

Data Analysis using Riemannian Geometry and Applications to Chemical Engineering

Alexander Smith[¶], Benjamin Laubach[†], Ivan Castillo[†], and Victor M. Zavala^{¶*}

[¶]Department of Chemical and Biological Engineering

University of Wisconsin-Madison, 1415 Engineering Dr, Madison, WI 53706, USA

[†]Chemometrics, AI and Statistics

The Dow Chemical Company, 332 SH 332 E, Lake Jackson, TX 77566,

Abstract

We explore the use of tools from Riemannian geometry for the analysis of symmetric positive definite matrices (SPD). An SPD matrix is a versatile data representation that is commonly used in chemical engineering (e.g., covariance/correlation/Hessian matrices and images) and powerful techniques are available for its analysis (e.g., principal component analysis). A key observation that motivates this work is that SPD matrices live on a Riemannian manifold and that implementing techniques that exploit this basic property can yield significant benefits in data-centric tasks such classification and dimensionality reduction. We demonstrate this via a couple of case studies that conduct anomaly detection in the context of process monitoring and image analysis.

1 Introduction

The assumption that data lies in a Euclidean space is pervasive throughout science and engineering and is the basis of diverse data analysis techniques used in these domains. Making this blanket assumption, however, is not always appropriate and can affect the accuracy/interpretability of such techniques or even break fundamental physical laws.

Recognizing that data can live in spaces that are governed by *non-Euclidean* geometry is critical to appropriately representing, manipulating, and analyzing certain data objects. A simple example of this arises when computing travel paths between a couple of points that are located on the surface of the Earth; when computing the distance between such points, the *elliptic* geometry of the Earth surface must be taken into consideration. If Euclidean geometry is assumed, travel paths between antipodal locations (e.g., United States and China) can require traversing the Earth (not physically-realizable paths).

Elliptic geometry is a non-Euclidean geometry in which one of the postulates of Euclidean geometry (sum of interior angles of a triangle is equal to 180°) no longer holds, because of the presence of positive curvature in a surface [7]. Elliptic geometry was first proposed by Bernhard Riemann

*Corresponding Author: victor.zavala@wisc.edu

in the 19th century, and he further developed these ideas into the field that is now known as Riemannian geometry [30].

An important example in which assuming Euclidean geometry can lead to spurious results is in the analysis of *symmetric positive definite* (SPD) matrices (e.g., covariance/correlation matrices). SPD matrices lie on a high-dimensional space which is governed by Riemannian geometry (known as a Riemannian manifold) [19]. Standard techniques for the analysis of SPD matrices (e.g., PCA or basic matrix norms) do not take this property into consideration and can lead to misleading results. For instance, computing the distance between SPD matrices in Euclidean geometry (e.g., via the Frobenius norm) ignores the fact that such matrices live on a Riemannian manifold, and this can yield misleading results [12, 2, 34]. Specifically, the so-called *swelling effect* can occur when applying operations in Euclidean geometry to SPD matrices [2]. This effect introduces spurious results by inflating the determinants of SPD matrices and can also distort the results of commonly used methods [34]. Computing interpolations and averages of SPD matrices, which is key in understanding physical systems (e.g. Brownian motion), can also break physical conservation laws if performed under Euclidean geometry [24].

In this work, we focus our attention on the use of techniques from Riemannian geometry for the analysis of data objects that can be represented as SPD matrices. An SPD matrix is a simple but versatile data representation that is widely used in multivariate analysis techniques such as PCA [29, 39, 42]. SPD representations are also used in process control, monitoring, and anomaly detection [11, 16, 20, 40, 31, 33], in the study of functional brain networks [28, 35, 38, 14, 33], in object detection [36, 25, 41], in biomedical image analysis [24, 27], in the analysis of Laplacian matrices in graph theory, and in the analysis of Hessian matrices in optimization [22]. In applications such as image analysis, an SPD representation can be obtained by applying transformations to a raw data object (e.g., smoothing via kernel functions and/or combination of image features). The defining feature of an SPD matrix is that all its eigenvalues are real and positive. In the context of optimization, it is well-known that an SPD Hessian matrix defines positive curvature of a multi-dimensional quadratic function and is key in defining the geometry of objective and constraint functions (e.g., convex or non-convex). In the context of statistics, it is well-known that an SPD covariance matrix defines a multi-dimensional ellipsoid (a surface with positive curvature) and that the level sets of a multivariate Gaussian probability density function are ellipsoids.

This paper provides a practical introduction to the Riemannian geometry of SPD matrices and demonstrates applications of interest to the process systems engineering community. Specifically, we illustrate the benefits of exploiting the Riemannian geometry of SPD matrices and discuss how these tools can be incorporated into common dimensionality reduction and classification workflows. We also provide example applications of interest in science and engineering. The first application focuses on analysis of covariance matrices derived from multivariate time series, which is a common task in process monitoring. Our application focuses on the so-called Tennessee Eastman Process (TEP) [9]. The TEP is a process where anomalies/faults are systematically introduced which shifts the relationships between the measured variables. Covariance matrices encode these changing relationships and are then used to predict what type of anomaly the process is experiencing. The second application is in defect/anomaly detection of textiles taken from the MVTEC AD dataset [4]. Here, grayscale images of textiles are represented as covariance matrices by characterizing the relationships between the original image and multiple transformations of the image

(e.g., smoothing with different kernels). The image transformations emphasize different features of the original image (e.g., edges and fibers). Subsequently, when an anomaly/defect is introduced (e.g., a cut or discoloration of the textile) these relationships will change, impacting the covariance matrix and allowing us to detect anomalies. This type of analysis can also be applied to other relevant image/field datasets such as those arising in microscopy and flow cytometry [33, 32]. All data and scripts needed to conduct such analyses is provided as open-source code in <https://github.com/zavalab/ML/tree/master/RiemannianSPD>. With this, we aim to provide a concise and easy introduction to non-experts to the field of Riemannian geometry.

