Gaussian distributions on Riemannian symmetric spaces : statistical learning with structured covariance matrices

Salem Said, Hatem Hajri, Lionel Bombrun, Baba C. Vemuri Fellow, IEEE

Abstract

The Riemannian geometry of covariance matrices has been essential to several successful applications, in computer vision, biomedical signal and image processing, and radar data processing. For these applications, an important ongoing challenge is to develop Riemannian-geometric tools which are adapted to structured covariance matrices. The present paper proposes to meet this challenge by introducing a new class of probability distributions, *Gaussian distributions of structured covariance matrices*. These are Riemannian analogs of Gaussian distributions, which only sample from covariance matrices having a preassigned structure, such as complex, Toeplitz, or block-Toeplitz. The usefulness of these distributions stems from three features: (1) they are completely tractable, analytically or numerically, when dealing with large covariance matrices, (2) they provide a statistical foundation to the concept of structured Riemannian barycentre (*i.e.* Fréchet or geometric mean), (3) they lead to efficient statistical learning algorithms, which realise, among others, density estimation and classification of structured covariance matrices spaces. Accordingly, it develops an original theory of Gaussian distributions on Riemannian symmetric spaces, of their statistical inference, and of their relationship to the concept of Riemannian barycentre. Then, it uses this original theory to give a detailed description of Gaussian distributions of structured covariance matrices, complex, Toeplitz, and block-Toeplitz. Finally, it describes algorithms for density estimation and classification mixture models.

Index Terms

Structured covariance matrix, Riemannian barycentre, Gaussian distribution, Riemannian symmetric space, Gaussian mixture model

I. INTRODUCTION

Data with values in the space of covariance matrices have become unavoidable in many applications (here, the term *covariance matrix* is taken to mean *positive definite matrix*). For example, these data arise as "tensors" in mechanical engineering and medical imaging [1]-[3], or as "covariance descriptors" in signal or image processing and computer vision [4]-[6][7]-[10]. Most often, the first step in any attempt to analyse these data consists in computing some kind of central tendancy, or average: the Riemannian barycentre [1][3][11] (*i.e.* Fréchet or geometric mean), the Riemannian median [5][6], the log-det or Stein centre [12][13], *etc.* Therefore, an important objective is to provide a rigorous statistical framework for the use of these various kinds of averages. In the case of the Riemannian barycentre, this objective was recently realised, through the introduction of Gaussian distributions on the space of covariance matrices [14][15].

Gaussian distributions on the space of covariance matrices were designed to extend the useful properties of normal distributions, from the usual context of Euclidean space, to that of the Riemannian space of covariance matrices equipped with its affine-invariant metric described in [3]. Accordingly, they were called "generalised normal distributions" in [14] and "Riemannian Gaussian distributions" in [15]. The most attractive feature of these Gaussians distributions is that they verify the following property, which may also be taken as their definition,

defining property of Gaussian distributions: maximum likelihood is equivalent to Riemannian barycentre

This is an extension of the defining property of normal distributions on Euclidean space, dating all the way back to the work of Gauss [16]. Indeed, in a Euclidean space, the concept of Riemannian barycentre reduces to that of arithmetic average. By virtue of Property (1), efficient algorithms were developed, which use Riemannian barycentres in carrying out prediction and regression [14][17], as well as classification and density estimation [18]–[20], for data in the space of covariance matrices.

The present paper addresses the problem of extending the definition of Gaussian distributions, from the space of covariance matrices to certain spaces of structured covariance matrices. The motivation for posing this problem is the following. The structure of a covariance matrix is a highly valuable piece of information. Therefore, it is important to develop algorithms which deal specifically with data in the space of covariance matrices, when these are known to possess a certain structure. To this end, the paper introduces a new class of probability distributions, *Gaussian distributions of structured covariance matrices*. These are Gaussian distributions which only sample from covariance matrices having a certain preassigned structure. Using these new distributions, it becomes possible to develop new versions of the algorithms from [14][17][18]–[20], which deal specifically with structured covariance matrices.

In the following, Gaussian distributions of structured covariance matrices will be given a complete description for three structures: complex, Toeplitz, and block-Toeplitz. This is achieved using the following observation, which was made in [14][15]. The fact that it is possible to use Property (1), in order to define Gaussian distributions on the space of covariance matrices, is a result of the fact that this space, when equipped with its affine-invariant metric, is a Riemannian symmetric space of non-positive curvature. Starting from this observation, the following new contributions are developed below, in Sections III, IV and V. These employ the mathematical background provided in Section II and Appendix I.

— In Section III: Gaussian distributions on Riemannian symmetric spaces of non-positive curvature are introduced. If \mathcal{M} is a Riemannian symmetric space of non-positive curvature, then a Gaussian distribution $G(\bar{x}, \sigma)$ is defined on \mathcal{M} by its probability density function, with respect to the Riemannian volume element of \mathcal{M} ,

p.d.f. of a Gaussian distribution:
$$p(x|\bar{x},\sigma) = \frac{1}{Z(\sigma)} \times \exp\left[-\frac{d^2(x,\bar{x})}{2\sigma^2}\right] \quad \bar{x} \in \mathcal{M}, \sigma > 0$$
 (2)

where $d: \mathcal{M} \times \mathcal{M} \to \mathbb{R}_+$ denotes Riemannian distance in \mathcal{M} . In all generality, as long as \mathcal{M} is a Riemannian symmetric space of non-positive curvature, Proposition 1 provides the exact expression of the normalising factor $Z(\sigma)$, and Proposition 2 provides an algorithm for sampling from the Gaussian distribution $G(\bar{x}, \sigma)$. Also, Propositions 3 and 4 characterise the maximum likelihood estimates of the parameters \bar{x} and σ . In particular, Proposition 3 shows that definition (2) implies Property (1).

S. Said, H. Hajri and L. Bombrun are with Laboratory IMS, CNRS- UMR-5218, University of Bordeaux, Bordeaux, France (email: { salem.said, hatem.hajri, lionel.bombrun}@ims-bordeaux.fr)

B. C. Vemuri is with the Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL, 32611 USA (e-mail: vemuri@cise.ufl.edu).

(1)

— In Section IV: as special cases of the definition of Section III, Gaussian distributions of structured covariance matrices are described for three spaces of structured covariance matrices: complex, Toeplitz, and block-Toeplitz. Each one of these three spaces of structured covariance matrices becomes a Riemannian symmetric space of non-positive curvature, when equipped with a so-called Hessian metric [21]–[24],

$$Riemannian metric = Hessian of entropy$$
(3)

where the entropy function, for a covariance matrix x, is $H(x) = -\log \det(x)$, as in the case of a multivariate normal model [25]. For each one of the three structures, the metric (3) is given explicitly, the normalising factor $Z(\sigma)$ is computed analytically or numerically, and an algorithm for sampling from a Gaussian distribution with known parameters is provided, (these algorithms are described in Appendix II).

— In Section V: efficient statistical learning algorithms are derived, for classification and density estimation of data in a Riemannian symmetric space of non-positive curvature, \mathcal{M} . These algorithms are based on the idea that any probability density on \mathcal{M} , at least if sufficiently regular, can be approximated to any desired precision by a finite mixture of Gaussian densities,

mixture of Gaussian densities:
$$p(x) = \sum_{\kappa=1}^{K} \omega_{\kappa} \times p(x | \bar{x}_{\kappa}, \sigma_{\kappa})$$
 (4)

where $\omega_1, \ldots, \omega_K > 0$ satisfy $\omega_1 + \ldots + \omega_K = 1$ and where each density $p(x | \bar{x}_{\kappa}, \sigma_{\kappa})$ is given by (2). A BIC criterion for selecting the mixture order K, and an EM (expectation-maximisation) algorithm for estimating mixture parameters $\{(\omega_{\kappa}, \bar{x}_{\kappa}, \sigma_{\kappa})\}$ are provided. Applied in the setting of Section IV, these lead to new versions of the algorithms in [15][18]–[20], which deal specifically with structured covariance matrices, and which perform, with a significant reduction in computational complexity, the same tasks as the recent non-parametric methods of [26]–[28].

Table 1. below summarises some of the new contributions of Sections III and IV. The table is intended to help readers navigate directly to desired results.

The main geometric ingredient in the defining Property (1) of Gaussian distributions is the Riemannian barycentre. Recall that the Riemannian barycentre of a set of points x_1, \ldots, x_N in a Riemannian manifold \mathcal{M} is any global minimiser \hat{x}_N of the variance function

variance function:
$$E_N(x) = \frac{1}{N} \sum_{n=1}^N d^2(x, x_n)$$
 (5)

where \hat{x}_N exists and is unique whenever \mathcal{M} is a Riemannian space of non-positive curvature [29][30]. The Riemannian barycentre of covariance matrices arises when $\mathcal{M} = \mathcal{P}_n$, the space of $n \times n$ covariance matrices, equipped with its affine-invariant metric. Its existence and uniqueness follow from the fact that \mathcal{P}_n is a Riemannian symmetric space of non-positive curvature [14][15].

Due to its many applications, the Riemannian barycentre of covariance matrices is the subject of intense ongoing research [11][31][32]. Recently, attention has been directed to the problem of computing the Riemannian barycentre of structured covariance matrices [23][24][33]. This problem arises because if $Y_1, \ldots, Y_N \in \mathcal{P}_n$ are covariance matrices which have some additional structure, then their Riemannian barycentre \hat{Y}_N need not have this same structure [33]. In order to resolve this issue, two approaches have been proposed, as follows.

Note that a space of structured covariance matrices is the intersection of \mathcal{P}_n with some real vector space V of $n \times n$ matrices. This is here denoted $V\mathcal{P}_n = V \cap \mathcal{P}_n$. The vector space V can be that of matrices which have complex, Toeplitz, or block-Toeplitz structure, *etc.* The first approach to computing the Riemannian barycentre \hat{Y}_N of $Y_1, \ldots, Y_N \in V\mathcal{P}_n$ was proposed in [33]. It defines \hat{Y}_N as a minimiser of the variance function (5) over $V\mathcal{P}_n$, rather than over the whole of \mathcal{P}_n . This is an extrinsic approach, since it considers $V\mathcal{P}_n$ as a subspace of \mathcal{P}_n . By contrast, the second approach to computing \hat{Y}_N is an intrinsic approach, proposed in [21]–[24]. It considers $V\mathcal{P}_n$ as a "stand-alone" Riemannian manifold $\mathcal{M} = V\mathcal{P}_n$ equipped with the Hessian metric (3). Then, \hat{Y}_N is given by the general definition of the Riemannian barycentre, applied to this Riemannian manifold \mathcal{M} .

For the present paper, the intrinsic approach to the definition of the Riemannian barycentre of structured covariance matrices is preferred. In this approach, a space $V\mathcal{P}_n$ of structured covariance matrices is equipped with the Hessian metric (3). If $V\mathcal{P}_n$ is the space of complex, Toeplitz, or block-Toeplitz covariance matrices, then it becomes a Riemannian symmetric space of non-positive curvature, when equipped with this metric [21]– [24]. In turn, this allows for the definition of Gaussian distributions, based on Section III. In general, note that $V\mathcal{P}_n$ is an open convex cone in the real vector space V. If $V\mathcal{P}_n$ is moreover a homogeneous and self-dual cone, then it becomes a Riemannian symmetric space of non-positive curvature, when equipped with a Hessian metric of the form (3), which is given in [21][34]. Thus, the definition of Gaussian distributions, based on Section III, extends to any space $V\mathcal{P}_n$ of structured covariance matrices, whenever $V\mathcal{P}_n$ is a homogeneous and self-dual cone.

Throughout the following, the normalising factor $Z(\sigma)$ appearing in (2) will play a major role. Proposition 1 establishes its two main properties : (i) although (2) contains two parameters \bar{x} and σ , this normalising factor does not depend on \bar{x} but only on σ , (ii) when expressed in terms of the "natural parameter" $\eta = -1/2\sigma^2$, this normalising factor is a strictly log-convex function. Proposition 3 shows these two properties are crucial ingredients in the definition of Gaussian distributions. Precisely, if x_1, \ldots, x_N are independent samples from a Gaussian distributions $G(\bar{x}, \sigma)$ given by (2), then the log-likelihood function for the parameters \bar{x} and σ is given by

log-likelihood function of a Gaussian distribution:
$$L(\bar{x},\sigma) = -N \log Z(\sigma) - \frac{1}{2\sigma^2} \sum_{n=1}^{N} d^2(\bar{x},x_n)$$
 (6)

Then, (i) guarantees that the defining Property (1) is verified. Indeed, since the first term in (6) does not depend on \bar{x} , maximisation of $L(\bar{x}, \sigma)$ with respect to \bar{x} is equivalent to minimisation of the sum of squared distances appearing in its second term. However, this is the same as the variance function (5). On the other hand, (ii) guarantees that maximisation of $L(\bar{x}, \sigma)$ with respect to σ is equivalent to a one-dimensional convex optimisation problem, (expressed as a Legendre transform in equation (28)). Note the exact expression of $Z(\sigma)$ is given by formulae (19) and (22) in the Proof of Proposition 1. In [35], a Monte Carlo integration technique was designed which uses these formulae to generate the graph of $Z(\sigma)$. In Section IV, this technique is shown to effectively compute $Z(\sigma)$ when the space \mathcal{M} is a space of large structured covariance matrices.

In conclusion, it is expected that Gaussian distributions on Riemannian symmetric spaces of non-positive curvature will be valuable tools in designing statistical learning algorithms which deal with structured covariance matrices. Such algorithms would combine efficiency with reduced computational complexity, since the main requirement for maximum likelihood estimation of Gaussian distributions is computation of Riemannian barycentres, a task for which there exists an increasing number of high-performance routines. As discussed in Section V, this is the essential advantage of the definition of Gaussian distributions adopted in the present paper, over other recent definitions, also used in designing statistical learning algorithms, such as those considered in [26]–[28].

Table 1.: Gaussian distributions on symmetric spaces and spaces of structured covariance matrices

	Notation	Riemannian metric	Normalising factor $Z(\sigma)$	Sampling algorithm
Section III: symmetric space M			(19) and (22) in Proposition 1	Algorithm 1. based on Proposition 2
Paragraph IV-A : $n \times n$ complex covariance matrices	$\mathcal{M} = \mathcal{H}_n$	Hessian metric (34b)	Monte Carlo integration from (34h)	Algorithm 2. in Appendix II
Paragraph IV-B : $n \times n$ Toeplitz covariance matrices	$\mathcal{M} = \mathcal{T}_n$	Hessian metric (37a)	Analytic expression (37h)	Algorithm 3. in Appendix II
Paragraph IV-C: $n \times n$ block-Toeplitz covariance matrices with $N \times N$ blocks	$\mathcal{M} = \mathcal{T}_n^N$	Hessian metric (40a)	Monte Carlo integration from (40i) and (40j)	Algorithm 4. in Appendix II

II. MATHEMATICAL BACKGROUND ON RIEMANNIAN SYMMETRIC SPACES

This section gives five facts about Riemannian symmetric spaces, Facts 1–5 below, which will be essential for the following. These are here stated in their general form, but three concrete examples are presented in Section IV.

A Riemannian symmetric space is a Riemannian manifold \mathcal{M} such that, for each point $x \in \mathcal{M}$, there exists an isometry transformation $I_x : \mathcal{M} \to \mathcal{M}$ whose effect is to reverse geodesic curves passing through x [36] (Chapter IV, Page 170). Any Riemannian symmetric space is a Riemannian homogeneous space, although the converse is not true. In particular, Riemannian symmetric spaces share in two important properties of Riemannian homogeneous spaces, given in Facts 1 and 2 below: invariance of distance, and invariance of integrals.

To say that \mathcal{M} is a Riemannian homogeneous space means a Lie group G acts transitively and isometrically on \mathcal{M} [36] (Chapter II, Page 113). Each element $g \in G$ defines a geometric transformation of \mathcal{M} which maps each point $x \in \mathcal{M}$ to its image $g \cdot x$. These transformations satisfy the group action property,

$$(g_1g_2) \cdot x = g_1 \cdot (g_2 \cdot x) \tag{7}$$

Then, transitivity of group action means that for any two points x and y of \mathcal{M} , there exists $g \in G$ such that $g \cdot x = y$.

When G acts on \mathcal{M} transitively, one says that \mathcal{M} is a homogeneous space. Homogeneous spaces admit a general description, as the quotient of G by a compact Lie subgroup H. Precisely, choose some point $o \in \mathcal{M}$ as origin, and let H be the subgroup of all $h \in G$ such that $h \cdot o = o$ [36] (Chapter II, Page 113). Then, \mathcal{M} is identified with the quotient G/H.

When G acts on \mathcal{M} isometrically, the transformations $x \mapsto g \cdot x$ preserve the Riemannian metric and Riemannian distance of \mathcal{M} . This is expressed in Fact 1 below, which uses the following notation [37]. The Riemannian metric of \mathcal{M} is denoted ds^2 . For $x \in \mathcal{M}$, this defines a quadratic form: $ds_x^2(dx) =$ squared length of a displacement dx away from x. The Riemannian distance corresponding to the metric ds^2 is denoted d(x, y) for $x, y \in \mathcal{M}$. Fact 1 is not a theorem, but really part of the definition of a Riemannian homogeneous space [36].

