

Riemannian statistics meets random matrix theory : towards learning from high-dimensional covariance matrices

Salem Said, Simon Heuveline, and Cyrus Mostajeran

Abstract

Riemannian Gaussian distributions were initially introduced as basic building blocks for learning models which aim to capture the intrinsic structure of statistical populations of positive-definite matrices (here called covariance matrices). While the potential applications of such models have attracted significant attention, a major obstacle still stands in the way of these applications: there seems to exist no practical method of computing the normalising factors associated with Riemannian Gaussian distributions on spaces of high-dimensional covariance matrices. The present paper shows that this missing method comes from an unexpected new connection, with random matrix theory. Its main contribution is to prove that Riemannian Gaussian distributions of real, complex, or quaternion covariance matrices are equivalent to orthogonal, unitary, or symplectic log-normal matrix ensembles. This equivalence yields a highly efficient approximation of the normalising factors, in terms of a rather simple analytic expression. The error due to this approximation decreases like the inverse square of dimension. Numerical experiments are conducted which demonstrate how this new approximation can unlock the difficulties which have impeded applications to real-world datasets of high-dimensional covariance matrices. The paper then turns to Riemannian Gaussian distributions of block-Toeplitz covariance matrices. These are equivalent to yet another kind of random matrix ensembles, here called “acosh-normal” ensembles. Orthogonal and unitary “acosh-normal” ensembles correspond to the cases of block-Toeplitz with Toeplitz blocks, and block-Toeplitz (with general blocks) covariance matrices, respectively.

Index Terms

Riemannian Gaussian distribution, random matrix theory, covariance matrix, block-Toeplitz covariance matrix, high-dimensional data

I. INTRODUCTION

The present paper aims to bring together the tools of Riemannian statistics and those of random matrix theory, in order to unlock the computational issues which have stood in the way of applying Riemannian learning models to datasets of high-dimensional covariance matrices.

S. Said is a CNRS research Scientist, at Laboratoire Jean Kuntzmann, Université Grenoble-Alpes

S. Heuveline is with the Centre for Mathematical Sciences of the University of Cambridge

C. Mostajeran is with the Department of Engineering of the University of Cambridge in the United Kingdom and the School of Physical and Mathematical Sciences of Nanyang Technological University (NTU) in Singapore

Over the past few years, Riemannian Gaussian distributions, and mixtures of these distributions, were introduced as a means of modeling the intrinsic structure of statistical populations of covariance matrices [1]–[3]. They have found successful applications in fields such as brain-computer interface analysis and artificial intelligence [4] [5]. However, such applications could not be pursued for covariance matrices of relatively larger dimension (for example, 50×50). Indeed, there seemed to exist no practical method of computing the normalising factors associated with Riemannian Gaussian distributions on spaces of high-dimensional covariance matrices.

The theoretical contribution of this paper is to show that this missing method arises quite naturally, as soon as one realises that a strong connection exists between Riemannian Gaussian distributions and random matrix theory. Roughly speaking, Riemannian Gaussian distributions on spaces of real, complex, or quaternion covariance matrices correspond to log-normal orthogonal, unitary or symplectic matrix ensembles. These are similar to the classical, widely-known Gaussian orthogonal, unitary and symplectic ensembles, but with the normal weight function replaced with a log-normal weight function. Thanks to this new connection, the powerful tools of random matrix theory can be employed to uncover several original properties of Riemannian Gaussian distributions, especially for higher-dimensional covariance matrices.

The present paper also extends the equivalence between Riemannian Gaussian distributions and random matrix ensembles to the case of block-Toeplitz matrices. Instead of a log-normal weight function $\exp(-\log^2(x))$, the corresponding weight function is “acosh-normal” $\exp(-\text{acosh}^2(x))$, where acosh is the inverse hyperbolic cosine. No attempt is made to develop these additional “acosh-normal” ensembles, at least not for now. In fact, there are already plenty of results to derive and apply with the log-normal ensembles, which cover the important cases of real and complex covariance matrices. This is in part because log-normal ensembles (also called Stieltjes-Wigert ensembles) are amenable to a direct analytic treatment, as will be seen below.

It is also because of a rather surprising connection also noted in [6]. The log-normal unitary ensemble had appeared in the theoretical physics literature about twenty years ago, as a random matrix model for the Chern-Simons quantum field theory [7]. As seen in the present paper, this ensemble corresponds to Riemannian Gaussian distributions on the space of complex covariance matrices. Of course, such Riemannian Gaussian distributions have nothing to do with quantum field theory, but some of their valuable properties can be obtained by carefully re-adapting already existing results, found within this theory.

The present paper was developed independently from a theoretical physics paper, published only very recently [8]. Both papers focus on the connection between Riemannian Gaussian distributions and log-normal matrix ensembles, but significantly differ in terms of their contribution and focus, as will be discussed below. The reader is also referred to [9] for previous results on asymptotic computations of log-partition functions from the random matrix theory literature, as well as the recent paper [10], which considers global and local scaling limits for the $\beta = 2$ Stieltjes-Wigert random matrix ensemble and associated physical interpretations.

The paper is organised as follows. Section II recalls basic background material on Riemannian Gaussian distributions. Section III presents the main original results, obtained by applying random matrix theory to the study of Riemannian Gaussian distributions. Section IV illustrates the importance of these original results to learning from datasets of high-dimensional covariance matrices. Finally, Section V considers Riemannian Gaussian distributions

on the space of block-Toeplitz covariance matrices, and introduces the corresponding “acosh-normal” ensembles. Proofs of the propositions stated in Section III and V are provided in Appendix A.

II. NOTATION AND BACKGROUND

Let \mathcal{P}_N^β denote the space of $N \times N$ covariance matrices which are either real ($\beta = 1$), complex ($\beta = 2$) or quaternion ($\beta = 4$). Precisely, the elements of \mathcal{P}_N^1 are real, symmetric positive-definite matrices, while those of \mathcal{P}_N^2 and \mathcal{P}_N^4 are respectively complex and quaternion, hermitian positive definite matrices.

Of course, \mathcal{P}_N^β is an open convex cone, sitting inside the real vector space \mathcal{S}_N^β of real, complex or quaternion self-adjoint matrices (the adjoint of a matrix being its conjugate transpose). Therefore, \mathcal{P}_N^β is a real differentiable manifold, whose tangent space $T_Y \mathcal{P}_N^\beta$ at any point $Y \in \mathcal{P}_N^\beta$ is isomorphic to \mathcal{S}_N^β . In particular, the dimension of \mathcal{P}_N^β is $\dim_N^\beta = NN_\beta$ where $N_\beta = \frac{\beta}{2}(N-1) + 1$.

The elements Y of \mathcal{P}_N^β are in one-to-one correspondence with the centred (zero-mean) N -variate normal distributions (real, complex circular, or quaternion circular, according to the value of β). Thus, \mathcal{P}_N^β can be equipped with the Rao-Fisher information metric of the centred N -variate normal model [11]. This is identical to the so-called affine-invariant metric introduced in [12],

$$\langle u, v \rangle_Y = \Re \operatorname{tr}((Y^{-1}u)(Y^{-1}v)) \quad u, v \in T_Y \mathcal{P}_N^\beta \simeq \mathcal{S}_N^\beta \quad (1)$$

where \Re denotes the real part and tr denotes the trace. The Riemannian geometry of the metric (1) is quite well-known in the geometric information science community (see the recent book [13]). Recall here the associated geodesic distance

$$d^2(X, Y) = \operatorname{tr} \left(\log^2 \left(X^{-\frac{1}{2}} Y X^{-\frac{1}{2}} \right) \right) \quad X, Y \in \mathcal{P}_N^\beta \quad (2)$$

where matrix logarithms and powers are understood as self-adjoint matrix functions, obtained by taking logarithms and powers of eigenvalues. The main advantage of the metric (1) is its invariance under the action of the group G_N^β of $N \times N$ invertible matrices with real, complex or quaternion entries (according to the value of β). For example, if $A \in G_N^\beta$ and X is replaced by $A \cdot X = AXA^\dagger$ while Y is replaced by $A \cdot Y = AYA^\dagger$, then the distance $d(X, Y)$ remains unchanged (note that † denotes the adjoint, or conjugate transpose).

