

Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces II: non-compact symmetric spaces

Iskander Azangulov*

ISKA.AZN@GMAIL.COM

St. Petersburg State University and University of Oxford

Andrei Smolensky*

ANDREI.SMOLENSKY@GMAIL.COM

St. Petersburg State University and Neapolis University Pafos

Alexander Terenin

AVT28@CORNELL.EDU

University of Cambridge and Cornell University

Viacheslav Borovitskiy

VIACHESLAV.BOROVITSKIY@GMAIL.COM

ETH Zürich

Editor: Philipp Hennig

Abstract

Gaussian processes are arguably the most important class of spatiotemporal models within machine learning. They encode prior information about the modeled function and can be used for exact or approximate Bayesian learning. In many applications, particularly in physical sciences and engineering, but also in areas such as geostatistics and neuroscience, invariance to symmetries is one of the most fundamental forms of prior information one can consider. The invariance of a Gaussian process' covariance to such symmetries gives rise to the most natural generalization of the concept of stationarity to such spaces. In this work, we develop constructive and practical techniques for building stationary Gaussian processes on a very large class of non-Euclidean spaces arising in the context of symmetries. Our techniques make it possible to (i) calculate covariance kernels and (ii) sample from prior and posterior Gaussian processes defined on such spaces, both in a practical manner. This work is split into two parts, each involving different technical considerations: part I studies compact spaces, while part II studies non-compact spaces possessing certain structure. Our contributions make the non-Euclidean Gaussian process models we study compatible with well-understood computational techniques available in standard Gaussian process software packages, thereby making them accessible to practitioners.

Keywords: Gaussian processes, kernels, geometric learning, stationarity, symmetries, Lie groups, homogeneous spaces, Riemannian manifolds.

1. Introduction

This work is the second part of a two-part series of papers that aims at bringing Gaussian process regression into the machine learning toolbox for a large class of Riemannian manifolds admitting analytical descriptions. To do so, in part I (Azangulov et al., 2024) we studied a rich class of compact manifolds, including compact *Lie groups* such as the special orthogonal group $SO(n)$ and the special unitary group $SU(n)$, as well as their *homogeneous spaces*, such as the sphere S_n and Stiefel manifold $V(k, n)$. We developed computational techniques

*Joint first author.

based on suitable generalizations of Bochner’s Theorem due to Yaglom (1961), which express stationary kernels on a compact homogeneous space $X = G/H$ as series in terms of *characters* of irreducible unitary representations of the group G , or the respective *zonal spherical functions*. These expressions generalize a number of previous works devoted to extending Gaussian processes to non-Euclidean spaces (Lindgren et al., 2011; Coveney et al., 2020; Borovitskiy et al., 2020, 2021, 2023; Hutchinson et al., 2021; Da Costa et al., 2023), and give a comprehensive mathematical treatment on how to implement stationary kernels and Gaussian processes on compact homogeneous spaces. These spaces occur in many application areas ranging from neuroscience (Jensen et al., 2020) to dynamical systems and robotics (Jaquier et al., 2020, 2022), as well as many others, and provide a rich setting for studying how geometry affects machine learning algorithms such as kernel methods (Rosa et al., 2023). For further motivation and a list of applications and examples, please see the introduction to part I.

Many application areas, such as for instance robotics and neuroscience, give rise to *non-compact* manifolds such as hyperbolic spaces, and spaces of positive definite matrices (Jaquier et al., 2022, 2024; Nejatbakhsh et al., 2024). For these, comprehensive computational techniques suitable for performing Gaussian process regression, and downstream techniques that rely on it such as Bayesian optimization, have yet to be developed. Our goal in this work is to develop them.

One of the key challenges when G is non-compact is that one must work with geometric analogs of the Fourier transform, rather than of Fourier series. Moreover, the setting is delicate: one can show that (i) stationary kernels on general non-compact non-abelian Lie groups usually do not exist, in an appropriate sense, and (ii) non-compact abelian Lie groups, provided they are connected, are isomorphic to products of Euclidean spaces with tori, which can already be handled by the combination of classical Euclidean techniques and compactness-based ideas developed in part I. Then, a product of kernels can be used to define a reasonable kernel on the corresponding product space, obtaining a reasonable Gaussian process model. This motivates us to work with a subset of homogeneous spaces, namely *symmetric spaces*, which are in some sense the right class of spaces on which one can reasonably expect Bochner-like results to hold in a concrete-enough manner to be amenable to numerical methods.

As in part I, we will begin with an abstract description of stationary kernels, and develop numerical techniques to make these kernels into concrete and practically viable machine learning models. Specifically, we develop the Bochner-type results of Yaglom (1961) into concrete numerical methods, using the *spherical Fourier transform* of Helgason (1962) for symmetric spaces. Our developments reveal that such kernels admit a maximum correlation value, meaning that their respective Gaussian processes are non-constant in the large length-scale limit—in sharp contrast with what happens in the Euclidean and compact settings.

For the first time, our work enables one to efficiently implement both the squared exponential (also known as heat, diffusion, RBF) and Matérn kernel classes in, for instance, hyperbolic space and the space of positive definite matrices. We study these cases in additional detail, and propose specific techniques that leverage their individual properties to enable more efficient approximations, compared to the general case. These developments reveal

unanticipated connections between geometric kernels and random matrix theory. As part of our work, we contribute implementations of these ideas to the `GEOMETRICKERNELS`¹ library.

Notation Following part I, we use bold italics (\mathbf{a}, \mathbf{b}) to simultaneously refer to finite-dimensional vectors in \mathbb{R}^n , and to elements of the product space X^n for a given set X . We use bold upface letters (\mathbf{A}, \mathbf{B}) to refer to matrices. If $f : X \rightarrow Y$ is a function, then for $\mathbf{x} \in X^n$ we write $f(\mathbf{x})$ to refer to the vector $(f(x_1), \dots, f(x_n))^\top$. For a two-argument function $k : A \times B \rightarrow C$, and $\mathbf{a} \in A^n$, $\mathbf{b} \in B^m$, the notation $\mathbf{K}_{\mathbf{a}\mathbf{b}}$ refers to the $n \times m$ matrix with entries $k(a_\ell, b_{\ell'})$ for $1 \leq \ell \leq n$ and $1 \leq \ell' \leq m$. We refer to conjugation of complex numbers by \bar{x} , and to the conjugate transpose of a matrix by $\mathbf{A}^* = \overline{\mathbf{A}^\top}$.

1.1 Gaussian Processes and Gaussian Process Regression

A *Gaussian process* f is a random function mapping some set X to the reals, whose finite dimensional distributions are all jointly Gaussian. One can show that every Gaussian process is uniquely determined by its mean function m and covariance kernel k given by

$$m(x) = \mathbb{E} f(x) \qquad k(x, x') = \text{Cov}(f(x), f(x')). \quad (1)$$

Gaussian processes are among the most commonly used priors for unknown functions in Bayesian learning. If $f \sim \text{GP}(0, k)$ and we observe the dataset \mathbf{x}, \mathbf{y} , where $y_i = f(x_i) + \varepsilon_i$ with $\varepsilon \sim \text{N}(\mathbf{0}, \Sigma_\varepsilon)$, then the posterior $f \mid \mathbf{y}$ is also a Gaussian process admitting the explicit form (Chiles and Delfiner, 2009; Wilson et al., 2020, 2021)

$$(f \mid \mathbf{y})(\cdot) = f(\cdot) - \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \Sigma_\varepsilon)^{-1}(\mathbf{y} - f(\mathbf{x}) - \varepsilon) \qquad \varepsilon \sim \text{N}(\mathbf{0}, \Sigma_\varepsilon) \quad (2)$$

to which we refer to as *pathwise conditioning*. From this expression, one can compute the posterior’s mean and covariance, obtaining predictions and their associated uncertainty.

By using (2), one can obtain samples from the posterior by generating samples from the prior and substituting them into (2) as $f(\cdot)$ and $f(\mathbf{x})$. Using this, if one is able to approximately sample from the prior f with a given computational complexity with respect to the number of evaluation locations, it follows one can approximately sample from the posterior with the same evaluation complexity.

1.2 Gaussian Processes That are Stationary Under Group Action

Suppose there is a group² (G, \cdot) acting on X , with group action $\triangleright : G \times X \rightarrow X$. A Gaussian process $f \sim \text{GP}(0, k)$ on X is called *stationary* if its distribution is invariant under group action. This means that for every finite set of points $x_1, \dots, x_m \in X$ and all $g \in G$, the random vectors $(f(x_1), \dots, f(x_m))$ and $(f(g \triangleright x_1), \dots, f(g \triangleright x_m))$ have the same distribution. It is easy to see that this is equivalent to

$$k(g \triangleright x, g \triangleright x') = k(x, x'). \quad (3)$$

When a kernel k satisfies the property above, it is termed *stationary* with respect to the group G and action \triangleright .

1. Available at [HTTPS://GEOMETRIC-KERNELS.GITHUB.IO](https://geometric-kernels.github.io). A prototypical software implementation, as well as our experiments, can be found at [HTTPS://GITHUB.COM/IMBIRIK/LIESTATIONARYKERNELS](https://github.com/imbirik/liestationarykernels).
 2. See Kondor (2008) or Robinson (2003) for a mild or more technical introduction to groups, respectively.

Stationarity under group action includes the classical case $X = \mathbb{R}^n$, where $G = \mathbb{R}^n$ and the group action is $g \triangleright x = g + x$, but is much richer as it allows one to consider general spaces and groups, far beyond Euclidean shift-invariance. For example, one obtains isotropic Gaussian processes by replacing the translation group $G = \mathbb{R}^n$ with the Euclidean group $G = E(n)$ acting on \mathbb{R}^n with rotations and reflections in addition to translations.

For stationary processes under the action of sufficiently regular and rich group of symmetries, one can show Bochner-type theorems. In the Euclidean case, this is

$$k(\mathbf{x}, \mathbf{x}') = \int_{\mathbb{R}^n} e^{2\pi i \boldsymbol{\lambda}^\top (\mathbf{x} - \mathbf{x}')} dS(\boldsymbol{\lambda}) = \int_{\mathbb{R}^n} a^{(\boldsymbol{\lambda})} e^{2\pi i \boldsymbol{\lambda}^\top (\mathbf{x} - \mathbf{x}')} d\boldsymbol{\lambda} \quad (4)$$

where, assuming S is absolutely continuous, $a^{(\boldsymbol{\lambda})}$ is the density of S , and $\int_{\mathbb{R}^n} a^{(\boldsymbol{\lambda})} d\boldsymbol{\lambda} < \infty$.