2 Riemannian Manifolds

The key observation driving this work is that SPD matrices lie on a Riemannian manifold; we thus begin our discussion by characterizing such manifolds. In this section, we aim to provide an intuitive understanding of manifolds and equip the reader with knowledge of their key properties.

An n -dimensional manifold is a topological space (a space where closeness and connectedness are defined but not directly measurable) that *locally* resembles an n -dimensional Euclidean space. Specifically, an n -dimensional manifold M is a set where every point $p \in M$ has an open neighborhood $U \subset M$ (known as a *chart*) that can be mapped to an open set of n -dimensional Euclidean space $V \subset \mathbb{R}^n$ via a one-to-one, onto, and continuous mapping $f : U \rightarrow V$ (i.e., a homeomorphism). The set of charts whose union covers the manifold is known as an *atlas* $\bigcup_{i=1}^n U_i = M$. The chart/atlas nomenclature is derived from navigation along the surface of the Earth (a 2D manifold); here, charts are flat (Euclidean) maps of the Earth that are collected in an atlas.

In Figure 1 we illustrate multiple topological spaces embedded in a 3D Euclidean space. The spiked (b) and smooth (c) hollow spheres are examples of 2D manifolds. The space in (a) represents a couple of cones that intersect at a single point and is not manifold; this is because a neighborhood drawn around the intersecting point will look like a smaller version of the intersecting cones, which cannot be mapped to 2D Euclidean space through a homeomorphism.




Figure 1: (a) Space composed of a couple of cones intersecting at a single point. This is a non-manifold space because any neighborhood formed around the intersecting point is not homeomorphic to 2D Euclidean space (the neighborhood is a smaller version of the two intersecting cones). (b) Represents a 2D manifold (all points and associated neighborhoods can be mapped to 2D Euclidean space) but is not a differentiable manifold because of the cusps occurring at the edges of the manifold (differential is not defined everywhere). (c) A smooth sphere is a 2D manifold that is also differentiable (curves on the surface can be differentiated everywhere).

Manifolds can also be endowed with geometric structure; for our analysis, we are particularly interested in whether or not a given manifold is *differentiable*. In simple terms, a differentiable

manifold is a manifold for which calculus (e.g., computing derivatives and integrals) can be performed on the charts that make up the manifold atlas [18]. This also means that *curves* on the surface of the manifold can be analyzed from a geometric perspective using calculus. A curve is defined as a continuous function $\gamma : [a, b] \rightarrow M$ mapping the interval $[a, b] \in \mathbb{R}$ to the manifold M . We will not cover the specific mathematical requirements that make a manifold differentiable in general, but refer interested readers to the following reference for details [18]. Examples of a differentiable and non-differentiable manifold are shown in Figure 1. We can see that the spiked sphere in (b) has multiple cusps where a derivative cannot be defined for a curve, making it non-differentiable. For the smooth sphere in (c), a derivative can be taken anywhere, allowing for more complex operations/transformations to be performed on the manifold [23].

We now restrict our attention to a special class of manifolds known as *Riemannian manifolds*. A Riemannian manifold is a differentiable manifold equipped with a defined *tangent space* at each point in the manifold $p \in M$, denoted as $T_p M$ [19]. The tangent space is the set of tangent vectors of all curves passing through point $p \in M$. The tangent space $T_p M$ is a vector (linear) space that is of the same dimension as the manifold itself. An illustration of a tangent space is shown in Figure 2 for a smooth sphere (a 2D Riemannian manifold). For a Riemannian manifold, the tangent space $T_p M$ is equipped with an inner product $g_p : T_p M \times T_p M \rightarrow \mathbb{R}$, along with a norm metric $|\cdot|_p : T_p M \rightarrow \mathbb{R}$ defined by $|v|_p = \sqrt{g_p(v, v)}$ for any vector $v \in T_p M$. These properties allow us to define the *length* of a curve on the manifold surface. A differentiable curve $\gamma : [a, b] \rightarrow M$ assigns to each $t \in (a, b)$ a tangent vector $\gamma'(t) \in T_{\gamma(t)} M$; thus, to obtain the length of the curve $L(\gamma)$, we integrate the norm of the tangent vectors along the curve (i.e., arc length):

$$L(\gamma) := \int_a^b |\gamma'(t)|_{\gamma(t)} dt \quad (2.1)$$

In our analysis, we are primarily interested in measuring the *shortest* curve between a couple points on the manifold (known as a *geodesic*). Given a couple of points on a Riemannian manifold $p, q \in M$ and the set of all curves $\gamma : [a, b] \rightarrow M$ such that $\gamma(a) = p$ and $\gamma(b) = q$, the geodesic $\bar{\gamma}$ is the curve with the shortest total length $L(\bar{\gamma})$:

$$L(\bar{\gamma}) := \inf\{L(\gamma) \mid \gamma : [a, b] \rightarrow M, \text{ with } \gamma(a) = p, \gamma(b) = q\} \quad (2.2)$$

An illustration of the geodesic between a couple of points on the smooth sphere is shown in Figure 2. Geodesics are a powerful tool for quantifying the relationship between points on a manifold surface and can be used to compute summarizing statistics for points on the surface (such as means and variances) [23].