In terms of the distance (2), we define Riemannian Gaussian distributions on the space \mathcal{P}_N^β to be given by parameterised probability density function $p(Y|\bar{Y}, \sigma)$, as in [2] [3],

$$p(Y|\bar{Y}, \sigma) = (Z(\sigma))^{-1} \exp \left[-\frac{d^2(Y, \bar{Y})}{2\sigma^2} \right] \quad (3)$$

where $\bar{Y} \in \mathcal{P}_N^\beta$ is the centre-of-mass parameter, and $\sigma > 0$ the dispersion parameter. The normalising factor $Z(\sigma)$ is given by the integral

$$Z(\sigma) = \int_{\mathcal{P}_N^\beta} \exp \left[-\frac{d^2(Y, \bar{Y})}{2\sigma^2} \right] dv(Y) \quad (4)$$

with respect to the Riemannian volume $dv(Y) = \det(Y)^{-N_\beta} \{dY\}$. Here,

$$\{dY\} = \prod_{i \leq j} dY_{ij} \quad (\beta = 1)$$

$$\{dY\} = \prod_{i \leq j} dY_{ij}^{(a)} \prod_{i < j} dY_{ij}^{(b)} \quad (\beta = 2)$$

$$\{dY\} = \prod_{i \leq j} dY_{ij}^{(a)} \prod_{i < j} dY_{ij}^{(b)} \prod_{i < j} dY_{ij}^{(c)} \prod_{i < j} dY_{ij}^{(d)} \quad (\beta = 4)$$

where $Y_{ij} = Y_{ij}^{(a)} + Y_{ij}^{(b)} i$ if Y_{ij} is complex and $Y_{ij} = Y_{ij}^{(a)} + Y_{ij}^{(b)} i + Y_{ij}^{(c)} j + Y_{ij}^{(d)} k$ if Y_{ij} is quaternion (here, i, j, k denote complex or quaternion imaginary units).

Both formulae (3) and (4) are greatly simplified by introducing “polar coordinates”. Each $Y \in \mathcal{P}_N^\beta$ can be diagonalised as $Y = Ue^rU^\dagger$ where e^r is a diagonal matrix, with diagonal elements e^{r_i} for $(r_1, \dots, r_N) \in \mathbb{R}^N$, and $U \in K_N^\beta$, which means $UU^\dagger = I_N$ (the $N \times N$ identity matrix). Note that K_N^1 is the orthogonal group $O(N)$, K_N^2 is the unitary group $U(N)$, and K_N^4 is the symplectic (quaternion unitary) group $\text{Sp}(N)$ (for deeper insight, see the recent review [14]).

Now, as in [3], let Y follow the Riemannian Gaussian density (3), and $Y' = \bar{Y}^{-\frac{1}{2}} Y \bar{Y}^{-\frac{1}{2}}$. If $Y' = Ue^rU^\dagger$, then U is uniformly distributed on K_N^β (for a rigorous definition of this “uniform distribution”, see [15]), while (r_1, \dots, r_N) have joint probability density function

$$p(r|\sigma) \propto \prod_i \exp \left[-\frac{r_i^2}{2\sigma^2} \right] \prod_{i < j} \sinh^\beta \left| \frac{r_i - r_j}{2} \right| \quad (5)$$

where \propto indicates proportionality and \sinh the hyperbolic sine. In addition, the normalising factor $Z(\sigma)$ in (4) reduces to a certain multiple integral $z_\beta(\sigma)$. Specifically,

$$Z(\sigma) = \Omega_N z_\beta(\sigma) \quad (6)$$

where Ω_N is a numerical constant, which appears after the uniformly distributed matrix U is integrated out of (4), and where

$$z_\beta(\sigma) = \frac{1}{N!} \int_{\mathbb{R}^N} \prod_{i < j} \sinh^\beta \left| \frac{r_i - r_j}{2} \right| \prod_i \exp \left[-\frac{r_i^2}{2\sigma^2} \right] dr_i \quad (7)$$

One of the main issues addressed in the present paper is the efficient approximation of the multiple integral $z_\beta(\sigma)$. Until now (in [2] [3]), this was done using a Monte Carlo technique which involved a smoothing method (containing certain arbitrarily fixed parameters) and which failed to produce coherent results when the dimension N increased beyond $N = 20$.

Before proceeding to present our main results, we should briefly note the existence of several alternative proposals for the extension of Gaussian distributions to Riemannian manifolds, including constructions based on heat flows and diffusion processes [16]. Further discussion of these alternative formulations is beyond the scope of this paper and the interested reader is referred to the literature on the topic for further information. See [13] (Section 3.4.3 and Chapter 10) and the references therein for a recent and comprehensive account.

III. MAIN RESULTS

The main results of the present paper stem from the equivalence between Riemannian Gaussian distributions on the spaces of real, complex, or quaternion covariance matrices and log-normal orthogonal, unitary, and symplectic matrix ensembles.

Let Y be a random matrix in \mathcal{P}_N^β which follows the Riemannian Gaussian density $p(Y|\bar{Y}, \sigma)$ of (3). It is always possible to assume that $\bar{Y} = I_N$, since this holds after replacing Y by $Y' = \bar{Y}^{-\frac{1}{2}} Y \bar{Y}^{-\frac{1}{2}}$. Then, the probability distribution of Y is described by the following proposition.

Proposition 1. *Let Y follow the Riemannian Gaussian density (3) with $\bar{Y} = \mathbf{I}_N$. If $X = e^{N\beta\sigma^2}Y$, then the probability distribution of X is given by*

$$\mathbb{P}(X \in B) \propto \int_B \text{etr} \left[-\frac{\log^2(X)}{2\sigma^2} \right] \{dX\} \quad (8)$$

for any measurable subset B of \mathcal{P}_N^β . Here, $\text{etr}(\cdot) = \exp(\text{tr}(\cdot))$ and the notation $\{dX\}$ was introduced after (4).

In plain words, this proposition states that X follows a log-normal matrix ensemble. If X is diagonalised as $X = UxU^\dagger$, then U is uniformly distributed on K_N^β , and the eigenvalues (x_1, \dots, x_N) , which are all positive, have joint probability density function

$$p(x|\sigma) \propto |V(x)|^\beta \prod_i \rho(x_i, 2\sigma^2) \quad (9)$$

where $V(x) = \prod_{i < j} (x_j - x_i)$ is the Vandermonde determinant, and $\rho(x, k) = \exp(-\log^2(x)/k)$ is the log-normal weight function.

In essence, Proposition 1 is already contained in [8]. As a consequence of this proposition, the multiple integral $z_\beta(\sigma)$ of (7) may be expressed as follows

$$z_\beta(\sigma) = (2\pi\sigma^2)^{N/2} \times \exp[-N\beta\sigma^2/2] \times \frac{1}{N!} \int_{\mathbb{R}_+^N} |V(x)|^\beta \omega(dx) \quad (10)$$

where $\omega(dx) = \prod \bar{\rho}(dx_i)$ (product over $i = 1, \dots, N$) with $\bar{\rho}(dx_i) = (2\pi\sigma^2)^{-1/2} \rho(x_i, 2\sigma^2) dx_i$. In [8], integrals as the one in (10) are expressed using the Andreev or De Bruijn identities, often employed in random matrix theory. These yield somewhat cumbersome formulae, involving determinants of size $N \times N$. On the other hand, the present paper focuses on an alternative approach, which turns out to be more suitable from a practical point of view. Instead of expressing $z_\beta(\sigma)$ exactly, by means of complicated formulae, the aim is to use a highly efficient approximation, which involves a single analytic expression. This is indicated by the following proposition.

Proposition 2. *In the limit where $N \rightarrow \infty$ and $\sigma \rightarrow 0$, while the product $t = N\sigma^2$ remains constant,*

$$\frac{1}{N^2} \log z_\beta(\sigma) \longrightarrow \frac{\beta}{2} \Phi\left(\frac{\beta}{2}t\right) \quad \Phi(\xi) = \frac{\xi}{6} - \frac{\text{Li}_3(e^{-\xi}) - \text{Li}_3(1)}{\xi^2} \quad (11)$$

where Li_3 is the trilogarithm function, $\text{Li}_3(\eta) = \sum_{k=1}^{\infty} \eta^k / k^3$.