Another form of Bochner's theorem is available in the case of a compact *homogeneous space* X . Deferring a formal definition, we will for now say that this is a compact smooth manifold with a sufficiently-rich group G of symmetries acting on it. If we fix an arbitrary point $x_0 \in X$, which is analogous to a sphere's pole, and define $g_1, g_2 \in G$ in a way that ensures $g_1 \triangleright x_0 = x$ and $g_2 \triangleright x_0 = x'$ —where such g_1, g_2 exist because G is sufficiently rich—the appropriate analog of Bochner's theorem (Azangulov et al., 2024; Yaglom, 1961) says

$$k(x, x') = \sum_{\lambda \in \Lambda} \sum_{j,k=1}^{r_\lambda} a_{jk}^{(\lambda)} \operatorname{Re} \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (5)$$

where the matrices $\mathbf{A}^{(\lambda)}$ with entries $a_{jk}^{(\lambda)}$ must be positive semi-definite and satisfy the condition $\sum_{\lambda \in \Lambda} \operatorname{tr} \mathbf{A}^{(\lambda)} < \infty$. The two key changes compared to the Euclidean case are: (i) instead of an integral, we have a sum over the countable set Λ of irreducible unitary representations of the group G , where the corresponding weights $\mathbf{A}^{(\lambda)}$ can be understood as a kind of spectral measure, which is now *discrete* and *matrix-valued*, and (ii) the complex exponentials are replaced with the *zonal spherical functions* $\pi_{jk}^{(\lambda)} : G \rightarrow \mathbb{C}$. The zonal spherical functions $\pi_{jk}^{(\lambda)}$ and the values r_λ are determined by the space X itself, while Λ is determined solely by the group G .

The expression (5) simplifies greatly when $r_\lambda \leq 1$, and the inner sum either collapses to one term, or vanishes. In this case, the matrices $\mathbf{A}^{(\lambda)}$ simplify to scalars $a^{(\lambda)}$, which now behave much like a discrete analog of the spectral measures in the Euclidean case. This simplification occurs for spaces termed *symmetric spaces* (Yaglom, 1961)—these include compact Lie groups acting on themselves from both sides (Terras, 2016), as well as homogeneous spaces of constant sectional curvature, such as for instance spheres or real projective spaces (Yaglom, 1961), or hyperbolic spaces in the non-compact setting.

In the non-compact case, general homogeneous spaces are much less well-behaved than in the compact case. Here, even in the seemingly-simplest case of X being a regular non-compact non-abelian Lie group, non-trivial stationary Gaussian processes do not exist.

Result 1. *Let G be a non-compact simple³ non-abelian Lie group. Then the only kernels on G stationary with respect to both left and right action of G on itself are identically constant.*

3. *Simple Lie groups* are the basic building blocks for general connected Lie groups. See the formal definition in Appendix A.

Proof sketch. Here, we sketch the argument, without introducing the rather heavy abstract machinery needed to make the statement mathematically rigorous: the omitted technicalities may be recovered from Yaglom (1961, pp. 601–602). Recall that every Lie group can be viewed as a homogeneous space, where $H = \{e\}$ is the trivial group. Under this viewpoint, one can show that every left-stationary kernel must satisfy an analog of (5), but where sums are replaced with abstract integrals, and $\mathbf{A}^{(\lambda)}$ is replaced with a suitable infinite-dimensional generalization. A mirrored expression holds for right-stationary kernels, which forces us to choose a form akin to $\mathbf{A}^{(\lambda)} = a^{(\lambda)}\mathbf{I}$, but this results in a statement similar to

$$\sum_{\lambda \in \Lambda} \text{tr } \mathbf{A}^{(\lambda)} = \sum_{\lambda \in \Lambda} a^{(\lambda)} r_{\lambda} < \infty. \quad (6)$$

When a homogeneous space is a group, each r_{λ} is equal to the dimension d_{λ} of the unitary representation corresponding to λ . Crucially, for non-compact simple non-abelian Lie groups, all unitary representations except the trivial one are infinite-dimensional (Segal and von Neumann, 1950). This implies $r_{\lambda} = d_{\lambda} = \infty$ for all λ corresponding to non-trivial representations, making it impossible for (6) to hold unless $a^{(\lambda)} = 0$ holds for them. Since the only remaining representation is the trivial representation, setting a non-zero $a^{(\lambda)}$ for this representation gives the class of constant kernels. \square

If we consider symmetric spaces, then $r_{\lambda} \leq 1$ holds even in the non-compact setting. This enables a clean generalization of Bochner’s theorem. Moreover, many interesting non-compact homogeneous spaces, such as hyperbolic space and certain spaces of matrices, are symmetric. This makes one wonder: are non-compact symmetric spaces a good setting in which to develop numerical techniques needed to use Gaussian-process-based models? We proceed to answer this question affirmatively in the sequel.

2. Stationary Gaussian Processes on Symmetric Spaces

To begin, we recall the formal definition of *homogeneous spaces*, which can be formulated in either geometric or algebraic terms. Geometrically, a homogeneous space is a smooth manifold⁴ X , equipped with a transitively acting Lie group G consisting of diffeomorphisms of X . Here, *transitive* means that for all $x, x' \in X$ there is a $g \in G$ with $g \triangleright x = x'$. We call G the group of symmetries of X . X can be represented as the quotient $X = G/H$ where $H = \{g \in G : g \triangleright x_0 = x_0\}$ for an arbitrary $x_0 \in X$ is termed the *isotropy group* (Arvanitogeorgos, 2003, Proposition 4.2). This implies that elements of X can be regarded as *cosets* $g \cdot H$, $g \in G$: we can express $x \in G/H$ as $g \cdot H$ where $g \in G$ satisfies $g \triangleright x_0 = x$. This is the algebraic definition of a homogeneous space: a quotient of a Lie group G by any of its subgroups $H \subseteq G$. There is a canonical smooth manifold structure on any such quotient (Arvanitogeorgos, 2003, Proposition 4.1) and one can show that G acts transitively on G/H .

A manifold X can potentially be given the structure of a homogeneous space in more than one way, by choosing different groups G . For example, consider the Euclidean group $E(n)$, acting transitively on \mathbb{R}^n by translations, rotations, and reflections. If we fix $\mathbf{0} \in \mathbb{R}^n$ as x_0 , its isotropy group will be exactly the group $O(n)$ generated by rotations and reflections, thus $\mathbb{R}^n = E(n)/O(n)$. On the other hand, the subgroup of $E(n)$ that consists of translations

4. Such a manifold often possesses a canonical (pseudo-)Riemannian structure, which we will discuss later.

alone already acts transitively on \mathbb{R}^n , and is isomorphic to the addition group of \mathbb{R}^n itself, meaning that $\mathbb{R}^n = \mathbb{R}^n / \{\mathbf{0}\}$ where $\{\mathbf{0}\}$ is the trivial subgroup of \mathbb{R}^n consisting of only the identity element.

We now define *symmetric spaces*. Geometrically, a symmetric space is a Riemannian manifold whose isometry group contains an inversion symmetry about every point. This means that for every $x \in X$ there is an isometry $g_x : X \rightarrow X$ which inverts the direction of all of the geodesics passing through x . From this definition, one can show that a symmetric space is necessarily also a homogeneous space (Arvanitogeorgos, 2003, Theorem 6.2). A symmetric space can also be defined algebraically as a homogeneous space G/H where G is a connected Lie group for which there exists an automorphism $s : G \rightarrow G$ which is an involution, namely $s(s(g)) = g$ for all $g \in G$, and H is the union of connected components of the subgroup $G^s = \{g \in G : s(g) = g\}$. These definitions are equivalent, provided the algebraically-defined symmetric space is equipped with a Riemannian metric which is left- G -invariant (Arvanitogeorgos, 2003, Theorem 6.3). In this work, we always assume conditions which imply $H = G^s$, or equivalently that $s(g) = g$ if and only if $g \in H$.⁵

We now consider the analog of Bochner’s Theorem for non-compact symmetric spaces. As noted before, if G is non-compact, one must work with general analogs of the Fourier transform instead of Fourier series. At best, one can therefore hope to represent stationary kernels through integrals over the index set Λ of irreducible unitary representations, which is uncountable. With these notions, one can obtain an integral expression for symmetric spaces which is similar to the Euclidean case, as long as $H = G^s$.

Theorem 2. *Let $X = G/H$ be a non-compact symmetric space, where G is a Lie group of type I.⁶ A Gaussian process $f \sim \text{GP}(0, k)$ is stationary on X if and only if k is of form*

$$k(x, x') = k(g_1 \cdot H, g_2 \cdot H) = \mathbb{k}(g_2^{-1} \cdot g_1) = \int_{\Lambda} \pi^{(\lambda)}(g_2^{-1} \cdot g_1) d\mu_k(\lambda) \quad (7)$$

for a function $\mathbb{k} : G \rightarrow \mathbb{R}$, where $g_1 \cdot H$ and $g_2 \cdot H$ are the cosets of x and x' , Λ is the set of all—not necessarily finite dimensional—irreducible unitary representations of G with $r_\lambda = 1$, $\pi^{(\lambda)}$ are the zonal spherical functions, and μ_k is a nonnegative finite measure over Λ .

Proof. Yaglom (1961, Theorem 6). □

Subsequently, when referring to irreducible unitary representations, we implicitly refer only to those for which $r_\lambda = 1$ —as opposed to those for which $r_\lambda = 0$. An illustration of a posterior Gaussian process, computed using approximations built from this result and presented in the sequel, is given in Figure 1. We proceed to study how to concretely calculate expressions of this type for various symmetric spaces of interest.

5. Note that Yaglom (1961) also works with this slightly more restrictive notion of a symmetric space.
 6. We do not dwell on the technical definition of a *type I* group, which involves C^* -algebraic considerations and implies that all irreducible unitary representations are uniquely characterized by characters, which is usually the case: our subsequent assumptions will imply G being type I. See Yaglom (1961, p. 599) and Naimark (1959) for details.