There are direct relationships between the tangent space of a Riemannian manifold and geodesics, such as the *exponential map* and the *logarithmic map*. For a tangent vector $v \in T_p M$ constructed at point $p \in M$, there exists a unique geodesic $\gamma : [0, 1] \rightarrow M$ such that $\gamma(0) = p$ and $\gamma'(0) = v$. The vector $v \in T_p M$ is mapped to the endpoint of the geodesic $\gamma(1) \in M$ through the exponential map (see Figure 2):

$$\exp_p(v) = \gamma(1). \quad (2.3)$$

The inverse of the exponential map is the logarithmic map, which maps the point $\gamma(1)$ in the neighborhood of $p \in M$ to a vector in the tangent space $v \in T_p M$ (see Figure 2):




Figure 2: (a) Illustration of a Riemannian manifold (M) and the associated tangent space at a point $p \in M$. (b) Illustration of a geodesic $\gamma(t)$ constructed between two points $p, q \in M$, along with the associated tangent space vector $v \in T_p M$. The exponential map ($\exp_p(v) : v \rightarrow q$) and the logarithmic map ($\log_p(q) : q \rightarrow v$) are also shown.

$$\log_p(\gamma(1)) = v \quad (2.4)$$

These functions provide a mapping from the surface of the manifold to the tangent space of a given point. The tangent space, which is a (linear) vector space, can be analyzed using standard techniques designed for Euclidean spaces (e.g., classification, regression, and dimensionality reduction) while properly capturing the relationships between points defined by the geometry of the manifold (e.g., geodesics). For instance, we can use these constructs to analyze and operate on the space of SPD matrices, as we discuss next.

3 Riemannian Geometry of SPD Matrices

This section will provide the mathematical background needed to understand the geometry of the space defined by SPD matrices. We first introduce the needed notation and define the properties of different matrix spaces that are used in our analysis. We then demonstrate how the properties of these matrix spaces can be used to quantify the geometry and structure of the space defined by SPD matrices. The main message is that SPD matrices lie on a Riemannian manifold and that important computations (e.g., matrix operations, summarizing statistics, classification, regression, and dimensionality reduction) can be performed by respecting the geometry of this manifold by conducting these on the tangent space; respecting such properties can lead to important improvements in efficiency and interpretability.

3.1 Matrix Spaces, Properties, and Notation

We define spaces and properties of matrices that reflect the structure of SPD matrices. We denote $\mathcal{S}(n) := \{\mathbf{S} \in \mathcal{M}(n), \mathbf{S} = \mathbf{S}^T\}$ as the set of symmetric $n \times n$ matrices in the space of square, real matrices $\mathcal{M}(n)$ and the set $\mathcal{P}(n) := \{\mathbf{P} \in \mathcal{S}(n), u^T \mathbf{P} u > 0, \forall u \in \mathbb{R}^n\}$ as the set of all $n \times n$ SPD matrices.

We also define the exponential and logarithmic mappings used in our analysis. The matrix exponential $\exp(\mathbf{P})$, where $\mathbf{P} \in \mathcal{S}(n)$, is defined as:

$$\exp(\mathbf{P}) := \mathbf{U} \text{diag}(\exp(\lambda_1), \dots, \exp(\lambda_n)) \mathbf{U}^T \quad (3.5)$$

where $\mathbf{U}$ represents the matrix of eigenvectors of $\mathbf{P}$ and $\lambda_1 > \dots > \lambda_n$ represent the eigenvalues of $\mathbf{P}$ (also denoted as $\lambda_i(\mathbf{P})$). We also define the inverse operation; the matrix logarithm $\log(\mathbf{P})$ as:

$$\log(\mathbf{P}) := \mathbf{U} \text{diag}(\log(\lambda_1), \dots, \log(\lambda_n)) \mathbf{U}^T. \quad (3.6)$$

The following properties should also be considered in the analysis [21]:

- $\forall \mathbf{P} \in \mathcal{P}(n)$ we have that $\det(\mathbf{P}) > 0$
- $\forall \mathbf{P} \in \mathcal{P}(n)$ we have that $\mathbf{P}^{-1} \in \mathcal{P}(n)$
- $\forall \mathbf{P} \in \mathcal{P}(n)$ we have that $\log(\mathbf{P}) \in \mathcal{S}(n)$
- $\forall \mathbf{S} \in \mathcal{S}(n)$ we have that $\exp(\mathbf{S}) \in \mathcal{P}(n)$

We also define the Frobenius inner product for matrices $\mathbf{A}$ and $\mathbf{B}$ as:

$$\langle \mathbf{A}, \mathbf{B} \rangle_F := \text{Tr}(\mathbf{A}^T \mathbf{B}), \quad (3.7)$$

where $\text{Tr}(\cdot)$ represents the matrix trace operator. The Frobenius norm for a matrix $\mathbf{A}$ is:

$$\|\mathbf{A}\|_F = \sqrt{\text{Tr}(\mathbf{A}^T \mathbf{A})} = \sqrt{\sum_{i=1}^n \lambda_i(\mathbf{A})}. \quad (3.8)$$

3.2 Manifold of SPD Matrices

To understand the Riemannian geometry of the space of SPD Matrices ($\mathcal{P}(n)$), we first need to construct a Riemannian *metric*, which will allow us to compute distances and other relationships between points in a manifold. The metric we consider for Riemannian manifolds is known as the *Affine Invariant Riemannian Metric (AIRM)*; a detailed derivation of the metric is found in the work of Bhatia [5]. We aim to provide an intuitive understanding of this metric, and its meaning.

Given SPD matrices $\mathbf{A} \in \mathcal{P}(n)$ and $\mathbf{B} \in \mathcal{P}(n)$, we construct a geodesic $\gamma(t) : [0, 1] \rightarrow \mathcal{P}(n)$ parameterized as:

$$\gamma(t) = \exp((1-t) \cdot \log(\mathbf{A}) + t \cdot \log(\mathbf{B})) \quad (3.9)$$

where we have $\log$ maps $\mathcal{P}(n) \rightarrow \mathcal{S}(n)$ and $\exp$ maps $\mathcal{S}(n) \rightarrow \mathcal{P}(n)$. Here, we are using the $n \times n$ identity matrix $\mathbf{I} \in \mathcal{P}(n)$ as our tangent space basis. Informally, we are mapping our matrices from our SPD manifold $\mathcal{P}(n)$ to the tangent vector space $\mathcal{S}(n)$ via the logarithmic map, constructing a line between these points in the tangent space, and then projecting the constructed line back to the manifold via the exponential map. This is guaranteed to be the shortest length path between the points in $\mathcal{P}(n)$. The corresponding geodesic distance between matrices $\mathbf{A}$ and $\mathbf{B}$ is given by:

$$d_g(\mathbf{A}, \mathbf{B}) := \|\log(\mathbf{A}) - \log(\mathbf{B})\|_F. \quad (3.10)$$