In [8], the limit in (11) is only mentioned in passing, and not stated under the same form. Here, this limit will be given centre stage. Proposition 2 states that (11) is valid in the “double-scaling regime” ($N \rightarrow \infty$ and $\sigma \rightarrow 0$), but numerical experiments have shown that $\log z_\beta(\sigma)$ can be replaced by the expression afforded by (11) without notable loss of accuracy, whenever σ is small in comparison with N (this is further illustrated below).

At least informally, this can be justified by appealing to arguments originating in theoretical physics [7]. Considered as a function of $t = N\sigma^2$, $F(t) = \log z_\beta(\sigma)$ is called the Free energy (log of partition function). This free energy can be expanded in an asymptotic series (see Section 1.3 of [7]),

$$F(t) \sim \sum_{g=0}^{\infty} F_g(t) \left(\frac{t}{N}\right)^{2g-2} \quad (12)$$

which is obtained by summing over Feynman diagrams, or so-called fatgraphs. Each coefficient $F_g(t)$ is itself a series $F_g(t) \sim \sum_h F_{g,h} t^h$, where $F_{g,h}$ counts fatgraphs which are said to have h holes and genus g (this means

that one thinks of a fatgraph as a graph with h loops, drawn on a surface of genus g , such as a sphere or torus, *etc*). Now, accepting (12), it follows that for each fixed value of t ,

$$\frac{1}{N^2} F(t) = F_0(t) + O\left(\frac{1}{N^2}\right) \quad (13)$$

so that $F_0(t)$ approximates the left-hand side up to an error of the order of $1/N^2$. Finally, recalling that $F(t) = \log z_\beta(\sigma)$, it is clear that $F_0(t)$ is the right-hand side of (11) — this follows by uniqueness of asymptotic expansions.

In addition to the asymptotic form of $\log z_\beta(\sigma)$, another quantity of interest is the asymptotic empirical distribution of eigenvalues, of a random matrix Y which follows the Riemannian Gaussian density $p(Y|\bar{Y}, \sigma)$ with $\bar{Y} = I_N$. Let (y_1, \dots, y_N) denote the eigenvalues of Y , and consider their empirical distribution

$$\hat{\nu}_\beta(I) = \mathbb{E} \left[\frac{|y_i \in I|}{N} \right] \quad (14)$$

for any open interval $I \subset \mathbb{R}_+$, where \mathbb{E} denotes expectation, and $|y_i \in I|$ the number of y_i which belong to I . In [8], the probability density function of $\hat{\nu}_2$ was expressed as a weighted sum of Gaussian distributions, by a direct application of the Christoffel-Darboux formula, as in [17]. It is possible to do so for any value of N , only because the $\beta = 2$ case can be studied using a well-known family of orthogonal polynomials, called Stieltjes-Wigert polynomials. No such analytic tool is available when $\beta = 1$ or 4.

To make up for this issue, the following proposition provides an asymptotic expression of the distribution $\hat{\nu}_\beta$, valid for all values $\beta = 1, 2, 4$.

Proposition 3. *In the limit where $N \rightarrow \infty$ and $\sigma \rightarrow 0$, while the product $t = N\sigma^2$ remains constant, the empirical distribution $\hat{\nu}_\beta$ converges weakly to a distribution with probability density function $n(y|\beta t/2)$, where*

$$n(y|\xi) = \frac{1}{\pi \xi y} \arctan \left(\frac{4e^\xi y - (y+1)^2}{y+1} \right) \quad (15)$$

on the interval $a(\xi) \leq y \leq b(\xi)$, where $a(\xi) = c(1 + \sqrt{1-c})^{-2}$ and $b(\xi) = c(1 - \sqrt{1-c})^{-2}$, with $c = e^{-\xi}$.

One hopes that, similar to the situation discussed after Proposition 2, the asymptotic density (15) approximates the finite- N empirical distribution $\hat{\nu}_\beta$ to such a good accuracy that one can replace $\hat{\nu}_\beta$ by this asymptotic density, for many practical purposes. This possibility will not be further investigated in the present paper.

IV. TOWARDS LEARNING APPLICATIONS

Riemannian Gaussian distributions were initially proposed as basic building blocks for learning models which aim to capture the intrinsic structure of statistical populations of covariance matrices. These include the mixture models and hidden Markov models, introduced in [2] [18] and further developed in [19]. In order to make use of these models in real-world applications, it is indispensable to know how to effectively compute the logarithm of the multiple integral $z_\beta(\sigma)$ of (7), for a dimension N which may be in the tens or hundreds.

Knowledge of $\log z_\beta(\sigma)$ is already crucial in the simplest situation, where one tries to fit a single Riemannian Gaussian density $p(Y|\bar{Y}, \sigma)$ (rather than a whole mixture) to data $Y_1, \dots, Y_M \in \mathcal{P}_N^\beta$. Indeed, this will require setting $\bar{Y} = \hat{Y}_M$ and $\sigma = \hat{\sigma}_M$, where \hat{Y}_M and $\hat{\sigma}_M$ are the maximum-likelihood estimates, given in [2] [3]. Specifically,

$$\hat{Y}_M = \operatorname{argmin}_{Y \in \mathcal{P}_N^\beta} \sum_{m=1}^M d^2(Y_m, Y) \quad (16)$$

is the Fréchet mean of the data Y_m , with respect to the distance (2), and $\hat{\sigma}_M$ is the solution of the nonlinear equation

$$\phi(\hat{\sigma}_M) = \frac{1}{M} \sum_{m=1}^M d^2(Y_m, \hat{Y}_M) \quad \phi(\sigma) = \sigma^3 \frac{d}{d\sigma} \log z_\beta(\sigma) \quad (17)$$

Therefore, it is already impossible to solve a toy problem, with a single Riemannian Gaussian density, without having some kind of hold on $\log z_\beta(\sigma)$.

Until now (in [2] [3]), $z_\beta(\sigma)$ was approximated using an *ad hoc* Monte Carlo technique, which failed to produce coherent results for a dimension N just above $N = 20$. In the present section, the aim will be to show that significantly improved results can be obtained by solving Equation (17) after approximating $\log z_\beta(\sigma)$ with the expression afforded by (11), according to Proposition 2.

First, numerical experiments were conducted to verify the validity of this new approximation. According to (13), the right-hand side of (11) should approximate $\log z_\beta(\sigma)/N^2$ up to an error of the order of $1/N^2$. To see that this is correct, the right-hand side of (11) was compared to certain exact expressions of $\log z_\beta(\sigma)$. Namely, for the $\beta = 2$ case, one has the following expression, obtained in [20],

$$\log z_2(\sigma) = \frac{N}{2} \log(2\pi\sigma^2) + N(N^2 - 1)(\sigma^2/6) + \sum_{n=1}^{N-1} (N - n) \log(1 - e^{-n\sigma^2}) \quad (18)$$

and for the $\beta = 1$ case, when the dimension N is even, the following expression, based on [8],

$$\log z_1(\sigma) = \frac{N}{2} \log(2\pi\sigma^2) - N(N + 1)^2(\sigma^2/8) + \log \text{Pf} [M(\sigma)] \quad (\text{for } N \text{ even}) \quad (19)$$

where Pf denotes the Pfaffian, equal to the square root of the determinant, and the matrix $M(\sigma)$ has entries

$$M_{ij}(\sigma) = \exp [(i^2 + j^2)(\sigma^2/2)] \text{erf}((j - i)(\sigma/2)) \quad i, j = 1, \dots, N \quad (20)$$

with erf the error function.

Numerical evaluation of (18) for large values of N or σ (up to $\sigma = 10$) is quite straightforward. Moreover, it immediately shows that (11) and (18) agree very closely when σ is smaller than N , and then gradually diverge away from one another as σ increases. Figure 1 provides graphical illustration for $N = 10$ and 20. Still larger values of N yield an even stronger match between (11) and (18).

