# Inductive Geometric Matrix Midranges ⋆

**Graham W. Van Goffrier** * **Cyrus Mostajeran** **
**Rodolphe Sepulchre** **

* *Department of Physics and Astronomy, University College London,
London, United Kingdom (e-mail: vangoffrier@gmail.com)*
** *Department of Engineering, University of Cambridge, Cambridge
CB2 1PZ, United Kingdom (e-mail: csm54@cam.ac.uk)*

**Abstract:** Covariance data as represented by symmetric positive definite (SPD) matrices are ubiquitous throughout technical study as efficient descriptors of interdependent systems. Euclidean analysis of SPD matrices, while computationally fast, can lead to skewed and even unphysical interpretations of data. Riemannian methods preserve the geometric structure of SPD data at the cost of expensive eigenvalue computations. In this paper, we propose a geometric method for unsupervised clustering of SPD data based on the Thompson metric. This technique relies upon a novel "inductive midrange" centroid computation for SPD data, whose properties are examined and numerically confirmed. We demonstrate the incorporation of the Thompson metric and inductive midrange into $X$-means and $K$-means++ clustering algorithms.

*Keywords:* Classification, Clustering, Covariance Matrices, Differential Geometry

## 1. INTRODUCTION

The arithmetic midrange of a finite collection of real numbers $\{y_i\}_{i \in I}$ is defined as the mean of the extremal values: $\frac{1}{2}(\min_i y_i + \max_i y_i)$. This number can be uniquely characterized as the solution $x^*$ to the optimization problem

$$\min_x \max_i |x - y_i|. \tag{1}$$

One can also characterize the midrange as the limit of an inductive procedure on the input data points $y_i$ in the following way. The merits of such a characterization will become clear in due course.

*Proposition 1.* For any $a, b \in \mathbb{R}$, define the curve $\gamma(a, b; \cdot) : [0,1] \to \mathbb{R}$ by $\gamma(a,b;t) = (1-t)a + tb$. For any initialization $x_1 \in \mathbb{R}$, we may generate a sequence $(x_k)$ as follows. Given a point $x_k$:

(1) choose a point $y_k^{\uparrow} \in \{y_i\}$ such that $|x_k - y_k^{\uparrow}| \geq |x_k - y_i|$ for all $i \in I$;

(2) set $x_{k+1} := \gamma(x_k, y_k^{\uparrow}; \frac{1}{k+1})$.

Any sequence $(x_k)$ generated as above converges to the midrange of $\{y_i\}$.

**Proof.** Denote by $y_{\min}$ and $y_{\max}$ the minimum and maximum values within $\{y_i\}$ and write $x^* = \frac{1}{2}(y_{\min} + y_{\max})$. For any sequence $(x_k)$ generated as above, there exists $m \in \mathbb{N}$ such that $x_m \in (y_{\min}, x^*)$ and $x_{m+1} \in (x^*, y_{\max})$. Such an $m$ can easily be found by taking a sufficiently large $m$. We then have

$$x_{m+1} = \frac{m}{m+1}x_m + \frac{1}{m+1}y_{\max}$$
$$x_{m+2} = \frac{m+1}{m+2}x_{m+1} + \frac{1}{m+2}y_{\min},$$

which combine to give

$$x_{m+2} = \frac{m}{m+2}x_m + \frac{1}{m+2}(y_{\min} + y_{\max})$$
$$= \frac{m}{m+2}x_m + \frac{2}{m+2}x^*.$$

Thus, $x_{m+2} - x_m = \frac{2}{m+2}(x^* - x_m) > 0$ and $x_{m+2} - x^* = \frac{m}{m+2}(x_m - x^*) < 0$, so that $x_m < x_{m+2} < x^*$. By symmetry in the reasoning, we also have $x^* < x_{m+3} < x_{m+1}$. By induction on $k$, we obtain

$$x_{m+2k} - x^* = \frac{m}{m+2k}(x_m - x^*) \leq \frac{m}{2k}(x_m - x^*) \tag{2}$$

for all $k \geq 1$. Thus, we see that the subsequence $(x_{m+2k})$ converges to $x^*$ at a rate of $\mathcal{O}(k^{-1})$. Since a similar result holds for the subsequence $(x_{m+2k-1})$, we conclude that $(x_k)$ converges to the midrange $x^*$ at the given rate. □