Here, note that we are simply projecting the matrices from the Riemannian manifold to the tangent vector space prior to measuring their distance using the Frobenius norm [5]. However, an important consideration must be made when using the $n \times n$ identity matrix $\mathbf{I}$ as the tangent space basis. In many applications, the data may lie within a particular neighborhood that is far from $\mathbf{I}$ on the manifold. Thus, projections to the tangent space $T_I M$ can result in distortions of the data [23]. Intuitively, one can think of this as similar to an analysis of a projection of the Earth surface onto a plane tangent to the $0'$ latitude and $0'$ longitude point (as many maps are represented). In this projection, landmasses near the edge of the projection are highly distorted, whereas points near $(0', 0')$ have almost no distortion. Thus, our aim is to identify a distance with respect to a tangent space defined by the matrices of concern. This distance can be constructed using a critical property of this Riemannian manifold and metric: *congruence invariance*.

Congruence invariance states that, for any $n \times n$ invertible matrix $\mathbf{X}$ and matrices $\mathbf{A}, \mathbf{B} \in \mathcal{P}(n)$:

$$d_g(\mathbf{X}^T \mathbf{A} \mathbf{X}, \mathbf{X}^T \mathbf{B} \mathbf{X}) = d_g(\mathbf{A}, \mathbf{B}) \quad (3.11)$$

We thus have that linear transformations of the given matrices do not impact the geodesic distance on the manifold. This property allows us to redefine the geodesic distance as:

$$d_g(\mathbf{A}, \mathbf{B}) = d_g(\mathbf{I}, \mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2}) \quad (3.12a)$$

$$= \|\log(\mathbf{I}) - \log(\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})\|_F \quad (3.12b)$$

$$= \|\log(\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})\|_F \quad (3.12c)$$

$$= \left(\sum_{i=1}^n \log^2 \lambda_i(\mathbf{A}^{-1} \mathbf{B}) \right)^{1/2} \quad (3.12d)$$

where $\mathbf{I}$ is the $n \times n$ identity matrix, $\mathbf{A} = \mathbf{A}^{1/2} \mathbf{A}^{1/2}$, and $\lambda_i(\mathbf{A}^{-1} \mathbf{B})$ are the eigenvalues of $\mathbf{A}^{-1} \mathbf{B}$.

We can simplify the geodesic distance further; we have that the geodesic between matrices $\mathbf{I}$ and $\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2}$ is:

$$\gamma_0(t) = \exp(\log(\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})t) = (\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})^t \quad (3.13)$$

We can then leverage congruence invariance to shift this geodesic to our matrices of interest as:

$$\gamma(t) = \mathbf{A}^{1/2} (\gamma_0(t)) \mathbf{A}^{1/2} = \mathbf{A}^{1/2} (\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})^t \mathbf{A}^{1/2} \quad (3.14)$$

where $\gamma(0) = \mathbf{A}$ and $\gamma(1) = \mathbf{B}$, which provides a geodesic that is independent of the tangent space basis. Essentially, we are leveraging congruence invariance to translate our points to a neighborhood of $\mathbf{I}$, which allows us to compute distances with minimal distortion, and then translate the points back through these linear transformations. We can apply this same logic to the exponential map (and the logarithmic map); for matrices $\mathbf{A}, \mathbf{B} \in \mathcal{P}(n)$ and $\mathbf{T}_B \in T_A M$ where $T_A M \subset \mathcal{S}(n)$:

$$\mathbf{B} = \exp_A(\mathbf{T}_B) = \mathbf{A}^{1/2} \exp(\mathbf{A}^{-1/2} \mathbf{T}_B \mathbf{A}^{-1/2}) \mathbf{A}^{1/2} \quad (3.15)$$

$$\mathbf{T}_B = \log_A(\mathbf{B}) = \mathbf{A}^{1/2} \log(\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2}) \mathbf{A}^{1/2} \quad (3.16)$$

Here, we are mapping the matrix $\mathbf{B} \in \mathcal{P}(n)$ to the tangent space vector $\mathbf{T}_B \in T_A M$ which is centered at $\mathbf{A} \in \mathcal{P}(n)$ via $\log_A(\mathbf{B})$ and inversely through $\exp_A(\mathbf{B})$. Thus, with these newly defined geodesics and mappings, we are able to compute relationships between SPD matrices with minimal distortions in the tangent space. However, when given a large dataset with multiple SPD matrices, the choice of a tangent space basis may not be immediately clear. In this case, the geometric mean of the matrices on the manifold is typically identified and used as reference point.

3.3 SPD Matrix Means and Tangent Spaces

In the analysis of a set of SPD matrices, we often need to identify a center point on the SPD manifold that will minimize the distortion of all geometric relationships between the matrices of the dataset when mapped to the tangent space. This matrix is the (Riemannian) geometric mean of the matrices [44]. For a set of SPD matrices $\mathbf{A}_i$, the geometric mean (see Figure 3) is the matrix $\bar{\mathbf{A}}$ that minimizes the sum of squared geodesic distances to all other matrices in the set:

$$\bar{\mathbf{A}} := \operatorname{argmin}_{\mathbf{M}} \sum_{i=1}^n \|\log(\mathbf{M}^{-1/2} \mathbf{A}_i \mathbf{M}^{-1/2})\|_F \quad (3.17)$$

We can see that the geometric mean is obtained by solving a matrix optimization problem. For the SPD manifold this problem is geodesically convex (similar to Euclidean convexity) [1]. This optimization problem can be solved by using classical optimization algorithms that have been adapted to geometric setting (e.g., gradient descent) [1]. A detailed review on these approaches can be found in the work by Absil, Mahony, and Sepulchre [1].