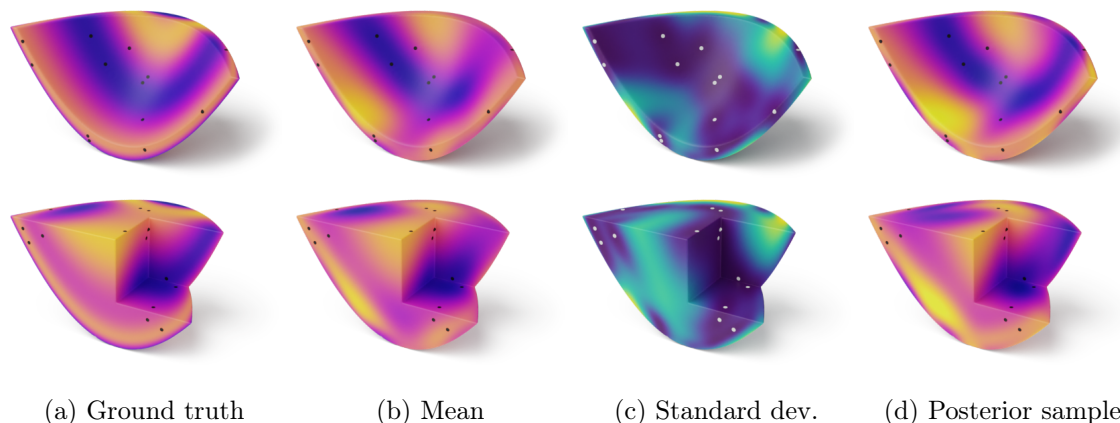


Figure 1: Posterior on the manifold $\text{SPD}(d)$ of symmetric positive definite $d \times d$ matrices, equipped with the affine-invariant metric. Here, $d = 2$ and $\text{SPD}(d)$ is represented as a cone: see Section 5.2. The top and the bottom rows correspond to two cross-sections thereof.

3. Computational Techniques for Stationary Gaussian Processes

We now develop computational techniques for working with the stationary Gaussian processes described in the preceding section. This involves the following two key computational primitives:

- (a) Evaluate the covariance kernel $k(\cdot, \cdot)$ pointwise.
- (b) Efficiently sample from the prior $f(\cdot)$.

Note that both (a) and (b) generally need to be performed in an automatically-differentiable manner, for instance to optimize model hyperparameters using maximum marginal likelihood. For a given set of evaluation locations, f can be sampled by forming and factorizing the kernel matrix at cubic cost with respect to the number of evaluation locations: we use the term *efficient sampling* to refer to sampling f approximately *without these costs*.

We focus on these primitives because Gaussian process priors are often used as building blocks for defining larger models. Here, most numerical methods require either (a) or (b) as part of their steps. For example, in Hamiltonian Monte Carlo, (a) is needed to compute the posterior density and evaluate the gradient of its logarithm. Similarly, in variational inference, (b) can be needed to compute stochastic estimators of expectations which appear in the variational objective. Computing Monte-Carlo-type acquisition functions (Wilson et al., 2018; Neiswanger et al., 2021) in Bayesian optimization, active learning, or other automated decision making applications further requires one to automatically differentiate both expressions—for instance, to perform Thompson sampling in an efficient and numerically stable manner using pathwise conditioning (Wilson et al., 2020, 2021; Terenin et al., 2024). These primitives therefore enable one to apply Gaussian processes as part of a wide class of learning and decision-making algorithms.

The general form of a stationary kernel on a symmetric space G/H under a non-compact group G is given by (7). From this, it is clear that to approximately compute such a kernel

pointwise it is enough to be able to sample from the measure μ_k and compute the zonal spherical functions $\pi^{(\lambda)}$. Letting $\sigma^2 = \mu_k(\Lambda)$, we obtain the random Fourier feature expansion

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{\sigma^2}{L} \sum_{l=1}^L \pi^{(\lambda_l)}(g_2^{-1} \cdot g_1) \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k. \quad (8)$$

To transform this into a bona-fide computational approach, we need to describe how to sample from μ_k , which is supported on the set of irreducible unitary representations Λ , and how to compute $\pi^{(\lambda)}$. Furthermore, since this expression does not directly factorize into an inner product, for efficient sampling and to guarantee positive semi-definiteness we will need to go beyond (8), leveraging specific algebraic properties of the zonal spherical functions $\pi^{(\lambda)}$.

The high-level idea will be to reduce the general case to the simplest possible non-compact setting, namely Euclidean space—mirroring how, in the compact case, the high-level idea was to reduce the general case to the simplest possible compact setting, namely the torus (Azangulov et al., 2024, Section 3.1). In this section, we assume henceforth the kernel and corresponding spectral measure μ_k are fixed: we study how to calculate spectral measures for specific kernel classes such as the Matérn and heat kernel classes in Section 4.

In Section 3.1, we formalize the intuition of the preceding paragraph. Then, in Section 3.2, we present a way to bijectively map a certain subset of Λ —which is sufficiently large to describe all square-integrable kernels—into a Euclidean space, and give an expression for the zonal spherical functions $\pi^{(\lambda)}$ —in total, this allows us to evaluate the kernel pointwise. Following this, in Section 3.3 we show how to leverage a certain very special property of the zonal spherical functions to obtain an efficient sampling strategy. As a byproduct, this allows us to build an alternative approximation to the kernel which, unlike the approximation we introduce in Section 3.2, is guaranteed to be positive semi-definite—an important property to ensure that downstream algorithms work well in practice.

Before starting this, we note that quotients G/H of a sufficiently regular Lie group G by a sufficiently nice subgroup H , including the symmetric spaces we study, admit a *canonical* Riemannian structure—something that will be essential for Section 4. To ensure this and to satisfy assumptions of the spherical Fourier transform theory of Helgason and Harish-Chandra—the theory that will enable us to make the abstract expressions we have seen so far into something concrete and amenable to computation—we assume the following technical conditions.

Assumption A. *In the symmetric space G/H , G is a SEMI-SIMPLE⁷ Lie group with FINITE CENTER, and H is a MAXIMAL COMPACT SUBGROUP of G .*

See Appendix A for definitions. In particular, hyperbolic spaces and the spaces of symmetric positive definite matrices with unit determinant arise in this way and satisfy the regularity assumptions. One can relax this assumption to situations where the symmetric space of interest is a Riemannian product of a compact space, Euclidean space, and symmetric space satisfying Assumption A: this factorization can be viewed as a crude and less-general analog of the basic classification results for symmetric spaces, see for instance Helgason (1962, Chapter V). We will use this factorization idea in our examples to handle the space of

7. Note that, if G is semi-simple, it is automatically of type I, see Harish-Chandra (1953) and Yaglom (1961).

symmetric positive definite matrices in Section 5.2, which does not satisfy Assumption A, but to ease presentation work with Assumption A otherwise.

3.1 Enumerating Representations of Symmetric Spaces

In the preceding section, we formulated a strategy for defining random Fourier features on symmetric spaces, where the high-level idea was to reduce the general case to the simplest possible non-compact setting—that is, Euclidean space, which is abelian. To actually carry this out, we will need a certain decomposition of G .

This is the *Iwasawa decomposition*, which can be viewed as a basis-free generalization of the well-known QR decomposition of a matrix. If $G = \text{GL}(\mathbb{R}^n)$ —we note that in this special case the result still holds, despite G not being semi-simple—then the Iwasawa decomposition can be reinterpreted as saying that the matrix corresponding to an arbitrary group element $g \in G$ can be represented as the product of a unit lower triangular matrix $N(g)$, positive diagonal matrix $A(g)$, and orthogonal matrix $H(g)$. The Iwasawa decomposition therefore generalizes the QR decomposition, up to transposition, with Q^\top corresponding to the product $N(g)A(g)$. A formal discussion of the Iwasawa decomposition can be found in Appendix A.

The key reason for introducing this decomposition in our setting is that it turns out, in order to compute a stationary kernel, it suffices—roughly speaking—to integrate only over the abelian, resp. diagonal, part of G in Iwasawa decomposition. For a Lie group G , define the *Lie algebra* \mathfrak{g} to be the tangent space at the identity element equipped with appropriate algebraic structure—see Appendix A for details. We now state the result.

Proposition 3. *Let G/H be a symmetric space satisfying Assumption A. Then there is an explicit surjective correspondence between a subset $\Lambda_A \subseteq \Lambda$ and the Euclidean space \mathfrak{a}^* , defined as the dual space of the Lie algebra \mathfrak{a} induced by the Lie group A of the abelian part of the Iwasawa decomposition of G . For any stationary k satisfying $\mathbb{k} \in L^2(G)$, we have $\text{supp } \mu_k \subseteq \Lambda_A$ and μ_k is absolutely continuous with respect to the Lebesgue measure on \mathfrak{a}^* .*

Proof. Follows from Proposition 4, which is presented subsequently. \square

This tells us what set to integrate over: to ease notation, where unambiguous we henceforth use λ to refer simultaneously to irreducible unitary representations $\lambda \in \Lambda$, and Lie algebra dual elements $\lambda \in \mathfrak{a}^*$. We proceed to understand how to carry out the necessary integration.

3.2 Pointwise Kernel Evaluation

We now leverage the Iwasawa decomposition to reduce the abstract integral of Theorem 2, which is supported on the set of irreducible unitary representations Λ , to a substantially more concrete integral over the space \mathfrak{a}^* introduced in Proposition 3. To do this, observe that the Lie algebra \mathfrak{a} is a vector space and hence can be equipped with an inner product. This inner product can be carried over to \mathfrak{a}^* , which can then be used to equip \mathfrak{a}^* with a Lebesgue measure. We can explicitly describe the integral of interest using this measure.

Proposition 4. *Let $\mathbb{k} \in L^2(G)$. Then there is a non-negative $a^{(\lambda)} \in L^2(\mathfrak{a}^*, |c(\lambda)|^{-2} d\lambda)$, where the L^2 -weight function is defined in terms of Harish-Chandra’s c -function, such that*

$$k(g_1 \cdot H, g_2 \cdot H) = \int_{\Lambda} \pi^{(\lambda)}(g_2^{-1} \cdot g_1) d\mu_k(\lambda) = \int_{\mathfrak{a}^*} a^{(\lambda)} \pi^{(\lambda)}(g_2^{-1} \cdot g_1) |c(\lambda)|^{-2} d\lambda \quad (9)$$




Figure 2: Values of the heat kernel $k(\bullet, \cdot)$ on \mathbb{H}_2 and $\text{SPD}(2)$.

where integration is defined with respect to the Lebesgue measure $d\lambda$, and $\pi^{(\lambda)}$ in the latter case is the zonal spherical function obtained via the bijective correspondence of Proposition 3.

Proof. This follows largely by specializing ideas of Helgason (2000) to the setting at hand: full details are given in Appendix B. \square

This substantially simplifies the matter: our kernel is now expressed as an ordinary integral over a Euclidean space. All that remains is to understand how to compute zonal spherical functions $\pi^{(\lambda)}$. One can show these functions are expressed by explicit integrals.