It is not equally straightforward to numerically evaluate (19). Even at $N = 10$, we begin encountering overflow problems for moderate values of σ when performing the computations in MATLAB. Still, as long as these overflow problems do not appear, it is possible to observe a close agreement between (11) and (19), as in the $\beta = 2$ case. This is shown in Figure 2 for $N = 6$ and 12.

Based on the numerical results summarised in Figures 1 and 2, it seems possible to use the right-hand side of (11), multiplied by N^2 , as a substitute for $\log z_\beta(\sigma)$. While this is only an approximation, it is a highly efficient one, and has the advantage of being given by a rather simple analytic expression. In the $\beta = 1$ case, direct numerical evaluation of $\log z_\beta(\sigma)$ is unstable for larger values of σ , and (11) offers a practical way out of this problem.

This can be verified by using symbolic computation (Mathematica), which allows us to extend the evaluation of (19) to larger values of σ for larger N . It is nonetheless associated with the drawback that the resulting curves tend to be artificially non-smooth as in Figure 3a, due to numerical artifacts. These non-smooth features also exist

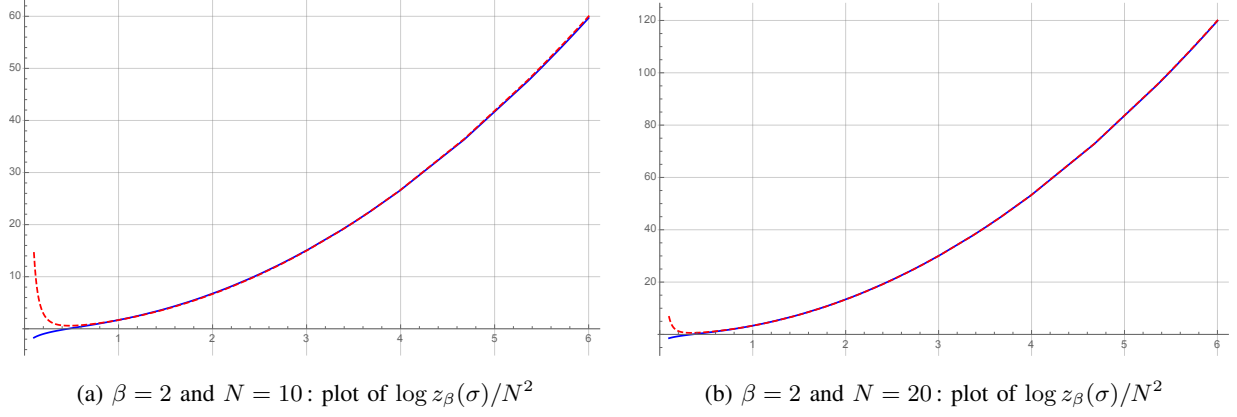


Fig. 1: Proposition 2 ($\beta = 2$): (11) in dashed red and (18) in solid blue

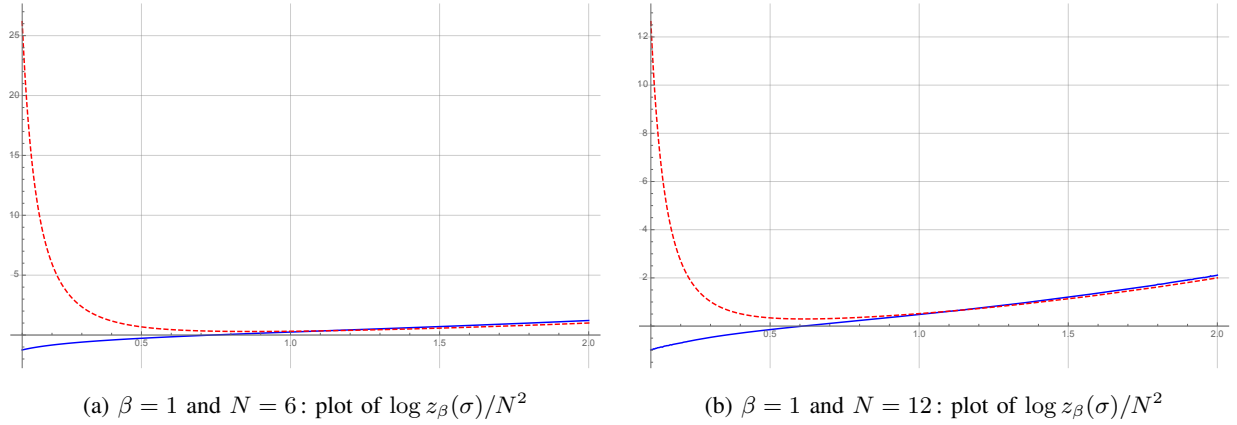


Fig. 2: Proposition 2 ($\beta = 1$): (11) in dashed red and (19) in solid blue

in the curve depicted in Figure 3b, but do not appear visible at the given resolution. This behaviour is particularly problematic in the context of Equation (17), due to the presence of the derivative of $\log z_\beta(\sigma)$. Furthermore, symbolic computation of (19) becomes exceedingly slow for sufficiently large N , and this can only be overcome by relying on approximations such as (11).

The second set of experiments directly addressed Equation (17), in the $\beta = 1$ case, for a dimension N ranging between 10 and 25. This equation was solved using the Newton method, after its right-hand side was approximated according to (11). Here, the solution obtained in this way will be denoted $\tilde{\sigma}_M$. This $\tilde{\sigma}_M$ is an approximation of the maximum-likelihood estimate $\hat{\sigma}_M$ (the exact solution of (17)). If this approximation is any good, $\tilde{\sigma}_M$ should approach the true value of σ for sufficiently large M (the number of data points Y_m). This can already be observed at $M = 10^3$ when $N = 10$, as reported in the following Table I. Each entry in this table gives the average value and standard deviation of $\tilde{\sigma}_M$, calculated over 20 independent trials (average \pm standard deviation). It is clear that increasing M from 10^3 to 10^4 only reduces the standard deviation of $\tilde{\sigma}_M$, without really affecting its average.

When $N = 10$, the right-hand side of (17) can still be approximated using the Monte Carlo technique mentioned




Fig. 3: Proposition 2 ($\beta = 1$): (11) in dashed red and (19) in solid blue for $N = 50$

true σ	1	2	3	4	5	6	7
$M = 10^3$	1.09 ± 0.00	2.01 ± 0.01	3.11 ± 0.30	3.56 ± 0.60	4.58 ± 0.60	5.13 ± 0.50	5.60 ± 0.75
$M = 10^4$	1.09 ± 0.00	2.01 ± 0.00	3.11 ± 0.15	3.80 ± 0.16	4.50 ± 0.23	5.37 ± 0.31	5.60 ± 0.37

TABLE I: The solution $\tilde{\sigma}_M$ of (17), for $\beta = 1$ and $N = 10$ (r.h.s. approximated using (11))

in [2] [3]. The solution obtained with this approximation will be denoted $[\tilde{\sigma}_M]_{\text{MC}}$. The following Table II shows that the estimation error from $\tilde{\sigma}_M$ is quite improved, in comparison with the estimation error from $[\tilde{\sigma}_M]_{\text{MC}}$. Moreover, $[\tilde{\sigma}_M]_{\text{MC}}$ seems to systematically overestimate the true value of σ .

true σ	1	2	3	4	5	6	7
$M = 10^3$	1.16 ± 0.01	2.90 ± 0.01	5.63 ± 0.93	7.38 ± 1.35	8.35 ± 1.50	9.40 ± 1.11	9.95 ± 1.12
$M = 10^4$	1.16 ± 0.01	2.90 ± 0.00	5.40 ± 0.37	7.00 ± 0.60	9.00 ± 0.50	9.50 ± 0.31	9.77 ± 0.35

TABLE II: The solution $[\tilde{\sigma}_M]_{\text{MC}}$ of (17), for $\beta = 1$ and $N = 10$ (MC approximation of r.h.s.)