Although the midrange is clearly sensitive to outliers, it can be a useful measure nonetheless. In particular, if the data points $y_i$ are uniformly spread within a domain, the mean coincides with the midrange and can thus be computed using only extremal values without sifting through all the points. Similarly, the midrange can be useful for analyzing data that is devoid of outliers and finds applications in clustering algorithms that rely on the isolation of outlying clusters (Steinley (2006); Carroll and Chaturvedi (1998)).

In this paper, we consider midrange statistics in the cone of symmetric positive definite (SPD) matrices. Data representations based on SPD matrices arise in an enormous range of applications as covariance matrices, including brain-computer interface (BCI) systems, radar data pro-

cessing, and diffusion tensor imaging (DTI). Let $\mathcal{P}_d$ denote the space of $d \times d$ real SPD matrices. The conic geometry of $\mathcal{P}_d$ often renders conventional Euclidean approaches to statistics and analysis on $\mathcal{P}_d$ ineffective. Indeed it is well-documented that using Euclidean algorithms on nonlinear spaces such as $\mathcal{P}_d$ often results in poor accuracy and undesirable effects, such as swelling phenomena in DTI (Arsigny et al. (2006)).

Several non-Euclidean geometries have been associated with $\mathcal{P}_d$ and successfully exploited in various applications. Geometries that have been studied in great detail include the log-Euclidean geometry and the affine-invariant Riemannian geometry. The log-Euclidean geometry is derived from the Lie group structure of $\mathcal{P}_d$ under the group operation $\Sigma_1 \circ \Sigma_2 = \exp(\log(\Sigma_1) + \log(\Sigma_2))$ for $\Sigma_1, \Sigma_2 \in \mathcal{P}_d$, where exp and log denote the usual matrix exponential and logarithm. The affine-invariant Riemannian geometry on the other hand is induced by the Riemannian metric

$$\langle V, W \rangle_\Sigma = \operatorname{tr}(\Sigma^{-1} V \Sigma^{-1} W), \qquad (3)$$

where $V$ and $W$ are vectors belonging to the tangent space $T_\Sigma \mathcal{P}_d$ at $\Sigma \in \mathcal{P}_d$. This smooth metric structure induces a well-defined distance function $d_2 : \mathcal{P}_d \times \mathcal{P}_d \to [0, \infty)$ given by

$$d_2(\Sigma_1, \Sigma_2) = \left[ \operatorname{tr}\left( \log^2(\Sigma_1^{-\frac{1}{2}} \Sigma_2 \Sigma_1^{-\frac{1}{2}}) \right) \right]^{\frac{1}{2}} = \left( \sum_{i=1}^d \log^2 \lambda_i \right)^{\frac{1}{2}}$$

where $\lambda_i$ denote the eigenvalues of $\Sigma_2 \Sigma_1^{-1}$. Moreover, the unique length-minimizing geodesic from $\Sigma_1$ to $\Sigma_2$ takes the form of the curve $\gamma(t) = \Sigma_1 \#_t \Sigma_2$, where

$$\Sigma_1 \#_t \Sigma_2 = \Sigma_1^{1/2} \left( \Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2} \right)^t \Sigma_1^{1/2}, \qquad (4)$$

for $t \in [0, 1]$ (Bhatia (2003)). The midpoint of this geodesic defines the geometric mean $\Sigma_1 \#_{\frac{1}{2}} \Sigma_2$ of $\Sigma_1$ and $\Sigma_2$. The geometric mean of $N$ SPD matrices $Y_1, \ldots, Y_N$ can be defined as the unique solution to

$$\operatorname{argmin}_{X \in \mathcal{P}_d} \sum_{i=1}^N d_2(X, Y_i)^2, \qquad (5)$$

which is also known as the Riemannian center of mass.