▶ *Fact* 1 (Invariant distance): For $g \in G$ and $x, y \in \mathcal{M}$,

p

reservation of metric:
$$z = g \cdot x \Rightarrow ds_z^2(dz) = ds_x^2(dx)$$
 (8a)

$$d(g \cdot x, g \cdot y) = d(x, y) \tag{8b}$$

Turning to Fact 2, recall that the Riemannian metric ds^2 defines a Riemannian volume element dv [37]. Fact 1 states that the transformations $x \mapsto g \cdot x$ preserve the metric ds^2 . Therefore, these transformations must also preserve the volume element dv. This leads to the invariance of integrals with respect to dv, as follows [36] (Chapter X, Page 361).

▶ Fact 2 (Invariant integrals): For any integrable function $f : \mathcal{M} \to \mathbb{R}$ and $g \in G$,

preservation of volume:
$$\int_{\mathcal{M}} f(g \cdot x) \, dv(x) = \int_{\mathcal{M}} f(x) \, dv(x) \tag{9}$$

Facts 1 and 2 above hold for any Riemannian homogeneous space. The following fact 3 requires that \mathcal{M} should have non-positive curvature. This guarantees the existence and uniqueness of Riemannian barycentres in \mathcal{M} [29].

Let π be a probability distribution on \mathcal{M} . The Riemannian barycentre of π is a point \hat{x}_{π} defined as follows. Consider $E_{\pi}: \mathcal{M} \to \mathbb{R}_+$,

variance function:
$$E_{\pi}(x) = \int_{\mathcal{M}} d^2(x, z) d\pi(z)$$
 (10)

The Riemannian barycentre \hat{x}_{π} is then a global minimiser of E_{π} . Fact 3 states that \hat{x}_{π} exists and is unique [29] (Theorem 2.1., Page 659). Fact 3 (*Riemannian barycentre*): Assume that \mathcal{M} has non-positive curvature. For any probability distribution π on \mathcal{M} , the Riemannian

barycentre \hat{x}_{π} exists and is unique. Moreover, \hat{x}_{π} is the unique stationary point of the variance function E_{π} . Fact 2 above states that integrals with respect to the Riemannian volume element dv are invariant. However, it does not show how these integrals can be computed. This is done by the following Fact 4, which requires that \mathcal{M} should be a Riemannian symmetric space of non-compact type. This means that the group G is a semisimple Lie group of non-compact type [36] (Chapter V). It should be noted that this property implies that \mathcal{M} has (strictly) negative curvature, so that Fact 3 holds true.

To state Fact 4, assume that G is a semisimple Lie group of non-compact type. Recall that H is a compact Lie subgroup of G. Let \mathfrak{g} and \mathfrak{h} denote the Lie algebras of G and H respectively, and let \mathfrak{g} have the following Iwasawa decomposition [36] (see details in Chapter VI, Page 219),

Iwasawa decomposition:
$$g = h + a + n$$
 (11)

where a is an Abelian subalgebra of \mathfrak{g} , and \mathfrak{n} is a nilpotent subalgebra of \mathfrak{g} . Each $x \in \mathcal{M}$ can be written in the following form [36] (Chapter X),

polar coordinates:
$$x = \exp(\operatorname{Ad}(h) a) \cdot o$$
 $a \in \mathfrak{a}, h \in H$ (12)

where exp : $\mathfrak{g} \to G$ is the Lie group exponential, and Ad is the adjoint representation. Here, it will be said that (a, h) are polar coordinates of x, which is then written x = x(a, h) (this is an abuse of notation, since (a, h) are not unique, but it will provide useful intuition).

The following integral formula holds [36] (Chapter X, Page 382). It makes use of the concept of roots. These are linear mappings $\lambda : \mathfrak{a} \to \mathbb{R}$, which characterise the Lie algebra \mathfrak{g} . For a minimal, introduction to this concept, see Appendix I.

▶ Fact 4 (Integration in polar coordinates): For any integrable function $f : \mathcal{M} \to \mathbb{R}$,

$$\int_{\mathcal{M}} f(x) \, dv(x) = \mathcal{C} \times \int_{H} \int_{\mathfrak{a}} f(a, h) \, D(a) \, da \, dh \tag{13}$$

where C is a constant which does not depend on the function f, dh is the normalised Haar measure on H and da is the Lebesgue measure on \mathfrak{a} . The function $D : \mathfrak{a} \to \mathbb{R}_+$ is given by

$$D(a) = \prod_{\lambda>0} \sinh^{m_{\lambda}}(|\lambda(a)|) \tag{14}$$

where the product is over positive roots $\lambda : \mathfrak{a} \to \mathbb{R}$, and m_{λ} is the dimension of the root space corresponding to λ .

Fact 5 provides a geometric interpretation of the polar coordinates (a, h) introduced in (12). Just like usual polar coordinates, these give a direct description of geodesics and distance, with respect to the origin o [36] (Chapter IV).

▶ Fact 5 (Geometric interpretation of polar coordinates): The two following properties hold true.

(i) Geodesic curves passing through the origin $o \in \mathcal{M}$ are exactly the curves $x : \mathbb{R} \to \mathcal{M}$ given by x(t) = x(t a, h) where a and h are constant. (ii) The Riemannian distance between the origin $o \in \mathcal{M}$ and x(a, h) is given by

distance to origin:
$$d^2(o, x) = B(a, a)$$
 (15)

where $B: \mathfrak{g} \times \mathfrak{g} \to \mathbb{R}$ is the Killing form, an Ad-invariant symmetric bilinear form on \mathfrak{g} , which is positive definite on \mathfrak{a} [36] (Chapter II).

III. GAUSSIAN DISTRIBUTIONS ON RIEMANNIAN SYMMETRIC SPACES

This section introduces Gaussian distributions on Riemannian symmetric spaces, and studies maximum likelihood estimation of their parameters. These distributions establish a connection between the statistical concept of maximum likelihood estimation, and the geometric concept of Riemannian barycentre.

Here, \mathcal{M} is a Riemannian symmetric space of non-positive curvature. Recall d(x, y) denotes Riemannian distance between $x, y \in \mathcal{M}$, and dv denotes the Riemannian volume element of \mathcal{M} . A Gaussian distribution is one that has the following probability density with respect to dv,

$$p(x|\bar{x},\sigma) = \frac{1}{Z(\sigma)} \times \exp\left[-\frac{d^2(x,\bar{x})}{2\sigma^2}\right]$$
(16)

where $\bar{x} \in \mathcal{M}$ and $\sigma > 0$. This distribution will be denoted $G(\bar{x}, \sigma)$.

The main feature of the probability density (16) is that the normalising factor $Z(\sigma)$ does not depend on the parameter \bar{x} . This will be proved as part of Proposition 1 below, and is the key ingredient in the relationship between Gaussian distributions and the concept of Riemannian barycentre. In [14][15], this relationship was developed in the special case where $\mathcal{M} = \mathcal{P}_n$, the symmetric space of real $n \times n$ covariance matrices. The present paper considers the general case where \mathcal{M} is any symmetric space of non-positive curvature. This allows for Gaussian distributions of the form (16) to be extended to spaces of structured covariance matrices, such as complex, Toeplitz, and block-Toeplitz covariance matrices, as done in Section IV below.

In the following, Paragraph III-A gives Proposition 1 which characterises the normalising factor $Z(\sigma)$, and Proposition 2 which explains how to sample from a Gaussian distribution with given parameters. Paragraph III-B is concerned with maximum likelihood estimation of the parameters of a Gaussian distribution. In particular, it contains Proposition 3 which shows that the maximum likelihood estimate \hat{x}_N of the parameter \bar{x} is precisely the Riemannian barycentre of available observations x_1, \ldots, x_N . Finally, Proposition 4 can be used to obtain the consistency and asymptotic normality of \hat{x}_N .

A. Definition : the role of invariance

This paragraph shows how the invariance properties of the Riemannian symmetric space \mathcal{M} , described in Section II, can be used to characterise Gaussian distributions on \mathcal{M} . Precisely, for a Gaussian distribution $G(\bar{x}, \sigma)$, it is possible to obtain a) an exact expression of the normalising factor $Z(\sigma)$ appearing in (16), and b) a generally applicable method for sampling from this distribution $G(\bar{x}, \sigma)$.

Proposition 1 is concerned with the normalising factor $Z(\sigma)$. A fully general expression of $Z(\sigma)$ is given by (19) and (22) in the proof of this proposition. To state Proposition 1, consider the following notation,

$$f(x|\bar{x},\sigma) = \exp\left[-\frac{d^2(x,\bar{x})}{2\sigma^2}\right] \qquad \qquad Z(\bar{x},\sigma) = \int_{\mathcal{M}} f(x|\bar{x},\sigma) \, dv(x) \tag{17}$$

In other words, $f(x|\bar{x}, \sigma)$ is the "profile" of the probability density $p(x|\bar{x}, \sigma)$.

Proposition 1 (Normalising factor): The two following properties hold true.

(i) $Z(\bar{x}, \sigma) = Z(o, \sigma)$ for any $\bar{x} \in \mathcal{M}$, where o is some point chosen as origin of \mathcal{M} .

(ii) $Z(\sigma) = Z(\sigma, \sigma)$ is a strictly log-convex function of the parameter $\eta = -1/2\sigma^2$.

Proof: the proof of (i) follows from Facts 1 and 2. Since G acts transitively on \mathcal{M} , any $\bar{x} \in \mathcal{M}$ is of the form $x = g \cdot o$ for some $g \in G$. Applying (8b) from Fact 2 to (17), it then follows that $f(x|\bar{x},\sigma) = f(g^{-1} \cdot x|o,\sigma)$. Accordingly,

$$Z(\bar{x},\sigma) = \int_{\mathcal{M}} f(x|\bar{x},\sigma) \, dv(x) = \int_{\mathcal{M}} f(g^{-1} \cdot x|o,\sigma) \, dv(x)$$

Now, by Fact 2, (and using g^{-1} instead of g in (9)),

$$\int_{\mathcal{M}} f(g^{-1} \cdot x | o, \sigma) \, dv(x) = \int_{\mathcal{M}} f(x | o, \sigma) \, dv(x) = Z(o, \sigma)$$

so that (i) is indeed true. The proof of (ii) is based on Facts 4 and 5. Assume first that \mathcal{M} is a Riemannian symmetric space of non-compact type. By (13) and (15),

$$Z(\sigma) = \int_{\mathcal{M}} f(x|o,\sigma) \, dv(x) = \mathcal{C} \times \int_{H} \int_{\mathfrak{a}} \exp\left[-\frac{B(a,a)}{2\sigma^2}\right] D(a) \, da \, dh \tag{18}$$

But the function under the integral does not depend on h. Since dh is the normalised Haar measure on H,

$$Z(\sigma) = C \times \int_{\mathfrak{a}} \exp\left[-\frac{B(a,a)}{2\sigma^2}\right] D(a) \, da \tag{19}$$

In general, a Riemannian symmetric space \mathcal{M} of non-positive curvature decomposes into a product Riemannian symmetric space [36] (Chapter V, Page 208),

$$\mathcal{M} = \mathcal{M}_1 \times \ldots \times \mathcal{M}_r \tag{20}$$

where each \mathcal{M}_p is either a Euclidean space or a Riemannian symmetric space of non-compact type. Now, each $x \in \mathcal{M}$ is an r-tuple $x = (x_1, \ldots, x_r)$ where $x_p \in \mathcal{M}_p$. Moreover, for $x, y \in \mathcal{M}$,

$$d^{2}(x,y) = \sum_{p=1}^{r} d_{p}^{2}(x_{p}, y_{p})$$
(21)

where $d_p(x_p, y_p)$ is the Riemannian distance in \mathcal{M}_p between x_p and y_p (where $y = (y_1, \ldots, y_r)$). Replacing (21) in (17), it follows that,

$$Z(\sigma) = Z_1(\sigma) \times \ldots \times Z_r(\sigma)$$
⁽²²⁾

where $Z_p(\sigma)$ is the normalising factor for the space \mathcal{M}_p . Recall that a product of strictly log-convex functions is strictly log-convex. If \mathcal{M}_p is a Euclidean space of dimension q, then $Z_p(\sigma) = (2\pi\sigma^2)^{-q/2}$, which is a strictly log-convex function of η . If \mathcal{M}_p is a Riemannian symmetric space of non-compact type, $Z_p(\sigma)$ is given by (19). Let $\rho = B(a, a)$ and $\mu(d\rho)$ the image of D(a)da under the mapping $a \mapsto \rho$. It follows that

$$Z(\sigma) = C \times \int_0^\infty e^{\eta \, \rho} \, \mu(d\rho) \tag{23}$$

As a function of η , this is the moment-generating function of the positive measure $\mu(d\rho)$, and therefore a strictly log-convex function [38], (Chapter 7, Page 103). Thus, $Z(\sigma)$ is a product (22) of strictly log-convex functions of η , and therefore itself strictly log-convex. This proves (ii).

The following Proposition 2 shows that polar coordinates (12) can be used to effectively sample from a Gaussian distribution $G(\bar{x}, \sigma)$. This proposition is concerned with the special case where \mathcal{M} is a Riemannian symmetric space of non-compact type. However, the general case where \mathcal{M} is a Riemannian symmetric space of non-positive curvature can then be obtained using decomposition (20). This is explained in the remark following the proposition.

In the statement of Proposition 2, the following usual notation is used. If X is a random variable and P is a probability distribution, $X \sim P$ means that X follows the probability distribution P.

Proposition 2 (Gaussian distribution via polar coordinates): Let \mathcal{M} be a Riemannian symmetric space of non-compact type. Let $\mathcal{M} = G/H$ where g has Iwasawa decomposition (11). Let h and a be independent random variables in H and a, respectively. Assume h is uniformly distributed on H, and a has probability density, with respect to the Lebesgue measure da on a,

$$p(a) \propto \exp\left[-\frac{B(a,a)}{2\sigma^2}\right] D(a)$$
 (24)

where \propto indicates proportionality. The following hold true.

(i) If x = x(a, h) as in (12), then $x \sim G(o, \sigma)$.

(ii) If
$$x \sim G(o, \sigma)$$
 and $\bar{x} = g \cdot o$, then $g \cdot x \sim G(\bar{x}, \sigma)$.

Proof: the proof of (ii) is a direct application of Fact 2, (compare to [15], proof of Proposition 5). Only the proof of (i) will be detailed. This follows using Fact 4. Recall that a uniformly distributed random variable h in H has constant probability density, identically equal to 1, with respect to the normalised Haar measure dh. Let h and a be as in the statement of the proposition. Since h and a are independent, (24) implies that for any function $\varphi : H \times \mathfrak{a} \to \mathbb{R}$, the expectation of $\varphi(h, a)$ is equal to

$$\int_{H} \int_{\mathfrak{a}} \varphi(h,a) \, p(a) \, da \, dh \, \propto \, \int_{H} \int_{\mathfrak{a}} \varphi(h,a) \, \exp\left[-\frac{B(a,a)}{2\sigma^{2}}\right] \, D(a) \, da \, dh$$

Then, if x = x(a, h) as in (12) and $\tilde{\varphi} : \mathcal{M} \to \mathbb{R}$, the expectation of $\tilde{\varphi}(x)$ is proportional to,

$$\int_{H} \int_{\mathfrak{a}} \tilde{\varphi}(x(h,a)) \exp\left[-\frac{B(a,a)}{2\sigma^{2}}\right] D(a) \, da \, dh \propto \int_{\mathcal{M}} \tilde{\varphi}(x) \, p(x|o,\sigma) \, dv(x)$$

as follows from (13), (15) and (16). Since the function $\tilde{\varphi}$ is arbitrary, this shows that the probability density of x with respect to dv is equal to $p(x|o, \sigma)$. In other words, $x \sim G(o, \sigma)$.

▶ Remark 1 : Proposition 2 only deals with the special case where \mathcal{M} is a Riemannian symmetric space of non-compact type. The general case where \mathcal{M} is a Riemannian symmetric space of non-positive curvature is obtained as follows. Let \mathcal{M} have decomposition (20) and $x = (x_1, \ldots, x_r)$ be a random variable in \mathcal{M} with distribution $G(\bar{x}, \sigma)$. If $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_r)$, then it follows from (16) and (21) that

$$p(x|\bar{x},\sigma) = p(x_1|\bar{x}_1,\sigma) \times \ldots \times p(x_r|\bar{x}_r,\sigma)$$
⁽²⁵⁾

Therefore, x_1, \ldots, x_r are independent and each x_p has distribution $G(\bar{x}_p, \sigma)$. Conversely, (25) implies that if x_1, \ldots, x_r are independent and each x_p has distribution $G(\bar{x}_p, \sigma)$, then x has distribution $G(\bar{x}, \sigma)$.

Proposition 2, along with the subsequent remark, provides a generally applicable method for sampling from a Gaussian distribution $G(\bar{x}, \sigma)$ on a Riemanniain symmetric space of non-positive curvature \mathcal{M} . This proceeds as follows, using decomposition (20).