Result 5. For $\lambda \in \mathfrak{a}^*$, we have that

$$\pi^{(\lambda)}(g) = \int_H e^{(i\lambda + \rho)^\top a(h \cdot g)} d\mu_H(h), \quad (10)$$

where $a(g) \in \mathfrak{a}$ satisfy $A(g) = \exp a(g)$, the term $\exp : \mathfrak{g} \rightarrow G$ is the exponential map, $\rho \in \mathfrak{a}^*$ is an explicit vector given in terms of representation-theoretic quantities of G which for most cases can be looked up in the literature,⁸ and μ_H is the normalized Haar measure on H .

Proof. Helgason (2000, Chapter IV, Theorem 4.3). \square

This completes the description, and establishes a way to evaluate the kernel pointwise. An illustration of kernel evaluations, for the heat kernel defined in Section 4 and computed using this method, can be seen in Figure 2. For every $g \in G$ and $\lambda \in \mathfrak{a}^*$, using Monte Carlo we have the following approximation

$$\pi^{(\lambda)}(g) \approx \frac{1}{L} \sum_{l=1}^L e^{(i\lambda + \rho)^\top a(h_l \cdot g)} \quad h_l \sim \mu_H \quad (11)$$

where we note that random sampling from the Haar measure on a compact group H is generally tractable (Mezzadri, 2007). From this, we have the random Fourier feature (Rahimi and Recht, 2008) expansion

$$\mathbb{k}(g) \approx \frac{\sigma^2}{L} \sum_{l=1}^L e^{(i\lambda_l + \rho)^\top a(h_l \cdot g)} \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k \quad h_l \sim \mu_H \quad (12)$$

where $\mathbb{k}(e) = \sigma^2$, and e is the identity of G . We now quantify its approximation error.

8. More precisely, $\rho = \frac{1}{2} \sum_{\alpha \in \Sigma^+} m_\alpha \alpha$ is the half-sum of all positive roots α weighted by the dimensions m_α of the corresponding root spaces \mathfrak{g}_α —see the definitions of *roots* and *root spaces* as well as the appropriate references in Appendix A.




Figure 3: Samples from a Gaussian process with heat kernel covariance on \mathbb{H}_2 and $\text{SPD}(2)$.

Proposition 6. *The standard deviation of the estimator (12) is uniformly bounded by $\frac{\sigma^2}{\sqrt{L}}$.*

Proof. Direct computation using the properties of $\pi^{(\lambda)}$, see Appendix C. \square

3.3 Efficient Sampling

The preceding expressions give us a way to evaluate the kernel pointwise. They do not, however, immediately give a way to efficiently sample from the prior. To do this, we will employ a different way of expressing the zonal spherical functions.

Result 7. *For any $\lambda \in \mathfrak{a}^*$ and $g_1, g_2 \in G$ we have*

$$\pi^{(\lambda)}(g_2^{-1} \cdot g_1) = \int_H e^{(i\lambda+\rho)^\top a(h \cdot g_1)} \overline{e^{(i\lambda+\rho)^\top a(h \cdot g_2)}} d\mu_H(h) \quad (13)$$

where a has the same meaning as in Result 5.

Proof. Helgason (1994, Chapter III, Theorem 1.1). \square

Note the resemblance between this expression given in part I, specifically Azangulov et al. (2024, equation (35)), from which one obtains generalized random phase Fourier features. The key property of (13) is that the right-hand side can be viewed as an inner product. Thanks to this property, the prior Gaussian process $f \sim \text{GP}(0, k)$ can be approximated as

$$f(x) \approx \frac{\sigma}{\sqrt{L}} \sum_{l=1}^L w_l e^{(i\lambda_l+\rho)^\top a(h_l \cdot x)} \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k \quad h_l \sim \mu_H \quad w_l \sim \text{N}(0, 1) \quad (14)$$

where μ_H is the normalized Haar measure over H and μ_k is the spectral measure. We therefore obtain a way to use random Fourier features for efficient sampling: an illustration of such samples, for a Gaussian process with heat kernel covariance defined in Section 4, is given in Figure 3. This is not the only use of (14): the corresponding kernel approximation is by construction guaranteed to be positive semi-definite—unlike (12). This is given by

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{\sigma^2}{L} \sum_{l=1}^L e^{(i\lambda_l+\rho)^\top a(h_l \cdot g_1)} \overline{e^{(i\lambda_l+\rho)^\top a(h_l \cdot g_2)}} \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k \quad h_l \sim \mu_H. \quad (15)$$

One key difference from (12), as well as from the standard Euclidean case, is that the convergence rate of the approximation given by (15), in the sense of the standard deviation of the estimated kernel, is not uniform, but rather only uniform on compact subsets.

Theorem 8. *For any compact subset $U \subset G/H$ there is a constant $C_U > 0$ such that, for any $g_1 \cdot H, g_2 \cdot H \in U$, the right-hand side of (15) is an unbiased estimator of $k(g_1 \cdot H, g_2 \cdot H)$, whose standard deviation is bounded by $C_U \frac{\sigma^2}{\sqrt{L}}$. The constant C_U , however, does not admit an upper bound: for any fixed L , the standard deviation on the full space G/H is unbounded.*

Proof. Appendix C. □

4. Heat and Matérn Kernels

In the previous section we obtained general computational recipes for working with stationary Gaussian processes on non-compact symmetric spaces. These depend on the specific kernel k only through its spectral measure μ_k , which by Theorem 2 in turn uniquely defines k . We now study specific families of stationary Gaussian processes on non-compact symmetric spaces, infer their associated spectral measures μ_k , and discuss general techniques for sampling μ_k .

To do this, we first need to define a Riemannian metric on $X = G/H$, a symmetric space satisfying Assumption A. We do so canonically as follows. Let \mathfrak{g} be the Lie algebra of G . Since G is semi-simple, it follows there is a canonical bilinear form on \mathfrak{g} , termed the *Killing form*, which is non-degenerate. By pushforward, the Killing form defines a bi-invariant pseudo-Riemannian metric on G . For non-compact semi-simple groups, this metric is never Riemannian, since the Killing form has both positive and negative eigenvalues (Arvanitogeorgos, 2003, Theorem 2.14). In spite of this, under Assumption A, the metric which the quotient G/H inherits from G is actually *Riemannian*. We assume henceforth that G/H is equipped with this metric.

We focus on Matérn kernels $k_{\nu, \kappa, \sigma^2}$, where σ^2, κ, ν are the amplitude, length scale, and smoothness parameters, respectively, as well as its infinite-smoothness limit, the squared exponential kernel $k_{\infty, \kappa, \sigma^2}$, which is also known as the heat, Gaussian, or RBF kernel. These are the most popular Euclidean Gaussian process kernels in machine learning and statistics.

To define their analogs in the setting at hand, we adopt the heat-equation-based approach used in the compact case (Azangulov et al., 2024). This works as follows. First, the Euclidean squared exponential kernel is reinterpreted as the fundamental solution of the heat equation. Next, we generalize the heat equation from Euclidean spaces to Riemannian manifolds, by replacing the Euclidean Laplacian with the Laplace–Beltrami operator, getting the following.

Definition 9. *Let X be a Riemannian manifold. Define the HEAT KERNEL to be*

$$k_{\infty, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\infty, \kappa}} \mathcal{P}(\kappa^2/2, x, x') \quad (16)$$

where $C_{\infty, \kappa}$ is the constant that ensures $\int_X k_{\infty, \kappa, \sigma^2}(x, x) dx = \sigma^2$, while \mathcal{P} is the solution of

$$\frac{\partial \mathcal{P}}{\partial t}(t, x, x') = \Delta_x \mathcal{P}(t, x, x') \quad \mathcal{P}(0, x, x') = \delta(x, x') \quad (17)$$

where Δ_x is the Laplace–Beltrami operator acting on the x variable, and $\delta(x, x')$ is the Dirac delta function, defined appropriately. The differential equation above is precisely the HEAT EQUATION, which in the given setting admits a unique solution in $L^2(X)$ (Grigoryan, 2009).




Figure 4: Samples from Matérn Gaussian processes on \mathbb{H}_2 , with varying smoothness.

This gives the appropriate analog of the squared exponential kernel. The final step is to observe that Euclidean Matérn kernels can equivalently be defined as integrals of squared exponential kernels, via a kernel analogue of the gamma mixture of Gaussians expression for the t -distribution, which is the spectral measure of a Matérn kernel. Using this observation, this expression can more generally be taken as the *definition* of a Matérn kernel.

Definition 10. Let X be a Riemannian manifold. Define the MATÉRN KERNEL to be

$$k_{\nu,\kappa,\sigma^2}(x, x') = \frac{\sigma^2}{C_{\nu,\kappa}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \mathcal{P}(u, x, x') du \quad (18)$$

where $C_{\nu,\kappa}$ is the constant that ensures $\int_X k_{\nu,\kappa,\sigma^2}(x, x) dx = \sigma^2$ and $n = \dim X$.

Note that both Definitions 9 and 10 give the classical expressions in the Euclidean case. We illustrate samples from Matérn Gaussian processes of varying smoothness in Figure 4.

It turns out that, in contrast with the Euclidean and compact Riemannian cases, the above generalizations are not the only possibility: to understand this, we proceed to introduce a slightly different notion of a Matérn Gaussian process, by considering a modified Laplace operator. We will prove both notions well-defined in Section 4.2.

4.1 Choice of Ordinary vs. Shifted Laplace–Beltrami Operator

The preceding definitions of $k_{\infty,\kappa,\sigma^2}$ and k_{ν,κ,σ^2} depend on the manifold $X = G/H$ through the Laplace–Beltrami operator Δ . One of the key ways in which non-compact symmetric spaces can differ from their compact counterparts is that the negated Laplacian $-\Delta$ can be strongly elliptic: $\langle -\Delta f, f \rangle_{L^2(X)} \geq s \|f\|_{L^2(X)}^2$ for some $s > 0$ and all $f \in L^2(X) \cap C^\infty(X)$. The maximal s such that the inequality holds is called the *spectral gap*.⁹

The spectral gap of the Laplace–Beltrami operator is always zero in the Euclidean and compact Riemannian cases.⁹ The fact that it can be non-zero on non-compact symmetric spaces can cause the differential equations such as the heat equation to behave in a qualitatively different manner: see, for instance, Vázquez (2022) and Anker et al. (2023).

In cases where a positive spectral gap $s > 0$ leads to undesired behavior, one can instead choose to work with the *shifted Laplacian* $\Delta + sI$ (Wildberger, 1995), which is defined by adding a scalar multiple of the identity map I to the Laplace–Beltrami operator Δ . Here, s is the spectral gap of Δ . Because of the strong ellipticity of $-\Delta$, the negated shifted

9. This notion of spectral gap is not to be confused with an alternative one, commonly used for compact Riemannian manifolds. This alternative notion is defined as the smallest *non-zero* eigenvalue of $-\Delta$.