An important property of the Riemannian geometry defined by (3) is affine-invariance, whereby congruence transformations form isometries of $\mathcal{P}_d$. In particular, for any matrix $A$ in the general linear group $GL(d)$, we have $d_2(A\Sigma_1 A^T, A\Sigma_2 A^T) = d_2(\Sigma_1, \Sigma_2)$ for all $\Sigma_1, \Sigma_2 \in \mathcal{P}_d$. Affine-invariance of algorithms at the level of SPD matrices corresponds to invariance under affine transformations of the underlying feature vectors, which is often a desirable property in applications involving covariance matrices.

A non-Riemannian affine-invariant geometry associated with $\mathcal{P}_d$ is that induced by the Thompson metric on the cone of $d \times d$ positive semidefinite matrices (Thompson (1963); Lemmens and Nussbaum (2012)). The Thompson metric $d_\infty$ on $\mathcal{P}_d$ takes the form

$$d_\infty(\Sigma_1, \Sigma_2) = \log \max\{\lambda_{\max}(\Sigma_1 \Sigma_2^{-1}), \lambda_{\max}(\Sigma_2 \Sigma_1^{-1})\},$$
$$(6)$$

where $\lambda_{\max}(\Sigma)$ denotes the maximum eigenvalue of $\Sigma$. Note that the Thompson metric can be rewritten as $d_\infty(\Sigma_1, \Sigma_2) = \max_{1 \le i \le d} |\log \lambda_i|$, which justifies the notation $d_\infty$. The pair $(\overline{\mathcal{P}}_d, d_\infty)$ constitutes a complete metric

space of non-positive curvature (Bhatia (2003)). While the Riemannian geodesic (4) is also a length-minimizing geodesic of the Thompson metric, it is known that $d_\infty$ generally admits infinitely many geodesics between a pair of points (Nussbaum (1994)). In particular, the curve $\gamma_\infty : [0, 1] \to \mathcal{P}_d$, $\gamma_\infty(t) = \Sigma_1 *_t \Sigma_2$ defined below is a geodesic of $d_\infty$ from $\Sigma_1$ to $\Sigma_2$, which satisfies a number of attractive properties. Let $\lambda_M$ and $\lambda_m$ denote the maximum and minimum eigenvalues of $\Sigma_2 \Sigma_1^{-1}$. If $\lambda_M \neq \lambda_m$, we define

$$\Sigma_1 *_t \Sigma_2 = \left( \frac{\lambda_M^t - \lambda_m^t}{\lambda_M - \lambda_m} \right) \Sigma_2 + \left( \frac{\lambda_M \lambda_m^t - \lambda_m \lambda_M^t}{\lambda_M - \lambda_m} \right) \Sigma_1, \; (7)$$

and $\Sigma_1 *_t \Sigma_2 = \lambda_m^t \Sigma_1$ otherwise. One of the attractive properties of (7) is that its midpoint $\Sigma_1 *_{1/2} \Sigma_2$ scales geometrically:

$$(a_1 \Sigma_1) *_{1/2} (a_2 \Sigma_2) = \sqrt{a_1 a_2}(\Sigma_1 *_{1/2} \Sigma_2), \quad a_1, a_2 > 0,$$

and coincides with the geometric mean when $d = 2$.

Computationally, the Thompson geodesic $\Sigma_1 *_t \Sigma_2$ (7) is considerably less expensive to construct than the Riemannian geodesic $\Sigma_1 \#_t \Sigma_2$ (4), particularly for high dimensional matrices. This is a consequence of the fact that $\Sigma_1 *_t \Sigma_2$ only relies on the computation of extremal generalized eigenvalues of the pair $(\Sigma_1, \Sigma_2)$, which can be computed efficiently by several algorithms (Golub and van der Vorst (2000); Mishra and Sepulchre (2016)). Similarly, the computation of the Thompson distance between a pair of SPD matrices only relies on the evaluation of extremal generalized eigenvalues, whereas the Riemannian distance involves the full generalized eigenspectrum of the pair of points.