When N is above 20, Monte Carlo approximation of the right-hand side of (17) is not feasible anymore, and one is left only with the approximation using (11). The solution $\tilde{\sigma}_M$ obtained from this approximation is shown in Table III, for $N = 20$ and $M = 10^4$. Here (contrary to Table II), $\tilde{\sigma}_M$ systematically underestimates the true value of σ . In fact, identical behavior was observed for N between 20 and 25, along with similar values of $\tilde{\sigma}_M$.

A practical means of overcoming this issue would be to include a penalty term into Equation (17), in order to enforce greater values of its solution. Then, the approximation using (11) can be successfully implemented for larger dimension N , where Monte Carlo approximation becomes useless (the present discussion stopped at $N = 25$, because the sampling algorithms used to generate the data points Y_m could not be taken any further).

In conclusion, the present section has demonstrated the new approximation of $\log z_\beta(\sigma)$, based on (11), significantly improves the existing Monte Carlo approximation, and also extends it to higher dimension. In upcoming

true σ	1	2	3	4	5	6	7
$M = 10^4$	1.02 ± 0.00	1.51 ± 0.00	1.87 ± 0.01	2.35 ± 0.02	2.82 ± 0.02	3.24 ± 0.03	3.63 ± 0.02

TABLE III: The solution $\tilde{\sigma}_M$ of (17), for $\beta = 1$ and $N = 20$ (r.h.s. approximated using (11))

work, the experiments conducted in this section will be generalised to more realistic learning models which may be applied to real-world data (for example, the mixture models considered in [4]).

V. BLOCK-TOEPLITZ COVARIANCE MATRICES

Similar to the spaces \mathcal{P}_N^β introduced in Section II, consider the spaces \mathcal{D}_N^β , defined as follows.

- \mathcal{D}_N^2 is the space of $N \times N$ matrices with complex entries, whose operator norm is < 1 .
- \mathcal{D}_N^1 is the space of $N \times N$ symmetric matrices with complex entries, whose operator norm is < 1 .

Here, operator norm means the largest singular value. The space \mathcal{D}_N^2 will be called the Hermitian Siegel domain, and \mathcal{D}_N^1 the symmetric Siegel domain. The focus will be restricted to these $\beta = 1$ or 2 cases, since they are closely related to time-series analysis and signal processing (the $\beta = 4$ case is similar, but involves quaternion matrices).

Specifically, assume a wide-sense stationary N -variate time series of length T is described by its autocovariance matrix Γ . Wide-sense stationarity implies Γ has a block-Toeplitz structure, with $T \times T$ blocks of size $N \times N$. When solving an optimal prediction problem, one may apply a multidimensional Szegő-Levinson algorithm to the autocovariance Γ [21] [22], and obtain a family of matrices $(\Gamma_0, \Omega_1, \dots, \Omega_{T-1})$, where Γ_0 is the zero-lag autocovariance of the original time series (this is a complex covariance matrix, $\Gamma_0 \in \mathcal{P}_N^2$), and each Ω_t is a so-called matrix reflection coefficient, which belongs to \mathcal{D}_N^2 . If the autocovariance Γ has Toeplitz blocks, then each Ω_t moreover belongs to \mathcal{D}_N^1 (which is a subspace of \mathcal{D}_N^2).

Riemannian geometry enters this picture in the following way [21] [22]. Consider an information metric on the space of block-Toeplitz autocovariance matrices Γ , equal to the Hessian of the entropy function $S(\Gamma) = \log \det(\Gamma)$. In terms of the new coordinates $(\Gamma_0, \Omega_1, \dots, \Omega_{T-1})$, this information metric is a direct product of the affine-invariant metric (1) on the first coordinate Γ_0 , and of a scaled copy of the Siegel domain metric on each of the remaining coordinates Ω_t (see formula (3.11) in [21]). Dropping the subscript t , this Siegel domain metric is given by,

$$\langle u, v \rangle_\Omega = \Re \operatorname{tr} \left[(\mathbf{I}_N - \Omega \Omega^\dagger)^{-1} u (\mathbf{I}_N - \Omega \Omega^\dagger)^{-1} v^\dagger \right] \quad u, v \in T_\Omega \mathcal{D}_N^\beta \quad (21)$$

where the tangent space $T_\Omega \mathcal{D}_N^\beta$ is isomorphic to the space of $N \times N$ complex matrices if $\beta = 2$, and to the space of $N \times N$ symmetric complex matrices if $\beta = 1$ (this isomorphism shows that the dimension of \mathcal{D}_N^β equals $2NN_\beta$). The geodesic distance associated with the Siegel metric (21) has the following expression [21] [22],

$$d^2(\Xi, \Omega) = \operatorname{tr} \left(\operatorname{arctanh}^2 \left(R^{\frac{1}{2}}(\Xi, \Omega) \right) \right) \quad \Xi, \Omega \in \mathcal{D}_N^\beta \quad (22)$$

where $R(\Xi, \Omega) = (\Xi - \Omega)(\mathbf{I}_N - \Omega^\dagger \Xi)^{-1}(\Xi^\dagger - \Omega^\dagger)(\mathbf{I}_N - \Omega \Xi^\dagger)^{-1}$ is called the matrix cross-ratio.

Riemannian Gaussian distributions on the space \mathcal{D}_N^β are given by their probability density function $p(\Omega|\bar{\Omega}, \sigma)$ which is of the same form as in (3), but with the distance $d(\Omega, \bar{\Omega})$ determined by (22), and the normalising factor

$$Z(\sigma) = \int_{\mathcal{D}_N^\beta} \exp \left[-\frac{d^2(\Omega, \bar{\Omega})}{2\sigma^2} \right] dv(\Omega) \quad (23)$$

with respect to the Riemannian volume $dv(\Omega) = \det(\mathbf{I}_N - \Omega\Omega^\dagger)^{-2N_\beta} \{d\Omega\}$. Here, $\{d\Omega\} = \prod_{ij} \Re d\Omega_{ij} \Im d\Omega_{ij}$ where \Re and \Im denote the real and imaginary parts (the product is over $i \leq j$ if $\beta = 1$ and over all i, j if $\beta = 2$).

As shown in [3], a Riemannian Gaussian distribution on the space of block-Toeplitz covariance matrices Γ is just a product of independent Riemannian Gaussian distributions, one for each coordinate $(\Gamma_0, \Omega_1, \dots, \Omega_{T-1})$. For Γ_0 , this is a Riemannian Gaussian distribution on \mathcal{P}_N^2 , already considered in Section II. Thus, to understand Riemannian Gaussian distributions of block-Toeplitz covariance matrices, it only remains to study Riemannian Gaussian distributions on the Siegel domain \mathcal{D}_N^β .

In Section III, it was seen that Riemannian Gaussian distributions on \mathcal{P}_N^β are equivalent to log-normal matrix ensembles. On the other hand, Riemannian Gaussian distributions on \mathcal{D}_N^β are equivalent to “acosh-normal” ensembles, which will be described in Proposition 4 below.

Note first that each matrix $\Omega \in \mathcal{D}_N^\beta$ can be factorised in the following way

$$\Omega = U\lambda V^\dagger \quad (\text{if } \beta = 2) \quad \Omega = U\lambda U^T \quad (\text{if } \beta = 1) \quad (24)$$

where U and V are unitary, T denotes the transpose, and λ is diagonal, with diagonal elements $\lambda_i = \tanh(r_i)$ for $(r_1, \dots, r_N) \in \mathbb{R}^N$. The $\beta = 2$ case follows from the singular value decomposition of Ω , and the $\beta = 1$ case follows from the Takagi decomposition of Ω [14]. In either case, λ_i is of the form $\tanh(r_i)$ because the singular values of Ω are all < 1 .

Proposition 4. *Let Ω follow a Riemannian Gaussian distribution on the Siegel domain \mathcal{D}_N^β with $\bar{\Omega} = 0_N$ ($N \times N$ zero matrix). If Ω is factorised as in (24), then U and V are uniformly distributed on the unitary group $U(N)$. Moreover, if $x_i = \cosh(2r_i)$ then (x_1, \dots, x_N) have joint probability density function*

$$p(x|\sigma) \propto |V(x)|^\beta \prod_i \rho(x_i, 8\sigma^2) \quad (x_1, \dots, x_N) \in (1, \infty)^N \quad (25)$$

where $V(x) = \prod_{i < j} (x_j - x_i)$, and where $\rho(x, k) = \exp(-\text{acosh}^2(x)/k)$ is the “acosh-normal” weight function.