11. for p = 1, ..., r do if \mathcal{M}_p is a Euclidean space 12. 13. sample x_p from a multivariate normal distribution $\mathcal{N}(\bar{x}_p, \sigma)$ 14 end if 15. if $\mathcal{M}_p = G/H$ is a Riemannian symmetric space of non-compact type sample h from a uniform distribution on H16. 17. sample a from the multivariate density (24) 18. put $x_p = x(a, h)$ as in (12) 19. reset x_p to $g \cdot x_p$ where $\bar{x} = g \cdot o$ as in (ii) of Proposition 2. 110. end if 111.end for

112. put $x = (x_1, ..., x_r)$

Algorithm 1. : sampling from a Gaussian distribution on a symmetric space \mathcal{M}

The above steps are completely general, and their concrete implementation depends on the specific space \mathcal{M} at hand. When $\mathcal{M} = \mathcal{P}_n$, the symmetric space of real $n \times n$ covariance matrices, this was described in [15] (Page 8). Three new implementations, for spaces of structured covariance matrices, will be given in Appendix II.

B. Maximum likelihood estimation

The current paragraph studies maximum likelihood estimation of the parameters \bar{x} and σ of a Gaussian distribution $G(\bar{x}, \sigma)$, on a Riemannian symmetric space of non-positive curvature \mathcal{M} . Proposition 3 characterises the resulting maximum likelihood estimates, and Proposition 4 yields their asymptotic behavior. Let \hat{x}_N and $\hat{\sigma}_N$ be maximum likelihood estimates, based on independent samples x_1, \ldots, x_N from $G(\bar{x}, \sigma)$. Proposition 3 shows that \hat{x}_N and $\hat{\sigma}_N$ provide a geometric description of these samples. Indeed, the proposition states that \hat{x}_N is the Riemannian barycentre of x_1, \ldots, x_N , and that $\hat{\sigma}_N$ measures the dispersion (that is, mean squared distance) of x_1, \ldots, x_N away from this Riemannian barycentre.

For Proposition 3, recall the Riemannian barycentre of x_1, \ldots, x_N is a global minimiser of the function $E_N : \mathcal{M} \to \mathbb{R}$, defined as follows [29],

empirical variance function:
$$E_N(x) = \frac{1}{N} \sum_{n=1}^N d^2(x, x_n)$$
 (26)

This is called an empirical variance function, since it is a special case of the variance function (10), corresponding to the empirical distribution $\pi = N^{-1} \sum_{n=1}^{N} \delta_{x_n}$, where δ_x denotes the Dirac distribution concentrated at $x \in \mathcal{M}$. Accordingly, existence and uniqueness of the Riemannian barycentre of x_1, \ldots, x_N follow from Fact 3.

Proposition 3 (MLE and Riemannian barycentre): Let x_1, \ldots, x_N be independent samples from $G(\bar{x}, \sigma)$. The maximum likelihood estimates \hat{x}_N and $\hat{\sigma}_N$, based on these samples, exist and are unique. Precisely,

(i) \hat{x}_N is the Riemannian barycentre of x_1, \ldots, x_N .

(ii) There exists a strictly increasing bijective function $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$, such that

$$\hat{\sigma}_N = \Phi\left(\frac{1}{N}\sum_{n=1}^N d^2(\hat{x}_N, x_n)\right)$$
(27)

Proof: from (16), the log-likelihood function for the parameters \bar{x} and σ is given by

$$L(\bar{x},\sigma) = -N \log Z(\sigma) - \frac{1}{2\sigma^2} \sum_{n=1}^{N} d^2(\bar{x},x_n)$$

By definition, \hat{x}_N and $\hat{\sigma}_N$ are found by maximising this function. The sum in the second term on the right hand side does not depend on σ . Therefore, \hat{x}_N can be found separately, by minimising this sum over the values of \bar{x} (minimisation instead of maximisation is due to the minus sign ahead of the sum). However, this is the same as minimising the empirical variance function (26). Now, the unique global minimiser of the empirical variance function is the Riemannian barycentre of x_1, \ldots, x_N , whose existence and uniqueness follow from Fact 3 as noted before the proposition. This proves (i). Now, assume \hat{x}_N has been found. Then, to find $\hat{\sigma}_N$, it is convenient to maximise $L(\hat{x}_N, \sigma)$ over the values of the parameter $\eta = -1/2\sigma^2$, which was already used in Proposition 1. Indeed, note that

$$rac{1}{N} L(\hat{x}_N,\sigma) = \eta \, \hat{
ho} - \psi(\eta)$$

where $\hat{\rho} = N^{-1} \sum_{n=1}^{N} d^2(\hat{x}_N, x_n)$ and $\psi(\eta) = \log Z(\sigma)$. Thus, the maximum likelihood estimate $\hat{\eta}_N$ is given by

$$\hat{\eta}_N = \operatorname{argmax}_n \left\{ \eta \,\hat{\rho} - \psi(\eta) \right\} \tag{28}$$

Recall that (ii) of Proposition 1 states that ψ is a strictly convex function, so its derivative ψ' is strictly increasing. But, from (28), $\hat{\eta}_N$ is given by $\psi'(\hat{\eta}_N) = \hat{\rho}$. It follows that $\hat{\eta}_N$ exists and is unique, whenever $\hat{\rho}$ is in the range of ψ' . It is possible to show, using (23), that this range is equal to \mathbb{R}_+ , which implies existence and uniqueness of $\hat{\eta}_N$ in general. Now, to complete the proof of (ii), it is enough to perform the change of parameter from $\hat{\eta}_N$ back to $\hat{\sigma}_N$.

Proposition 3 shows how the maximum likelihood estimates \hat{x}_N and $\hat{\sigma}_N$ can be computed, based on the samples x_1, \ldots, x_N . Since \hat{x}_N is the Riemannian barycentre of these samples, \hat{x}_N can be computed using a variety of existing algorithms for computation of Riemannian barycentres. These include gradient descent [39][40], Newton method [41][42], recursive estimators [43][44], and stochastic methods [45][46]. Some of these algorithms are specifically adapted to Riemannian symmetric spaces of non-positive curvature (that is, to the context of the present paper)[42]–[44]. Least computationally expensive are the recursive estimators [43][44]. On the other hand, the most computationally expensive, involving exact calculation of the Hessian of squared Riemannian distance, is the Newton method of [42]. Computation of $\hat{\sigma}_N$ amounts to solving the equation $\psi'(\hat{\eta}_N) = \hat{\rho}$. This is a scalar non-linear equation, which can be solved using standard methods, such as the Newton-Raphson algorithm.

Remark 2: the entropy of a Gaussian distribution $G(\bar{x}, \sigma)$ is directly related to the function ψ introduced in the proof of Proposition 1. By definition, the entropy of $G(\bar{x}, \sigma)$ is [25],

$$h(\bar{x}, \sigma) = \int_{\mathcal{M}} \log p(x|\bar{x}, \sigma) \times p(x|\bar{x}, \sigma) \, dv(x)$$

Using Fact 2 as in the proof of (i) in Proposition 1, it follows $h(\bar{x}, \sigma) = h(o, \sigma)$, so the entropy does not depend on \bar{x} . Evaluation of $h(o, \sigma)$ gives, in the notation of (23),

$$h(o,\sigma) = C \times \int_0^\infty (\eta \rho - \psi(\eta)) \times e^{\eta \rho - \psi(\eta)} \mu(d\rho)$$

Following [38] (Chapter 9), let ψ^* be the Legendre transform of ψ and $\bar{\rho} = \psi'(\eta)$. Then, the entropy $h(o, \sigma)$ is equal to

entropy of a Gaussian distribution: $\psi^*(\bar{\rho}) = \eta \bar{\rho} - \psi(\eta)$ (29)

Essentially, this is due to the fact that, when \bar{x} is fixed, $G(\bar{x}, \sigma)$ takes on the form of an exponential distribution with natural parameter η and sufficient statistic ρ .

The following Proposition 4 clarifies the significance of the parameters \bar{x} and σ of a Gaussian distribution $G(\bar{x}, \sigma)$. This Proposition is an asymptotic counterpart of Proposition 3. In particular, it can be used to obtain the consistency and asymptotic normality of \hat{x}_N . This is explained in the remark after the proposition. For the statement of Proposition 4, recall the definition (10) of the variance function $E_{\pi} : \mathcal{M} \to \mathbb{R}_+$ of a probability distribution π on \mathcal{M} . If $\pi = G(\bar{x}, \sigma)$, denote $E_{\pi}(x) = E(x|\bar{x}, \sigma)$.

Proposition 4 (significance of the parameters): The parameters \bar{x} and σ of a Gaussian distribution $G(\bar{x}, \sigma)$ are given by,

$$\bar{x}$$
 is the Riemannian barycentre of $G(\bar{x}, \sigma)$: $\bar{x} = \operatorname{argmin}_{x \in \mathcal{M}} E(x|\bar{x}, \sigma)$ (30a)

$$\sigma = \Phi\left(\int_{\mathcal{M}} d^2(\bar{x}, z) \, p(z|\bar{x}, \sigma) \, dv(z)\right) \tag{30b}$$

where $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ is the strictly increasing function introduced in Proposition 3.

Proof: the proof can be carried out by direct generalisation of the one in [15] (proof of Proposition 9). An alternative, shorter proof of (30a), based on the definition of a Riemannian symmetric space, can be obtained by generalising [14] (proof of Theorem 2.3.).

Remark 3: Proposition 4 yields the consistency and asymptotic normality of the maximum likelihood estimate \hat{x}_N . Precisely,

consistency of
$$\hat{x}_N$$
: $\lim_{N} \hat{x}_N = \bar{x}$ (31a)

asymptotic normality of
$$\hat{x}_N$$
: $\sqrt{N} \operatorname{Log}_{\bar{x}}(\hat{x}_N) \Rightarrow N_d(0, C^{-1})$ (31b)

where Log denotes the Riemannian logarithm mapping [37], and N_d a normal distribution in d dimensions, (d being the dimension of \mathcal{M}). Here,

$$C = \int_{\mathcal{M}} \left(\sigma^{-2} \operatorname{Log}_{\bar{x}}(z) \right) \otimes \left(\sigma^{-2} \operatorname{Log}_{\bar{x}}(z) \right) \, p(z|\,\bar{x},\sigma) \, dv(z) \tag{32}$$

where \otimes denotes exterior product of tangent vectors to \mathcal{M} . Indeed, (31a) follows from Propositions 3 and 4. According to [47] (Theorem 2.3., Page 8), the Riemannian barycentre \hat{x}_N of independent samples x_1, \ldots, x_N from $G(\bar{x}, \sigma)$ converges to the Riemannian barycentre of $G(\bar{x}, \sigma)$. But this is exactly \bar{x} . As for (31b), it can be proved by direct generalisation of the proof in [15] (Proposition 11).

IV. APPLICATION TO SPACES OF STRUCTURED COVARIANCE MATRICES

The previous Section III introduced the notion of a Gaussian distribution on a Riemannian symmetric space of non-positive curvature, \mathcal{M} . In the present section, the development of Section III is applied to the special case where \mathcal{M} is a space of structured covariance matrices. Three variants are considered: $\mathcal{M} =$ space of complex covariance matrices, $\mathcal{M} =$ space of Toeplitz covariance matrices, and $\mathcal{M} =$ space of block-Toeplitz covariance matrices. In [14][15], Gaussian distributions were studied for $\mathcal{M} =$ space of real covariance matrices. Here, they are extended to spaces of covariance matrices with additional structure.

It seems that [21][22], in general, a space of structured covariance matrices becomes a Riemannian symmetric space of non-positive curvature, when it is equipped with a so-called Hessian metric. Roughly, a Hessian metric is a Riemannian metric arising from the entropy function,

$$Riemannian metric = Hessian of entropy$$
(33)

where, if \mathcal{M} is a space of structured covariance matrices, the entropy of $x \in \mathcal{M}$ is minus its log-determinant, $H(x) = -\log \det(x)$ [25] (Chapter 8, Page 254), as in the usual formula for the entropy of a multivariate normal distribution, with covariance x.

For each of the three variants mentioned above, the following development will be carried out. First, (33) is used to obtain a Riemannian metric ds^2 on \mathcal{M} . It turns out that, as intended, this metric makes \mathcal{M} into a Riemannian symmetric space of non-positive curvature. Second, Gaussian distributions on \mathcal{M} are characterised, as in Section III. Precisely, the graphs of the functions $\log Z : \mathbb{R}_+ \to \mathbb{R}$ and $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ are given based on Propositions 1 and 3. Further, the reader is referred to Appendix II, for algorithms which provide samples from a Gaussian distribution on \mathcal{M} .

A. Complex covariance matrices

Here, $\mathcal{M} = \mathcal{H}_n$, the space of $n \times n$ complex covariance matrices. Precisely, each element $Y \in \mathcal{H}_n$ is an $n \times n$ Hermitian positive definite matrix. This space \mathcal{H}_n is equipped with a Hessian metric as follows.

- Hessian metric: application of the general prescription (33) leads to a Riemannian metric on \mathcal{H}_n which is identical to the well-known affine-invariant metric, used in [48][3]. Indeed, the Hessian of the entropy function $H(Y) = -\log \det(Y)$ has the following expression,

Hessian of entropy:
$$D^2 H(Y)[v, w] = \operatorname{tr} \left[Y^{-1} v Y^{-1} w \right] \quad Y \in \mathcal{H}_n$$

for any complex matrices v and w, where tr denotes the trace. This expression can be found using the matrix differentiation formulae, as in [49] (Appendix A, Page 196),

$$D\log \det(Y)[w] = \operatorname{tr} \left[Y^{-1}w \right] \qquad DY^{-1}[v] = -Y^{-1}vY^{-1}$$
(34a)

Replacing in (33), gives the affine-invariant metric of \mathcal{H}_n ,

$$ds_Y^2(dY) = D^2 H(Y)[dY, dY] = \operatorname{tr}\left[Y^{-1}dY Y^{-1}dY\right] \quad Y \in \mathcal{H}_n$$
(34b)

The Riemannian distance corresponding to the metric (34b) has the following expression [48][3],

$$d^{2}(X,Y) = \operatorname{tr}\left[\log\left(X^{-1/2}YX^{-1/2}\right)\right]^{2} \quad X,Y \in \mathcal{H}_{n}$$
(34c)

where all matrix functions, (matrix power and logarithm), are Hermitian matrix functions [50]. Moreover, the Riemannian volume element dv is given by, (this can be checked using [51], Lemma 2.1),

$$dv(Y) = \det(Y)^{-n} \prod_{i \le j} \operatorname{Re} dY_{ij} \prod_{i < j} \operatorname{Im} dY_{ij}$$
(34d)

where subscripts denote matrix elements, and Re and Im denote real and imaginary parts.

- Symmetric space: equipped with the Riemannian metric (34b), \mathcal{H}_n is a Riemannian symmetric space of non-positive curvature [36]. Accordingly, it is possible to apply Facts 1–5. To do so, note that $G = GL(n, \mathbb{C})$, the Lie group of $n \times n$ invertible complex matrices, acts on \mathcal{H}_n transitively and isometrically by congruence transformations [48],

$$g \cdot Y = g Y g^H \qquad g \in G, \, Y \in \mathcal{H}_n \tag{34e}$$

where ^{*H*} denotes the conjugate transpose. Then, Facts 1 and 2 state that definitions (34b)–(34d) verify the general identities (8) and (9). Furthermore, Fact 3 states the existence and uniqueness of Riemannian barycentres in \mathcal{H}_n .

The application of Facts 4 and 5 to \mathcal{H}_n leads to the two following formulae, which are proved in Example 1 of Appendix I,

$$\int_{\mathcal{H}_n} f(Y) \, dv(Y) = \mathcal{C} \times \int_{U(n)} \int_{\mathbb{R}^n} f(r, U) \prod_{i < j} \sinh^2(|r_i - r_j|/2) \, dr \, dU \tag{34f}$$

$$l^{2}(I,Y) = ||r||^{2} = \sum_{j=1}^{n} r_{j}^{2}$$
 ($I = n \times n$ identity matrix) (34g)



Fig. 1: $\log Z$ and Φ for 20×20 complex covariance matrices

For \mathcal{H}_n , formulae (34f) and (34g) play the same role as formulae (13) and (15), for a general symmetric space \mathcal{M} . Here $Y = Ue^r U^H$ is the spectral decomposition of Y, with $U \in U(n)$, the group of $n \times n$ unitary matrices, and $e^r = \text{diag}(e^{r_1}, \ldots, e^{r_n})$ for $r = (r_1, \ldots, r_n) \in \mathbb{R}^n$. Moreover, dr denotes the Lebesgue measure on \mathbb{R}^n and dU denotes the normalised Haar measure on U(n).

- Gaussian distributions: the functions $\log Z : \mathbb{R}_+ \to \mathbb{R}$ and $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$, for Gaussian distributions on \mathcal{H}_n , can be obtained using formulae (34f) and (34g). These yield the following expression for $Z(\sigma)$, similar to (18) in the proof of Proposition 1,

$$Z(\sigma) = \int_{\mathcal{H}_n} \exp\left[-\frac{d^2(I,Y)}{2\sigma^2}\right] dv(Y) = \mathcal{C} \times \int_{U(n)} \int_{\mathbb{R}^n} \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i < j} \sinh^2(|r_i - r_j|/2) dr dU$$

Since the function under the double integral does not depend on U, this integral simplifies to a formula similar to formula (19) in the proof of Proposition 1,

$$Z(\sigma) = \mathcal{C} \times \int_{\mathbb{R}^n} \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i < j} \sinh^2(|r_i - r_j|/2) dr$$
(34h)

For any moderate value of n, (in practice, up to $n \approx 40$ has been considered), this formula can be numerically evaluated on a desktop computer, using a specifically designed Monte Carlo technique [35], leading to the graph of $Z(\sigma)$. Then, obtaining the graphs of $\log Z : \mathbb{R}_+ \to \mathbb{R}$ and $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ is straightforward. In particular, it follows from (28) in the proof of Proposition 3 that the function Φ is obtained by solving, for each $\rho \in \mathbb{R}_+$, the equation $\rho = \psi'(\eta)$. Denoting $\hat{\eta}$ the unique solution of this equation, it follows that $\Phi(\rho) = \hat{\sigma}$ where $\hat{\eta} = -1/2\hat{\sigma}^2$.