In the case of positive scalars $(d = 1)$, affine-invariance simply reduces to invariance under scaling in the cone of positive real numbers. The affine-invariant midrange of a collection of $N$ positive scalars $y_i$ is then the geometric mean of the extremal values, which can be formulated as the unique solution of the optimization problem

$$\min_{x > 0} \max_{1 \le i \le N} \left| \log \frac{x}{y_i} \right|. \qquad (8)$$

The affine-invariant SPD matrix analogue of the above optimization formulation of the geometric midrange takes the form

$$\min_{X \in \mathcal{P}_d} \max_{1 \le i \le N} \| \log(Y_i^{-1/2} X Y_i^{-1/2}) \|_\infty, \qquad (9)$$

which can be recast as the convex optimization problem

$$\begin{cases} \min\limits_{\tau, \xi, X \succeq 0} \xi \\ \tau Y_i \preceq X \preceq \xi Y_i \\ 1/\xi - \tau \le 0 \end{cases} \qquad (10)$$

where $\preceq$ denotes the matrix Löwner order (Mostajeran et al. (2020)).

Although (10) can be solved using standard convex optimization packages, the problem does not scale well with the dimension $d$ or the number of data points $N$. Thus, the aim of this work is to develop an alternative route to a notion of a geometric midrange of SPD matrices that has favorable computational properties that scale with the dimension $d$. It is in this context that we investigate the natural generalization of the inductive procedure for computing the midrange of scalars outlined in Proposition

1 to the matrix setting as an alternative definition of the geometric midrange of a collection of points in $\mathcal{P}_d$. By using the Thompson metric $d_\infty$ to compute the distances in the first condition of Proposition 1, and the Thompson geodesics $\gamma_\infty$ as the interpolating curves in the second condition, we arrive at an algorithm that relies only on the computation of generalized extremal eigenvalues. This algorithm and its computational properties are investigated in the following sections.

## 2. INDUCTIVE MIDRANGE

The inductive midrange (IMR) algorithm takes as input a set of symmetric positive definite (SPD) matrix data $\{Y_i\}_{i \in I}$, as well as an initialization point which may or may not be a member of the data set, and iterates towards a representative midrange centroid of the data set. Given an initial point $X_1$, a sequence $(X_k)$ is generated according to the following process.

(1) Choose a point $Y_k^\uparrow \in \{Y_i\}$ such that $d_\infty(X_k, Y_k^\uparrow) \geq d_\infty(X_k, Y_i)$ for all $i \in I$.
(2) Set $X_{k+1} := X_k *_{\frac{1}{k+1}} Y_k^\uparrow$.

As shown in Lim (2013), the point $X = A *_{1/2} B$ is a Thompson metric midpoint of $A$ and $B$. That is,

$$d_\infty(A, X) = d_\infty(X, B) = \frac{1}{2}d_\infty(A, B).$$

The effect of the weighted midrange in the IMR algorithm is to produce step sizes that decrease as $\frac{1}{k}$. As outlined above, these steps are restricted to the direction of the furthest data point from each successive IMR candidate, reflecting the original intent of midrange statistics to primarily account for data outliers.

We emphasize that the convergence point of the IMR *is not equivalent* to the optimization midrange discussed in Section 1. As a demonstrative example, we evaluate both midranges on the data set:

$$\{Y_i\} = \left\{ \begin{pmatrix} 0.95 & -0.6 \\ -0.6 & 1.1 \end{pmatrix}, \begin{pmatrix} 1.0 & 0.5 \\ 0.5 & 2.1 \end{pmatrix}, \begin{pmatrix} 2.5 & -0.2 \\ -0.2 & 1.2 \end{pmatrix} \right\}.$$

The optimization midrange ($M_{OPT}$) and IMR midrange ($M_{IMR}$) are evaluated to be:

$$M_{OPT} = \begin{pmatrix} 1.32 & -0.53 \\ -0.53 & 1.62 \end{pmatrix}, \quad M_{IMR} = \begin{pmatrix} 1.14 & -0.25 \\ -0.25 & 1.25 \end{pmatrix}.$$

While the $M_{OPT}$ midrange attains the minimum cost-function value of 0.790, the IMR has a nearby cost-function evaluation of 0.811, which represents a less than 3% increase. The Thompson distance between these two midrange matrices is 0.33, which is modest in comparison to the $\sim 1.4$ average separation between data set matrices. Thus, we observe an example of the interesting phenomenon that equivalent characterizations of a mathematical object in a linear space can generalize to distinct notions in nonlinear spaces.