Note that the assumption that the centre-of-mass parameter $\bar{\Omega}$ is equal to 0_N does not entail any loss of generality. The transformation [22]

$$\Omega \mapsto \Psi(\Omega) = (\mathbf{I}_N - \bar{\Omega}\bar{\Omega}^\dagger)^{-\frac{1}{2}} (\Omega - \bar{\Omega})(\mathbf{I}_N - \bar{\Omega}^\dagger\Omega)^{-1} (\mathbf{I}_N - \bar{\Omega}^\dagger\bar{\Omega})^{\frac{1}{2}} \quad (26)$$

maps $\bar{\Omega}$ to 0_N , while preserving the Siegel domain metric (21) and the associated distance and Riemannian volume. Thus, if Ω follows a Riemannian Gaussian density $p(\Omega|\bar{\Omega}, \sigma)$, it is enough to replace Ω by $\Psi(\Omega)$, which will have a Riemannian Gaussian density with $\bar{\Omega} = 0_N$ and with the same σ .

Proposition 4 implies that the transformed singular values $x_i = \cosh(2r_i)$ follow the classical eigenvalue distribution of an orthogonal ($\beta = 1$) or unitary ($\beta = 2$) matrix ensemble with “acosh-normal” weight function. Therefore, in particular, the normalising factor $Z(\sigma)$ of (23) reduces to a multiple integral (compare to (6) and (10))

$$z_\beta(\sigma) = \frac{1}{N!} \int_{(1,\infty)^N} |V(x)|^\beta \omega(dx) \quad (27)$$

where $\omega(dx) = \prod \rho(dx_i)$, with $\rho(dx_i) = \rho(x_i, 8\sigma^2)dx_i$ for the weight function $\rho(x, k) = \exp(-\text{acosh}^2(x)/k)$.

Thus, Riemannian Gaussian distributions on the Siegel domain \mathcal{D}_N^β can be studied by applying the tools of random matrix theory to new “acosh-normal” matrix ensembles. Hopefully, in the near future, this will lead to similar results to the ones obtained with log-normal ensembles in Section III, paving the way to learning from datasets of high-dimensional block-Toeplitz covariance matrices.

APPENDIX A

PROOFS OF PROPOSITIONS 1 TO 4

A. Proof of Proposition 1

If Y follows the Riemannian Gaussian density (3) with $\bar{Y} = \mathbf{I}_N$, then

$$\mathbb{P}(Y \in B) = (Z(\sigma))^{-1} \int_B \exp \left[-\frac{d^2(Y, \mathbf{I}_N)}{2\sigma^2} \right] dv(Y) \quad (28)$$

Let (y_1, \dots, y_N) denote the eigenvalues of Y . Using (2) and the fact that $dv(Y) = \det(Y)^{-N_\beta} \{dY\}$, (28) becomes

$$\mathbb{P}(Y \in B) = (Z(\sigma))^{-1} \int_B \left(\prod_{i=1}^N \exp \left[-\frac{\log^2(y_i)}{2\sigma^2} \right] y_i^{-N_\beta} \right) \{dY\} \quad (29)$$

Recall that $X = e^{N_\beta \sigma^2} Y$. Accordingly, if (x_1, \dots, x_N) are the eigenvalues of X , an elementary calculation yields

$$\exp \left[-\frac{\log^2(y_i)}{2\sigma^2} \right] y_i^{-N_\beta} = \exp[N_\beta^2(\sigma^2/2)] \times \exp \left[-\frac{\log^2(x_i)}{2\sigma^2} \right]$$

Therefore, (29) can be written

$$\mathbb{P}(Y \in B) = \exp[NN_\beta^2(\sigma^2/2)] \times (Z(\sigma))^{-1} \int_B \text{etr} \left[-\frac{\log^2(X)}{2\sigma^2} \right] \{dY\} \quad (30)$$

To conclude, it is enough to use once more the definition $X = e^{N_\beta \sigma^2} Y$. This implies

$$\{dY\} = \exp[-2NN_\beta^2(\sigma^2/2)] \{dX\}$$

Thus, using the fact that

$$\mathbb{P}(X \in B) = \mathbb{P}(e^{N_\beta \sigma^2} Y \in B)$$

and changing the variable of integration from Y to X in (30), it follows that

$$\mathbb{P}(X \in B) = \exp[-NN_\beta^2(\sigma^2/2)] \times (Z(\sigma))^{-1} \int_B \text{etr} \left[-\frac{\log^2(X)}{2\sigma^2} \right] \{dX\}$$

as required in (8).

B. Proof of Proposition 2

The $\beta = 2$ case follows from (18), by an elementary calculation, after noting that

$$\frac{1}{N^2} \sum_{n=1}^{N-1} (N-n) \log(1 - e^{-n\sigma^2})$$

is a Riemann sum for the improper integral

$$\int_0^1 (1-x) \log(1 - e^{-tx}) dx = \frac{\text{Li}_3(1) - \text{Li}_3(e^{-t})}{t^2}$$

For other values of β , the result will be obtained from the reasoning presented in the proof of Proposition 3, based on the scaling equation (33).

C. Proof of Proposition 3

The $\beta = 2$ case has already been proved in [20]. To deal with the general case, write (10) under the form $z_\beta(\sigma) = c_\beta(\sigma) I_\beta(\sigma)$, where

$$c_\beta(\sigma) = \frac{1}{N!} \exp[-N N_\beta^2(\sigma^2/2)]$$

and $I_\beta(\sigma)$ is the multiple integral

$$I_\beta(\sigma) = \int_{\mathbb{R}_+^N} \exp \left[-\frac{1}{2\sigma^2} \sum_{i=1}^N \log^2(x_i) + \beta \sum_{i < j} \log|x_j - x_i| \right] dx$$

Now, if $\mu_N = (1/N) \sum_{i=1}^N \delta_{x_i}$ is the empirical distribution of x_i (where δ_{x_i} denotes the Dirac measure at x_i), then

$$I_\beta(\sigma) = \int_{\mathbb{R}_+^N} \exp[-N^2 E_\beta(\mu_N, t)] dx \quad (31)$$

where $E_\beta(\mu, t)$ is the so-called energy functional

$$E_\beta(\mu, t) = \frac{1}{2Nt} \int_{\mathbb{R}_+} \log^2(x) \mu(dx) - \beta \int_{\mathbb{R}_+} \int_{\mathbb{R}_+} \log|y - x| \mu(dx) \mu(dy) \quad (32)$$

defined for any probability distribution μ on \mathbb{R}_+ . This energy functional satisfies the following scaling equation

$$E_\beta(\mu, t) = (\beta/2) E_2(\mu, (\beta/2)t) \quad (33)$$

which is easily obtained after dividing (32) by $\beta/2$. The proof can now be completed by applying to (31) and (33) the arguments in [17]. First [17] (Corollary 6.90, Page 155),

$$\frac{1}{N^2} \log I_\beta(\sigma) \longrightarrow -E_\beta(\mu^*, t)$$

where $\mu^* = \mu_\beta^*(t)$ is the so-called equilibrium distribution, the unique minimiser of the energy (32) among probability distributions on \mathbb{R}_+ . From the scaling equation (33), it is now clear

$$\frac{1}{N^2} \log I_\beta(\sigma) \longrightarrow -(\beta/2) E_2(\mu^*, (\beta/2)t) \quad (34)$$

Therefore, by adding the limit of $\log c_\beta(\sigma)/N^2$,

$$\frac{1}{N^2} \log z_\beta(\sigma) \longrightarrow -(\beta/2) ((\beta t/4) + E_2(\mu^*, (\beta/2)t)) \quad (35)$$

However, when $\beta = 2$, it is already known that this limit is $\Phi(t)$, where Φ is defined in (11). This provides $E_2(\mu^*, t) = -t/2 - \Phi(t)$, which can be replaced back into (35), yielding the general case of Proposition 2.