Figure 1 shows the graphs of $\log Z$ and Φ when n = 20. The constant factor C appearing in (34h), which depends on n through an involved expression (compare to [15], Proposition 4), is ignored in these graphs. This has no influence on the graph of Φ .

► Remark 4 : the space \mathcal{H}_2 of 2 × 2 complex covariance matrices plays a central role in polarimetric synthetic aperture radar (PolSAR) image analysis [52][53]. Each pixel in a PolSAR image is given by a covariance matrix $Y \in \mathcal{H}_2$, which is the coherency matrix of a polarised lightwave. For the special case of the space \mathcal{H}_2 , the analytic expression of the function $Z(\sigma)$ can be found by integrating (34h),

$$Z(\sigma) = \mathcal{C} \times \pi \sigma^2 \left(e^{\sigma^2} - 1 \right)$$
(34i)

Formula (34i) gives the same behavior as the function represented in Figure 3a. Indeed, it is clear from this formula that $\log Z(\sigma) \to -\infty$ when $\sigma \to 0$ and that $\log Z(\sigma) \to +\infty$ when $\sigma \to +\infty$.

B. Toeplitz covariance matrices

Here, $\mathcal{M} = \mathcal{T}_n$, the space of $n \times n$ Toeplitz covariance matrices. Precisely, each element $T \in \mathcal{T}_n$ is an $n \times n$ Hermitian positive definite matrix, which also has Toeplitz structure. Toeplitz covariance matrices are of fundamental importance in signal processing, since they are the autocovariance matrices of wide-sense stationary random signals [54]. A Hessian metric on \mathcal{T}_n was introduced in [21][22]. This metric is expressed in terms of so-called reflection or partial autocorrelation coefficients, as given by (37a) below.

It should be noted that \mathcal{T}_n is a subspace of the space \mathcal{H}_n of complex covariance matrices, studied in IV-A. As suggested in [33], this means \mathcal{T}_n can be equipped with a Riemannian metric which is simply the restriction of the Riemannian metric (34b) from \mathcal{H}_n . Unlike the Hessian metric of [21][22], this metric does not make \mathcal{T}_n into a Riemannian symmetric space. Therefore, it cannot be used in the present paper.

▶ Remark 5 : in signal processing, reflection coefficients arise naturally in linear prediction problems [54] (Chapter 8). Equivalently, they also arise from the study of orthogonal polynomials on the unit circle [55]. Here, it is suitable to recall their definition from this analytical point of view [55] (see Chapter 1, Page 55). Recall each element $T \in \mathcal{T}_n$ has Toeplitz structure. That is, there exist $r_0 > 0$ and $r_1, \ldots, r_{n-1} \in \mathbb{C}$ such that $T_{ij} = T_{ji}^* = r_{i-j}$, (where * denotes the complex conjugate). Since T is positive definite, it can be used to define a scalar product on polynomials of degree $\leq n - 1$, which depend on a complex variable z on the unit circle, $|z|^2 = zz^* = 1$. This scalar product is given by

$$(z^{i}, z^{j}) = (z^{j}, z^{i})^{*} = r_{i-j}$$
(35a)

Let $\phi_0, \ldots, \phi_{n-1}$ be orthogonal monic polynomials with respect to this scalar product (monic means the leading degree coefficient is equal to 1). Reflection coefficients are the complex numbers $\alpha_1, \ldots, \alpha_{n-1}$ defined by the Szegö-Levinson recursion,

$$\alpha_{j} = -(1, z\phi_{j-1})/\delta_{j-1} \quad \phi_{j} = z\phi_{j-1} + \alpha_{j}\phi_{j-1} \quad \delta_{j} = \delta_{j-1}(1 - |\alpha_{j}|^{2})$$
(35b)

where $\phi_0 = 1$ and $\delta_0 = r_0$, and $\tilde{\phi}$ is the polynomial obtained from ϕ by conjugating and reversing the order of its coefficients. The significance of the numbers δ_j is that $\delta_j = (\phi_j, \phi_j)$, so δ_j must be strictly positive. From the third formula in (35b), this implies that each reflection coefficient

 α_j verifies $|\alpha_j| < 1$. Reflection coefficients allow a useful expression of the determinant of the matrix T. Indeed, since the polynomials ϕ_j are orthogonal with respect to the scalar product defined by T,

$$\det(T) = \prod_{j=0}^{n-1} (\phi_j, \phi_j) = r_0 \prod_{j=1}^{n-1} \delta_j = r_0^n \prod_{j=1}^{n-1} (1 - |\alpha_j|^2)^{n-j}$$
(35c)

This will be useful in obtaining the Hessian metric (37a) below.

- Hessian metric: consider now the Hessian metric on \mathcal{T}_n , introduced in [21][22]. This is based on the following identification,

$$\mathcal{T}_n = \mathbb{R}_+ \times \mathbb{D} \times \ldots \times \mathbb{D} \quad (n \text{ factors}) \tag{36}$$

where \mathbb{D} denotes the set of complex α which verify $|\alpha| < 1$. Precisely [56][57], \mathcal{T}_n is diffeomorphic to the product space on the right hand side, and the required diffeomorphism is given in the notation of (35) by $T \mapsto (r_0, \alpha_1, \ldots, \alpha_{n-1})$. The advantage of parameterising T by means of $(r_0, \alpha_1, \ldots, \alpha_{n-1})$ instead of $(r_0, r_1, \ldots, r_{n-1})$ is that the latter are subject to a complicated set of non-linear constraints, expressing the fact that T is positive definite, whereas the former are subject to no constraints, except $r_0 \in \mathbb{R}_+$ and $\alpha_j \in \mathbb{D}$. In the following, r_0 is denoted r, since there is no risk of confusion.

Using (36), the Hessian metric on \mathcal{T}_n can be found directly. Note first the entropy function $H(T) = -\log \det(T)$ follows from (35c),

entropy function:
$$H(T) = -n \log r - \sum_{j=1}^{n-1} (n-j) \log(1 - |\alpha_j|^2)$$

Note the following elementary calculations,

$$\frac{\partial^2 H}{\partial r^2} = n \frac{1}{r^2} \qquad \frac{\partial^2 H}{\partial \alpha_j \partial \alpha_j^*} = (n-j) \left[1 - |\alpha_j|^2\right]^{-2}$$

These immediately yield the Riemannian metric [21][22],

$$ds_{T}^{2}(dr, d\alpha) = n \left(\frac{dr}{r}\right)^{2} + \sum_{j=1}^{n-1} (n-j) ds_{\alpha_{j}}^{2}(d\alpha_{j})$$
(37a)

where
$$ds_{\alpha}^{2}(d\alpha) = |d\alpha|^{2} / [1 - |\alpha|^{2}]^{2}$$
 $\alpha \in \mathbb{D}$

Here, ds_{α}^2 is recognised as the Poincaré metric on \mathbb{D} [58] (Chapter III). Accordingly, the Riemannian distance and Riemannian volume element on \mathcal{T}_n have the following expressions,

$$d^{2}(T^{(1)}, T^{(2)}) = n \left| \log \left(r^{(2)} \right) - \log \left(r^{(1)} \right) \right|^{2} + \sum_{j=1}^{n-1} (n-j) d_{\mathbb{D}}^{2}(\alpha_{j}^{(1)}, \alpha_{j}^{(2)})$$
(37b)
where $d_{\mathbb{D}}(\alpha, \beta) = \operatorname{atanh} \left| \frac{\alpha - \beta}{1 - \alpha^{*}\beta} \right| \quad \alpha, \beta \in \mathbb{D}$
 $dv(T) = \sqrt{n} \prod_{j=1}^{n-1} (n-j) \times \frac{dr}{r} \prod_{j=1}^{n-1} dv(\alpha_{j})$
where $dv(\alpha) = \frac{\operatorname{Re} d\alpha \operatorname{Im} d\alpha}{[1 - |\alpha|^{2}]^{2}} \quad \alpha \in \mathbb{D}$

Here, $d_{\mathbb{D}}^2(\alpha,\beta)$ and $dv(\alpha)$ denote the Riemannian distance and Riemannian volume element on \mathbb{D} , corresponding to the Poincaré metric ds_{α}^2 appearing in (37a) [58].

- Symmetric space: to see that the Riemannian metric (37a) makes T_n into a Riemannian symmetric space of non-positive curvature, it is enough to notice that decomposition (36) is of the form (20), as in the proof of Proposition 1 in III-A. Indeed, in this decomposition, \mathbb{R}_+ appears as a one-dimensional Euclidean space, by identification with $\mathbb R$ through the logarithm function, while each copy of $\mathbb D$ appears as a scaled version of the Poincaré disc, with the scale factor $(n-j)^{1/2}$, and therefore as a Riemannian symmetric space of non-compact type [58] (Chapter III), (the Poincaré disc is the set \mathbb{D} with the Poincaré metric ds_{α}^2 , and is a space of constant negative curvature). To apply Facts 1–5 to \mathcal{T}_n , consider the group $G = \mathbb{R}_+ \times SU(1,1) \times \ldots \times SU(1,1)$, where SU(1,1) is the group of 2 × 2 complex

matrices u such that

$$u = \begin{pmatrix} a & b \\ c & d \end{pmatrix}; \quad u^{T} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} u = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} , \quad u^{H} \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} u = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$$
(37d)

with $a, b, c, d \in \mathbb{C}$ and T denotes the transpose. This group acts transitively and isometrically on \mathcal{T}_n in the following way,

$$g \cdot T = (s \, r, u_1 \cdot \alpha_1, \dots, u_{n-1} \cdot \alpha_{n-1}) \quad , \quad g = (s, u_1, \dots, u_{n-1}) \tag{37e}$$

where
$$u \cdot \alpha = rac{a \, \alpha + b}{c \, \alpha + d}$$
 $u \in SU(1, 1), \, \alpha \in \mathbb{D}$

Then, Facts 1 and 2 state that definitions (37a)-(37c) verify the general identities (8) and (9). Furthermore, Fact 3 states the existence and uniqueness of Riemannian barycentres in \mathcal{T}_n .

The application of Facts 4 and 5 to T_n leads to the two following formulae, which are proved in Example 2 of Appendix I,

$$\int_{\mathcal{T}_n} f(T) \, dv(T) = \mathbf{C} \times \int_{-\infty}^{+\infty} \int_{\mathbb{D}} \dots \int_{\mathbb{D}} f(t, \alpha_1, \dots, \alpha_{n-1}) \, dv(\alpha_1) \dots dv(\alpha_{n-1}) \, dt$$

$$\int_{\mathbb{D}} f(\alpha) \, dv(\alpha) = \mathbf{C} \times \int_{0}^{2\pi} \int_{-\infty}^{+\infty} f(\rho, \theta) \, \sinh(|\rho|) d\rho \, d\theta$$
(37f)



(a) Graph of $\log Z : \mathbb{R}_+ \to \mathbb{R}$

Fig. 2: $\log Z$ and Φ for 20×20 Toeplitz covariance matrices

$$d^{2}(I,T) = n t^{2} + \sum_{n=1}^{n-1} (n-j)\rho_{j}^{2}$$
 $(I = n \times n \text{ identity matrix})$ (37g)

For \mathcal{T}_n , formulae (37f) and (37g) play the same role as formulae (13) and (15), for a general symmetric space \mathcal{M} . Here $r = e^t$ and $\alpha_j =$ $\tanh(\rho_i) e^{i\theta_j}$ for $j = 1, \dots, n-1$, where $i = \sqrt{-1}$.

- Gaussian distributions: using formulae (37f) and (37g), it is possible to find the exact analytic expression of the function $\log Z : \mathbb{R}_+ \to \mathbb{R}$ for Gaussian distributions on \mathcal{T}_n . From this expression, which is given in (37i) below, the function $\Phi: \mathbb{R}_+ \to \mathbb{R}_+$ is found in the same way discussed in IV-A.

Unlike Gaussian distributions on \mathcal{H}_n , for which the function $\log Z$ cannot be found analytically, but only using the Monte Carlo technique of [35], Gaussian distributions on \mathcal{T}_n admit an exact analytic expression of $\log Z$ for any value of n. This means there is no computational limitation on the size (given by n) of Toeplitz covariance matrices which can be modelled using Gaussian distributions.

Consider the calculation of $Z(\sigma)$ using (37f) and (37g). Recall from (18) in the proof of Proposition 1,

$$Z(\sigma) = \int_{\mathcal{T}_n} \exp\left[-\frac{d^2(I,T)}{2\sigma^2}\right] dv(T)$$

Upon replacing (37f) and (37g), this becomes,

$$Z(\sigma) = \mathcal{C} \times \int_{-\infty}^{+\infty} \exp\left[-\frac{nt^2}{2\sigma^2}\right] dt \times \prod_{j=1}^{n-1} \int_0^{2\pi} \int_{-\infty}^{+\infty} \exp\left[-\frac{(n-j)\rho_j^2}{2\sigma^2}\right] \sinh(|\rho_j|) d\rho_j \, d\theta_j$$

by computing each integral in the product, it follows that

$$Z(\sigma) = C \times \sqrt{2\pi/n} \sigma \times \prod_{j=1}^{n-1} Z_{\mathbb{D}}\left(\sigma / \sqrt{n-j}\right)$$

$$Z_{\mathbb{D}}(\sigma) = (2\pi)^{2/3} \sigma e^{\frac{\sigma^2}{2}} \times \operatorname{erf}\left(\frac{\sigma}{\sqrt{2}}\right)$$
(37h)

where erf denotes the error function [59]. Finally, this yields the analytic expression of $\log Z$,

$$\log Z(\sigma) = \mathcal{C} + \log \left(\sigma / \sqrt{n} \right) + \sum_{j=1}^{n-1} \log Z_{\mathbb{D}} \left(\sigma / \sqrt{n-j} \right)$$
(37i)

Using (37h) and (37i), Figure 2 shows the graphs of $\log Z$ and Φ when n = 20. These graphs ignore the constant C appearing in (37i).

C. Block-Toeplitz covariance matrices

Here, $\mathcal{M} = \mathcal{T}_n^N$, the space of $n \times n$ block-Toeplitz covariance matrices which have $N \times N$ Toeplitz blocks. Precisely, each element $T \in \mathcal{T}_n^N$ is an $nN \times nN$ Hermitian positive definite matrix, which has a block structure of the following form

$$T = \begin{pmatrix} T_0 & T_1^* & \dots & T_{n-1}^* \\ T_1 & T_0 & T_1^* & \dots & T_{n-2}^* \\ & T_1 & T_0 & & & \\ \vdots & \vdots & & & \vdots \\ T_{n-1} & T_{n-2} & & \dots & T_0 \end{pmatrix} \qquad T_0, \dots, T_{n-1} \in \mathcal{T}_N$$

Block-Toeplitz covariance matrices which have Toeplitz blocks play a fundamental role in signal and image processing, where they arise in multichannel and two-dimensional linear prediction and filtering problems [60][61]. A Hessian metric on \mathcal{T}_n^N was considered in [23], see (40a) below. This generalises the Hessian metric (37a) on the space \mathcal{T}_n of Toeplitz covariance matrices, considered in the previous Paragraph IV-B. Precisely, in the same way as the metric (37a) was expressed in terms of scalar reflection coefficients $\alpha_1, \ldots, \alpha_{n-1}$ which belong to the complex unit disc \mathbb{D} , the metric (40a) will be expressed in terms of matrix reflection coefficients $\Omega_1, \ldots, \Omega_{n-1}$ which belong to the matrix unit disc \mathbb{D}_N : the set of $N \times N$ symmetric complex matrices Ω which verify $I - \Omega^* \Omega \succ 0$, where I denotes the $N \times N$ identity matrix and \succ the Loewner order, $(X \succ Y \text{ means that } X - Y \text{ is positive definite})$ [48].

