φ ,
which we can use to minimize the mean square error.

Note that in general, one may certainly assume a non-*i.i.d* Gaussian noise model (6) and arrive at another probabilistic model for the principal components, other than (8).
 110 However, it is not guaranteed that this alternately derived model will take the form of a matrix Bingham density, which has nice sampling properties, as we will later see.

Now, let $\pi_{\mathcal{V}_{R,m}}(\Phi)$ be the uniform density on the Stiefel manifold for $\Phi \in \mathcal{V}_{R,m}$. Bayes' rule then gives

$$p(\Phi|x, \sigma) \propto p(x|\Phi, \sigma)\pi_{\mathcal{V}_{R,m}}(\Phi), \quad (9)$$

where the proportionality is up to a constant, which only depends on the data, and the
 115 noise term, $\sigma \in \mathbb{R}^+$, is fixed for now. If we let X be the data matrix where the columns are i.i.d. realizations of x , then (9) can be written more explicitly as

$$p(\Phi|X, \sigma) \propto \prod_{i=1}^n p(x_i|\Phi, \sigma)\pi_{\mathcal{V}_{R,m}}(\Phi), \quad (10)$$

Plugging (8) into (10), gives

$$p(\Phi|X, \sigma) \propto \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n \|x_i - \Phi\Phi^T x_i\|^2\right) \pi_{\mathcal{V}_{R,m}}(\Phi). \quad (11)$$

Finally, after some algebra, one can show that

$$p(\Phi|X, \sigma) \propto \text{etr}(\Phi^T(nS)\Phi/2\sigma^2)\pi_{\mathcal{V}_{R,m}}(\Phi), \quad (12)$$

where etr is the exponential trace of a matrix, $S \doteq n^{-1}XX^T$, and the prior probability
 120 distribution is over the Stiefel manifold, $\mathcal{V}_{R,m}$. The density in (12) is referred to as the matrix Bingham density, which is a density on $\mathcal{S}^{m-1}$ (see [8] for more details on the normalization constant).

Sampling from this density can be tricky because samples from the matrix Bingham
 density must be $(m \times R)$ -dimensional orthonormal matrices. To sample from (12) we propose
 125 breaking up the density into conditional densities over the individual columns of Φ , utilizing the chain rule for probability densities. Thus, we can write the posterior as

$$p(\Phi|X, \sigma) \propto p(\phi_1|X, \sigma)p(\phi_2|\phi_1, X, \sigma) \cdots p(\phi_R|\phi_1, \dots, \phi_{R-1}, X, \sigma), \quad (13)$$

where $\Phi = [\phi_1 \cdots \phi_R] \in \mathbb{R}^{m \times R}$, subject to orthonormality between $\phi_1, \dots, \phi_R$. To be clear,

the conditional densities in (13) are as follows:

$$\begin{aligned}
p(\phi_1|X, \sigma) &\propto \exp(\phi_1^T(nS)\phi_1/2\sigma^2), \quad \|\phi_1\|_2 = 1 \\
p(\phi_2|\phi_1, X, \sigma) &\propto \exp(\phi_2^T(nS)\phi_2/2\sigma^2), \quad \text{s.t. } \phi_2 \perp \phi_1, \|\phi_2\|_2 = 1 \\
&\vdots \\
p(\phi_R|\phi_1, \dots, \phi_{R-1}, X, \sigma) &\propto \exp(\phi_R^T(nS)\phi_R/2\sigma^2), \quad \text{s.t. } \phi_R \perp \phi_2, \dots, \phi_{R-1}, \|\phi_R\|_2 = 1.
\end{aligned} \tag{14}$$

For each conditional density we will use a Gibbs sampling procedure. Each conditional density is not independent of its predecessors so some maneuvering must be taken in order to sample over the correct orthogonal space, which is explained in the next section. In short, we will first show how to sample $p(\phi_1|X, \sigma)$, then use a simple linear transformation using the left null space of ϕ_1 to generate samples from $p(\phi_1|\phi_2, X, \sigma)$, and so on.

4. Sampling the Posterior Density

In this section, we will show how to sample from (13) where each conditional density is defined by (14) using Hoff's Gibbs sampling algorithm introduced in [7]. Note that the algorithm presented here is a slight modification of the aforementioned algorithm, which will give us a set of ordered vectors, analogous to retrieving principal components in order of decreasing variance. Without this modification, we lose the ordering of the principal components in order of decreasing variance. Nonetheless, both algorithms will produce principal components from the matrix Bingham density defined in (12).

First, in Section 4.1 we will describe how to sample the vector Bingham density, i.e. (12) for $R = 1$. Then, in Section 4.2 we will introduce a fairly simple trick to sample from the full matrix Bingham density using a left null space transformation.