Laplacian $-\Delta - sI$ is still positive semi-definite on $L^2(X)$ —but its spectral gap vanishes, making it more similar to its Euclidean and compact Riemannian counterparts. This operator is sometimes used in defining analogues of the heat equation (Maher, 2014) as well the analogues of wave equation (Helgason, 1992) and Helmholtz equation (Livshits, 2015) in the non-compact symmetric space setting.

Since heat and Matérn kernels can be defined just-as-easily with the shifted Laplacian as with the ordinary Laplacian, the results we develop can be used in both settings with minor modifications. We will examine practical differences that result from working with the shifted Laplacian vs. the ordinary Laplacian in the sequel. First, however, we study these kernels' basic properties, which will already reveal some of the consequences of this choice.

4.2 Stationarity and Spectral Densities

We now deduce a more explicit form for the kernels defined in Definitions 9 and 10, proving them positive semi-definite along the way. This entails computing the solution of the heat equation in the setting where $X = G/H$ is a symmetric space. The standard way of doing this in the Euclidean case is by applying the Fourier transform, solving the resulting algebraic equation, then applying the inverse Fourier transform to map the solution back to the original space. To proceed, we apply this argument adapted to the setting of non-compact symmetric spaces, using the *spherical Fourier transform* of Helgason (1994).

Proposition 11. *The heat and Matérn kernels on a non-compact symmetric space $X = G/H$ are positive semi-definite, stationary, and given by*

$$k_{\nu,\kappa,\sigma^2}(g_1 \cdot H, g_2 \cdot H) = \int_{\mathfrak{a}^*} \pi^{(\lambda)}(g_2^{-1} \cdot g_1) d\mu_{\nu,\kappa,\sigma^2}(\lambda), \quad (19)$$

where the respective spectral measures $\mu_{\nu,\kappa,\sigma^2}$ are

$$\mu_{\nu,\kappa,\sigma^2} = \begin{cases} \frac{\sigma^2}{C'_{\infty,\kappa}} e^{-\frac{\kappa^2}{2}(\|\lambda\|^2 + \|\rho\|^2)} |c(\lambda)|^{-2} d\lambda & \nu = \infty \\ \frac{\sigma^2}{C'_{\nu,\kappa}} \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda & \nu < \infty \end{cases} \quad (20)$$

where ρ is defined in Result 5, $c(\lambda)$ is Harish-Chandra's c -function, and $C'_{\nu,\kappa}$ are constants which ensure that $\int_X k_{\nu,\kappa,\sigma^2}(x, x) dx = \sigma^2$ which in this case reduces to $k_{\nu,\kappa,\sigma^2}(x, x) = \sigma^2$. Moreover, the kernels $\widehat{k}_{\nu,\kappa,\sigma^2}$ defined via the shifted Laplacian admit analogous expressions for their spectral measures $\widehat{\mu}_{\nu,\kappa,\sigma^2}$, but with $\|\rho\|^2$ removed and with different constants $C'_{\nu,\kappa}$.

Proof. A proof for the heat kernel ($\nu = \infty$) can be found in Gangolli (1968, Proposition 3.1), also see Appendix B for a proof sketch. From this, the Matérn kernel's form follows by applying its definition, changing the order of integration, and appropriate algebra, which works the same way as in the compact case—see Azangulov et al. (2024).

The proof applies mutatis mutandis for $\widehat{k}_{\nu,\kappa,\sigma^2}$. The fact that both k_{ν,κ,σ^2} and $\widehat{k}_{\nu,\kappa,\sigma^2}$ are positive semi-definite and stationary follows from their form in (19) and Theorem 2. \square

We now state an important relation between the solution \mathcal{P} of the heat equation for the Laplace–Beltrami operator and the solution $\widehat{\mathcal{P}}$ for its shifted counterpart which follows from the proof of Proposition 11.

Corollary 12. *We have*

$$\widehat{\mathcal{P}}(t, x, x') = e^{t\|\rho\|^2} \mathcal{P}(t, x, x'). \quad (21)$$

It follows that $k_{\infty, \kappa, \sigma^2} = \widehat{k}_{\infty, \kappa, \sigma^2}$.

Note that this does *not* imply the similar relation for $\nu < \infty$: the kernels $k_{\nu, \kappa, \sigma^2}$ and $\widehat{k}_{\nu, \kappa, \sigma^2}$ will typically still be different.

All that remains to compute the heat and Matérn kernels numerically using random Fourier features is to develop techniques for sampling from the explicit spectral measures. The main obstacle that seemingly remains is Harish-Chandra's c -function, which produces densities that do not have a known form. Before handling this, we verify that these kernels' hyperparameters correspond to the same regularity classes as they do in the Euclidean case.

Theorem 13. *For all $\nu, \kappa, \sigma^2 > 0$ and all $g_2 \in G$, the Matérn kernels $k_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$ and $\widehat{k}_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$ lie in the Sobolev space $H^{\nu+n/2}$ if $\nu < \infty$, and in all Sobolev spaces H^α , $\alpha > 0$, if $\nu = \infty$. Thus, they are continuous and possess continuous derivatives of all integer orders strictly less than ν .*

Proof. The first part follows from a certain spectral characterization of $H^{\nu+n/2}$. The second part follows from an appropriate Sobolev embedding theorem. See Appendix C. \square

4.3 General Computation Techniques

We now study how to sample from the spectral measures $\mu_{\nu, \kappa, \sigma^2}$, which is all that remains for applying the techniques of Section 3. Here, we present general techniques that do not depend on the space being studied: Section 5 gives specialized techniques for hyperbolic space and the space of positive definite matrices, leveraging their special structure to improve efficiency compared to the techniques we now describe. We work with the ordinary Laplacian: the same techniques also apply to the shifted Laplacian, provided that one sets $\rho = 0$ and redefines the resulting normalizing constants appropriately. Observe that the densities

$$p_{\nu, \kappa}(\lambda) \propto \begin{cases} e^{-\frac{\kappa^2}{2}(\|\lambda\|^2 + \|\rho\|^2)} & \nu = \infty \\ \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 + \|\rho\|^2\right)^{-\nu-n/2} & \nu < \infty \end{cases} \quad (22)$$

are precisely the usual multivariate Gaussian and (generalized) t -densities. Using these densities, the respective spectral measures can be expressed as

$$\mu_{\nu, \kappa, \sigma^2} \propto \sigma^2 p_{\nu, \kappa}(\lambda) |c(\lambda)|^{-2} d\lambda. \quad (23)$$

The analogs $\widehat{p}_{\nu, \kappa}$ of $p_{\nu, \kappa}$ for the shifted Laplacian and $\widehat{\mu}_{\nu, \kappa, \sigma^2}$ are obtained by setting $\|\rho\|^2 = 0$.

We can therefore apply importance sampling to obtain

$$\mathbb{k}(g) \approx \frac{\sigma^2}{C_{\nu, \kappa}'' L} \sum_{l=1}^L |c(\lambda_l)|^{-2} e^{(i\lambda_l + \rho)^\top a(h_l \cdot g)} \quad \lambda_l \sim p_{\nu, \kappa} \quad h_l \sim \mu_H \quad (24)$$

$$f(g \cdot H) \approx \frac{\sigma^2}{\sqrt{C_{\nu, \kappa}'' L}} \sum_{l=1}^L w_l |c(\lambda_l)|^{-1} e^{(i\lambda_l + \rho)^\top a(h_l \cdot g)} \quad \lambda_l \sim p_{\nu, \kappa} \quad h_l \sim \mu_H \quad w_l \sim \mathcal{N}(0, 1) \quad (25)$$

where μ_H is the normalized Haar measure on H , and $C''_{\nu,\kappa}$ are normalization constants chosen to ensure the approximation has variance σ^2 .

Note that the kernel approximation corresponding to the right-hand side of (25)—the kernel approximation which is necessarily positive semi-definite—is given by

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{\sigma^2}{C''_{\nu,\kappa} L} \sum_{l=1}^L |c(\lambda_l)|^{-2} e^{(i\lambda_l + \rho)^\top a(h_l \cdot g_1)} \overline{e^{(i\lambda_l + \rho)^\top a(h_l \cdot g_2)}} \quad \begin{array}{l} \lambda_l \sim p_{\nu,\kappa} \\ h_l \sim \mu_H. \end{array} \quad (26)$$

Since these estimators are based on importance sampling, it is easy to see that they are less efficient compared to direct Monte Carlo sampling. We will demonstrate in Section 5 that for certain spaces it is possible to instead derive efficient rejection samplers, as long as one is willing to work with space-specific techniques.

5. Examples

We now specialize the techniques developed to the hyperbolic space \mathbb{H}_n and the space of symmetric positive-definite matrices $\text{SPD}(d)$ —arguably the two most important examples of non-compact symmetric spaces—and propose efficient approximations for kernels on those spaces, leveraging their specific structure. The key idea is that densities of the measures $\mu_{\nu,\kappa,\sigma^2}$, and respectively $\hat{\mu}_{\nu,\kappa,\sigma^2}$, can be upper-bounded by tractable densities, allowing efficient rejection sampling. For \mathbb{H}_n these are the simple mixtures of chi or beta-prime distributions, while for $\text{SPD}(d)$ these are the densities arising as the distributions of eigenvalues of certain random matrices. Moreover, in this section, we will also discuss the Riemannian metric used on $\text{SPD}(d)$, and handle the fact that this space does not satisfy the formal assumptions of Section 3 by representing it as a suitable product of spaces that do or are otherwise well-behaved, following the strategy discussed alongside Assumption A.