Figure 3: (a) Illustration of the geometric mean $\bar{\mathbf{A}}$ of a set of matrices $\{\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3, \mathbf{A}_4\} \in \mathcal{P}(n)$. The geometric mean represents a point on the manifold that minimizes the geodesic to all other matrices in the set. (b) Representation of the tangent space $T_{\bar{\mathbf{A}}}M$ at the geometric mean. The set of matrices $\mathbf{A}_i$ are projected (through the logarithmic map) onto this tangent space with minimal geometric distortion $\mathbf{T}_{A_i} \in T_{\bar{\mathbf{A}}}M$.

Given a set of SPD matrices $\mathbf{A}_i$ and a geometric mean $\bar{\mathbf{A}}$, we can construct a tangent space at the geometric mean $T_{\bar{\mathbf{A}}}M$, and project the SPD matrices onto the tangent space $\mathbf{T}_{A_i} = \log_{\bar{\mathbf{A}}}(\mathbf{A}_i)$, as shown in Figure 3. The matrices are now represented in a vector (linear) space that reflects the geometry of the SPD manifold. Projecting the data into a vector space allows us to apply common matrix analysis methods such as PCA.

4 Case Study : Process Monitoring




Figure 4: Simplified illustration of the Tennessee Eastman Process (TEP) and resulting multivariate time series sensor data. Process sensors measure values such as temperature, pressure, flow, and level. From the multivariate time series data it is difficult to distinguish whether there is a fault occurring, or what type of fault may be occurring. Thus, a simplified and informative representation of the data is needed to be able to distinguish when the process is behaving normally or is experiencing a particular fault.

We focus on data obtained from a simulated industrial process known as the Tennessee Eastman Process (TEP) [9]. This dataset is a widely used benchmark dataset for testing and comparing various anomaly (i.e. fault) detection methods [17, 43, 8]. Figure 4 provides a high-level illustration of the process along with the multivariate time series data that is produced by the sensors monitoring the process. The process has a total of 52 measurements, 41 are process variables, 11 are manipulated variables. There are 20 different potential faults, which are defined in Table 1 (for further details see Appendix A in [9]). Our aim is to use geometric methods to detect and classify the presence and type of fault using only the multivariate sensor data.

The difficulty in distinguishing potential faults in the TEP is illustrated in Figure 4, where multivariate time series sensor signals for the 52 monitored variables are shown for the TEP without faults, and in the presence of fault 7 (step change in the component "C" header pressure) and fault 12 (random variation in the condenser cooling water inlet). The complexity of the multivariate process dynamics and of the number of sensor measurements make it difficult to reliably identify if a fault is occurring and to distinguish between fault types. Our method focuses on simplifying the data by quantifying the relationships between the 52 measured variables through covariance matrices and leveraging the geometry of if covariance matrices to detect and distinguish faults in the TEP.

4.1 Data Pre-Processing

To begin our analysis of the TEP, we must first pre-process the TEP data into multiple covariance matrices. To accomplish this, we represent each of the 52 measured variables as a univariate random variable x_i where $i = 1, 2, \dots, 52$ and we denote the collection of signals as a multivariate random vector $\mathbf{X} = (x_1, \dots, x_n)$ where $n = 52$. We denote the observations of each signal at time $t = 1, 2, \dots, m$ as $x_i(t) \in \mathbb{R}^m$. We use this representation to construct the sample covariance matrix for the process data $\mathbf{P} \in \mathcal{P}(n)$ as:

Table 1: Types of Faults for Tennessee Eastman Process [9].

Fault ID	Fault Name	Type
Fault 1	A/C feed ratio, B composition constant (stream 4)	Step
Fault 2	B composition, A/C ratio constant (stream 4)	Step
Fault 3	D feed temperature (stream 2)	Step
Fault 4	Reactor cooling water inlet temperature	Step
Fault 5	Condenser cooling water inlet temperature	Step
Fault 6	A feed loss (stream 1)	Step
Fault 7	C header pressure loss - reduced availability (stream 4)	Step
Fault 8	A, B, C feed composition (stream 4)	Random variation
Fault 9	D feed temperature (stream 2)	Random variation
Fault 10	C feed temperature (stream 4)	Random variation
Fault 11	Reactor cooling water inlet temperature	Random variation
Fault 12	Condenser cooling water inlet temperature	Random variation
Fault 13	Reaction kinetics	Slow drift
Fault 14	Reactor cooling water valve	Sticking
Fault 15	Condenser cooling water valve	Sticking
Fault 16-20	Unknown	Unknown




Figure 5: Representation of the data pre-processing workflow for the geometric analysis of the TEP data. (a) Covariance matrices are constructed from the sensor data multivariate time series forming a set of SPD matrices ($P_i \in \mathcal{P}(n)$). (b) The geometric mean of the matrices ($\bar{P} \in \mathcal{P}(n)$) is identified as the point that minimizes the squared geodesic distance to all other points. (c) All derived covariance matrices are mapped to the tangent space $T_{\bar{P}}M$ through the logarithmic mapping. This maps the matrices into a vector space that reflects the manifold geometry. The mapped data can then be analyzed with commonly defined dimensionality reduction and classification/regression methods.

$$P := \frac{1}{m-1} XX^T \quad (4.18)$$

The TEP dataset consists of multiple separate simulations of the process, both with and without faults. Thus, for each simulation we construct a sample covariance matrix P_i . Our goal is to pair each simulation sample covariance matrix with the fault occurring in the simulation. Figure

5 provides examples of the sample covariance matrices that are constructed from simulations containing no faults, fault 7, and fault 12. We note that there are no obvious differences between the covariance matrices that would identify a given fault. These covariance matrices lie on the SPD manifold, and can be integrated into our geometric framework. An illustration of this computational workflow is found in Figure 5. We compute the geometric mean $\bar{\mathbf{P}}$ for our set of covariance matrices $\mathbf{P}_i \in \mathcal{P}(n)$, the matrices are then projected to the tangent space $T_{\bar{\mathbf{P}}}M$ centered at the geometric mean via the logarithmic mapping $\mathbf{T}_{\mathbf{P}_i} = \log_{\bar{\mathbf{P}}}(\mathbf{P}_i)$. The data is now projected into a vector space that retains the geometric characteristics of the SPD matrices with minimal distortion, and can be integrated in dimensionality reduction and classification algorithms to perform analysis.