4.1. Sampling the vector Bingham density

Consider the vector Bingham density on the m -dimensional sphere with respect to the uniform distribution over the unit sphere,

$$p(\phi|A) \propto \exp(\phi^T A \phi) \pi_{\mathcal{V}_{R,m}}(\Phi). \tag{15}$$

Without loss of generality, let us assume that $A \in \mathbb{R}^{m \times m}$ is symmetric. For the purposes of this paper, A is in fact proportional to the sample covariance matrix, i.e., $A = n\sigma^{-2}S/2$,

150 which is always semi-positive definite. Since this matrix A is symmetric, it always admits an eigenvector/eigenvalue decomposition

$$A = E\Lambda E^T, \quad (16)$$

where $E \in \mathbb{R}^{m \times m}$ is unitary and $\Lambda \in \mathbb{R}^{m \times m}$ is diagonal and has non-negative values. If we transform our density under the isometric mapping E , letting $y = E^T \phi$, then we can write (15) under the new variable y as

$$p(y|E, \Lambda) \propto \exp\left(\sum_{i=1}^m \lambda_i y_i^2\right) s(y), \quad (17)$$

155 where λ_i 's are the eigenvalues of S and $s(y)$ is the uniform density on $\mathcal{V}_{1,R}$. Note that since the change of variables is given by a linear, unitary mapping E , the determinant of the Jacobian is 1. Furthermore, we can write the probability density function for $s(y)$ w.r.t. to cartesian coordinates as

$$s(y_1, \dots, y_{m-1}) \propto \left(1 - \sum_{i=1}^{m-1} y_i^2\right)^{-1/2}, \quad \text{and} \quad y_m^2 = 1 - \sum_{i=1}^{m-1} y_i^2. \quad (18)$$

The derivation of this density is not trivial and is explained in [11, see §1.1, Eq(1)]. The basic idea is the following. Due to the radial symmetry of the standard multivariate normal density, if we let $y \sim \mathcal{N}(0, I)$, then

$$\hat{y} = \frac{y}{\|y\|}, \quad (19)$$

where $\|y\| = \sqrt{y_1^2 + \dots + y_m^2}$, is a uniform random variable on $\mathcal{S}^{m-1}$. Consequently, the density of $\hat{y}$ in (19) is exactly the density defined in (18) via a multivariate change of variables. To be clear, if $\{(y_{1,j}, \dots, y_{m-1,j})\}_{j=1}^n$ denotes n samples from $s(y_1, \dots, y_{m-1})$, then

$$\{(y_{1,j}, \dots, y_{m-1,j}, 1 - \sum_{i=1}^{m-1} y_{i,j}^2)\}_{j=1}^n,$$

represents n samples on the m -dimensional unit sphere. Thus, the density in (17) can be explicitly written as

$$p(y|E, \Lambda) \propto \exp\left(\sum_{i=1}^m \lambda_i y_i^2\right) \left(1 - \sum_{i=1}^{m-1} y_i^2\right)^{-1/2}, \quad \text{s.t.} \quad y_m^2 = 1 - \sum_{i=1}^{m-1} y_i^2. \quad (20)$$

A Gibbs sampling procedure can be performed to sample from (20), which means that we need to derive the one-dimensional conditional densities for (20). Hoff suggests performing

170 a simple transformation before deriving the conditional densities for $p(y|E, \Lambda)$, in order to improve the mixing of the Markov Chain [7, 11]. We briefly go over the suggested transformation.

In a straightforward Gibbs sampling procedure for (20), we need to sample from $p(y_i|y_{-i}, E, \Lambda)$ where

$$y_{-i} = (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_m) \in \mathbb{R}^{m-1}.$$

175 This conditional density is hard to sample from in practice, so we perform the following transformation. Let $\theta \doteq y_i^2$ and define

$$q_{-i} \doteq \frac{1}{1-\theta}(y_1^2, \dots, y_{i-1}^2, y_{i+1}^2, \dots, y_m^2),$$

so that $\{y_i^2, y_{-i}^2\} = \{\theta, (1-\theta)q_{-i}\}$. Then, after some calculation, Hoff shows that the conditional density, $p(y_i|y_{-i}, E, \Lambda)$, in terms of $\theta \doteq y_i^2$

$$p(\theta|q_{-i}, E, \Lambda) = \exp(\theta[\lambda_i - q_{-i}^T \lambda_{-i}])\theta^{1/2}(1-\theta)^{(m-3)/2}, \quad (21)$$

where λ_i is the i^{th} diagonal element of Λ . Since we are making a change of variables, we
180 will also need the determinant of the Jacobian which is given by the following:

$$\left| \frac{d\theta}{dy_i} \right| = 2|y_i| = 2\theta^{1/2} \quad \left| \frac{dq_j}{dy_j} \right| = 2 \frac{|y_j|}{1-y_{j2}} = 2q_j^{1/2}(1-\theta)^{1/2}, \quad j \neq i. \quad (22)$$