The primary observed features of the IMR algorithm are a universal $\frac{1}{k}$ convergence rate, initialization-invariance, and dependence exclusively on a subset of the input data called *active data*. We will demonstrate each of these features in turn through numerical studies on a broad range of data set sizes $N$ and matrix dimensions $d$.

### 2.1 Numerical Results

*Claim 1.* The IMR algorithm converges at a rate of $\frac{1}{k}$ regardless of the size or matrix dimensionality of the data set.

We observed that the IMR algorithm converges in all of our numerical experiments. The convergence rate of the IMR algorithm is assessed by measuring the Thompson distance between the IMR estimate at each step and the final IMR value after a prescribed number of steps *numiters*. Fig. 1 (a) shows a typical example of the convergence measure throughout an IMR run for $d = 2$, $N = 5$, where $N$ is the number of data points in $\mathcal{P}_d$. After $10^3$ iterations, the error inherent in our estimate of the true convergence point leads to nonlinearity in this measure. Therefore the range $10^3$ to $10^4$ of iterations is excluded from fitting, and an averaged trend of $\sim \frac{1}{k}$ convergence is demonstrated. Different colors in the plot correspond to different input data taken to initialize the algorithm. Fig. 1 (b) shows a typical example of the same kind for $d = 100$, $N = 5$, with only one initialization included.

Table 1 aggregates observed convergence rates for setups with larger input data sets or greater matrix dimensionality. Average convergence rates for 10 runs with *numiters* $= 10^4$ iterations are given, with the fit excluding nonlinearity past $10^3$ iterations. Random SPD data are generated via the transpose-product method, as are all other SPD data in this paper unless otherwise specified.

Table 1. Average convergence rates for several sample SPD data set configurations

| $(d, N)$ | (5,5) | (5,20) | (50,5) | (50,20) |
|---|---|---|---|---|
| Rate | -0.9942 | -0.9932 | -0.9965 | -1.0019 |

*Claim 2.* For a given data set, the IMR algorithm converges to the same SPD matrix regardless of initialization.

The input data and IMR trajectories of $2 \times 2$ SPD matrices can be visualized in a cone in $\mathbb{R}^3$ as described by Mostajeran and Sepulchre (2018). Fig. 2 shows a successful IMR run for $d = 2$ and $N = 3$ depicted in a 2D projection of the cone in $\mathbb{R}^3$. Although the figure depicts a 2D projection for simplicity, the convergence is observed in the full three-dimensional space. The top and right-hand input data points move towards each other to their mutual midpoint during the first IMR step, and therefore trivially converge to a single result; however, it is less trivial that the bottom input data point converges to the same result.

In Fig. 1 (c) the pairwise Thompson distances between an arbitrary reference initialization and all others are plotted throughout a $d = 2$, $N = 10$ IMR run. Since different initializations lead to IMR paths which are drawn to some data points more often than others, quasiperiodic features dominate the pairwise distances. Despite this, all pairwise distances appear to converge to zero with a $\sim \frac{1}{k}$ rate.