4.2 Principal Geodesic Analysis

Mapping the process data covariance matrices to the tangent space provides an avenue for the application of common dimensionality reduction techniques. Here, we apply PCA to the matrices mapped to the tangent (vector) space. PCA applied on the tangent space of the SPD manifold is commonly known as Principal Geodesic Analysis (PGA), as it identifies the geodesics that capture the most variance in the data [15]. An example comparison of PGA versus PCA (directly on the covariance matrices) is presented in Figure 6. The simulations with no faults are colored in red and the faulty simulations are represented by different grayscale values. We can see that using only the first two components in PGA, we are able to perfectly separate the faulty and non-faulty simulations; on the other hand, when applying PCA directly on the covariance matrices we can see that there is significant overlap between the faulty and non-faulty simulations. The comparison of these projections demonstrates that capturing the geometry of the Riemannian manifold in the analysis of covariance matrices can improve the performance with minimal added complexity.

This improvement in separation of the data through the geometric approach is due in part to the congruence invariance of our defined metric on the Riemannian manifold. As previously stated, congruence invariance means that any $n \times n$ invertible matrix $\mathbf{X}$ applied to a set of covariance matrices $\mathbf{P}_i \in \mathcal{P}(n)$ does not impact the geodesic distance between the two matrices:

$$d_g(\mathbf{X}^T \mathbf{P}_i \mathbf{X}, \mathbf{X}^T \mathbf{P}_j \mathbf{X}) = d_g(\mathbf{P}_i, \mathbf{P}_j) \quad (4.19)$$

Operations such as re-scaling and normalization, which can be represented algebraically as invertible square matrices, have no impact on the geodesic distance between the matrices [3]. Therefore, there is no need to select specific scaling or normalization strategies when applying our geometry based analysis of the data (e.g., PGA). However, this is not true when ignoring the data geometry, making methods such as PCA susceptible to the chosen framework (or lack of) for normalization/scaling. We perform no scaling prior to PCA in this case to ensure a direct comparison between PCA and PGA.

4.3 Classification and Clustering Results

PGA analysis of the covariance matrices also reveals that there is definite clustering of the data with respect to the different fault types within the TEP dataset. This suggests that classification of the different fault types can be done directly using the tangent space of the SPD manifold. To investigate this we use a simple linear (ridge) classifier. We compare the prediction accuracy of the linear classifier using coefficients of the tangent space projected matrices versus the coefficients of





Figure 6: Comparison of PCA applied to the raw covariance matrices and PCA applied to data mapped onto the tangent space (PGA). The red points represent the simulations where no fault is occurring and the grayscale points represent simulations with different faults. (a) PCA on the raw covariance matrices shows minimal separation in the data; faultless simulations are overlapped with faulty simulations. (b) PCA performed in the tangent space provides perfect separation between the faulty and faultless simulations, and also shows clustering of the faulty systems into separate groups. This demonstrates that simple considerations for the geometry of the data can yield improved results.


the non-transformed covariance matrices as input. In the analysis, we perform a simple train-test split of the data, where 30% of the data is used for testing and 70% of the data is use for training. Figure 7 illustrates the dramatic increase in accuracy when the model incorporates geometric information, which is reflected in the normalized confusion matrices. Here, a value $x \in [0, 1]$ on the diagonal indicates an accuracy of $x * 100\%$ when classifying a particular fault. All values in the off diagonal (e.g., row i , column j , where $i \neq j$) represent the percentage of covariance matrices associated with fault i that have been incorrectly labeled as experiencing fault j . When the SPD manifold is accounted for via the tangent space projection, there is perfect classification of the data (with the exception of faults 3, 4, 9 and 15). When the manifold geometry is ignored, there are few instances where high classification accuracy is achieved. The faults 3, 4, 9 and 15 have been shown in prior work to be difficult to classify [43]. We also note that these faults are only misclassified within their group (are never classified as having no fault), which suggests that there is limited quantifiable difference in the covariance matrices for these faults. The inclusion of more information around these particular faults may correct this issue and further increase accuracy.

5 Case Study - Image Analysis

Another important application of covariance matrices is in the analysis of images [26, 10, 6]. Here, we focus on an analysis of real images from textile manufacturing that contain non-defective and defective woven textiles taken. The images were obtained from the public MVTEC AD dataset [4]. Example images of non-defective and defective textiles are found in Figure 8. Our goal is to classify textiles using a linear classifier and Riemannian geometry.




(a) Classification raw covariance matrices.



(b) Classification in tangent space ($T_{\bar{P}}M$).

Figure 7: Comparison of linear classification on the raw covariance matrices versus matrices mapped to the tangent space $T_{\bar{P}}M$. (a) Classification of the covariance matrices without regard for the data geometry results in poor classification accuracy. (b) Simple mapping of the data to the appropriate tangent space provides a dramatic improvement in classification accuracy.



(a) Normal textiles.

(b) Defective textiles.

Figure 8: Example images from the woven textile dataset [4]. (a) Representative sample of woven textiles with no defects. (b) Representative sample of woven textiles that are considered defective.