In order to sample (21) over $\theta \in (0, 1)$, we can proceed in either of two directions. The first is the most straightforward, but not the most efficient. In the first approach, we can either build the inverse cumulative distribution function (CDF) via interpolation and then sample based on the inverse CDF method, or evaluate $p(\theta|q_{-i}, E, \Lambda)$ at a set of uniform grid
185 points, weight them according to their PDF value, and then draw samples from this discrete density. A more efficient alternative is to use a rejection sampler. The target density (21) is of the form

$$p(\theta|q_{-i}, E, \Lambda) \propto \theta^{-1/2}(1-\theta)^{k-1}e^{\theta a}, \quad (23)$$

where $k = (m-3)/2$ and $a = \lambda_i - q_{-i}^T \lambda_{-i}$. In order to obtain a proper rejection sampler, we need a proposal density $f(\theta)$, also known as an envelope function, s.t. $Mf(\theta) >$
190 $p(\theta|q_{-i}, E, \Lambda)$ for some fixed constant $M > 0$. Since (23) is very similar to a beta density, Hoff proposes using a $\text{beta}(1/2, 1+k \wedge [(k-a) \vee -1/2])$ envelope which works well for many

choices of k and a [7]. Note that choosing the constant M is not trivial in practice. Please see the companion R implementation to [7] for a proper choice of M .

Under the Gibbs sampling approach the above procedure generates a Markov chain in
195 $\{y_1^2, \dots, y_m^2\}$ with a stationary distribution equal to $p(y_1^2, \dots, y_m^2 | E, \Lambda)$ [7]. The sign of y_i does not actually effect the density so it can be randomly assigned. The algorithm is summarized below.

Input: $A, \phi^{(0)} = e_1$, where e_1 is the canonical unit vector

Output: Gibbs sample from $p(\phi | A)$

Let $A = E^T \Lambda E$ and set $y = E^T \phi$;

for $i = 1, \dots, m$, *in random order* **do**

Set $\{q_1, \dots, q_m\} = \{y_1^2/(1 - y_i)^2, \dots, y_m^2/(1 - y_i)^2\}$;

Sample $\theta \in (0, 1)$ from $p(\theta | q_{-i}, E, \Lambda) \propto \theta^{-1/2}(1 - \theta)^k e^{\theta a}$ with $k = (m - 3)/2$ and

$a = \lambda_i - q_{-i}^T \lambda_{-i}$;

Sample s_i on $\{-1, +1\}$ using a binomial with $p = .5$;

Transform θ back to y : $y_i = s_i \theta^{1/2}$ and $y_j^2 = (1 - \theta)q_j$ for $j \neq i$;

end

Transform y back to ϕ : $\phi = Ey$;

Add new sample to the Markov chain: $\phi^{(1)} = \phi$

Algorithm 1: Gibbs sampler for the vector Bingham density.

Algorithm 1 can be repeated to obtain a Markov chain of samples from the vector
200 Bingham density. In practice, the mixing seems to work rather quickly, usually with a burn-in of only about five to ten samples. Figure 1 shows samples from a three-dimensional vector Bingham density using this Gibbs sampling procedure outlined above. Notice that the distribution is bi-modal because the vector Bingham density is antipodally symmetric, i.e., the density is invariant under a scalar multiplication by -1 . In the next section, we
205 explore how one obtains samples from the individual conditional densities in (14), i.e., a set of orthogonal vectors.

4.2. Sampling the matrix Bingham Density

In the previous section we discussed how to sample from the vector Bingham density.
210 Now, we will explain how to sample from the conditional densities given in (14). In general,

Figure 1: Samples of a three-dimensional vector Bingham density. Arrows indicate the direction of the principal eigenvector. Points in green indicate the samples from the vector Bingham density, while points in blue indicate samples from the Stiefel manifold, illustrating the surface of the sphere, $\mathcal{S}^2$.

suppose we want to sample from the vector Bingham density (15) subject to $\phi \perp \Psi$ where $\Psi \in \mathbb{R}^{m \times k}$ is some set of k orthonormal columns, with $k < R$. This density can be written as

$$p(\phi|A, \phi \perp \Psi) \propto \exp(\phi^T A \phi) \pi_{\mathcal{V}_{R,m}}(\phi|\phi \perp \Psi)(\phi) \quad (24)$$

In short, to sample from this density, we will write ϕ as a function of the left null space of Ψ and show that under this linear, isometric mapping, the density is again a vector Bingham density.