To prove Proposition 3, it is possible to use the scaling equation (33), once more. From [17] (Section 6.4), if $\hat{\mu}_N$ is defined as in (14), but with the x_i instead of the y_i , then $\hat{\mu}_N$ converges weakly to the equilibrium distribution $\mu^* = \mu_\beta^*(t)$. On the other hand, the scaling equation (33) implies $\mu^* = \mu_2^*(\beta t/2)$, and $\mu_2(t)$ is already known to be the image of the distribution with density $n(y|t)$, defined as in (15), under the change of variables $x = e^{\frac{\beta}{2}t}y$. This provides μ^* for any value of β , and the Proposition then follows by changing the variables back from x to y .

D. Proof of Proposition 4

The proof relies on the general theory of Gaussian distributions on Riemannian symmetric spaces, as outlined in [3] [20]. Here, to keep the proof self-contained, it will be helpful to briefly recall certain aspects of this theory (for a more detailed, in-depth discussion, the reader is referred to [20], Sections 1.9 and 3.3).

A Gaussian distribution on a Riemannian symmetric space M is defined by its probability density function

$$p(x|\bar{x}, \sigma) = (Z(\sigma))^{-1} \exp \left[-\frac{d^2(x, \bar{x})}{2\sigma^2} \right] \quad x \in M \quad (36)$$

with respect to the Riemannian volume $dv(x)$ on M , with $\bar{x} \in M$ and $\sigma > 0$.

It is always assumed M is a Riemannian symmetric space of non-positive curvature, associated to a symmetric pair (G, K) . Precisely, G is a connected Lie group which acts transitively and isometrically on M (this action is denoted $x \mapsto g \cdot x$, for $g \in G$ and $x \in M$), and K is a compact subgroup of G , made up of those elements $k \in G$ which fix a certain point $o \in M$ (that is, $k \in K$ if and only if $k \cdot o = o$).

For a concrete understanding of Gaussian distributions on M , it is necessary to take a closer look at the Lie algebras of G and K , denoted \mathfrak{g} and \mathfrak{k} , respectively. These are related together by the so-called Cartan decomposition, $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$ (direct sum), where the subspace \mathfrak{p} of \mathfrak{g} can be identified with the tangent space $T_o M$ (tangent space to M at o). In terms of this decomposition, the main construction needed for the proof can be described as follows.

Without any loss of generality, G and K are taken to be matrix Lie groups, and \mathfrak{g} and \mathfrak{k} their matrix Lie algebras. Let \mathfrak{a} be a maximal abelian subspace of \mathfrak{p} (that is, all the matrices $a \in \mathfrak{a}$ commute with each other). Then, each matrix $u \in \mathfrak{p}$ can be written under the form $u = kak^{-1}$, where $k \in K$ and $a \in \mathfrak{a}$ (this factorisation is the general template for the fifty three matrix factorisations, outlined in [14]). Moreover, any $x \in M$ admits a representation

$$x = \exp(kak^{-1}) \cdot o \quad k \in K \text{ and } a \in \mathfrak{a} \quad (37)$$

where \exp denotes the matrix exponential. Incidentally, this representation is not unique, but for almost all $x \in M$ (for all $x \in M$, except a subset of zero volume) x has exactly $|W|$ couples (k, a) which satisfy (37), where $|W|$ is the number of elements of the Weyl group of the symmetric couple (G, K) .

It is now possible to state the following Lemma 1 [3] [20], which will yield the entire proof, by direct application.

Lemma 1. *Let x follow the Gaussian density (36) on M , with $\bar{x} = o$. If x is represented as in (37), then k is uniformly distributed on the compact group K . Moreover, a has the following probability density function on \mathfrak{a} (recall that \mathfrak{a} is a real vector space, so the density is with respect to the usual Lebesgue measure on \mathfrak{a})*

$$p(a|\sigma) \propto \exp \left[-\frac{\|a\|_o^2}{2\sigma^2} \right] \prod_{\rho \in \Delta_+} (\sinh |\rho(a)|)^{m_\rho} \quad (38)$$

where $\|a\|_o$ is the Riemannian norm of a (since $a \in \mathfrak{p}$, it can be identified with a vector in $T_o M$), and Δ_+ is a set of positive roots ρ on \mathfrak{a} , with respective multiplicities m_ρ (each $\rho \in \Delta_+$ is a certain linear function $\rho : \mathfrak{a} \rightarrow \mathbb{R}$).

Recall that positive roots $\rho : \mathfrak{a} \rightarrow \mathbb{R}$ are any set of linear functions on \mathfrak{a} such that, for any $a \in \mathfrak{a}$, the eigenvalues of the linear operator $\text{ad}_a : \mathfrak{g} \rightarrow \mathfrak{g}$, given by $\text{ad}_a(\xi) = [a, \xi]$ (this is $a\xi - \xi a$), are equal to $\pm\rho(a)$ with respective multiplicities m_ρ .

The proof may now begin in earnest. It merely consists in identifying G , K , \mathfrak{g} , \mathfrak{k} , \mathfrak{p} , and \mathfrak{a} , for each Siegel domain \mathcal{D}_N^β , and then writing down the corresponding version of Lemma 1, which directly yields Proposition 4. Fortunately, all the necessary information can be found in [23] [24].

The symmetric pair: \mathcal{D}_N^β is associated to the symmetric pair (G, K) , where G and K are groups of $2N \times 2N$ complex matrices g , defined in the following way (here, $*$ denotes the complex conjugate).

$$G_{\beta=2} = \left\{ g = \begin{pmatrix} A & B \\ C & D \end{pmatrix} : gPg^\dagger = P \right\} \quad ; \quad G_{\beta=1} = \left\{ g = \begin{pmatrix} A & B \\ C & D \end{pmatrix} : gPg^\dagger = P \text{ and } gSg^T = S \right\}$$

$$K_{\beta=2} = \left\{ k = \begin{pmatrix} U & \\ & V \end{pmatrix} : U, V \in U(N) \right\} \quad ; \quad K_{\beta=1} = \left\{ k = \begin{pmatrix} U & \\ & U^* \end{pmatrix} : U \in U(N) \right\}$$

where P and S denote the following matrices

$$P = \begin{pmatrix} I_N & 0_N \\ 0_N & -I_N \end{pmatrix} \quad S = \begin{pmatrix} 0_N & I_N \\ -I_N & 0_N \end{pmatrix}$$

These groups G and K act on \mathcal{D}_N^β by matrix fractional transformations

$$g \cdot \Omega = (A\Omega + B)(C\Omega + D)^{-1} \quad g \in G \text{ and } \Omega \in \mathcal{D}_N^\beta$$

The Lie algebras: these are given by

$$\mathfrak{g}_{\beta=2} = \{(\gamma, \delta, \varepsilon) : \gamma + \gamma^\dagger = \varepsilon + \varepsilon^\dagger = 0_N\} \quad ; \quad \mathfrak{g}_{\beta=1} = \{(\gamma, \delta, \varepsilon) : \gamma + \gamma^\dagger = \varepsilon + \varepsilon^\dagger = \delta - \delta^T = \gamma + \varepsilon^T = 0_N\}$$

$$\mathfrak{k}_{\beta=2} = \{(\gamma, 0_N, \varepsilon) : \gamma + \gamma^\dagger = \varepsilon + \varepsilon^\dagger = 0_N\} \quad ; \quad \mathfrak{k}_{\beta=1} = \{(\gamma, 0_N, \gamma^*) : \gamma + \gamma^\dagger = 0_N\}$$

in terms of the notation

$$(\gamma, \delta, \varepsilon) = \begin{pmatrix} \gamma & \delta \\ \delta^\dagger & \varepsilon \end{pmatrix}$$

The subspaces \mathfrak{p} and \mathfrak{a} : these are given by

$$\mathfrak{p}_{\beta=2} = \{(0_N, \delta, 0_N)\} \quad ; \quad \mathfrak{p}_{\beta=1} = \{(0_N, \delta, 0_N) : \delta - \delta^T = 0_N\}$$

$$\mathfrak{a}_{\beta=2} = \mathfrak{a}_{\beta=1} = \{a = (0_N, r, 0_N) : r \text{ real diagonal}\}$$