▶ Remark 6 : in signal and image processing, matrix reflection coefficients arise naturally in multi-channel and two-dimensional linear prediction problems [60][61]. Here, they are introduced through the formalism of orthogonal matrix polynomials on the unit circle [62][63] (compare to Remark 5, Paragraph IV-B). A matrix polynomial P on the unit circle is an expression $P(z) = P_0 + ... + P_m z^m$, where z is a complex variable on the unit circle, and $P_0, ..., P_m$ are complex $N \times N$ matrices (here, $m \le n - 1$). Following [62] (Section 3, Page 43), consider for a fixed $T \in \mathcal{T}_n^N$ the so-called left and right matrix scalar products of matrix polynomials. These are given by,

$$\langle P z^i, Q z^j \rangle_L = P T_{i-j} Q^H \qquad \langle P z^i, Q z^j \rangle_R = P^H T_{i-j} Q$$
(38a)

for any complex $N \times N$ matrices P and Q, where $T_{i-j} = T_{j-i}^*$. Let $\varphi_0^L, \ldots, \varphi_{n-1}^L$ and $\varphi_0^R, \ldots, \varphi_{n-1}^R$ be orthonormal matrix polynomials with respect to these matrix scalar products. Matrix reflection coefficients $\Omega_1, \ldots, \Omega_{n-1}$ are given by the generalised Szegö-Levinson recursion,

$$z\varphi_{j-1}^{L} - \rho_{j}\varphi_{j}^{L} = \widehat{\varphi_{j-1}^{R}}\Omega_{j} \qquad z\varphi_{j-1}^{R} - \varphi_{j}^{R}\rho_{j}^{*} = \widehat{\Omega_{j}}\varphi_{j-1}^{L}$$
(38b)

where $\rho_j = (I - \Omega_j^* \Omega_j)^{1/2}$, and where \tilde{P} , for a matrix polynomial P, denotes the polynomial obtained by taking the conjugate transposes of the coefficients of P and reversing their order. It follows from [63] (Definition 2.3., Page 1615), that $\Omega_1, \ldots, \Omega_{n-1} \in \mathbb{D}_N$. Moreover, $\Omega_1, \ldots, \Omega_{n-1}$ allow a useful expression of the determinant of the matrix T [23] (Page 10),

$$\det(T) = \det(T_0)^n \prod_{j=1}^{n-1} \det(I - \Omega_j^* \Omega_j)^{n-j}$$
(38c)

This will be useful in obtaining the Hessian metric (40a) below.

- Hessian metric: consider now the Hessian metric on \mathcal{T}_n^N , introduced in [23]. This is based on the following identification,

$$\mathcal{T}_n^N = \mathcal{T}_N \times \mathbb{D}_N \times \ldots \times \mathbb{D}_N \quad (n \text{ factors})$$
(39)

where, as already stated, \mathbb{D}_N denotes the set of $N \times N$ symmetric complex matrices Ω which verify $I - \Omega^* \Omega \succ 0$. Precisely [23], \mathcal{T}_n^N is diffeomorphic to the product space on the right hand side, and the required diffeomorphism is given by $T \mapsto (T_0, \Omega_1, \ldots, \Omega_{n-1})$, in the notation of (38). In the following, T_0 is denoted P, in order to avoid the unnecessary subscript.

Using (39), the Hessian metric on \mathcal{T}_n can be found directly [23]. Note first the entropy function $H(T) = -\log \det(T)$ follows from (38c),

entropy function:
$$H(T) = -n \log \det(P) - \sum_{j=1}^{n-1} (n-j) \log \det(I - \Omega_j^* \Omega_j)$$

Then [23] (Page 10), applying the matrix differentiation formulae (34a) yields the required Hessian metric

$$ds_{T}^{2}(dP, d\Omega) = n \, ds_{P}^{2}(dP) + \sum_{j=1}^{n-1} (n-j) \, ds_{\Omega_{j}}^{2}(d\Omega_{j})$$
(40a)
where
$$ds_{\Omega}^{2}(d\Omega) = \operatorname{tr}\left[(I - \Omega\Omega^{*})^{-1} \, d\Omega \, (I - \Omega^{*}\Omega)^{-1} \, d\Omega^{*} \right] \quad \Omega \in \mathbb{D}_{N}$$

Here, ds_P^2 is the Riemannian metric (37a) of the space \mathcal{T}_N , evaluated at $P \in \mathcal{T}_N$. On the other hand, ds_Ω^2 is recognised as the Siegel metric on \mathbb{D}_N [64]. Accordingly, the Riemannian distance and Riemannian volume element on \mathcal{T}_n have the following expressions,

$$d^{2}(T^{(1)}, T^{(2)}) = n d^{2}_{\mathcal{T}_{N}}(P^{(1)}, P^{(2)}) + \sum_{j=1}^{n-1} (n-j) d^{2}_{\mathbb{D}_{N}}(\Omega^{(1)}_{j}, \Omega^{(2)}_{j})$$
(40b)
where $d^{2}_{\mathbb{D}_{N}}(\Phi, \Psi) = \text{tr} \left[\operatorname{atanh}^{2} \left(R^{1/2} \right) \right] \qquad R = (\Psi - \Phi)(I - \Phi^{*}\Psi)^{-1}(\Psi^{*} - \Phi^{*})(I - \Phi\Psi^{*})^{-1} \quad \Phi, \Psi \in \mathbb{D}_{N}$

Here, $d_{\mathcal{T}_N}^2(P^{(1)}, P^{(2)})$ is the Riemannian distance (37b) between $P^{(1)}, P^{(2)} \in \mathcal{T}_N$.

$$dv(T) = \sqrt{n} \prod_{j=1}^{n-1} (n-j) \times dv_{\mathcal{T}_N}(P) \prod_{j=1}^{n-1} dv(\Omega_j)$$
(40c)
where $dv(\Omega) = \frac{\prod_{i \le j} \operatorname{Re} d\Omega_{ij} \operatorname{Im} d\Omega_{ij}}{\det (I - \Omega\Omega^*)^{N+1}} \quad \Omega \in \mathbb{D}_N$

Here, $dv_{\mathcal{T}_N}(P)$ is the Riemannian volume element (37c) computed at $P \in \mathcal{T}_N$. In expressions (40b) and (40c), $d^2_{\mathbb{D}_N}(\Phi, \Psi)$ and $dv(\Omega)$ denote the Riemannian distance and Riemannian volume element on \mathbb{D}_N , corresponding to the Siegel metric ds^2_{Ω} appearing in (40a). These may be found in [23] (Page 10) and [65] (Chapter IV, Page 84), respectively.

- Symmetric space: to see that the Riemannian metric (40a) makes \mathcal{T}_n^N into a Riemannian symmetric space of non-positive curvature, it is enough to notice that decomposition (39) identifies \mathcal{T}_n^N with a product of Riemannian symmetric spaces of non-positive curvature. The first factor in this decomposition is the space \mathcal{T}_N of Toeplitz covariance matrices. This was seen to be a Riemannian symmetric space of non-positive curvature in the previous Paragraph IV-B. Each of the remaining factors is a copy of \mathbb{D}_N , equipped with a scaled version of the Siegel metric ds_{Ω}^2 , as in (40a). With this metric, \mathbb{D}_N is a Riemannian symmetric space of non-compact type, known as the Siegel disc [36] (Exercice D.1., Page 527).

as in (40a). With this metric, \mathbb{D}_N is a Riemannian symmetric space of non-compact type, known as the Siegel disc [36] (Exercice D.1., Page 527). To apply Facts 1–5 to \mathcal{T}_n^N , consider the group $G = G_{\mathcal{T}_N} \times Sp(n, \mathbb{R}) \times \ldots \times Sp(n, \mathbb{R})$, where $G_{\mathcal{T}_N} = \mathbb{R}_+ \times SU(1, 1) \times \ldots \times SU(1, 1)$, as in (37d)–(37e) of Paragraph IV-B, and where $Sp(n, \mathbb{R})$ is the group of $2N \times 2N$ complex matrices U such that [64] (Page 10),

$$U = \begin{pmatrix} A & B \\ C & D \end{pmatrix}; \quad U^T \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix} U = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix} , \quad U^H \begin{pmatrix} -I & 0 \\ 0 & I \end{pmatrix} U = \begin{pmatrix} -I & 0 \\ 0 & I \end{pmatrix}$$
(40d)

with A, B, C, D complex $N \times N$ matrices. This group G acts transitively and isometrically on \mathcal{T}_n^N in the following way,

$$g \cdot T = (S \cdot P, U_1 \cdot \Omega_1, \dots, U_{n-1} \cdot \Omega_{n-1}) \quad , \quad g = (S, U_1, \dots, U_{n-1})$$
(40e)
where $U \cdot \Omega = (A \Omega + B)(C \Omega + D)^{-1} \quad U \in Sp(n, \mathbb{R}) , \ \Omega \in \mathbb{D}_N$

where $S \cdot P$ is given by (37e) for $S \in G_{\mathcal{T}_N}$ and $P \in \mathcal{T}_N$. Then, Facts 1 and 2 state that definitions (40a)–(40c) verify the general identities (8) and (9). Furthermore, Fact 3 states the existence and uniqueness of Riemannian barycentres in \mathcal{T}_n^N .

The application of Facts 4 and 5 to \mathcal{T}_n^N leads to the two following formulae, which are proved in Example 3 of Appendix I,

$$\int_{\mathcal{T}_{n}^{N}} f(T) \, dv(T) = \mathcal{C} \times \int_{\mathcal{T}_{N}} \int_{\mathbb{D}_{N}} \dots \int_{\mathbb{D}_{N}} f(P, \Omega_{1}, \dots, \Omega_{n-1}) \, dv(\Omega_{1}) \dots dv(\Omega_{n-1}) \, dv_{\mathcal{T}_{N}}(P) \tag{40f}$$

$$\int_{\mathbb{D}_{N}} f(\Omega) \, dv(\Omega) = \mathcal{C} \times \int_{SU(N)} \int_{\mathbb{R}^{N}} f(r, \Theta) \prod_{i < j} \sinh(|r_{i} - r_{j}|) \prod_{i \le j} \sinh(|r_{i} + r_{j}|) \, dr \, d\Theta$$

$$d^{2}(\mathcal{I}, T) = n \, d^{2}_{\mathcal{T}_{N}}(I, P) + \sum_{j=1}^{n-1} (n-j) \, d^{2}_{\mathbb{D}_{N}}(I, \Omega_{j}) \quad (\mathcal{I} = nN \times nN \text{ identity matrix}) \tag{40g}$$

$$d^{2}_{\mathbb{D}_{N}}(I, \Omega) = ||r||^{2} = \sum_{j=1}^{n} r_{j}^{2}$$

For \mathcal{T}_n^N , formulae (40f) and (40g) play the same role as formulae (13) and (15), for a general symmetric space \mathcal{M} . Here, $dv_{\mathcal{T}}(P)$ and $d^2_{\mathcal{T}_N}(I, P)$ are given by (37f) and (37g). Moreover, $\Omega \in \mathbb{D}_N$ is written in the form,

$$\Omega = \Theta \tanh(r) \Theta^T \qquad \tanh(r) = \operatorname{diag} \left(\tanh(r_1), \dots, \tanh(r_n) \right) \tag{40h}$$

where $\Theta \in SU(N)$, the group of unitary $N \times N$ matrices which have unit determinant, and $r \in \mathbb{R}^N$ (this is called the Takagi factorisation in [66]). In (40f), dr denotes the Lebesgue measure on \mathbb{R}^N and $d\Theta$ denotes the normalised Haar measure on SU(N).

- Gaussian distributions: using formulae (40f) and (40g), it is possible to construct the functions $\log Z : \mathbb{R}_+ \to \mathbb{R}$ and $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ for Gaussian distributions on \mathcal{T}_n^N . Computation of $\log Z$, based on these two formulae, leads to a similar structure to the one found in (37i) of IV-B, for the case of Gaussian distributions on \mathcal{T}_n .

To see this, recall according to (18) in the proof of Proposition 1,

$$Z(\sigma) = \int_{\mathcal{T}_n^N} \exp\left[-\frac{d^2(\mathcal{I},T)}{2\sigma^2}\right] dv(T)$$

Upon replacing (40f) and (40g), this becomes,

$$Z(\sigma) = \mathcal{C} \times \int_{\mathcal{T}_N} \exp\left[-\frac{n d_{\mathcal{T}_N}^2(I, P)}{2\sigma^2}\right] dv_{\mathcal{T}_N}(P) \times \prod_{j=1}^{n-1} \int_{\mathbb{D}_N} \exp\left[-\frac{(n-j) d_{\mathbb{D}_N}^2(I, \Omega_j)}{2\sigma^2}\right] dv(\Omega_j)$$

This formula states that $Z(\sigma)$ is the product of n integrals. The first of these integrals, which extends over \mathcal{T}_N , is of the form $Z_{\mathcal{T}_N}(\sigma/\sqrt{n})$, where $Z_{\mathcal{T}_N}(\sigma)$ is the function given by (37h) of IV-B. On the other hand, each one of the remaining n-1 integrals, which extend over \mathbb{D}_N , is of the form $Z_{\mathbb{D}_N}(\sigma/\sqrt{n-j})$, where $Z_{\mathbb{D}_N}(\sigma)$ is given by

$$Z_{\mathbb{D}_N}(\sigma) = \int_{\mathbb{D}_N} \exp\left[-\frac{d_{\mathbb{D}_N}^2(I,\Omega)}{2\sigma^2}\right] dv(\Omega)$$
(40i)

Using these definitions of $Z_{\mathcal{T}_N}(\sigma)$ and $Z_{\mathbb{D}_N}(\sigma)$, it follows that,

$$Z(\sigma) = C \times Z_{\mathcal{T}_N} \left(\sigma / \sqrt{n} \right) \times \prod_{j=1}^{n-1} Z_{\mathbb{D}_N} \left(\sigma / \sqrt{n-j} \right)$$
(40j)

Then, by taking logarithms, the following expression of $\log Z$ is found,

$$\log Z(\sigma) = C + \log Z_{\mathcal{T}_N}\left(\sigma/\sqrt{n}\right) + \sum_{j=1}^{n-1} \log Z_{\mathbb{D}_N}\left(\sigma/\sqrt{n-j}\right)$$
(40k)

Now, this expression indeed has a similar structure to the one in (37i) of IV-B. Formally, it is obtained from (37i) after two substitutions: $\log(\sigma) \rightsquigarrow \log Z_{\mathcal{T}_N}(\sigma)$ and $\log Z_{\mathbb{D}}(\sigma) \rightsquigarrow \log Z_{\mathbb{D}_N}(\sigma)$. Referring to decompositions (36) and (39), it is seen that, on a geometric level, these correspond to replacing the Euclidean space \mathbb{R}_+ by the space \mathcal{T}_N , and replacing the Poincaré disc \mathbb{D} by the Siegel disc \mathbb{D}_N .

The function $\log Z$ lends itself to numerical evaluation from expression (40k). In this expression, the function $\log Z_{\mathcal{T}_N}$ can be computed analytically, using (37h) and (37i) of IV-B. On the other hand, the function $\log Z_{D_N}$ can be evaluated from (40i) using (40f) and (40g). Precisely, by replacing the second line of each one of these formulae in (40i), it follows that,

$$Z_{\mathbb{D}_N}(\sigma) = \mathcal{C} \times \int_{SU(N)} \int_{\mathbb{R}^N} \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i < j} \sinh(|r_i - r_j|) \prod_{i \le j} \sinh(|r_i + r_j|) \, dr \, d\Theta$$

and since the function under the integral does not depend on Θ ,

$$Z_{\mathbb{D}_N}(\sigma) = \mathcal{C} \times \int_{\mathbb{R}^N} \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i < j} \sinh(|r_i - r_j|) \prod_{i \le j} \sinh(|r_i + r_j|) dr$$
(401)

Using this result, the function $\log Z_{\mathbb{D}_N}$ can be evaluated by means of the Monte Carlo technique of [35]. Then, complete evaluation of $\log Z$ is achieved by substitution in (40k).

In addition to being useful in the evaluation of $\log Z$, the function $Z_{\mathbb{D}_N}$ is interesting in its own right, since it gives the normalising factor for Gaussian distributions on the Siegel disc \mathbb{D}_N . Figure 3 presents the graph of the function $\log Z_{\mathbb{D}_N}$ and of the corresponding function $\Phi_{\mathbb{D}_N}$, (defined as in (27) from Proposition 3 of III-B), for N = 20.

▶ Remark 7 : the Siegel disc \mathbb{D}_N is isometric to the Siegel half-plane \mathbb{H}_N , the set of $N \times N$ symmetric complex matrices Z such that Im $Z \succ 0$ [64][67]. The Siegel half-plane \mathbb{H}_N , (in addition to its usefulness in analytical mechanics and number theory), provides a geometric setting which leads to deep results with regard to linear filtering and linear control theory [68]–[70]. In particular, both the discrete and the continuous forms of the Riccati equation are equivalent to invariant dynamical systems on \mathbb{H}_N , which have remarkable geometric properties. This provides a geometric means of studying the asymptotic behavior of the Riccati equation, and therefore of linear filters and control systems.



Fig. 3: $\log Z_{\mathbb{D}_N}$ and $\Phi_{\mathbb{D}_N}$ for the Siegel disc \mathbb{D}_N with N = 20

V. DENSITY ESTIMATION WITH FINITE MIXTURES OF GAUSSIAN DENSITIES

The present section introduces finite mixtures of Gaussian densities, as a tool for dealing with the problem of density estimation in a Riemannian symmetric space of non-positive curvature, \mathcal{M} . A finite mixture of Gaussian densities on \mathcal{M} is a probability density of the form

$$p(x) = \sum_{\kappa=1}^{K} \omega_{\kappa} \times p(x|\bar{x}_{\kappa}, \sigma_{\kappa})$$
(41)

where the weights $\omega_1, \ldots, \omega_K > 0$ are such that $\omega_1 + \ldots + \omega_K = 1$, and where each density $p(x | \bar{x}_{\kappa}, \sigma_{\kappa})$ is a Gaussian density given by (16) of Section III. Definition (41) is directly based on the standard definition of a finite mixture of probability densities [71][72].