To further demonstrate invariance beyond only initializations from the input data set, IMR runs were performed with randomly generated SPDs as initializations for several data set configurations. In all cases, a maximum Thompson distance separation could be identified within which all IMR results could be bounded for the chosen *numiters* and initializations. By increasing *numiters*,
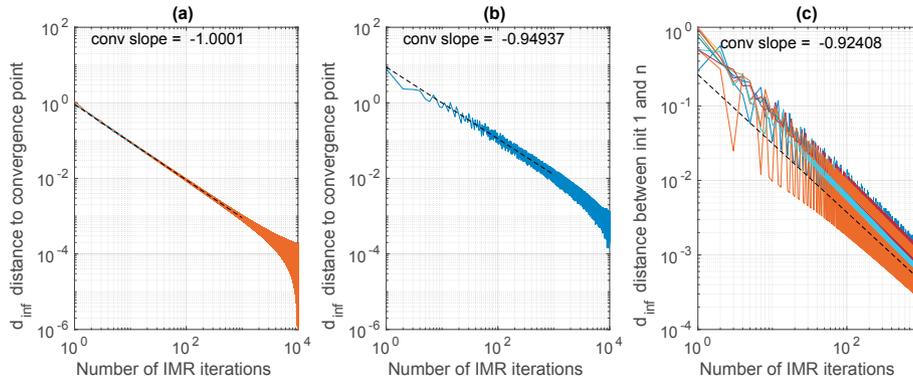
Fig. 1. IMR convergence plot for (a) $d = 2$, $N = 5$ and (b) $d = 100$, $N = 5$. (c) Convergence plot showing the evolution of the distances between IMR trajectories with $d = 2$ and $N = 10$. The Thompson distances are computed from an arbitrary reference initialization. The plot shows that all trajectories contract towards each other at approximately the specified rate.
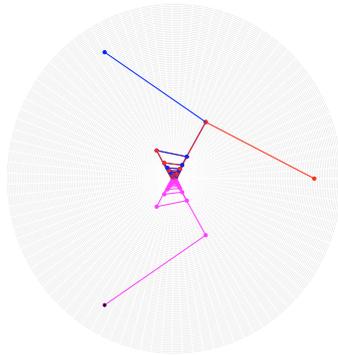


Fig. 2. Invariance of the point of convergence of the IMR algorithm with respect to the choice of initialization. The plot depicts an example with $d = 2$ and $N = 3$.

the maximum separation could be reduced arbitrarily. The average separation between the trajectories after $numiters = 10^4$ steps was significantly smaller than the maximum separation in all cases, as shown in Table 2.

Table 2. Initialization invariance results for the IMR algorithm with random initializations.

| $(d, N)$ | $(2,5)$ | $(5,5)$ | $(20,5)$ | $(100,5)$ |
|---|---|---|---|---|
| Max separation | .0018 | .0310 | .0864 | .1453 |
| Average separation | .0008 | .0072 | .0188 | .0277 |
| Convergence count (/100) | 100 | 100 | 100 | 100 |

*Claim 3.* The IMR convergence point depends exclusively on a subset of the initial data called the *active data*.

This property is essential for the IMR to be interpreted as a midrange-type centroid for a data set. For a data set $D \subset \mathcal{P}_d^N$, we define the "external" data $E \subset D$ as the subset of all data points $D_i \in D$ that for some choice of initialization $X_1$ and $k \geq 1$,

$$X_{k+1} = X_k *_{\frac{1}{k+1}} D_i. \tag{11}$$

Only the external data $E$ are available for the IMR algorithm to target and step towards on each iteration. As the size of an SPD data set increases, it becomes more difficult to arrange them in such a way that $|E|$ increases. However, simulations demonstrate that the IMR need not be dependent on all data in $E$ in the limit of infinite time, and in fact are generally dependent on a subset $A \subseteq E$
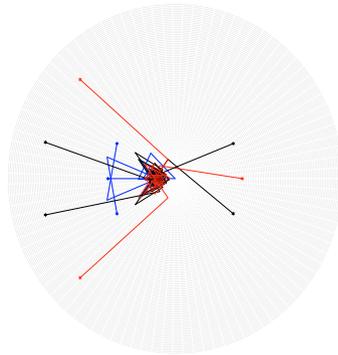


Fig. 3. Active (red), inactive external (black), and internal data (blue) for an example with $d = 2$, $N = 10$. The IMR trajectories are only influenced by the external data, while the point of convergence itself is only dependent on the active data.

which may be significantly smaller – we term this subset the active data.