Lemma 1: the only information still needed is the positive roots. Matrix multiplication shows that for any $a \in \mathfrak{a}$, ad_a has the following eigenvalues and eigenvectors in $\mathfrak{g}_{\beta=2}$.

$$\xi = (\omega_{ij}, \tau_{ij}, \omega_{ij}) \implies \text{ad}_a(\xi) = (r_i - r_j)\xi \quad \xi = \mathfrak{i}(\tau_{ij}, \omega_{ij}, \tau_{ij}) \implies \text{ad}_a(\xi) = (r_i - r_j)\xi$$

$$\xi = \mathfrak{i}(-\tau_{ij}, \tau_{ij}, \tau_{ij}) \implies \text{ad}_a(\xi) = (r_i + r_j)\xi \quad \xi = (\omega_{ji}, \omega_{ij}, \omega_{ij}) \implies \text{ad}_a(\xi) = (r_i + r_j)\xi$$

where $\tau_{ij} = e_{ij} + e_{ji}$ and $\omega_{ij} = e_{ij} - e_{ji}$, with e_{ij} a matrix all of whose entries are zero, except the one on line i and column j , which is equal to 1, and where $r = \text{diag}(r_1, \dots, r_N)$.

This shows that the positive roots are $\rho(a) = r_i - r_j$ where $i < j$ and $\rho(a) = r_i + r_j$ where $i \leq j$, which all have multiplicity $m_\rho = 2$ (this is β), except for $\rho(a) = 2r_i$ (this is $r_i + r_j$ when $i = j$), which has multiplicity $m_\rho = 1$. Moreover, of the above eigenvectors, only the ones in the left column belong to $\mathfrak{g}_{\beta=1}$. Thus, in the $\beta = 1$ case, all the multiplicities m_ρ are equal to 1. Finally, note that, from the power series of the matrix exponential,

$$\exp(a) = \begin{pmatrix} \cosh(r) & \sinh(r) \\ \sinh(r) & \cosh(r) \end{pmatrix} \quad a = (0_N, r, 0_N) \quad (39)$$

From the above form of the action of G on \mathcal{D}_N^β , it is then straightforward that the representation (37) is the same as (24). This implies the first part of Proposition 4 (U, V are uniformly distributed on $U(N)$). For the second part, it follows from (38) that, in the present case,

$$p(r|\sigma) \propto \prod_i \exp\left[-\frac{r_i^2}{2\sigma^2}\right] \sinh(2r_i) \prod_{i < j} (\sinh|r_i - r_j| \sinh|r_i + r_j|)^\beta \quad (40)$$

This yields (25), by using $2 \sinh|r_i - r_j| \sinh|r_i + r_j| = \cosh(2r_i) - \cosh(2r_j)$, and introducing the change of variables $x_i = \cosh(2r_i)$.

ACKNOWLEDGMENT

S.H. is supported by the Science and Technology Facilities Council (STFC) and St. John’s College, Cambridge. S.H. also benefited from partial support from the Cambridge Mathematics Placement (CMP) Programme and the European Research Council under the Advanced ERC Grant Agreement Switchlet n.670645. C.M. was supported by a Henslow Fellowship from the Cambridge Philosophical Society as well as a Presidential Postdoctoral Fellowship at NTU and is grateful for support from Fitzwilliam College, Cambridge.

REFERENCES

- [1] 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.
- [2] S. Said, L. Bombrun, Y. Berthoumieu, and J. H. Manton, “Riemannian Gaussian distributions on the space of symmetric positive definite matrices,” *IEEE Transactions on Information Theory*, vol. 63, no. 4, pp. 2153–2170, 2017.
- [3] S. Said, H. Hajri, L. Bombrun, and B. C. Vemuri, “Gaussian distributions on Riemannian symmetric spaces: statistical learning with structured covariance matrices,” *IEEE Transactions on Information Theory*, vol. 64, no. 2, pp. 752–772, 2018.
- [4] P. Zanini, M. Congedo, C. Jutten, S. Said, and Y. Berthoumieu, “Transfer learning: A Riemannian geometry framework with applications to brain-computer interfaces,” *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 5, pp. 1107–1116, 2018.
- [5] E. Mathieu, C. Le Lan, C. Maddison, R. Tomioka, and Y. W. Teh, “Continuous hierarchical representations with Poincaré variational auto-encoders,” in *Neural Information Processing Systems (NIPS)*, 2019.
- [6] S. Heuveline, S. Said, and C. Mostajeran, “Gaussian distributions on Riemannian symmetric spaces in the large N limit,” in *Geometric Science of Information*, F. Nielsen and F. Barbaresco, Eds. Cham: Springer International Publishing, 2021, pp. 20–28.
- [7] M. Mariño, *Chern-Simons theory, matrix models, and topological strings*. Oxford University Press, 2005.
- [8] L. Santilli and M. Tierz, “Riemannian Gaussian distributions, random matrix ensembles and diffusion kernels,” *Nucl. Phys. B*, vol. 973, p. 115582, 2021.
- [9] P. J. Forrester, “Vicious random walkers in the limit of a large number of walkers,” *Journal of Statistical Physics*, vol. 56, no. 5, pp. 767–782, 1989. [Online]. Available: <https://doi.org/10.1007/BF01016779>
- [10] —, “Global and local scaling limits for the $\beta=2$ stieljes–wigert random matrix ensemble,” *Random Matrices: Theory and Applications*, p. 2250020, 2021.
- [11] S. I. Amari, *Information geometry and its applications*. Springer Publishing Company, 2016.
- [12] X. Pennec, “Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements,” *Journal of Mathematical Imaging and Vision*, vol. 25, no. 1, pp. 127–154, 2006.
- [13] X. Pennec, S. Sommer, and T. Fletcher, *Riemannian geometric statistics in medical image analysis*. Elsevier, 2020.
- [14] A. Edelman and S. Jeong, “Fifty three matrix factorizations: A systematic approach,” *ArXiv*, vol. abs/2104.08669, 2021.
- [15] E. S. Meckes, *The random matrix theory of the classical compact groups*. Cambridge University Press, 2019.
- [16] S. Sommer, “Anisotropic distributions on manifolds: Template estimation and most probable paths,” in *Information Processing in Medical Imaging*, S. Ourselin, D. C. Alexander, C.-F. Westin, and M. J. Cardoso, Eds. Cham: Springer International Publishing, 2015, pp. 193–204.
- [17] P. Deift, *Orthogonal polynomials and random matrices: a Riemann-Hilbert approach*, American Mathematical Society, 1998.
- [18] S. Said, N. L. Bihan, and J. H. Manton, “Hidden Markov chains and fields with observations in Riemannian manifolds,” in *Mathematical Theory of Networks and Systems*, 2020.
- [19] Q. Tupker, S. Said, and C. Mostajeran, “Online learning of Riemannian hidden Markov models in homogeneous Hadamard spaces,” in *Geometric Science of Information*, F. Nielsen and F. Barbaresco, Eds. Cham: Springer International Publishing, 2021, pp. 37–44.
- [20] S. Said, “Statistical models and probabilistic methods on Riemannian manifolds,” 2021, Université de Bordeaux, Habilitation thesis (HDR).
- [21] B. Jeuris and R. Vandebril, “The Kähler mean of block Toeplitz matrices with Toeplitz structured blocks,” *SIAM Journal on matrix analysis and applications*, vol. 37, pp. 1151–1175, 2016.
- [22] Y. Cabanes, “Multidimensional complex stationary centred Gaussian autoregressive model classification: Applications for audio and radar machine learning in hyperbolic and Siegel spaces,” Ph.D. dissertation, Institut de Mathématiques de Bordeaux, 2022, M. Arnaudon, F. Barbaresco, J. Bigot (supervisors).

- [23] I. I. Piatetskii-Shapiro, *Automorphic functions and the geometry of classical domains*. Gordon and Breach Science Publishers, 1969.
- [24] S. Helgason, *Differential geometry and symmetric spaces*. New York and London: Academic Press, 1962.