5.1 Hyperbolic Space

Hyperbolic space \mathbb{H}_n is defined as the unique, up to an isometry, simply connected n -dimensional Riemannian manifold with sectional curvature -1 . More concretely, it is usually represented by *models* which are concrete manifolds that are isometric to the general definition. For instance, the *Poincaré ball model* identifies \mathbb{H}_n with the open unit ball $B(n) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1\} \subseteq \mathbb{R}^n$ endowed with the non-Euclidean Riemannian metric

$$g_{\mathbf{x}}(\mathbf{u}_{\mathbf{x}}, \mathbf{v}_{\mathbf{x}}) = \frac{4\langle \mathbf{u}_{\mathbf{x}}, \mathbf{v}_{\mathbf{x}} \rangle}{(1 - \|\mathbf{x}\|^2)^2}, \quad (27)$$

where $\mathbf{x} \in B(n)$ and $\mathbf{u}_{\mathbf{x}}, \mathbf{v}_{\mathbf{x}}$ are tangent vectors (Wolf, 2011, Corollary 2.4.13). Alternatively, \mathbb{H}_n can be represented by the *hyperboloid model*. To construct this model, one starts with generalized Minkowski space, defined as \mathbb{R}^{n+1} with the bilinear form

$$\langle \mathbf{v}, \mathbf{u} \rangle_M = -v_0 u_0 + v_1 u_1 + \dots + v_n u_n \quad (28)$$

in place of the standard inner product. This is a pseudo-Euclidean space, and hence a pseudo-Riemannian manifold, the $n = 3$ version of which models spacetime in special relativity. One then considers the unit sphere $\{\mathbf{v} \in \mathbb{R}^{n+1} : \langle \mathbf{v}, \mathbf{v} \rangle_M = 1\}$, which is a two-sheeted hyperboloid.




Figure 5: Gaussian process regression on hyperbolic space \mathbb{H}_2 driven by the heat kernel.

Hyperbolic space is realized as its forward sheet, that corresponds to $v_0 > 0$, which is a Riemannian manifold. We show a simple Gaussian process regression model on \mathbb{H}_n , computed using the techniques of this section, in Figure 5.

The hyperboloid model helps identify that \mathbb{H}_n is a homogeneous space. Let $\mathrm{SL}(1+n)$ be the group of $(1+n) \times (1+n)$ matrices with unit determinant. Define the *proper Lorentz group* $\mathrm{SO}(1, n)$ as the subgroup of $\mathrm{SL}(1+n)$ that preserves $\langle \cdot, \cdot \rangle_M$, namely

$$\mathrm{SO}(1, n) = \{ \mathbf{M} \in \mathrm{SL}(1+n) : \langle \mathbf{M}\mathbf{v}, \mathbf{M}\mathbf{u} \rangle_M = \langle \mathbf{v}, \mathbf{u} \rangle_M, \forall \mathbf{v}, \mathbf{u} \in \mathbb{R}^{n+1} \}. \quad (29)$$

The proper Lorentz group's action on \mathbb{H}_n is simply the restriction onto the unit sphere of the natural action of $\mathrm{SL}(1+n)$ on \mathbb{R}^{n+1} . The group $\mathrm{SO}(1, n)$ is not connected and has two connected components. Let $\mathrm{SO}_0(1, n)$ be the connected component that contains the identity matrix. It consists (Lester, 1993) of $\mathbf{M} \in \mathrm{SO}(1, n)$ such that

$$\mathbf{M} = \begin{pmatrix} a & \mathbf{b}^\top \\ \mathbf{c} & \mathbf{D} \end{pmatrix} \quad a > 0 \quad \mathbf{b}, \mathbf{c} \in \mathbb{R}^n \quad \det \mathbf{D} > 0. \quad (30)$$

Let $\mathbf{x}_0 = (0, 1/\sqrt{n}, \dots, 1/\sqrt{n})$. Then the subgroup of $\mathrm{SO}_0(1, n)$ which preserves \mathbf{x}_0 , in the sense $\mathbf{M}\mathbf{x}_0 = \mathbf{x}_0$, will be the group $H \subseteq \mathrm{SO}_0(1, n)$ that is isomorphic to the group $\mathrm{SO}(n)$ of orthogonal matrices with unit determinant. With these definitions, we have the following.

Result 14. *The quotient manifold $\mathrm{SO}_0(1, n)/\mathrm{SO}(n)$ is a symmetric space satisfying Assumption A. Moreover, when equipped with the canonical metric induced by the Killing form, this space is isometric to \mathbb{H}_n .*

Proof. See Wolf (2011, Theorem 8.12.2(iii.b')). □

5.1.1 EXPLICIT FORMULAS FOR SPHERICAL FUNCTIONS AND SPECTRAL DENSITIES

On \mathbb{H}_n , the set \mathfrak{a}^* that parametrizes the irreducible unitary representations of interest corresponding to \mathbb{H}_n is one-dimensional: $\mathfrak{a}^* = \mathbb{R}$. Moreover, $\rho = \frac{n-1}{2}$. The integral form of the zonal spherical functions on \mathbb{H}_n ostensibly involves integration over $\mathrm{SO}(n)$: however, this integral actually simplifies further. The basic reason for this is that, for a from Result 5, there are distinct values $\mathbf{V}, \mathbf{U} \in H \subseteq \mathrm{SO}_0(1, n)$ for which $a(\mathbf{V}\mathbf{M}) = a(\mathbf{U}\mathbf{M})$ for all $\mathbf{M} \in \mathrm{SO}_0(1, n)$, which allows one to instead integrate over a smaller set. We formally state the underlying symmetry in terms of the diagonal part $A(\mathbf{M})$ of the Iwasawa decomposition of elements $\mathbf{M} \in \mathrm{SO}_0(1, n)$ —recall that A and a are connected by $A(\mathbf{M}) = \exp a(\mathbf{M})$.

Lemma 15. *We have $A(\mathbf{UM}) = A(\mathbf{M})$ for all $\mathbf{M} \in \text{SO}_0(1, n)$ and all $\mathbf{U} \in H \subseteq \text{SO}_0(1, n)$ of the form*

$$\mathbf{U} = \begin{pmatrix} 1 & 0 & \\ 0 & 1 & \\ & & \tilde{\mathbf{U}} \end{pmatrix} \quad \tilde{\mathbf{U}} \in \text{SO}(n-1). \quad (31)$$

Proof. This follows from the Iwasawa decomposition for $\text{SO}_0(1, n)$, see Appendix C. \square

Consequently, the same holds for $a(\cdot)$, making it possible to use this symmetry to replace integration over $\text{SO}(n)$ with integration over $\text{SO}(n)/\text{SO}(n-1) \cong \mathbb{S}_n$, the n -dimensional sphere. In the Poincaré disk model, the resulting zonal spherical functions admit a particularly simple form (Sun, 1994). Here, the sphere \mathbb{S}_n appears as boundary of the unit ball $B(n)$ that models hyperbolic space, and the zonal spherical functions are given by

$$\pi^{(\lambda)}(\mathbf{x}) = \int_{\mathbb{S}_n} \left(\frac{1 - \|\mathbf{x}\|^2}{\|\mathbf{x} - \mathbf{b}\|^2} \right)^{i|\lambda| + (n-1)/2} d\mu_{\mathbb{S}_n}(\mathbf{b}) \quad \mathbf{x} \in B(n) \subseteq \mathbb{R}^n \quad \lambda \in \mathbb{R} \quad (32)$$

where $\mu_{\mathbb{S}_n}$ is the Haar measure on the n -dimensional sphere. For \mathbb{H}_n , Harish-Chandra's c -function depends on the parity of the space's dimension. It is given by Sun (1994) as

$$|c(\lambda)|^{-2} \propto \begin{cases} |\lambda| \tanh(\pi|\lambda|) \prod_{j=2}^m \left(|\lambda|^2 + \frac{(2j-3)^2}{4} \right) & n = 2m \\ \prod_{j=0}^{m-1} (|\lambda|^2 + j^2) & n = 2m + 1. \end{cases} \quad (33)$$

With this, applying the general theory of Section 3 allows one to approximate heat and Matérn kernels numerically, and perform efficient sampling of the respective Gaussian processes.

5.1.2 EFFICIENT SAMPLING FROM SPECTRAL DENSITIES

In Section 3, we showed that the problem of computing heat and Matérn kernels can ultimately be reduced to the problem of sampling from their respective spectral measures. We now develop specific techniques for doing so on \mathbb{H}_n . We concentrate on the ordinary Laplace–Beltrami operator, remarking on how to apply the same techniques to the shifted Laplacian at the end of this section. First, let us expand the polynomials in (33):

$$|c(\lambda)|^{-2} \propto \begin{cases} \tanh(\pi|\lambda|) \sum_{j=0}^{m-1} a_{n,j} |\lambda|^{2j+1} & n = 2m \\ \sum_{j=0}^m a_{n,j} |\lambda|^{2j} & n = 2m + 1 \end{cases} \quad (34)$$

where $a_{n,j}$ are the coefficients of the respective polynomials. Using this, and $\rho = (n - 1)/2$, we can express heat kernels' spectral measure as

$$\mu_{\infty,\kappa,\sigma^2} \propto \begin{cases} \tanh(\pi|\lambda|) \sum_{j=0}^{m-1} a_{n,j} |\lambda|^{2j+1} e^{-\frac{\kappa^2}{2}|\lambda|^2} d\lambda & n = 2m \\ \sum_{j=0}^m a_{n,j} |\lambda|^{2j} e^{-\frac{\kappa^2}{2}|\lambda|^2} d\lambda & n = 2m + 1. \end{cases} \quad (35)$$

Note that we dropped the $(\frac{n-1}{2})^2$ term from the exponent: we can do this because it only changes the normalizing constant. This density takes the form of a Gaussian times a polynomial, which is an analytically tractable form. Since the spectral densities and spherical functions depend only on the absolute value of λ , we may also restrict support of λ onto $[0, \infty)$.

Proposition 16. *For any $m \in \mathbb{N}$, and any coefficients $\alpha_j \geq 0$, the density*

$$p_{\infty,\kappa}(\lambda) \propto \sum_{j=0}^m \alpha_j \lambda^j e^{-\frac{\kappa^2}{2}\lambda^2} \quad \lambda \in \mathbb{R}_{>0} \quad (36)$$

is a mixture of scaled chi-distributed random variables. Specifically, we have

$$p_{\infty,\kappa}(\lambda) = \sum_{j=0}^m c_j p_{y_j}(\lambda) \quad c_j = \frac{\alpha_j/\beta_j}{\sum_{j=0}^m \alpha_j/\beta_j} \quad \beta_j = 2^{\frac{1-j}{2}} \Gamma\left(\frac{j+1}{2}\right)^{-1} \kappa \quad (37)$$

where $\Gamma(\cdot)$ is the gamma function, and $p_{y_j}(\lambda)$ is the density of the random variable $y_j = x_j/\kappa$ where $x_j \sim \chi_{j+1}$ is a chi-distributed random variable with $j + 1$ degrees of freedom.