5.1 Data Pre-Processing

Covariance matrices are useful data representations for images because they are invariant to translations and rotations [36]. These matrices can also be used to combine multiple image features that can be quantified through filters and kernel methods [37]. Here, we select nine image features, as shown in Figure 9. We implement both the Frangi and Hessian filters, which are designed to detect edges and fiber structures of the woven textures [13]. We apply these filters to the image; multiple transformed images that are smoothed through a Gaussian filter of varying strength are shown in

Figure 9. We use Gaussian filters to emphasize features of different scale within the image [45]. This yields nine total feature images (including the original image), which we use to construct a sample covariance matrix. We do this by taking each feature image (which we treat as a 64×64 matrix) and vectorizing each image $x_i \in \mathbb{R}^m$, where $i = 1, 2, \dots, 9$ and $m = 4096$. Each feature image can therefore be represented as a realization of a multivariate random vector $\mathbf{X} = (x_1, x_2, \dots, x_9)$. We use this observation to construct a sample covariance matrix $\mathbf{P} \in \mathcal{P}(n)$ (see Figure 9) where $n = 9$:

$$\mathbf{P} := \frac{1}{m-1} \mathbf{X} \mathbf{X}^T. \quad (5.20)$$

The covariance matrices are SPD, and can be directly integrated into our geometric analysis framework.




Figure 9: Workflow for the construction of an image covariance matrix. (left) Image filters and transformations emphasize specific characteristics of an image. Gaussian filters emphasize features of different scale within an image, and the Frangi and Hessian filters capture important fiber and edge features of an image. (right) The covariance between each image representation can be computed and used to form a covariance matrix. Importantly, the covariance matrix representation is invariant to transformations such as rotation and translation which are present in the textile images.

5.2 Classification Results

We first apply PGA to the image covariance matrices (see Figure 10). PGA reveals a distinct clustering and separation of the data with only two components. The PGA suggests that a simple classifier can be used to separate the defective and non-defective textile samples (after projecting to the tangent space). Thus, we construct a simple linear support vector machine classifier, which we train with 70% of the image data and test our trained model on the remaining 30% of the data. The covariance representation of the images provides a simple characterization of the data and its various features, while also imbuing the data with the inherent geometry of SPD matrices. This results in the trained model being able to separate the defective and non-defective samples in the testing data with 92% accuracy. We also compare our PGA analysis of the covariance matrix representation to PCA applied to the raw image data in Figure 10. The PCA analysis reveals almost

no separation of the data, and results in poor classification accuracy when used as input to a linear SVM (50 % accuracy). This is likely due to sensitivity of the analysis on the raw images to rotations and translations of the textiles.




Figure 10: (a) PGA of the covariance matrices derived from the set of woven textile images. The analysis reveals clustering of the data into groups representing the defective and non-defective textiles. A simple linear SVM classifier is able to separate the defective and non-defective textile images with 92% accuracy. (b) PCA analysis of the raw images; without the covariance matrix representation and considerations for the data geometry, we see there is almost no separation in the data.

6 Conclusions and Future work

We have presented an introduction to concepts of Riemannian geometry for symmetric positive definite (SPD) matrices and show how these concepts can be used in applications of interest to chemical engineers. Specifically, we discussed approaches to capture the geometry of the SPD manifold (a Riemannian manifold) in dimensionality reduction and classification tasks. Through a couple of case studies, we demonstrated that capturing such geometry can lead to significant improvements in accuracy and interpretability. This work focused primarily on the manifold of SPD matrices, but we note that similar concepts can be applied to other matrix representations such as fixed-rank matrices, orthogonal matrices, and shape spaces as well as to other datasets where a Riemannian manifold is known to exist (e.g., points on a sphere). Furthermore, geometry can be leveraged in optimization and control to relax certain constraints that are already enforced by the underlying system geometry (e.g., optimization on the SPD manifold ensures that all solutions are SPD).

Acknowledgments

We acknowledge funding from the U.S. National Science Foundation (NSF) under BIGDATA grant IIS-1837812.

References

- [1] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [2] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM journal on matrix analysis and applications*, 29(1):328–347, 2007.
- [3] Alexandre Barachant and Marco Congedo. A plug&play p300 bci using information geometry. *arXiv preprint arXiv:1409.0107*, 2014.
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9592–9600, 2019.
- [5] Rajendra Bhatia. *Positive definite matrices*. Princeton university press, 2009.
- [6] Charlene E Caefer, Jerry Silverman, Oded Orthal, Dani Antonelli, Yaron Sharoni, and Stanley R Rotman. Improved covariance matrices for point target detection in hyperspectral data. *Optical Engineering*, 47(7):076402, 2008.
- [7] Isaac Chavel. *Riemannian geometry: a modern introduction*, volume 98. Cambridge university press, 2006.
- [8] Leo H Chiang, Evan L Russell, and Richard D Braatz. Fault diagnosis in chemical processes using fisher discriminant analysis, discriminant partial least squares, and principal component analysis. *Chemometrics and intelligent laboratory systems*, 50(2):243–252, 2000.
- [9] James J Downs and Ernest F Vogel. A plant-wide industrial process control problem. *Computers & chemical engineering*, 17(3):245–255, 1993.
- [10] Ian L Dryden, Alexey Koloydenko, and Diwei Zhou. Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *The Annals of Applied Statistics*, 3(3):1102–1123, 2009.
- [11] Bo Feng, Mengyin Fu, Hongbin Ma, Yuanqing Xia, and Bo Wang. Kalman filter with recursive covariance estimation—sequentially estimating process noise covariance. *IEEE Transactions on Industrial Electronics*, 61(11):6253–6263, 2014.
- [12] P Thomas Fletcher, Conglin Lu, Stephen M Pizer, and Sarang Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE transactions on medical imaging*, 23(8):995–1005, 2004.
- [13] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *International conference on medical image computing and computer-assisted intervention*, pages 130–137. Springer, 1998.
- [14] Joaquín Goñi, Martijn P Van Den Heuvel, Andrea Avena-Koenigsberger, Nieves Velez De Mendizabal, Richard F Betzel, Alessandra Griffa, Patric Hagmann, Bernat Corominas-Murtra, Jean-Philippe Thiran, and Olaf Sporns. Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences*, 111(2):833–838, 2014.
- [15] Mehrtaash T Harandi, Mathieu Salzmann, and Richard Hartley. From manifold to manifold: Geometry-aware dimensionality reduction for spd matrices. In *European conference on computer vision*, pages 17–32. Springer, 2014.