The problem of density estimation in \mathcal{M} is the following [73][74]. Assume a population of data points $x_1, \ldots, x_N \in \mathcal{M}$ is under study. Assume further this population was generated from some unknown probability density q(x). The problem is to approximate this unknown density using the data x_1, \ldots, x_N .

Concretely, this problem arises in the following manner [18]–[20] [26]–[28]. As in Section IV, the space \mathcal{M} is a space of structured covariance matrices. For example, $\mathcal{M} = \mathcal{H}_n$, $\mathcal{M} = \mathcal{T}_n$, or $\mathcal{M} = \mathcal{T}_n^N$, (recall these notations from Table 1. in the introduction). A database of signals or images is under study, and each object in this database is represented by a covariance descriptor. This descriptor is taken to be a structured covariance matrix $x \in \mathcal{M}$. The database of signals or images is thus mapped into a population $x_1, \ldots, x_N \in \mathcal{M}$. In a real-world setting, it is hopeless to attempt any simplifying hypothesis with regard to the probability density q(x) of this population. Rather, one attempts to approximate q(x) within a sufficiently rich class of candidate densities p(x).

The use of mixtures of Gaussian densities as a tool for solving this problem is founded on the idea that the probability density q(x), at least if sufficiently regular, can be approximated to any precision by a density p(x) of the form (41). Accordingly, consider the mathematical problem of finding a density p(x) of the form (41) which best approximates the unknown density q(x), in the sens of Kullback-Leibler divergence. Recall this is given by [25],

$$\hat{p} = \operatorname{argmin}_{p} D(q \| p) \qquad D(q \| p) = \int_{\mathcal{M}} q(x) \log \left[\frac{q(x)}{p(x)}\right] dv(x)$$
(42)

where dv is the Riemannian volume element of \mathcal{M} . Here, the minimum is over all densities p(x) of the form (41). Following the general idea of "empirical risk minimisation" [73] (Page 32), the above integral with respect to the unknown density q(x) is replaced by an empirical average. The minimisation problem (42) is then replaced by

$$\hat{p}_N = \operatorname{argmin}_p \frac{1}{N} \sum_{n=1}^N \log \left[\frac{q(x_n)}{p(x_n)} \right] = \operatorname{argmin}_p - \frac{1}{N} \sum_{n=1}^N \log p(x_n)$$
(43)

This new minimisation problem is equivalent to the problem of maximum likelihood estimation of the parameters $\{(\omega_{\kappa}, \bar{x}_{\kappa}, \sigma_{\kappa}); \kappa = 1, ..., K\}$ of the finite mixture of Gaussian densities p(x), under the assumption that the data $x_1, ..., x_N$ are independently generated from p(x).

In [15], an original EM (expectation maximisation) algorithm was proposed for solving this problem of maximum likelihood estimation. This algorithm was derived in the special case where $\mathcal{M} = \mathcal{P}_n$, the space of $n \times n$ covariance matrices. It was successfully applied to real data, mainly from remote sensing, with the aim of carrying out density estimation as well as classification [18]–[20]. This algorithm generalises directly from $\mathcal{M} = \mathcal{P}_n$ to any \mathcal{M} which is a Riemannian symmetric space of non-positive curvature, and its derivation remains formally the same as in [15] (Paragraph IV.A., Page 13).

In the present general setting, it can be described as follows. The algorithm assumes that K, the number of mixture components in (41), has already been selected. It iteratively updates an approximation $\hat{\theta} = \{(\hat{\omega}_{\kappa}, \hat{x}_{\kappa}, \hat{\sigma}_{\kappa})\}$ of the maximum likelihood estimates of the parameters $\{(\omega_{\kappa}, \bar{x}_{\kappa}, \sigma_{\kappa})\}$. The required update rules (45a)–(45c), as given below, are repeated as long as this leads to a sensible increase in the joint likelihood function of $(\hat{\omega}_{\kappa}, \hat{x}_{\kappa}, \hat{\sigma}_{\kappa})$. These update rules involve the two quantities $\pi_{\kappa}(x_n, \hat{\theta})$ and $N_{\kappa}(\hat{\theta})$, defined as follows,

$$\pi_{\kappa}(x_n,\hat{\theta}) \propto \hat{\omega}_{\kappa} \times p(x_n | \hat{x}_{\kappa}, \hat{\sigma}_{\kappa}) \qquad N_{\kappa}(\hat{\theta}) = \sum_{n=1}^{N} \pi_{\kappa}(x_n, \hat{\theta})$$
(44)

where the constant of proportionality, corresponding to " \propto " in the first expression, is chosen so $\sum_{\kappa} \pi_{\kappa}(x_n, \hat{\theta}) = 1$. Each iteration of the EM algorithm consists in an application of the updatte rules,

EM update rules for mixture estimation:

$$\hat{\omega}_{\kappa}^{\text{new}} = N_{\kappa}(\hat{\theta}) / N \tag{45a}$$

$$\hat{x}_{\kappa}^{\text{new}} = \operatorname{argmin}_{x \in \mathcal{M}} E_{\kappa}(x) \qquad E_{\kappa}(x) = N_{\kappa}^{-1}(\hat{\theta}) \sum_{n=1}^{N} \pi_{\kappa}(x_n, \hat{\theta}) d^2(x, x_n)$$
(45b)

$$\hat{\sigma}_{\kappa}^{\text{new}} = \Phi(E_{\kappa}(\hat{x}_{\kappa}^{\text{new}})) \quad \Phi \text{ given by (27) of Proposition 3}$$
 (45c)

Realisation of update rules (45a) and (45c) is straightforward. On the other hand, realisation of update rule (45b) requires minimisation of the function $E_{\kappa}(x)$ over $x \in \mathcal{M}$. It is clear this function is a variance function of the form (10), corresponding to the probability distribution $\pi \propto \sum_{n=1}^{N} \pi_{\kappa}(x_n, \hat{\theta}) \times \delta_{x_n}$ where δ_x denotes the Dirac distribution concentrated at $x \in \mathcal{M}$. Therefore, existence and uniqueness of $\hat{x}_{\kappa}^{\text{new}}$ follow from Fact 3, and computation of $\hat{x}_{\kappa}^{\text{new}}$ can be carried out using existing routines for the computation of Riemannian barycentres [40][42][44][45].

The EM algorithm just described requires prior selection of the number K of mixture components. When this EM algorithm was applied in [18]–[20], (for the special case $\mathcal{M} = \mathcal{P}_n$), K was selected by maximising a BIC criterion. In general, this criterion is given following [75],

BIC criterion:
$$K^* = \operatorname{argmax}_K BIC(K) \quad BIC(K) = \ell(K) - \frac{1}{2}DF \times \ln(N)$$
 (46)

Here, $\ell(K)$ is the maximum value of the joint likelihood function of $(\omega_{\kappa}, \bar{x}_{\kappa}, \sigma_{\kappa})$ for given K, and DF is the number of parameters to be estimated, or number of degrees of freedome. This is $DF = K \times [2 + \dim(\mathcal{M})] - 1$. For the three spaces of structured covariance matrices considered in Section IV, the dimension of \mathcal{M} is given by,

$$\dim(\mathcal{H}_n) = n^2 \quad \dim(\mathcal{T}_n) = 2n - 1 \quad \dim(\mathcal{T}_n^N) = 2N - 1 + (n - 1) \times (N^2 + N) \tag{47}$$

The EM algorithm given by update rules (45a)–(45c), and the BIC criterion (46), together provide a complete programme for solving the problem of density estimation in \mathcal{M} , using finite mixtures of Gaussian densities. The output of this programme is an approximation of the finite mixture of Gaussian densities \hat{p}_N defined by (43). Then, with a slight abuse of notation, it is possible to write,

$$\hat{p}_N(x) = \sum_{\kappa=1}^K \hat{\omega}_\kappa \times p(x|\hat{x}_\kappa, \hat{\sigma}_\kappa)$$
(48)

In [15][18]–[20], an additional problem of classification was addressed, whose aim is to use (48) as an explicative model of the data x_1, \ldots, x_N , revealing the fact that these data arise from K disjoint classes, where the κ th class is generated from a single Gaussian distribution $G(\hat{x}_{\kappa}, \hat{\sigma}_{\kappa})$. In [15] (Paragraph IV.B., Page 15), this problem was solved in the special case $\mathcal{M} = \mathcal{P}_n$. In general, the proposed solution remains valid as long as \mathcal{M} is a Riemannian symmetric space of non-positive curvature. It is based on a Bayes optimal classification rule. This is a function $\kappa^* : \mathcal{M} \to \{1, \ldots, K\}$, obtained by minimising the *a posteriori* probability of classification error [74] (Section 2.4., Page 21),

$$\kappa^*(x) = \operatorname{argmin}_{\kappa=1,\dots,K} \left[1 - P(\kappa|x) \right] \tag{49}$$

where $P(\kappa|x)$ is found from Bayes rule, $P(\kappa|x) \propto p(x|\kappa) \times P(\kappa)$, after replacing $p(x|\kappa) = p(x|\hat{x}_{\kappa}, \hat{\sigma}_{\kappa})$ and $P(\kappa) = \hat{\omega}_{\kappa}$, where $p(x|\hat{x}_{\kappa}, \hat{\sigma}_{\kappa})$ is given by (16). Carrying out the minimisation (49) readily gives the expression,

$$\kappa^*(x) = \operatorname{argmin}_{\kappa=1,\dots,K} \left[-\log \hat{\omega}_{\kappa} + \log Z(\hat{\sigma}_{\kappa}) + \frac{d^2(x, \hat{x}_{\kappa})}{2\hat{\sigma}_{\kappa}^2} \right]$$
(50)

This is a generalisation of the so-called minimum-distance-to-mean classification rule, which corresponds to all $\hat{\omega}_{\kappa}$ and all $\hat{\sigma}_{\kappa}$ being constant, (in the sense that they do not depend on κ). In [4], the minimum-distance-to-mean classification rule was shown to give very good performance in the area brain-computer interface decoding. Ongoing work, started in [35], examines the application of (50) in this same area.

Comparison to kernel density estimation

In recent literature [26]–[28], a fully non-parametric method for density estimation in \mathcal{M} was developed. This is based on the method of kernel density estimation on Riemannian manifolds, first outlined in [76]. This non-parametric method is alternative to the method of the present section, based on finite mixtures of Gaussian distributions. Currently, no practical comparison of these two methods is available, as they have not yet been tested side-by-side on real or simulated data. However, it is interesting to compare them from a theoretical point of view.

The main idea of kernel density estimation on Riemannian manifolds is the following [76]: classical performance bounds for kernel density estimation on a Euclidean space can be immediately transposed to any Riemannian manifold \mathcal{M} , provided the kernel function $K : \mathbb{R} \to \mathbb{R}_+$ is adjusted to the Riemannian volume element of \mathcal{M} .

In the present context, where \mathcal{M} is a Riemannian symmetric space of non-positive curvature, consider the case of a Gaussian kernel $K(x) = (2\pi)^{-1/2} \exp(-x^2/2)$. When this Gaussian kernel is adjusted to the Riemannian volume element of \mathcal{M} , the result is a family of Riemannian Gaussian kernels $K(x|\bar{x},\sigma)$ on \mathcal{M} , which provides the following method.

Kernel density estimation on \mathcal{M} :

$$K(x|\bar{x},\sigma) = (2\pi\sigma^2)^{-\dim(\mathcal{M})/2} \times J(x,\bar{x}) \times \exp\left[-\frac{d^2(x,\bar{x})}{2\sigma^2}\right]$$
(51a)

where the "adjusting factor" $J(x, \bar{x})$ is expressed in [27] (Page 11), The kernel density estimate $\hat{q}_N(x)$ of the unknown density q(x) of the data x_1, \ldots, x_N is then defined by [27] (Formula (7), Page 8)

$$\hat{q}_N(x) = \frac{1}{N} \sum_{n=1}^{N} K(x | x_n, \sigma)$$
(51b)

and the integrated mean square error between $\hat{q}_N(x)$ and q(x) can be made to decrease as a negative power of N, by a suitable choice of $\sigma = \sigma(N)$, under the condition that q(x) is twice differentiable [76] (Theorem 3.1., Page 302).

The EM algorithm (45a)–(45c) and the kernel density estimation method (51a)–(51b) provide two approaches to the problem of density estimation in \mathcal{M} , which differ in the following ways,

A-Cost in memory: for the EM algorithm, evaluation of the output \hat{p}_N from (48) requires bounded memory, essentially independent of N. For

kernel density estimation, evaluation of \hat{q}_N from (51b) requires storing N data points, each of dimension dim(\mathcal{M}), which can be found in (47). B-Computational cost: each iteration of the EM algorithm involves computing K Riemannian barycentres of N data points. For kernel density estimation, computing $\hat{q}_N(x)$ at a single point $x \in \mathcal{M}$ requires computing N factors $J(x, x_n)$ and N Riemannian distances $d(x, x_n)$. Concretely, Riemannian distances on \mathcal{M} are found from highly non-linear formulae, such as (34c), (37b) and (40b).

C-Rate of convergence: EM algorithms are known to be slow, and to get trapped in local minima or saddle points. However, stochastic versions of EM, such as the SEM algorithm [77], effectively remedy these problems. For kernel density estimation, (51b) is an exact formula which can be computed directly. However, it should be kept in mind, the precision obtained using this formula suffers from the curse of dimensionality.

D-Free parameters: running the EM algorithm requires fixing the mixture order K. The BIC criterion (46) proposed to this effect may be difficult to evaluate. Kernel density estimation requires choosing the kernel width $\sigma = \sigma(N)$. This also seems to be a difficult task, especially when the dimension of \mathcal{M} is high.

In conclusion, the EM algorithm (45a)–(45c) provides an approach to the problem of density estimation in \mathcal{M} , which can be expected to offer a suitable rate of convergence and which is not greedy in terms of memory. The main computational requirement of this algorithm is the ability to find Riemannian barycentres, a task for which there exists an increasing number of high-performance routines [42][44][45][11][31][32]. The fact that the EM algorithm reduces the problem of probability density estimation in \mathcal{M} to one of repeated computation of Riemannian barycentres is due to the unique connection which exists between Gaussian distributions in \mathcal{M} and the concept of Riemannian barycentre. This connection is stated in the defining Property (1) of these Gaussian distributions, and was a guiding motivation for the present paper.

► **Remark 8** : a Riemannian Gaussian kernel $K(x|\bar{x},\sigma)$ is a probability density with respect to the Riemannian volume element of \mathcal{M} . That is, it is positive and its integral with respect to this volume element is equal to one. Precisely, $K(x|\bar{x},\sigma)$ is the probability density of a random point x in \mathcal{M} , given by $x = \exp_{\bar{x}}(v)$ where v is a tangent vector to \mathcal{M} at \bar{x} with distribution $v \sim \mathcal{N}(0,\sigma)$, and where $\exp_{\bar{x}}$ is the Riemannian exponential mapping [37]. This probability density $K(x|\bar{x},\sigma)$ is not a Gaussian density and therefore is not compatible with Property (1).

ACKNOWLEDGEMENTS

This research was in part funded by the NSF grant IIS-1525431 to B.C. Vemuri.