To show this, let $N \in \mathbb{R}^{m \times m-k}$ be an orthonormal basis for the left null space of $\Psi \in \mathbb{R}^{m \times k}$. That is, for $l \in \mathbb{R}^m$ in the span of the columns of N , we have $l^T \Psi = 0$. Now, if ϕ must be orthogonal to Ψ , then ϕ must be a linear combination of the columns of N , i.e., $\phi = Nz$ for some $z \in \mathbb{S}^{m-k}$, i.e., a vector on the $(m-k)$ -dimensional unit sphere. Then we can perform the following change of variables for (24)

$$\begin{aligned} p(\phi = Nz|A) &\propto \exp((Nz)^T A (Nz)) \pi_{\mathcal{V}_{1,m-k}}(z) \\ &= \exp(z^T \tilde{A} z) \pi_{\mathcal{V}_{1,m-k}}(z), \end{aligned} \quad (25)$$

where $\tilde{A} \doteq N^T A N$. Thus, (25) is again a vector Bingham density. This means we can use

Algorithm 1 again. In particular, we can use Algorithm 1 in an iterative fashion to sample from the conditional densities defined in (14), with A defined as $n\sigma^{-2}S/2$.

225 In summary, consider matrix Bingham density

$$p(\Phi|A) \propto \exp(\Phi^T A \Phi) \pi_{\mathcal{V}_{R,m}}(\Phi), \quad (26)$$

subject to $\Phi \in \mathcal{V}_{R,m}$, where $\Phi = [\phi_1, \dots, \phi_R]$. We can now sample from this density using the following algorithm.

Input: A , $\Phi^{(0)} = I$, where I is the $m \times R$ identity matrix.

Output: Set of Gibbs samples from $p(\Phi|A)$

Using Algorithm 1 generate M samples from $p(\phi|A)$, denoted by $\{\phi_1^{(1)}, \dots, \phi_1^{(M)}\}$;

for $r = 2, \dots, m$ **do**

for $j = 1, \dots, M$ **do**

 Let $\Psi_j \doteq [\phi_1^{(j)} \dots \phi_r^{(j)}]$, s.t. $\Psi \in \mathcal{V}_{R,r-1}$;

 Compute the Null space of Ψ_j , $N_j \in \mathbb{R}^{m \times m-r+1}$;

 Set $\tilde{A}_j = N_j^T A N_j$;

 Use Algorithm (1) to get sample of $z \sim \exp(z^T \tilde{A} z)$;

 Transform z to get sample of ϕ_r : $\phi_r^{(j)} = N_j z$;

end

end

Algorithm 2: Gibbs sampler for the matrix Bingham density using conditional densities.

Since the Gibbs sampling procedure for the vector Bingham density converges to the
 230 unique stationary distribution, giving us exact samples from each conditional density in
 (14), Algorithm (2) will converge to the target matrix Bingham density as well (see [12, 7]
 for more details about convergence of the Gibbs sampler).

In summary, Algorithm (2) obtains samples of the matrix Bingham density by breaking
 up the joint density into conditional densities, using the chain rule, for which each conditional
 235 density can be sampled via the vector Bingham algorithm. This allows flexibility in obtaining
 the orthonormal vectors Φ in two ways. First, by obtaining samples of the conditional
 densities, one can choose to increase R , i.e., the number of basis elements of Φ , adaptively
 if more basis terms are needed. Secondly, if one is probabilistically certain of a particular
 subset of the columns of Φ , i.e., one might know the first few columns exactly, then Algorithm
 240 (2) can be used to sample over the space orthogonal to the know subspace. This allows one
 to obtain a probability distribution on select columns of Φ only.

Hoff proposes a more classical Gibbs procedure, which is perfectly valid for a fixed choice of R [7]. Essentially, Hoff's Gibbs sampling algorithm runs over each individual column of Φ , while fixing all other columns simultaneously. He shows that this algorithm indeed generates a reversible, irreducible, aperiodic Markov chain for $R < m$. The difference between the algorithms introduced in this paper and Hoff's algorithm is that the principal components obtained in the latter formulation will not necessarily be ordered in decreasing projection error. In contrast, Algorithm (2) will indeed give us a set of ordered principal directions. This is more consistent with computational techniques involving SVD used to compute principal component vectors, which return the orthonormal vectors in order of decreasing variance. This ordering is useful in determining which basis terms to keep in the Karhunen-Loève expansion. Typically, the basis terms are chosen so that the cumulative energy, given by $\sum_{i=1}^j \lambda_i^2 / \sum_{i=1}^m \lambda_i^2$ where $j \leq m$ is the number of basis terms retained and λ_i 's are the variances along the respective ϕ_i directions, is above a prescribed threshold, i.e. 90%. In order to illustrate this point, Figure 2 shows samples of a three-dimensional matrix Bingham density with $m = 3$ and $R = 2$. Both algorithms provide samples from the same density, but Algorithm (2) provides samples in the *correct* ordering. Note that both algorithms provide subspaces which are more-or-less equivalent under rotation.