The properties of active and external data are observed even in very small data sets. If the right-hand point in Fig. 2 is moved leftward, a sharp transition occurs beyond which the IMR is dependent only on the remaining two data points, and can be shown to occur at their geometric midrange. Fig. 3 shows a more typical example, in which any or all of three input data internal to the data set (blue) may be removed without affecting the IMR convergence algorithm, and four external but non-active data points (black) may also be removed without affecting the IMR convergence point. In this example, there are only three active (red) data points.

The IMR algorithm may therefore be made more efficient by advance identification of all active points and restricting the max distance search to only target this subset. Data sets with extreme outliers will tend to have fewer active points, and the IMR would accordingly experience a dramatic speedup in such cases.

## 3. MIDRANGE CLUSTERING ON MATRIX DATA

Clustering methodologies using nonlinear geometries related to eigenspectra of SPD matrices have shown promise.

For example, see the adaptation of $K$-means clustering to several similarity measures considered by Stanitsas et al. (2017), which include the affine-invariant Riemannian (AIRM) and log-Euclidean metrics. See also Nielsen and Sun (2017) for an account of clustering in Hilbert simplex geometry. In this work, we consider $K$-means clustering with the Thompson distance $d_\infty$ as the similarity measure, and employ the IMR to calculate centroids for clusters. We further employ $X$-means and $K$-means++ algorithms with geometric midrange statistics and evaluate their performance.

The $K$-means clustering algorithm may be summarized as follows, where the inputs are a set of SPD data $\bar{Y}$ and a desired quantity of clusters $K$ (for a more thorough review, see Steinley (2006)):

(1) Assign an initial cluster label $l_i \in \{1, ..., k\}$ to each $Y_i$. The set of points with label $L$ define a cluster $C_L \subset \bar{Y}$.
(2) For each cluster $C_L$, a centroid $\mu_L$ is calculated (here via IMR).
(3) For each data point $Y_i$, the nearest centroid $\mu_{L'}$ is identified, and the cluster label is reassigned $l_i \leftarrow L'$.
(4) Repeat steps 2 and 3 until the cluster labels do not change between iterations.

Two straightforward initialization methods are random assignment of labels, and random selection of $K$ data points as initial centroids, where labels are then assigned as in step 3. Both of these methods can have severe limitations due to false merging of clusters, where the severity decreases with the connectedness of the data set. In a highly disconnected data set, a true cluster which does not initially contain at least one centroid is unlikely to be accurately identified during the course of the algorithm. Oversplitting of clusters is a less severe but still significant source of error. Near-perfect accuracy can be achieved given advance knowledge of approximate cluster locations, so that one initial centroid can be placed in each; however this is not the usual case in applications, where the optimal value of $K$ is usually not even known.

### 3.1 X-means Clustering

$X$-means is an iterative extension of $K$-means, first proposed by Pelleg and Moore (2000), which resolves the issue of erroneously merged clusters by repeated binary splits. In this work the authors make use of the Bayesian information criterion (BIC) to accept or reject these binary splits. In our $X$-means implementation, a low (empirical) limit is placed on the number of attempted splits that may be performed on a particular cluster. BIC scoring is modified to include a Thompson-based maximum likelihood estimator. We also split centroids following Thompson geometry, by randomly selecting one matrix on a fixed-radius $d_\infty$-sphere from the unsplit centroid, and taking this SPD and its antipodal point on the $d_\infty$-sphere as the initial split centroids before performing the intermediary $K = 2$ clustering. $d_\infty$-sphere sampling is performed by generating SPDs in a near-radially symmetric distribution around the identity. As the $d_\infty$ geodesic given by (7) is parameterized proportionally to Thompson distance, a retraction (or extension) along the geodesic is then made such that the generated point has the desired radius from

the identity. By affine invariance of the metric, any point $\Sigma$ generated on the $d_\infty$ sphere may then be transported to the desired centroid $C$ by a congruence transformation $\Sigma \mapsto C^{\frac{1}{2}} \Sigma C^{\frac{1}{2}}$.