Proof. Direct computation, see Appendix C. \square

For odd n , this immediately yields a way to sample exactly from the spectral density. For even n , this is *almost* the heat kernel's spectral density: once constants are dropped: all that remains is the hyperbolic tangent term. Since $\tanh(x) < 1$, we can devise an efficient rejection sampler. Before stating the algorithm, we extend this idea to the Matérn class. Using (20) we can express Matérn kernel's spectral density as

$$\mu_{\nu,\kappa,\sigma^2} \propto \begin{cases} \tanh(\pi|\lambda|) \sum_{j=0}^{m-1} a_{n,j} |\lambda|^{2j+1} \left(\frac{2\nu}{\kappa^2} + |\lambda|^2 + \left| \frac{n-1}{2} \right|^2 \right)^{-\nu-n/2} d\lambda & n = 2m \\ \sum_{j=0}^m a_{n,j} |\lambda|^{2j} \left(\frac{2\nu}{\kappa^2} + |\lambda|^2 + \left| \frac{n-1}{2} \right|^2 \right)^{-\nu-n/2} d\lambda & n = 2m + 1. \end{cases} \quad (38)$$

In this case the respective part of the spectral density, excluding the hyperbolic tangent, is again tractable.

Proposition 17. For any $m \in \mathbb{N}$, $m \leq n - 1$, and any coefficients $\alpha_j \geq 0$, the density

$$p_{\nu, \kappa}(\lambda) \propto \sum_{j=0}^m \alpha_j \lambda^j \left(\frac{2\nu}{\kappa^2} + \lambda^2 + \left(\frac{n-1}{2} \right)^2 \right)^{-\nu - \frac{n}{2}} \quad \lambda \in \mathbb{R}_{>0} \quad (39)$$

is a mixture of transformed beta-prime distributed random variables. Specifically, we have

$$p_{\nu, \kappa}(\lambda) = \sum_{j=0}^m c_j p_{y_j}(\lambda) \quad c_j = \frac{\alpha_j / \beta_j}{\sum_{j=0}^m \alpha_j / \beta_j} \quad \beta_j = 2 \mathbf{B} \left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2} \right)^{-1} \gamma^{\nu + \frac{n-j-1}{2}} \quad (40)$$

where \mathbf{B} is the beta function, $\gamma = \frac{2\nu}{\kappa^2} + \left(\frac{n-1}{2} \right)^2$, and p_{x_j} is the density of the random variable $y_j = \sqrt{\gamma x_j}$ where $x_j \sim \text{Beta}' \left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2} \right)$ is beta-prime distributed.

Proof. Direct computation, see Appendix C. \square

Using this, for both heat and Matérn kernels, we obtain a rejection-sampling-based strategy for sampling from the respective spectral measure.

Algorithm 18. Rejection sampler for $\mu_{\nu, \kappa, \sigma^2}$ on \mathbb{H}_n .

1. Compute a_j by expanding (33).
2. Sample the mixture component index j with probability c_j from Proposition 16 or Proposition 17, for the heat or Matérn cases, respectively.
3. Sample λ from p_{Y_j} where p_{Y_j} is given in Proposition 16 or Proposition 17, for the heat or Matérn cases, respectively.
4. If n is even, reject λ with probability $1 - \tanh(\pi\lambda)$ returning to step 2, otherwise proceed.
5. Uniformly sample a sign $s \in \{+1, -1\}$ and output $s\lambda$.

Remark 19. Setting $\gamma = \frac{2\nu}{\kappa}$ instead of $\gamma = \frac{2\nu}{\kappa^2} + \left(\frac{n-1}{2} \right)^2$ results in an algorithm for sampling the measure $\hat{\mu}_{\nu, \kappa, \sigma^2}$ that corresponds to the shifted Laplacian, instead of the measure $\mu_{\nu, \kappa, \sigma^2}$.

This completes our development: we can now compute heat and Matérn kernels, and perform efficient sampling from the prior. Note that, here, λ is one-dimensional, which follows from \mathbb{H}_n being a rank-1 symmetric space. This is a key reason for the relative simplicity of these results.

5.1.3 CLOSED FORM EXPRESSIONS AND RECURRENCE RELATIONS FOR HEAT KERNELS

It is known (Jaquier et al., 2020; Grigoryan, 1998) that in odd dimensions heat kernels on hyperbolic spaces admit closed-form expressions, while in even dimensions the general formulas simplify to certain one-dimensional integrals. For $x, x' \in \mathbb{H}_n$ denote the geodesic distance between x and x' by $\rho = \text{dist}_{\mathbb{H}_n}(x, x')$. Then, in two and three dimensions, we have

$$k_{\infty, \kappa, \sigma^2}^{\mathbb{H}_2}(x, x') = \frac{\sigma^2}{C_{\infty, \kappa}} \int_{\rho}^{\infty} \frac{s \exp(-s^2 / (2\kappa^2))}{(\cosh(s) - \cosh(\rho))^{1/2}} ds \quad (41)$$

$$k_{\infty, \kappa, \sigma^2}^{\mathbb{H}_3}(x, x') = \frac{\sigma^2}{C_{\infty, \kappa}} \frac{\rho}{\sinh \rho} e^{-\frac{\rho^2}{2\kappa^2}} \quad (42)$$

and other cases may be obtained by the following recurrence relation, known as Millson’s formula, for $d > 3$:

$$k_{\infty, \kappa, \sigma^2}^{\mathbb{H}_d}(x, x') = -\frac{\sigma^2}{C_{\infty, \kappa} \sinh \rho} \frac{\partial}{\partial \rho} k_{\infty, \kappa, \sigma^2}^{\mathbb{H}_{d-2}}(x, x'). \quad (43)$$

Note that, unfortunately, for even dimensions the obvious Monte Carlo approximations of these integrals are not guaranteed to be positive semi-definite.

5.2 The Space of Symmetric Positive-definite Matrices $\text{SPD}(d)$

Symmetric positive-definite matrices arise in many real-world applications (Arnaudon et al., 2013; Dong and Kuang, 2015; Jaquier et al., 2020; Jayasumana et al., 2013). We denote the set of such matrices of size $d \times d$ by $\text{SPD}(d)$. This is an open subset (Terras, 2016, Section 1.1.2) of the $\frac{d(d+1)}{2}$ -dimensional vector space of symmetric matrices, from which it inherits the structure of a smooth manifold. There are different Riemannian metrics on $\text{SPD}(d)$ (Jayasumana et al., 2013; Lin, 2019; Dryden et al., 2009): we will concentrate on the *affine-invariant metric*, which makes $\text{SPD}(d)$ into a symmetric space of negative curvature.

The smooth manifold $\text{SPD}(d)$ has the same dimension as the space of symmetric matrices. In the case $d = 2$, this is a three-dimensional manifold, which can be understood as the set of all matrices $\begin{pmatrix} a & b \\ b & c \end{pmatrix}$ satisfying $a > 0$ and $ac - b^2 > 0$. Geometrically, this space is a cone: to illustrate how Matérn kernels reflect the geometry of $\text{SPD}(d)$, we visualize them for various length scales in Figure 6. We show cross-sections that correspond, respectively, to unit-determinant matrices, and central slices that correspond to entries equal to $a = 1$, $c = 1$, and $b = 0$. The kernel’s first argument is fixed at the identity matrix.

The geodesic distance induced by the affine-invariant metric (Bhatia, 2009, Theorem 6.1.6) is given by

$$\text{dist}_{\text{SPD}(d)}(\mathbf{S}_1, \mathbf{S}_2) = \left\| \log \left(\mathbf{S}_1^{-1/2} \mathbf{S}_2 \mathbf{S}_1^{-1/2} \right) \right\|_F \quad (44)$$

where $\|\cdot\|_F$ is the Frobenius norm of a matrix. This distance is invariant under linear transformations, namely

$$\text{dist}_{\text{SPD}(d)}(\mathbf{A} \mathbf{S}_1 \mathbf{A}^\top, \mathbf{A} \mathbf{S}_2 \mathbf{A}^\top) = \text{dist}_{\text{SPD}(d)}(\mathbf{S}_1, \mathbf{S}_2). \quad (45)$$

It can therefore be extended to the space of Gaussian measures: in doing so, the above invariance transforms into an invariance with respect to affine transformations—this property gives it the name *affine-invariant*.

5.2.1 THE GEOMETRY OF $\text{SPD}(d)$

Consider first the case $d = 1$: in this case, symmetric positive definite matrices are merely positive reals $\mathbb{R}_{>0}$, which form an abelian Lie group under multiplication. Every abelian Lie group is isomorphic to a product of an Euclidean addition group with a torus: in this particular case, $\text{SPD}(1)$ is isomorphic to \mathbb{R} via the logarithm map $\log : \mathbb{R}_{>0} \rightarrow \mathbb{R}$, which preserves the group structure. This allows us to identify kernels on $\text{SPD}(1)$ as kernels on \mathbb{R} composed with the logarithm. From now on we study the higher-dimensional case, where the group $\mathbb{R}_{>0}$ will play an important role: assume henceforth that $d > 1$.




Figure 6: Values of Matérn-5/2 kernels $k(\bullet, \cdot)$ on $\text{SPD}(2)$, with various length scales.

We now proceed to develop a more precise understanding of the geometric structure of $\text{SPD}(d)$, including how to give it the structure of a symmetric space, in order to apply the preceding ideas. The following proposition shows how one can view $\text{SPD}(d)$ as a homogeneous space, which in this form is *almost*, but not quite, symmetric.

Proposition 20. *The group $G = \text{GL}(d)$ acts on $\text{SPD}(d)$ transitively by conjugation as*

$$\triangleright : \text{GL}(d) \times \text{SPD}(d) \rightarrow \text{SPD}(d) \qquad \mathbf{M} \triangleright \mathbf{S} = \mathbf{M}\mathbf{S}\mathbf{M}^\top \quad (46)$$

where $H = \text{O}(d) \subseteq G$ is the isotropy subgroup corresponding to the element $\mathbf{I}_d \in \text{SPD}(d)$. $\text{SPD}(d)$ can thus be regarded as the homogeneous space $\text{GL}(d)/\text{O}(d)$. Moreover, the map

$$s : \text{GL}(d) \rightarrow \text{GL}(d) \qquad s(\mathbf{M}) = \mathbf{M}^{-\top} \quad (47)$$

is an involution, namely $s(s(\mathbf{M})) = \mathbf{M}$, and $s(\mathbf{M}) = \mathbf{M}$ if and only if $\mathbf{M} \in \text{O}(d)$. Using this representation, $\text{SPD}(d)$ satisfies all aspects of the definition of a symmetric space, except for connectedness of the group G .