Figure 2: The samples on the left provide samples around the vectors determined by PCA, while Hoff's algorithm on the right provide equivalent samples, but under different rotations.

5. Random Karhunen-Loève Expansions

The Karhunen-Loève theorem states that one can represent any square integrable stochastic process as a linear combination of deterministic, orthogonal vectors, whose coefficients are uncorrelated, but not necessarily independent, random variables [13, 14]. To be precise, for a zero-mean, m -dimensional stochastic process $W \in \mathbb{R}^m$, and any set of m orthonormal basis functions $\Psi = \{\psi_1, \dots, \psi_m\}$ where $\psi_i \in \mathbb{R}^m$, one can write

$$W \stackrel{d}{=} \sum_{i=1}^m \alpha_i \psi_i, \quad (27)$$

where $\{\alpha_i = \langle W, \psi_i \rangle\}_{i=1}^m$ constitute a set of uncorrelated random variables induced by the projection of W onto each ψ_i , and equality is given in distribution. Typically, one chooses the basis functions, ψ_i , to be the eigenvectors of the covariance matrix of W (or the sample covariance matrix if one is only given samples of the stochastic process), where λ_i are the corresponding eigenvalues. In this way, the expansion (27) can be optimally chosen in the L^2 sense if one chooses eigenvectors in decreasing order of eigenvalues, i.e. ψ_i are ordered such that $\lambda_1 \geq \dots \geq \lambda_m$. In particular, Parseval's theorem gives

$$\mathbb{E} \left[\left(W - \sum_{i=1}^R \alpha_i \psi_i \right)^2 \right] = \sum_{i=R+1}^m \lambda_i^2,$$

where $R \leq m$. Thus, for $R < m$, if the residual eigenvalues are small, the k -dimensional approximation, $\sum_{i=1}^R \alpha_i \psi_i$, to W may be a sufficient representation of the original stochastic process.

In the present context, we no longer have a single set of principal components, but rather a set of M random subspaces $\{\Phi^{(1)}, \dots, \Phi^{(M)}\}$, where $\Phi^{(i)} \in \mathcal{V}_{R,m}$ are sampled from (5) via the Gibbs sampling procedure outlined in Algorithms (1) and (2). Thus, each set of R orthonormal basis functions, $\Phi^{(i)}$, admits an associated Karhunen-Loève approximate to the stochastic process W , such that we have, in aggregate, the M KLEs:

$$\left\{ \sum_{j=1}^R \alpha_j^{(i)} \phi_j^{(i)} \right\}_{i=1}^M, \quad (28)$$

where $\Phi^{(i)} = \{\phi_1^{(i)}, \dots, \phi_R^{(i)}\}$ contains R orthonormal m -dimensional columns. For each fixed i , the distribution on coefficients, $\{\alpha_1^{(i)}, \dots, \alpha_R^{(i)}\}$ can be determined by projecting samples of W onto each basis vector $\phi_j^{(i)}$. That is, samples of $\{\alpha_1^{(i)}, \dots, \alpha_R^{(i)}\}$ are given by

$$\{\langle \phi_1^{(i)}, x_j \rangle, \dots, \langle \phi_R^{(i)}, x_j \rangle\}_{j=1}^M. \quad (29)$$

It is important to note that even though the individual $\alpha_j^{(i)}$'s are uncorrelated random variables, they may not be independent. Thus, the full joint density must be determined in most cases, unless independence is known. For a Gaussian process W , uncorrelated, in fact, implies independence, and, moreover, one can show that the $\alpha_j^{(i)}$'s are independent, zero-mean normal random variables with variance $\mathbb{E}[\langle \phi_j^{(i)}, W \rangle^2]$. In all other cases, in order to obtain the full joint density on $\{\alpha_1^{(i)}, \dots, \alpha_R^{(i)}\}$ from the projection samples (29), one could use kernel density estimation (KDE) along with the inverse Rosenblatt transformation, to produce a polynomial chaos expansion (PCE) for the $\alpha_j^{(i)}$'s [15, 16]. A thorough discussion of alternative methods for inferring the full joint density on the coefficients goes beyond the scope of this paper, so we refer the readers to the references.