REFERENCES

- [1] M. Moakher, "On the averaging of symmetric positive-definite tensors," Journal of Elasticity, vol. 82, no. 3, pp. 273–296, 2006.
- [2] P. J. Basser, J. Mattiello, and D. Le Bihan, "MR diffusion tensor spectroscopy and imaging," Biophys. J., vol. 66, no. 1, pp. 259–267, 1994.
- [3] X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," Int. J. Comput. Vision, vol. 66, no. 1, pp. 41-66, 2006.
- [4] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Multiclass brain-computer interface classification by Riemannian geometry," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 4, pp. 920–928, 2012.
- [5] M. Arnaudon, F. Barbaresco, and L. Yang, "Riemannian medians and means with applications to Radar signal processing," IEEE J. Sel. Topics Signal Process, vol. 7, no. 4, pp. 595–604, 2013.
- [6] L. Yang, "Medians of probability measures in Riemannian manifolds and applications to radar target detection," Ph.D. dissertation, Université de Poitiers, 2011, arnaudon, M. and Barbaresco, F. (supervisors). [Online]. Available: https://tel.archives-ouvertes.fr/tel-00664188
- [7] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi, "Kernel methods on the Riemannian manifold of symmetric positive definite matrices," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 73–80.
- [8] L. Zheng, G. Qiu, J. Huang, and J. Duan, "Fast and accurate nearest neighbor search in the manifolds of symmetric positive definite matrices," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 3804–3808.
- [9] O. Tuzel, F. Porikli, and P. Meer, "Pedestrian detection via classification on Riemannian manifolds," IEEE Trans. Pattern Anal. and Mach. Intell., vol. 30, no. 10, pp. 1713–1727, 2008.
- [10] R. Caseiro, J. F. Henriques, P. Martins, and J. Batista, "A nonparametric Riemannian framework on tensor field with application to foreground segmentation," *Pattern Recognition*, vol. 45, no. 11, pp. 3997 – 4017, 2012.
- [11] B. Jeuris, R. Vandebril, and B. Vandereycken, "A survey and comparison of contemporary algorithms for computing the matrix geometric mean," *Electron. Trans. Numer. Anal.*, vol. 39, pp. 379–402, 2012. [Online]. Available: http://etna.mcs.kent.edu/volumes/2011-2020/vol39/ abstract.php?vol=39&pages=379-402
- [12] H. Salehian, G. Cheng, B. C. Vemuri, and J. Ho, "Recursive estimation of the stein center of spd matrices and its applications," in *IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 1793–1800.
- [13] A. Cherian, "Similarity search in visual data," Ph.D. dissertation, University of Mnnesota, 2013, Papanikolopoulos, N. (supervisor). [Online]. Available: https://lear.inrialpes.fr/people/cherian/papers/
- [14] G. Cheng and B. C. Vemuri, "A novel dynamic system in the space of SPD matrices with applications to appearance tracking," SIAM J. Imaging Sci., vol. 6, no. 1, pp. 592-615, 2013.
- [15] S. Said, L. Bombrun, Y. Berthoumieu, and J. H. Manton, "Riemannian Gaussian distributions on the space of symmetric positive definite matrices (under review)," *IEEE Trans. Inf. Theory*, 2016. [Online]. Available: http://arxiv.org/abs/1507.01760
- [16] S. Stahl, "The Evolution of the Normal Distribution," Mathematics Magazine, vol. 79, pp. 96–113, 2006.
- [17] M. Banerjee, R. Chakraborty, E. Ofori, D. Vaillancourt, and B. C. Vemuri, "Nonlinear regression on Riemannian manifolds and its applications to neuro-image analysis," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 719–727.
- [18] R. Rosu, M. Donias, L. Bombrun, S. Said, O. Regniers, and J. P. Da Costa, "Structure tensor Riemannian statistical models for CBIR and classification of remote sensing images (accepted)," *IEEE Trans. Geosci. Remote Sens.*, 2016.
- [19] I. Ilea, H. Hajri, S. Said, L. Bombrun, C. Germain, and Y. Berthoumieu, "An M-estimator for robust centroid estimation on the manifold of covariance matrices: performance analysis and application to image classification," in *European Signal Processing Conference (EUSIPCO)*, 2016, pp. 2196–2200.
- [20] R. Rosu, L. Bombrun, S. Said, O. Regniers, M. Donias, and J. P. Da Costa, "Modèles probabilistes Riemanniens de tenseurs de structure pour le traitement d'images texturées," in Groupe d'Etude du Traitement du Signal et des Images (GRETSI), 2015, pp. 1–4.
- [21] F. Barbaresco, "Interactions between symmetric cone and information geometries: Bruhat-Tits and Siegel spaces models for high resolution autoregressive Doppler imagery," in *Emerging trends in Visual Computing*, ser. Lecture notes in Computer Science, F. Nielsen, Ed. Springer Berlin Heidelberg, 2009, vol. 5416, pp. 124–163.
- [22] —, "Information geometry of covariance matrix: Cartan-Siegel homogeneous bounded domains, Mostow/Berger fibration and Fréchet median," in *Matrix Information Geometry*, F. Nielsen and R. Bhatia, Eds. Springer Berlin Heidelberg, 2013, pp. 199–255.
- [23] B. Jeuris and R. Vandebril, "The Kähler mean of Block-Toeplitz matrices with Toeplitz structured blocks," SIAM J. Matrix Anal. Appl., 2016. [Online]. Available: https://lirias.kuleuven.be/handle/123456789/497758

- [24] B. Jeuris, "Riemannian Optimization for Averaging Positive Definite Matrices," Ph.D. dissertation, Department of Computer Science, KU Leuven, 2015, vandebril, R. and Nicaise, J. (supervisors). [Online]. Available: https://lirias.kuleuven.be/handle/123456789/496544
- [25] T. M. Cover and J. A. Thomas, Elements of Information Theory (Wiley series in telecommunications). Wiley-Interscience, 2006.
- [26] E. Chevallier, F. Barbaresco, and J. Angulo, "Non-parametric density estimation on Siegel space applied to radar processing (accepted)," *Entropy*, 2016. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01344910/
- [27] E. Chevallier, E. Kalunga, and J. Angulo, "Kernel density estimation on spaces of Gaussian distributions and symmetric positive definite matrices (accepted)," SIAM J. Imaging Sci., 2015. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01245712/
- [28] E. Chevallier, F. Barbaresco, and J. Angulo, "Probability density estimation on the hyperbolic space applied to radar processing," in *Geometric Science of Information (GSI)*, 2015, pp. 753–761.
- [29] B. Afsari, "Riemannian L^p center of mass: existence, uniqueness and convexity," Proc. Amer. Math. Soc., vol. 139, no. 2, pp. 655–673, 2011.
- [30] M. Moakher, "A differential geometric approach to the geometric mean of symmetric positive-definite matrices," SIAM J. Matrix Anal. and Appl., vol. 26, no. 3, pp. 735–747, 2005.
- [31] M. Congedo, B. Afsari, A. Barachant, and M. Moakher, "Approximate joint diagonalisation and geometric mean of symmetric positive definite matrices," *PLoS ONE, Public Library of Science*, vol. 10, no. 4, pp. 379–402, 2015.
- [32] M. Congedo, R. Phlypo, and A. Barachant, "A fixed-point algorithm for estimating power means of symmetric positive definite matrices," in *European Signal Processing Conference (EUSIPCO)*, 2016.
- [33] D. A. Bini, B. Iannazzo, B. Jeuris, and R. Vandebril, "Geometric means of structured matrices," BIT., vol. 54, no. 1, pp. 55-83, 2014.
- [34] J. Faraut and A. Korányi, Analysis on symmetric cones. Claredon Press, Oxford, 1994.
- [35] P. Zanini, S. Said, M. Congedo, Y. Berthoumieu, and C. Jutten, "Parameter estimates of Riemannian Gaussian distributions in the manifold of covariance matrices," in Sensor Array and Multichannel Signal Processing Workshop (SAM), 2016.
- [36] S. Helgason, Differential geometry and symmetric spaces. New York and London: Academic Press, 1962.
- [37] M. Berger, A panoramic view of Riemannian geometry. Berlin-Heidelberg: Springer-Verlag, 2002.
- [38] O. Barndorff-Nielsen, Information and exponential families in statistical theory. John Wiley & Sons, 1978.
- [39] X. Pennec, "Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements," J. Math. Imaging Vis., vol. 25, no. 1, pp. 127–154, 2006.
- [40] C. Lenglet, M. Rousson, R. Deriche, and O. Faugeras, "Statistics on the manifold of multivariate normal distributions," J. Math. Imaging Vis., vol. 25, no. 3, pp. 423–444, 2006.
- [41] D. Groisser, "Newton's method, zeroes of vector fields, and the Riemannian center of mass," Adv. in Appl. Math., vol. 33, no. 1, pp. 95 135, 2004.
- [42] R. Ferreira, J. Xavier, J. P. Costeira, and V. Barroso, "Newton algorithms for Riemannian distance related problems on connected locally symmetric manifolds," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 4, pp. 634–645, 2013.
- [43] K. T. Sturm, "Probability measures on metric spaces of nonpositive curvature," in *Heat kernels and analysis on manifolds, graphs and metric spaces*, ser. Contemporary Mathematics, P. Auscher, T. Coulhon, and A. Grigoryan, Eds. American Mathematical Society, 2002, vol. 338, pp. 357–391.
- [44] J. Ho, G. Cheng, H. Salehian, and B. C. Vemuri, "Recursive Karcher expectation estimators and geometric law of large numbers." in Artificial Intelligence and Statistics (AISTATS), ser. JMLR Workshop and Conference Proceedings, vol. 31, 2013, pp. 325–332.
- [45] S. Bonnabel, "Stochastic gradient descent on Riemannian manifolds," IEEE Trans. Autom. Control., vol. 122, no. 4, pp. 2217–2229, 2013.
- [46] M. Arnaudon and L. Miclo, "Means in complete manifolds : uniqueness and approximation," ESAIM Probability and statistics, vol. 18, pp. 185–206, 2014.
- [47] R. Bhattacharya and V. Patrangenaru, "Large sample theory of intrinsic and extrinsic sample means on manifolds. I," Ann. Stat., vol. 31, no. 1, pp. 1–29, 2003.
- [48] R. Bhatia, Positive definite matrices. Princeton University Press, 2007.
- [49] P. A. Absil, R. Mahony, and R. Sepulchre, Optimization algorithms on matrix manifolds. Princeton University Press, 2008.
- [50] N. J. Higham, Functions of matrices, Theory and computation. Society for Industrial and Applied Mathematics, 2008.
- [51] D. Nagar, A. K. Gupta, and L. E. Sanchez, "A class of integral identities with hermitian matrix argument," Proc. Amer. Math. Soc., vol. 134, no. 11, pp. 3329–3341, 2006.
- [52] Y. Wang, C. Han, and F. Tupin, "PolSAR data segmentation by combining tensor space cluster analysis and Markovian framework," *IEEE Geosci. Remote Sens. Lett.*, vol. 7, no. 1, pp. 210–214, 2010.
- [53] O. D'Hondt, S. Guillaso, and O. Hellwich, "Bilateral filtering of PolSAR data based on Riemannian metrics," in International Geoscience and Remote Sensing Symposium (IGARSS), 2012, pp. 5153–5156.
- [54] C. W. Therrien, Discrete random signals and statistical signal processing. Prentice-Hall Inc., 1992.
- [55] B. Simon, Orthogonal polynomials on the unit circle. Part 1: classical theory. American Mathematical Society, 2005.
- [56] O. Barndorff-Nielsen and G. Schou, "On the parameterization of autoregressive models by partial autororrelations," J. Multivariate Anal., vol. 3, no. 4, pp. 408–419, 1973.
- [57] F. L. Ramsey, "Characterization of the partial autocorrelation function," Ann. Statist., vol. 2, no. 6, pp. 1296–1301, 1974.
- [58] A. Terras, Harmonic analysis on symmetric spaces and applications, Vol. I. Springer-Verlag, 1985.
- [59] M. Abramowitz and I. A. Stegun, Handbook of mathematical functions with formulas, graphs, and mathematical tables. U.S. Government Printing Office, Washington, D.C., 1964.
- [60] C. Therrien, "Relations between 2-D and multichannel linear prediction," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 29, no. 3, pp. 454–456, 1981.
- [61] A. Jakobsson, S. L. Marple, and P. Stoica, "Computationally efficient two-dimensional Capon spectrum analysis," *IEEE Trans. Signal Process.*, vol. 48, no. 9, pp. 2651–2661, 2000.
- [62] D. Damanik, A. Pushnitski, and B. Simon, "The analytic theory of matrix orthogonal polynomials," Surveys in approximation theory, vol. 4, pp. 1 – 85, 2008.
- [63] H. Dette and J. Wagener, "Matrix measures on the unit circle, moment spaces, orthogonal polynomials and the geronimus relations," *Linear Algebra Appl.*, vol. 432, no. 7, pp. 1609 1626, 2010.
- [64] C. L. Siegel, "Symplectic geometry," Am. J. Math., vol. 65, no. 1, pp. 1-86, 1943.
- [65] L. K. Hua, Harmonic analysis of functions of several complex variables in the classical domains. American Mathematical Society, 1963.
- [66] R. A. Horn and C. R. Johnson, Matrix analysis (second edition). Cambridge University Press, 2012.
- [67] A. Terras, Harmonic analysis on symmetric spaces and applications, Vol. II. Springer-Verlag, 1988.
- [68] M. I. Zelikin, Control theory and optimization I, (Encyclopedia of Mathematical Sciences, Vol. 86). Springer-Verlag, 1999.
- [69] P. Bougerol, "Filtre de kalman-Bucy et exposants de Lyapounov," in Lyapunov Exponents, ser. Lecture Notes in Mathematics 1186, L. Arnold, H. Crauel, and J. P. Eckmann, Eds., 1991, pp. 112–122.
- [70] —, "Kalman filtering with random coefficients and contractions," SIAM J. Control Optim., vol. 31, no. 4, pp. 942–959, 1993.
- [71] D. M. Titterington, A. F. M. Smith, and U. E. Makov, Statistical analysis of finite mixture distributions. John Wiley & Sons Ltd., 1985.
- [72] S. Fruhwirth-Schnatter, Finite mixture and Markov switching models. Springer, 2006.
- [73] V. N. Vapnik, Statistical learning theory. John Wiley and Sons, 1998.

- [74] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction.* Springer, 2009.
 [75] B. G. Leroux, "Consistent estimation of a mixing distribution," *Ann. Statist.*, vol. 20, no. 3, pp. 1350–1360, 1992.
 [76] B. Pelletier, "Kernel density estimation on riemannian manifolds," *Statistics and Probability Letters*, vol. 73, no. 3, pp. 297–304, 2005.
 [77] G. Celeux, D. Chauveau, and J. Diebolt, "On Stochastic Versions of the EM Algorithm," INRIA, Research Report RR-2514, 1995. [Online]. Available: https://hal.inria.fr/inria-00074164
- [78] M. L. Eaton, Multivariate statistics, a vector space approach (Monograph series, Vol. 53). Institute of Mathematical Statistics, 2007.

Section II

APPENDIX I

THE CARTAN DECOMPOSITION, KILLING FORM, AND ROOTS

This appendix provides additional details on the mathematical facts, concerning Riemannian symmetric spaces, which were introduced in Section II. It also justifies the application of these facts to spaces of structured covariance matrices, which was made in Section IV.

In Section II, Facts 4 and 5 introduced roots in formula (13) and the Killing form in formula (15), in order to express i) integrals of functions over \mathcal{M} , and ii) distances between pairs of points in \mathcal{M} , where \mathcal{M} is a Riemannian symmetric space of non-compact type. In Section IV, formulae (13) and (15) were specified to spaces of structured covariance matrices, according to the following scheme,

Facto	IV-A : Complex covariance matrices	formulae (34f) and (34g)
4 & 5	IV-B : Toeplitz covariance matrices	formulae (37f) and (37g)
	IV-C : Block-Toeplitz covariance matrices	formulae (40f) and (40g)

Section IV

Here, the definitions of roots and of the Killing form are briefly recalled, and the specific formulae of Section IV are accordingly proved. The starting point will be the Cartan decomposition.

THE CARTAN DECOMPOSITION AND KILLING FORM : A Riemannian symmetric space of non-compact type \mathcal{M} may always be expressed in the form $\mathcal{M} = G/H$, where G is a semisimple Lie group of non-compact type and H a compact Lie subgroup of G [36] (Chapter VI). Let \mathfrak{g} and \mathfrak{h} denote the Lie algebras of G and H, respectively. The local geometric properties of \mathcal{M} are equivalent to algebraic properties of the pair $(\mathfrak{g}, \mathfrak{h})$. This equivalence is expressed in terms of the Cartan decomposition, the Killing form, and roots. The Cartan decomposition is the following [36] (Chapter III, Page 156),

Cartan decomposition:
$$g = h + p$$
 (52)

where \mathfrak{p} is a subspace of \mathfrak{g} , which is complementary to \mathfrak{h} , and such that : $[\mathfrak{h}, \mathfrak{p}] \subset \mathfrak{p}$ and $[\mathfrak{p}, \mathfrak{p}] \subset \mathfrak{h}$. For the geometric meaning of these conditions, see [37] (Chapter 4, Page 186). In particular, for any point $o \in \mathcal{M}$ there exists a Cartan decomposition (52) such that $T_o \mathcal{M}$ is identified with \mathfrak{p} , and the isotropy algebra of $T_o \mathcal{M}$ is identified with \mathfrak{h} . Here, $T_o \mathcal{M}$ denotes the tangent space to \mathcal{M} at o. The Killing form $B : \mathfrak{g} \times \mathfrak{g} \to \mathbb{R}$ is an Ad-invariant symmetric bilinear form on \mathfrak{g} , given by

lling form:
$$B(v, w) = \operatorname{tr} (\operatorname{ad}(v) \operatorname{ad}(w)) \qquad v, w \in \mathfrak{g}$$
 (53)

where $ad(v) : \mathfrak{g} \to \mathfrak{g}$ is the linear mapping ad(v)u = [v, u]. The Killing form is negative definite on \mathfrak{h} and positive definite on \mathfrak{p} [36] (Proposition 7.4., Page 158).

The Cartan decomposition and Killing form provide an algebraic expression for the local Riemannian geometry of \mathcal{M} . Precisely, when $T_o \mathcal{M}$ is identified with p, the Riemannian metric and the Riemannian curvature tensor of \mathcal{M} can be written as follows [36] (Theorem 4.2., Page 180),

Riemanniain metric:
$$\|v\|_{c}^{2} = B(v, v)$$
 $v \in T_{o} \mathcal{M} \cong \mathfrak{p}$ (54a)

Riemanniain curvature:
$$R_o(v, w)u = -[[v, w], u]$$
 $v, w, u \in T_o \mathcal{M} \cong \mathfrak{p}$ (54b)

where $||v||_o$ denotes the Riemannian length of $v \in T_o \mathcal{M}$ and R_o the Riemannian curvature tensor at $o \in \mathcal{M}$.