- [16] M Willjuice Iruthayarajan and S Baskar. Covariance matrix adaptation evolution strategy based design of centralized pid controller. *Expert systems with Applications*, 37(8):5775–5781, 2010.
- [17] Wenfu Ku, Robert H Storer, and Christos Georgakis. Disturbance detection and isolation by dynamic principal component analysis. *Chemometrics and intelligent laboratory systems*, 30(1):179–196, 1995.
- [18] Serge Lang. *Fundamentals of differential geometry*, volume 191. Springer Science & Business Media, 2012.
- [19] John M Lee. *Riemannian manifolds: an introduction to curvature*, volume 176. Springer Science & Business Media, 2006.
- [20] M Mansouri, MZ Sheriff, R Baklouti, M Nounou, H Nounou, A Ben Hamida, and N Karim. Statistical fault detection of chemical process-comparative studies. *Journal of Chemical Engineering & Process Technology*, 7(1):282–291, 2016.
- [21] Maher Moakher and Philipp G Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In *Visualization and Processing of Tensor Fields*, pages 285–298. Springer, 2006.
- [22] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [23] Xavier Pennec. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- [24] Xavier Pennec. Manifold-valued image processing with spd matrices. In *Riemannian geometric statistics in medical image analysis*, pages 75–134. Elsevier, 2020.
- [25] Fatih Porikli. Achieving real-time object detection and tracking under extreme conditions. *Journal of Real-Time Image Processing*, 1(1):33–40, 2006.
- [26] Fatih Porikli and Oncel Tuzel. Fast construction of covariance matrices for arbitrary size image windows. In *2006 International Conference on Image Processing*, pages 1581–1584. IEEE, 2006.
- [27] Nara M Portela, George DC Cavalcanti, and Tsang Ing Ren. Semi-supervised clustering for mr brain image segmentation. *Expert Systems with Applications*, 41(4):1492–1497, 2014.
- [28] Anqi Qiu, Annie Lee, Mingzhen Tan, and Moo K Chung. Manifold learning on brain functional networks in aging. *Medical image analysis*, 20(1):52–60, 2015.
- [29] Gregory C Reinsel. *Elements of multivariate time series analysis*. Springer Science & Business Media, 2003.
- [30] Bernhard Riemann. *On the Hypotheses which lie at the Bases of Geometry*. Birkhäuser, 2016.
- [31] Evan L Russell, Leo H Chiang, and Richard D Braatz. Fault detection in industrial processes using canonical variate analysis and dynamic principal component analysis. *Chemometrics and intelligent laboratory systems*, 51(1):81–93, 2000.
- [32] Alexander D Smith, Paweł Dłotko, and Victor M Zavala. Topological data analysis: concepts, computation, and applications in chemical engineering. *Computers & Chemical Engineering*, 146:107202, 2021.
- [33] Alexander D Smith and Victor M Zavala. The euler characteristic: A general topological descriptor for complex data. *Computers & Chemical Engineering*, 154:107463, 2021.

- [34] Stefan Sommer, François Lauze, Søren Hauberg, and Mads Nielsen. Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations. In *European conference on computer vision*, pages 43–56. Springer, 2010.
- [35] Olaf Sporns. Network analysis, complexity, and brain function. *Complexity*, 8(1):56–60, 2002.
- [36] Oncel Tuzel, Fatih Porikli, and Peter Meer. Region covariance: A fast descriptor for detection and classification. In *European conference on computer vision*, pages 589–600. Springer, 2006.
- [37] Oncel Tuzel, Fatih Porikli, and Peter Meer. Pedestrian detection via classification on Riemannian manifolds. *IEEE transactions on pattern analysis and machine intelligence*, 30(10):1713–1727, 2008.
- [38] Gaël Varoquaux, Flore Baronnet, Andreas Kleinschmidt, Pierre Fillard, and Bertrand Thirion. Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 200–208. Springer, 2010.
- [39] Larry Wasserman. *All of statistics: a concise course in statistical inference*, volume 26. Springer, 2004.
- [40] Barry M Wise and Neal B Gallagher. The process chemometrics approach to process monitoring and fault detection. *Journal of Process Control*, 6(6):329–348, 1996.
- [41] Yang Xu, Zebin Wu, Jun Li, Antonio Plaza, and Zhihui Wei. Anomaly detection in hyperspectral images based on low-rank and sparse representation. *IEEE Transactions on Geoscience and Remote Sensing*, 54(4):1990–2000, 2015.
- [42] Jianfeng Yao, Shurong Zheng, and ZD Bai. *Sample covariance matrices and high-dimensional data analysis*. Cambridge University Press Cambridge, 2015.
- [43] Shen Yin, Steven X Ding, Adel Haghani, Haiyang Hao, and Ping Zhang. A comparison study of basic data-driven fault diagnosis and process monitoring methods on the benchmark tennessee eastman process. *Journal of process control*, 22(9):1567–1581, 2012.
- [44] Kisung You and Hae-Jeong Park. Re-visiting Riemannian geometry of symmetric positive definite matrices for the analysis of functional connectivity. *Neuroimage*, 225:117464, 2021.
- [45] Ian T Young and Lucas J Van Vliet. Recursive implementation of the Gaussian filter. *Signal processing*, 44(2):139–151, 1995.