6. Examples

In this section, we illustrate the approaches introduced in the previous sections on a discretized version of a continuous time, square integrable stochastic process, given by

$$W_t = \sum_{k=1}^3 \frac{\xi_k}{\sqrt{\pi(k - \frac{1}{2})}} \sqrt{2} \sin\left(\left(k - \frac{1}{2}\right) \pi t\right), \quad (30)$$

for $t \in [0, 1]$, $\xi_k \sim \mathcal{N}(0, 1)$, where $\{\sqrt{2} \sin((k - \frac{1}{2}) \pi t)\}_{k=1}^3$ are the first three eigenfunctions of the covariance function for standard Brownian motion, and $\{1/\sqrt{\pi(k - \frac{1}{2})}\}_{k=1}^3$ are the corresponding eigenvalues. We discretize in time to obtain an m -dimensional approximation, $W_m = (w_1, \dots, w_m) \in \mathbb{R}^m$, to (30), where

$$w_j = \sum_{k=1}^3 \frac{\xi_k}{\sqrt{\pi(k - \frac{1}{2})}} \sqrt{2} \sin\left(\left(k - \frac{1}{2}\right) \pi \frac{j}{N}\right), \quad (31)$$

for $j = 1, \dots, m$. Figure 3 shows the the first three PCA modes, in absolute value, of W_m with $m = 100$, alongside realizations of this finite-dimensional stochastic process.

6.1. Samples from the matrix Bingham density

Suppose we are given n realizations of our stochastic process, where $n \ll m$ (this is not a requirement but is indicative of a scenario in which we have very few samples of a discrete random process relative to the dimensionality). Assuming a fixed noise parameter

Figure 3: Left: first three PCA modes for W_m . Absolute value is shown since the principal vectors are invariant under scalar multiplication by -1. Right: realizations of the stochastic process W_m .

value, σ , we use the Bayesian KLE approach outlined in Algorithm (2) to obtain samples of the principal components. Figure 4 shows the spread of samples from the matrix Bingham density (12) for $n = 10$, illustrated by the shaded regions, color coded for each principal component. The dashed lines show the exact principal components and the solid lines show the mean. Figure 4 also displays a single realization from the matrix Bingham density, which lives on the Stiefel manifold $\mathcal{V}_{3,100}$.

Figure 4: Left: shaded regions represent $\pm 2\sigma$ error bars for principal components sampled from the matrix Bingham density. The dashed lines represent the exact PCA modes, which are the eigenvectors of the true covariance matrix. The solid lines represent the mean values. Right: one sample from the matrix Bingham density. In this example, $n = 10$, $m = 100$, and $\sigma = .1$. Again, absolute values are shown since the density is antipodally symmetric.

If the noise parameter, σ , is not known, the Gibbs sampling procedure makes it fairly easy to obtain a posterior on the noise, given an appropriate choice of a prior. In fact, by

choosing the conjugate prior on $1/\sigma^2$ to be $\text{gamma}(\alpha, \beta)$, then

$$p(1/\sigma^2 | \Phi, X) \sim \text{gamma} \left(\alpha + \frac{n}{2}, \beta + \frac{1}{2} \sum_{i=1}^n \|x_i - \Phi \Phi^T x_i\|^2 \right), \quad (32)$$

where $p(\Phi | X, \sigma)$ is defined in (12). Figure 5 shows samples of the matrix Bingham density for $n = 10$ when we impart a gamma prior on $1/\sigma^2$. Note the similarity to the results shown in Figure 4. In general, if one does not choose a conjugate prior on σ , one can perform a Metropolis-Hastings accept/reject scheme for σ , for every sample of Φ , and then iterate between the two. This approach is known as Metropolis-within-Gibbs sampling or block MCMC [17].

Figure 5: Left: shaded regions represent $\pm 2\sigma$ error bars at each point in time for samples from the matrix Bingham density. The dashed lines represent the exact PCA modes, which are the eigenvectors of the true covariance matrix. The solid lines represent the mean values. Right: Samples from the matrix Bingham density. In this example, $n = 10$, $m = 100$, and the prior on $1/\sigma^2$ is given by $\text{beta}(100, .1)$. Again, absolute values are shown since the density is antipodally symmetric.