ROOTS : It is clear that application of expressions (54a) and (54b) requires an efficient means of computing the linear mapping $\operatorname{ad}(v)$ for $v \in \mathfrak{p}$. This is nicely done using the roots of the Lie algebra \mathfrak{g} . A root λ of \mathfrak{g} is a linear mapping $\lambda : \mathfrak{a} \to \mathbb{R}$, where \mathfrak{a} is a maximal Abelian subspace of \mathfrak{p} (Abelian means commutative, so $[a_1, a_2] = 0$ for $a_1, a_2 \in \mathfrak{a}$). Precisely, $\lambda : \mathfrak{a} \to \mathbb{R}$ is called a root of \mathfrak{g} if there exists a non-trivial subspace $\mathfrak{g}_{\lambda} \subset \mathfrak{g}$ such that $\operatorname{ad}(a)w = \lambda(a)w$ for all $w \in \mathfrak{g}_{\lambda}$, see [36] (Chapter III).

First, roots provide an efficient means of computing ad(a) when $a \in \mathfrak{a}$. Note that the subspace \mathfrak{g}_{λ} is called the root space corresponding to λ . Its dimension is denoted m_{λ} , (as in (14) of Section II). The various subspaces \mathfrak{g}_{λ} are in direct sum, so that the linear mappings ad(a) for $a \in \mathfrak{a}$ are jointly diagonalisable, and equal to $\lambda(a)$ times the identity mapping on each subspace \mathfrak{g}_{λ} . Second, any $v \in \mathfrak{p}$ is of the form $v = \mathrm{Ad}(h) a$ for some $a \in \mathfrak{a}$ and $h \in H$. Then, $\mathrm{Ad}(v) = \mathrm{Ad}(h) \circ \mathrm{Ad}(h^{-1})$, so that the linear mappings $\mathrm{ad}(a)$ are the same up to similarity.

Note finally that the system of roots λ of \mathfrak{g} possesses a distinct symmetry, which is described by a discrete group of permutations, called the Weyl group. Here, it is only mentioned that if λ is a root, then so is $-\lambda$. Moreover, the set of non-zero roots can be split into two disjoint subsets, called positive roots and negative roots, in such a way that the positive roots span the dual space of \mathfrak{a} [36] (Chapter VI). To indicate that a root λ is positive, one writes $\lambda > 0$, (again, as in (14) of Section II).

Example 1: here, it is shown that formulae (34f) and (34g) follow from Facts 4 and 5, applied to the case where $\mathcal{M} = \mathcal{H}_n$ as in IV-A. Precisely, (34f) and (34g) can be obtained from Facts 4 and 5, by substitution from the following table [36][67].

Table 2.: Data for the symmetric space \mathcal{H}_n

Group action:	$G = \{n \times n \text{ complex invertible matrices}\}$	action given by (34e)
Quotient space:	$\mathcal{H}_n = G/H$	$H = \{n \times n \text{ unitary matrices}\}$
Lie algebra of G :	$\mathfrak{g} = \{n \times n \text{ complex matrices}\}$	
Lie algebra of H :	$\mathfrak{h} = \{n \times n \text{ skew-Hermitian matrices}\}$	
Cartan decomposition (52):	$\mathfrak{p} = \{n \times n \text{ Hermitian matrices}\}$	$T_I \mathcal{H}_n \cong \mathfrak{p}, (I = n \times n \text{ identity matrix})$
Abelian sub- space a:	$\mathfrak{a} = \{n \times n \text{ real diagonal matrices}\}$	$r \in \mathfrak{a}$ is written $r = \operatorname{diag}(r_1, \ldots, r_n)$
Positive roots :	$\lambda_{ij}(r) = r_i - r_j$; $i < j$	with $m_{\lambda_{ij}} = 2$
Killing form :	$B(r,r) = \frac{1}{2n} \sum_{ij} (r_i - r_j)^2$	for $r \in \mathfrak{a}$

The data in this table can be substituded into (13) and (15), to obtain (34f) and (34g). While (34f) follows immediately, additional care is required for (34g). This is because \mathcal{H}_n is not a Riemannian symmetric space of non-compact type, but only a Riemannian symmetric space of non-positive curvature [67] (Page 262). Therefore, (34g) does not follow directly from (15), but rather from (15) and (21), after noticing $\mathcal{H}_n = \mathbb{R} \times S\mathcal{H}_n$, where $S\mathcal{H}_n$ is the subset of $Y \in \mathcal{H}_n$ such that $\det(Y) = 1$. This implies that $r \in \mathfrak{a}$ should be considered as a couple $r = (\operatorname{tr}(r), \bar{r})$ where $\bar{r} = r - \frac{1}{n}\operatorname{tr}(r) I \in \mathfrak{a}$. Then, using the notation of (54a), it follows from (21) that $||r||_I^2 = \frac{1}{n}\operatorname{tr}^2(r) + B(\bar{r}, \bar{r})$, which is the same as the right hand side of (34g).

1 .

- >

► Example 2: the space \mathcal{T}_n of Toeplitz covariance matrices, treated in IV-B, is a special case of the space \mathcal{T}_n^N of block-Toeplitz covariance matrices, treated in IV-C. Precisely, this is the special case obtained when N = 1. Therefore, in order to show that formulae (37f) and (37g) follow from Facts 4 and 5, applied to the case where $\mathcal{M} = \mathcal{T}_n$, it is enough to replace N = 1 in the argument of the following example.

Example 3: here, it is shown that formulae (40f) and (40g) follow from Facts 4 and 5, applied to the case where $\mathcal{M} = \mathcal{T}_n^N$ as in IV-C. Recall that \mathcal{T}_n^N admits the decomposition (39). Through (40b) and (40c), this decomposition implies the first line in each of the two formulae (40f) and (40g). The second line in each of these two formulae can be found from Facts 4 and 5, by substitution from the following table [36][67]. In this table, U, J, C, w, and v, denote $2N \times 2N$ complex matrices, where

- \

$$U = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \quad J = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix}, \quad C = \begin{pmatrix} -I & 0 \\ 0 & I \end{pmatrix} \quad (I = N \times N \text{ identity matrix})$$

$$\underline{\text{Table 3.: Data for the Siegel disc } \mathbb{D}_N$$
Group action:
$$G = \{U \mid U^T J U = J, U^H C U = C\} \quad \text{action given by (40e)}$$
Quotient space:
$$\mathbb{D}_N = G/H \quad H = \{U \mid A = D^H, A \text{ unitary}\}$$
Lie algebra of G :
$$\mathfrak{g} = \{w \mid w^T J + J w = 0, w^H C + C w = 0\}$$
Lie algebra of H :
$$\mathfrak{h} = \left\{w = \begin{pmatrix} \omega & 0 \\ 0 & \omega^H \end{pmatrix} \mid \omega \text{ skew-Hermitian}\right\}$$
Cartan decomposition (52):
$$\mathfrak{g} = \left\{v = \begin{pmatrix} 0 & s \\ s^H & 0 \end{pmatrix} \mid s \text{ complex symmetric}\right\} \quad T_0 \mathbb{D}_N \cong \mathfrak{p}, (0 = N \times N \text{ zero matrix})$$
Abelian subspace \mathfrak{a} :
$$\mathfrak{a} = \left\{\begin{pmatrix} 0 & r \\ r & 0 \end{pmatrix} \mid r \text{ real diagonal}\right\} \quad r = \operatorname{diag}(r_1, \dots, r_N)$$
Positive roots :
$$\lambda_{ij}(r) = r_i - r_j ; i < j \quad \text{with } m_{\lambda_{ij}} = 1 \quad (\text{identically zero when } N = 1)$$

$$\psi_{ij}(r) = r_i + r_j ; i \leq j \quad \text{with } m_{\nu_{ij}} = 1 \quad (\text{positive for any value of } N)$$
Killing form :
$$B(r, r) = \sum_i r_i^2$$

Indeed, the data in this table can be substituded into formulae (13) and (15), as these directly apply to the Siegel disc \mathbb{D}_N , which is a Riemannian symmetric space of non-compact type, in order to obtain the second line in each of (40f) and (40g).

APPENDIX II

SAMPLING GAUSSIAN STRUCTURED COVARIANCE MATRICES

Algorithm 1. of III-A gave a generally applicable method for sampling from a Gaussian distribution on a Riemannian symmetric space of non-positive curvature M. Here, Algorithms 2., 3., and 4. are specific versions of Algorithm 1., which deal with the three spaces of structured covariance matrices considered in Section IV. Precisely, Algorithms 2., 3., and 4. sample Gaussian structured covariance matrices, by implementing the instructions of Algorithm 1., labeled 11 to 112. Accordingly, in the description below, each instruction in Algorithms 2., 3., and 4. will be given the same label as the corresponding instruction in Algorithm 1.

Example 1: here, Algorithm 2. samples from a Gaussian distribution on the space $\mathcal{M} = \mathcal{H}_n$, which was considered in IV-A. Precisely, the output of Algorithm 2. is a sample Y from a Gaussian distribution $G(\bar{Y}, \sigma)$ with parameters $\bar{Y} \in \mathcal{H}_n$ and $\sigma > 0$. In reference to Algorithm 1., only instructions 16 to 19 are required. In these instructions, the role of the compact group H is played by the group U(n) of $n \times n$ unitary matrices, and the role of (12) and (24) is played by the following (55a) and (55b), respectively,

spectral decomposition:
$$Y(r, U) = U e^r U^H$$
 (55a)

$$p(r) \propto \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i < j} \sinh^2(|r_i - r_j|/2) \quad \text{for } r = (r_1, \dots, r_n) \in \mathbb{R}^n$$
(55b)

In instruction 19 below, the notation $\overline{Y}^{1/2} \cdot Y$ is that of (34e), so that $\overline{Y}^{1/2} \cdot I = \overline{Y}$, where I is the $n \times n$ identity matrix, (compare to (ii) of Proposition 2).

16. san	pple U from a uniform distribution on U	T(n)	
16-a.	generate Z with $n \times n$ i.i.d. complex	normal entries	
16-b.	perform QR decomposition : $Z = UT$	% T upper triangular with unit diagonal	
		% see [78] (Proposition 7.3., Page 234)	
17. san	pple r from the multivariate density (55b))	Algorithm 2. : Sampling from a
18. put	Y = Y(r, U) as in (55a)		GAUSSIAN DISTRIBUTION ON \mathcal{H}_n
19. res	et Y to $\bar{Y}^{1/2} \cdot Y$	% $\bar{Y}^{1/2}$ = Hermitian square root of \bar{Y}	

1)

Example 2: here, Algorithm 3. samples from a Gaussian distribution on the space $\mathcal{M} = \mathcal{T}_n$, which was considered in IV-B. The output of this algorithm is a sample T from a Gaussian distribution $G(\bar{T}, \sigma)$ with parameters $\bar{T} \in \mathcal{T}_n$ and $\sigma > 0$. This output sample T is obtained via decomposition (36), in the form

$$T = \Psi(r, \alpha_1, \dots, \alpha_{n-1}) \quad r \in \mathbb{R}_+ \text{ and } \alpha_1, \dots, \alpha_{n-1} \in \mathbb{D}$$
(56a)

where the mapping $\Psi(r_0, \alpha_1, \ldots, \alpha_{n-1}) = (r_0, r_1, \ldots, r_{n-1})$ yields the matrix entries of T, given by $T_{ij} = T_{ji}^* = r_{i-j}$. Concretely, Ψ is computed by means of Verblunsky's formula, which maps $\alpha_1, \ldots, \alpha_j$ to r_1, \ldots, r_j , for $j = 1, \ldots, n-1$ [55] (Theorem 1.5.5., Page 60).

For Algorithm 3., let $\overline{T} = \Psi(\overline{r}, \overline{\alpha}_1, \dots, \overline{\alpha}_{n-1})$. Instruction 13 of this algorithm samples r, given \overline{r} and σ . Then, instructions 16 to 19 sample each one of the α_j , given $\overline{\alpha}_j$ and σ . Here, in reference to instructions 16 to 19 of Algorithm 1., the role of the compact group H is played by the circle group SU(1), and the role of (12) and (24) is played by the following (56b) and (56c).

polar coordinates on
$$\mathbb{D}$$
: $\alpha(\rho, \theta) = \tanh(\rho) e^{i\theta}$ (56b)

$$p(\rho|\sigma) \propto \exp\left[-\frac{\rho^2}{2\sigma^2}\right] \sinh(|\rho|)$$
 (56c)

In instruction 19, the notation $u \cdot \alpha$ is that of (37e), and $u_j = u(\alpha_j)$, where

$$u(\alpha) = \begin{pmatrix} e^{i\frac{\theta}{2}}\cosh(\rho) & e^{i\frac{\theta}{2}}\sinh(\rho) \\ e^{-i\frac{\theta}{2}}\sinh(\rho) & e^{-i\frac{\theta}{2}}\cosh(\rho) \end{pmatrix} \quad \text{for } \alpha = \alpha(\rho, \theta)$$
(56d)

so that $u(\alpha) \cdot 0 = \alpha$.

13. sample r from a log-normal distribution with median \bar{r} and scale σ/\sqrt{n}	% $r = e^t$ where $t \sim \mathcal{N}\left(\log \bar{r}, \sigma/\sqrt{n}\right)$
. for $j=1,\ldots,n-1$ do	
16. sample θ_j from a uniform distribution on $[0, 2\pi]$	% so $e^{\mathrm{i} \theta_j} \sim$ uniform distribution on $SU(1)$
17. sample ρ_j from the density $p(\rho_j \sigma / \sqrt{n-j})$ given by (56c)	
18. put $\alpha_j = \alpha(\rho, \theta)$ as in (56b)	
19. reset α_j to $u_j \cdot \alpha_j$ where u_j is given by (56d)	Algorithm 3.: Sampling from a
111.end for	GAUSSIAN DISTRIBUTION ON \mathcal{T}_n
112. put $T = \Psi(r, \alpha_1,, \alpha_{n-1})$ as in (56a)	

Example 3: here, Algorithm 4. samples from a Gaussian distribution on the space $\mathcal{M} = \mathcal{T}_n^N$, which was considered in IV-C. The algorithm obtains a sample T from a Gaussian distribution $G(\bar{T}, \sigma)$, where $\bar{T} \in \mathcal{T}_n^N$ and $\sigma > 0$, via decomposition (39). This decomposition is expressed in the form,

$$\Gamma = \Psi(P, \Omega_1, \dots, \Omega_{n-1}) \quad P \in \mathcal{T}_N \text{ and } \Omega_1, \dots, \Omega_{n-1} \in \mathbb{D}_N$$
(57a)

where the mapping $\Psi(P, \Omega_1, \dots, \Omega_{n-1}) = (T_0, T_1, \dots, T_{n-1})$ yields the $N \times N$ blocks of T. Concretely, Ψ is computed as in [23] (formula (3.4), Page 6), which recursively gives the blocks T_1, \dots, T_j in terms of $\Omega_1, \dots, \Omega_{n-1}$, for $j = 1, \dots, n-1$. For Algorithm 4., let $\overline{T} = \Psi(\overline{P}, \overline{\Omega}_1, \dots, \overline{\Omega}_{n-1})$. Instruction 13 of this algorithm samples P, given \overline{P} and σ . Then, instructions 16 to 19

For Algorithm 4., let $\overline{T} = \Psi(\overline{P}, \Omega_1, \dots, \Omega_{n-1})$. Instruction 13 of this algorithm samples P, given \overline{P} and σ . Then, instructions 16 to 19 sample each one of the Ω_j , given $\overline{\Omega}_j$ and σ . Here, in reference to instructions 16 to 19 of Algorithm 1., the role of the compact group H is played by the group SU(N) of unitary $N \times N$ matrices with unit determinant, and the roles of (12) and (24) are respectively played by (40h) and by the following (57b),

$$p(r|\sigma) \propto \exp\left[-\frac{\|r\|^2}{2\sigma^2}\right] \prod_{i< j} \sinh(|r_i - r_j|) \prod_{i \le j} \sinh(|r_i + r_j|) \quad \text{for } r = (r_1, \dots, r_N) \in \mathbb{R}^N$$
(57b)

In instruction 19, the notation $U \cdot \Omega$ is that of (40e), and $U_i = U(\Omega_i)$, where

$$U(\Omega) = \begin{pmatrix} \Theta \cosh(r) & \Theta \sinh(r) \\ \Theta^* \sinh(r) & \Theta^* \cosh(r) \end{pmatrix} \quad \text{for } \Omega = \text{ right-hand side of (40h)}$$
(57c)

so that $U(\Omega) \cdot 0 = \Omega$, where 0 denotes the $N \times N$ zero matrix.

13. sample P from a Gaussian distribution $G\left(\bar{P},\sigma/\sqrt{n} ight)$ on \mathcal{T}_N	% use Algorithm 3.
. for $j=1,\ldots,n-1$ do	
16. sample Θ_j from a uniform distribution on $SU(N)$	% compare to 16 in Algorithm 2.
16-a. generate Z with $n \times n$ i.i.d. complex normal entries	
16-b. perform QR decomposition: $Z = \Theta T$	
16-c. reset Θ to $(\det \Theta)^{1/N} \Theta$	
17. sample r_j from the multivariate density $p(r_j \sigma / \sqrt{n-j})$ given by	(57b)
18. put $\Omega_j = \Omega(r_j, \Theta_j)$	$\Re \ \Omega(r_j, \Theta_j) =$ right-hand side of (40h)
19. reset Ω_j to $U_j \cdot \Omega_j$ where U_j is given by (57c)	ALGORITHM 4.: SAMPLING FROM A
111.end for	GAUSSIAN DISTRIBUTION ON \mathcal{T}_n^N
112. put $T = \Psi(P, \Omega_1, \dots, \Omega_{n-1})$ as in (56a)	