Fig. 4 shows representative $X$-means results for 200 $d = 2$ SPD data in 10 disjoint clusters, visualized in the cone projection. These disjoint clusters are created by generating 10 cluster centers with the requirement that all pairwise Thompson distances are at least 1. For each cluster center, 20 data points are uniformly generated on a $d_\infty$-sphere of radius 0.2 around the cluster center. Although oversplitting can still occur, as exemplified by the center-left cluster in the plot, $X$-means massively reduces the risk of falsely merged clusters. Only one falsely merged cluster appears: the center-most cluster in this figure, where the bounding boxes of the two true clusters overlap.

For higher-dimensional SPDs, clustering results are best visualized in tabular form; for a selection of matrix dimensionalities of up to 20, $X$-means accuracy with $10 \times 20$ constructed data sets are summarized in Table 3, averaged over 20 runs. The low error rate obtained confirms that midrange statistics provides a viable setting for clustering SPD matrix data. Due to volumetric scaling that naturally occurs with increased dimensionality, more accurate clustering is observed for high dimensions. Significant false-merging of clusters in the $d = 2$ case appears to be attributable to poor performance of the BIC for small matrices, which gives a lower likelihood of splitting during step 4 of the algorithm.

Table 3. X-means clustering accuracy in high dimensions

| $d$ | 2 | 5 | 10 | 20 |
|---|---|---|---|---|
| Points identified (/200) | 109.2 | 172.0 | 170.0 | 200.0 |
| Clusters identified (/10) | 4.3 | 8.2 | 8.3 | 10.0 |
| Clusters lost (/10) | 4.5 | 1.4 | 1.5 | 0.0 |

### 3.2 K-means++ Clustering

A perhaps less invasive approach, rather than iterating the entire $K$-means procedure to reassign relatively few cluster labels, is to directly modify the initialization procedure which is the source of many oversplitting and false-merging errors. Arthur and Vassilvitskii (2007) achieve this by constructing a new initialization procedure from the ground up, which relies on no prior knowledge of clusters, other than the desired quantity of clusters. This procedure gives a vanishing likelihood that more than one initial centroid from the same cluster will be selected, and also can give such a strong initial guess at the clustering that $K$-means requires far fewer iterations to converge. In our implementation, the Thompson distance is employed for weight assignments.

Fig. 4 shows a typical $K$-means++ result for the same 200 $d = 2$ SPD data as in the above $X$-means example. No oversplitting or false merging occurs. The center-most pair of clusters, which were falsely merged in the $X$-means run, are now distinguished into two clusters in a qualitatively sensible way despite the overlap.

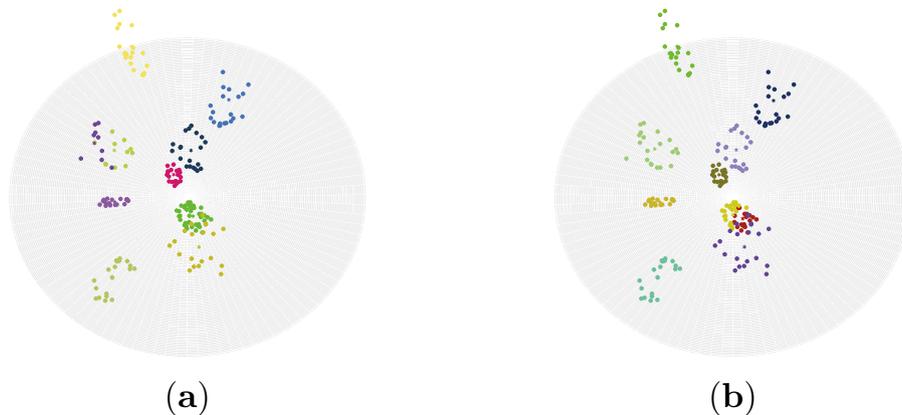Clustering results are again obtained for a selection of matrix dimensionalities up to 100 and averaged over 20