An important observation from these figures is that the samples from the matrix Bingham density exhibit a large amount of fluctuation, compared with the analytical modes, or the vectors obtained via traditional PCA, which, in general, seems far smoother as a function of t . This is not a consequence of the Gibbs sampling algorithm, but rather a consequence of the assumption on the data likelihood in (6), which assumes that the projection error is i.i.d Gaussian white noise. This assumption is equivalent to having a matrix Bingham density for the principal directions (12), and, furthermore, indicates that the manifold defined by the matrix Bingham density does not impose any smoothness constraint, and, in fact, gives more weight to *noisy* realizations. In other words, samples from the matrix Bingham density are inherently *noisy* as illustrated by Figures 5, 6, and 7. If one does require some degree

of regularity on the principal component samples, then one option is to use less noisy data, or additional measurements, if available. Alternatively, one can impose a prior on Φ in (12) which tends to smooth the samples. For example, if D is the m -dimensional, first-order, finite difference operator, then one might consider

$$p(\Phi|X, \sigma) \propto \text{etr}(\Phi^T n(S - \delta D^T D)\Phi/2\sigma^2), \quad (33)$$

where $\delta > 0$ is a tunable parameter which penalizes the columns of Φ for having a high total variation or squared difference. This is by no means the only prior that imposes regularity, but, rather we want to emphasize that the Bayesian framework allows for regularization via an appropriately chosen prior.

Before we move on to the random Karhunen-Loeve expansions, we illustrate how the uncertainty about the principal components decreases as the sample size, n , increases. In particular, Figures 6 and 7 show the uncertainty in the principal vectors for higher n values. Note the reduction in the spread of the samples, shown in shaded regions of color, compared with Figures 5 and 4.

Figure 6: Left: shaded regions represent $\pm 2\sigma$ error bars at each point in time for samples from the matrix Bingham density. The dashed lines represent the exact PCA modes, which are the eigenvectors of the true covariance matrix. The solid lines represent the mean values. Right: samples from the matrix Bingham density. In this example, $n = 25$, $m = 100$, and the prior on $1/\sigma^2$ is given by $\text{gamma}(100, .1)$. Compare this with Figures 4 and 5, where $n = 10$.

6.2. Random KLE

Each sample from the matrix Bingham density, $\Phi^{(i)}$, for $i = 1, \dots, M$, admits a Karhunen-Loeve expansion, where the coefficients are independent, zero-mean Gaussians (this is only

Figure 7: Left: shaded regions represent $\pm 2\sigma$ error bars at each point in time for samples from the matrix Bingham density. The dashed lines represent the exact PCA modes, which are the eigenvectors of the true covariance matrix. The solid lines represent the mean values. Right: samples from the matrix Bingham density. In this example, $n = 250$, $m = 100$, and the prior on $1/\sigma^2$ is given by $\text{gamma}(100, .1)$. Compare this with Figures 4, 5, and 6 where $n = 25$.

valid because W_m is a Gaussian process). To determine the variance for each projection coefficient, we use (29) to obtain samples and then compute the empirical variance. The set of random approximates (28) can now be used to generate new sample data, where each set lives on the subspace defined by $\Phi^{(i)}$. Figure 8 shows the original 25 samples versus a set of samples generated from (28).

Figure 8: (left) $n = 25$ samples from the stochastic process, W . (right) Samples from the random Karhunen-Loeve approximates. We take $m = 100$ and impose a gamma prior on the noise (See Figure 5).

Figure 9 shows samples from the Karhunen-Loeve approximates when we increase the number of samples by a factor of ten. Note that realizations exhibit slightly less variation due to higher certainty in the principal components (see Figure ??).

As per the discussion in the previous section, because the principal components exhibit

Figure 9: (left) $n = 250$ samples from the stochastic process, W . (right) Samples from the random Karhunen-Loève approximates. We take $m = 100$ and, again, impose a gamma prior on the noise (See caption of Figure 5).

variability (See Figures 4 and 5 (right)), the samples from the KLE, which are linear combinations of the principal components, exhibit similar variation (see Figure 8 (right)). These fluctuations would be smaller if the data had less noise, with a smaller associated σ^2 in (12), or if there is additional data. Alternatively, they can be reduced by imposing a smoothing prior for Φ (33). However, as previously discussed, this is a natural consequence of the matrix Bingham density, which stems from assumption that our projection error is *i.i.d.* white noise (6). Regardless, summary statistics such as means, variances, and correlations can still be computed and can reasonably approximate the statistics of the true, underlying stochastic process. In fact, for most types of summary statistics, for which *smoothness* is not a necessity, e.g., $P(W(t) \in [a, b] | t \in [t_1, t_2])$, one can interpret these *noisy* realizations in Figures 4 and 5 (right) as samples of the stochastic process from non-contiguous realizations.

7. Conclusion

In this paper we formulated and demonstrated a Bayesian procedure to obtain the basis functions for the Karhunen-Loève Expansion of a square-integrable stochastic process, which allowed for uncertainty in the principal components as a function of the sample size and data noise. We derived the matrix Bingham density on the Stiefel manifold as the posterior density in this Bayesian approach in the case where the projection error is, or may be reasonably modeled as, *i.i.d.* Gaussian white noise. We also introduced a modified Gibbs sampling procedure, based on work by Hoff [7], to sample from this density. The modification

380 allowed us to obtain an orthonormal basis in order of decreasing variance, similar to how the
 principal components are computed numerically via classical Singular Value Decomposition
 (SVD). Moreover, the Bayesian framework allowed flexibility in the use of regularization
 in the form of priors on the principal components. After samples were obtained from the
 matrix Bingham density, we computed random Karhunen-Loeve expansions to generate
 realizations of the original stochastic process. This probabilistic characterization of the
 principal components is important in the context of uncertainty quantification so that one
 385 can accurately predict the effects of sample size for any quantities of interest which depend
 on these stochastic processes.

It is reasonable to consider, of course, the utility of the construction in a context where
 the Gaussian assumption on the projection error is *not* justified. In this case, the above
 matrix Bingham density is not useful, although the general framework we present is valid.
 390 Application to the non-Gaussian case would, indeed, require formulation of alternate means
 for efficient MCMC sampling of orthonormal vectors, but otherwise the construction is
 applicable in general. Alternatively, one possible means of resolving this difficulty would
 be to employ approximate Bayesian computation methods, also known as Likelihood-free
 methods [10], where a kernel density is used as a pseudo-likelihood to enforce the constraint
 395 minimizing the weighted L2 norm of the projection error.

8. Acknowledgement

Support for this work was provided through the Scientific Discovery through Advanced
 Computing (SciDAC) program funded by the U.S. Department of Energy, Office of Science,
 Advanced Scientific Computing Research. Sandia National Laboratories is a multi-program
 400 laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of
 Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Secu-
 rity Administration under contract DE-AC04-94-AL85000.

References

- [1] O. P. L. Maitre, O. M. Knio, Spectral Methods for Uncertainty Quantification with
 405 applications to computational fluid dynamics, Springer, New York, 2010.

- [2] O. Moklyachuk, Simulation of random processes with known correlation function with the help of Karhunen-Love decomposition, *Theory of Stochastic Processes* 13 (29) (2008) 163–169.
- [3] I. T. Jolliffe, *Principal Component Analysis*, Springer, New York, 2002.
- 410 [4] T. Needham, A visual explanation of Jensen’s inequality, *The American Mathematical Monthly* 100 (8) (1993) 768–771.
- [5] M. Bouriga, O. Fron, Estimation of covariance matrices based on hierarchical inverse-wishart priors, *Journal of Statistical Planning and Inference* 143 (4) (2013) 795 – 808.
- 415 [6] G. Cao, L. R. Bachega, C. A. Bouman, The sparse matrix transform for covariance estimation and analysis of high dimensional signals, *Image Processing: IEEE Transactions* 64 (1) (2011) 27–43.
- [7] P. D. Hoff, Simulation of the matrix Bingham-von Mises-Fisher distribution with applications for multivariate and relational data, *Journal of Computational and Graphical Statistics* 18 (2) (2013) 438–456.
- 420 [8] Y. Chikuse, *Statistics on special manifolds*, Springer, New York, 2003.
- [9] M. E. Tipping, C. M. Bishop, Probabilistic principal component analysis, *Journal of the Royal Statistical Society* 61 (3) (1999) 611–622.
- [10] B. M. Turner, T. V. Zandt, A tutorial on Approximate Bayesian Computation, *Journal of Mathematical Psychology* 56 (2) (2012) 69 – 85.
- 425 [11] A. Kume, S. G. Walker, Sampling from compositional and directional distributions, *Statistics and Computing* 16 (3) (2006) 261–265.
- [12] G. Casella, G. E. I., Explaining the gibbs sampler, *The American Statistician* 46 (3) (1992) 167–174.
- 430 [13] R. G. Jaimez, J. C. Ruiz, M. J. Valderrama, On the numerical expansion of a second order stochastic process, *Applied stochastic models and data analysis* 8 (2) (1992) 67–77.
- [14] R. G. Jaimez, M. J. Bonnet, On the Karhunen-Love expansion for transformed processes, *Trabajos de Estadistica* 2 (2) (1987) 81–90.

- [15] R. G. Jaimez, M. J. Bonnet, On the Karhunen-Loève expansion for transformed processes, *Trabajos de Estadística* 2 (2) (1987) 81–90.
- [16] K. Sargsyan, B. Debusschere, H. Najm, O. L. Matre, Spectral representation and reduced order modeling of the dynamics of stochastic reaction networks via adaptive data partitioning, *SIAM Journal on Scientific Computing* 31 (6) (2010) 4395–4421.
- [17] S. Chib, E. Greenberg, Understanding the metropolis-hastings algorithm, *The american statistician* 49 (4) (1995) 327–335.