Proof. This is a standard result: we present its proof in Appendix C. □

Unfortunately, the resulting quotient $\text{SPD}(d) \cong \text{GL}(d)/\text{O}(d)$ is not a symmetric space: the first problem is that $\text{GL}(d)$ is not connected. This is relatively easy to fix, by restricting from $\text{GL}(d)$ to its identity-connected-component subgroup $\text{GL}_+(d)$, and, using the same argument as above, representing $\text{SPD}(d) = \text{GL}_+(d)/\text{SO}(d)$. This shows that $\text{SPD}(d)$ is actually a symmetric space, but is still not good enough for our purposes since this representation fails to satisfy Assumption A, as $\text{GL}_+(d)$ is not semi-simple. To see this, note that a semi-simple Lie group cannot contain a normal abelian subgroup, like the subgroup of $\text{GL}_+(d)$ formed by the matrices $a\mathbf{I}_d$, $a > 0$.

To sidestep these issues, the idea will be to follow the strategy previously mentioned in the context of Assumption A, and view $\text{SPD}(d)$ as the product of the space of unit-determinant

positive definite matrices $\text{SSPD}(d) \subseteq \text{SPD}(d)$ with the multiplicative group of positive reals $\mathbb{R}_{>0}$, representing the determinant, which already appeared when considering $\text{SPD}(1)$. To proceed, we show how $\text{SSPD}(d)$ can be regarded as a symmetric space.

Proposition 21. *We have $\text{SSPD}(d) \cong \text{SL}(d)/\text{SO}(d)$. Moreover, this space is symmetric and satisfies Assumption A.*

Proof. This is a standard result: we present its proof in Appendix C. \square

Since $\text{SSPD}(d)$ satisfies Assumption A, it possesses a canonical Riemannian metric induced by its Killing form. The set $\mathbb{R}_{>0}$ of positive reals also possesses a canonical metric arising from pulling back the Euclidean metric on \mathbb{R} through the logarithm map. Thus, we can identify $\text{SPD}(d)$ with the product of two Riemannian manifolds $\mathbb{R}_{>0} \times \text{SSPD}(d)$, which we equip with the product metric to define a metric on $\text{SPD}(d)$. This metric is precisely the *affine-invariant* metric we began with, which therefore arises from the symmetric space structure.

Proposition 22. *Let $\mathbb{R}_{>0}$ and $\text{SSPD}(d)$ be equipped with their respective canonical metrics, and let $\text{SPD}(d)$ be equipped with the affine-invariant metric. Then*

$$\mathcal{T} : \mathbb{R}_{>0} \times \text{SSPD}(d) \rightarrow \text{SPD}(d) \qquad \mathcal{T}(a, \mathbf{S}) = a^{\sqrt{d}} \mathbf{S} \qquad (48)$$

is a Riemannian isometry.

Proof. This is a known result, but we present its proof in Appendix C. \square

5.2.2 EXPLICIT FORMULAS FOR SPHERICAL FUNCTIONS AND SPECTRAL MEASURES

We have established that $\text{SPD}(d) = \mathbb{R}_{>0} \times (\text{SL}(d)/\text{SO}(d))$ is the product of the abelian group $\mathbb{R}_{>0}$ with a symmetric space. To define the kernels of interest, formally speaking, we now need to (1) apply Section 3 together with standard Euclidean theory to define and compute the heat kernels on $\text{SSPD}(d)$ and $\mathbb{R}_{>0}$, then (2) use the fact that the heat kernel on a product of Riemannian manifolds is the product of the heat kernels corresponding to its components (Grigoryan, 2009, Theorem 9.11), and finally (3) extend this to Matérn kernels by leveraging their connection to the heat kernel.

Terras (2016, Section 1.3) shows that the above calculation produces the same results that one would obtain if one were to simply *pretend* that $\text{GL}(d)$ is semi-simple, consider $\text{SPD}(d) = \text{GL}(d)/\text{O}(d)$, and use the RQ decomposition¹⁰ in lieu of the Iwasawa decomposition. We will present the obtained results in these terms. The vector ρ from Result 5 is given by Terras (2016, Section 1.3.1)¹¹ as

$$\rho = \left(1 - \frac{d+1}{2}, \dots, j - \frac{d+1}{2}, \dots, d - \frac{d+1}{2} \right). \qquad (49)$$

The zonal spherical functions are

$$\pi^\lambda(\mathbf{M}) = \int_{\text{O}(d)} p_\lambda(\mathbf{H} \cdot \mathbf{M}) \, d\mu_{\text{O}(d)}(\mathbf{H}) \qquad \mathbf{M} \in \text{GL}(d) \qquad (50)$$

10. The RQ decomposition is defined similarly to the QR decomposition, but the two are distinct.

11. Note that in Terras (2016) ρ is two times smaller, this is because $\lambda = 2i \cdot \mathbf{r}$, where vector \mathbf{r} in the book is called *properly normalized variables in the power function*.

where, denoting by u_j the j -th diagonal entry of the R part of the RQ decomposition of a matrix $\mathbf{U} \in \text{GL}(d)$, we have

$$p_{\boldsymbol{\lambda}}(\mathbf{U}) = \prod_{j=1}^d e^{(i \cdot \lambda_j + j - \frac{d+1}{2}) \log u_j} = \prod_{j=1}^d u_j^{i \cdot \lambda_j + j - \frac{d+1}{2}}. \quad (51)$$

Finally, Harish-Chandra's c -function is given by

$$|c(\boldsymbol{\lambda})|^{-2} \propto \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \tanh(\pi |\lambda_i - \lambda_j|). \quad (52)$$

5.2.3 SAMPLING FROM SPECTRAL MEASURES

We now turn attention to sampling from the spectral measure, which is the final missing piece on $\text{SPD}(d)$. Here, again, we concentrate on the ordinary Laplace–Beltrami operator, commenting on the changes required for the shifted Laplacian at the end of the section. The respective spectral measures of heat and Matérn kernels on $\text{SPD}(d)$ are given by

$$\mu_{\nu, \kappa, \sigma^2} \propto \begin{cases} \exp\left(-\frac{\kappa^2}{2}(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2)\right) \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \tanh(\pi |\lambda_i - \lambda_j|) d\boldsymbol{\lambda} & \nu = \infty \\ \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2\right)^{-\nu - n/2} \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \tanh(\pi |\lambda_i - \lambda_j|) d\boldsymbol{\lambda} & \nu < \infty \end{cases} \quad (53)$$

where $\|\boldsymbol{\rho}\|^2 = \frac{d^3 - d}{12}$ and $n = d(d+1)/2$. We now observe these expressions are very closely related to the eigenvalue distribution of the following classical random matrix ensemble.

Lemma 23. *Define the GAUSSIAN ORTHOGONAL ENSEMBLE to be $\mathbf{M} = (\mathbf{X} + \mathbf{X}^\top)/2$, where \mathbf{X} is a matrix with independent standard Gaussian entries. Let $\boldsymbol{\lambda}_{\text{GOE}}$ be a random vector consisting of eigenvalues of \mathbf{M} and denote its density by $p_{\boldsymbol{\lambda}_{\text{GOE}}}$. Then*

$$p_{\boldsymbol{\lambda}_{\text{GOE}}}(\boldsymbol{\lambda}) \propto \exp\left(-\frac{\|\boldsymbol{\lambda}\|^2}{2}\right) \prod_{1 \leq i < j \leq d} |\lambda_i - \lambda_j| \quad (54)$$

Proof. Mehta (1991, Theorem 3.3.1).¹² □

An immediate corollary of Lemma 23 is the following.

Corollary 24. *Consider the density defined by*

$$p_{\infty, \kappa}(\boldsymbol{\lambda}) \propto \exp\left(-\frac{\kappa^2}{2}(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2)\right) \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \quad \boldsymbol{\lambda} \in \mathbb{R}^d. \quad (55)$$

If $\boldsymbol{\lambda}^\infty \sim p_{\boldsymbol{\lambda}_{\text{GOE}}}$ and $\boldsymbol{\lambda}^{\infty, \kappa} = \boldsymbol{\lambda}^\infty / \kappa$, then $\boldsymbol{\lambda}^{\infty, \kappa} \sim p_{\infty, \kappa}$.

Proof. Let $p_{\boldsymbol{\lambda}^{\infty, \kappa}}$ be the density of $\boldsymbol{\lambda}^{\infty, \kappa}$. Using the change of variable formula for densities

$$p_{\boldsymbol{\lambda}^{\infty, \kappa}}(\boldsymbol{\ell}) = p_{\boldsymbol{\lambda}^\infty}(\kappa \boldsymbol{\ell}) \kappa^d. \quad (56)$$

Hence $p_{\boldsymbol{\lambda}^{\infty, \kappa}}(\boldsymbol{\ell}) \propto p_{\infty, \kappa}(\boldsymbol{\ell})$, and since both are densities, they are equal. □

¹² We use a slightly different normalization compared to Mehta (1991).

This is *almost* the density which occurs in the heat kernel's spectral measure: once the length scale κ is standardized and normalizing constants are dropped: all that remains is the hyperbolic tangent term. Since $\tanh(|x|) < 1$, this connection can be used to devise an efficient rejection sampler, which behaves very similarly to the one previously developed for hyperbolic space. As before, we first extend this to the Matérn class.

Proposition 25. *Consider the density defined by*

$$p_{\nu,\kappa}(\boldsymbol{\lambda}) \propto \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2 \right)^{-\nu-n/2} \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \quad \boldsymbol{\lambda} \in \mathbb{R}^d. \quad (57)$$

Let $\boldsymbol{\lambda}^{\infty,\gamma} \sim p_{\infty,\gamma}$, where $\gamma = \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2 \right)^{-1/2}$ and the right-hand side is defined in Corollary 24, and let $y \sim \chi_{2\nu}$ be a chi-distributed random variable with $2\nu \in \mathbb{R}_{>0}$ degrees of freedom. If $\boldsymbol{\lambda}^{\nu,\kappa} = \boldsymbol{\lambda}^{\infty,\gamma}/y$, then $\boldsymbol{\lambda}^{\nu,\kappa} \sim p_{\nu,\kappa}$.

Proof. Direct computation, see Appendix C. \square

Using this, we obtain an analogous rejection-sampling-based strategy for sampling from the Matérn spectral measure.

Algorithm 26. *Rejection sampler for $\mu_{\nu,\kappa,\sigma^2}$ on $\text{SPD}(d)$.*

1. Sample a matrix \mathbf{X} whose entries are IID $\mathcal{N}(0, 1)$ and compute $\mathbf{M} = (\mathbf{X} + \mathbf{X}^\top)/\sqrt{2}$.
2. Compute the eigenvalues $\boldsymbol{\lambda}^\infty$ of \mathbf{M} .
3. This step depends on whether or not ν is equal to ∞ .
 - (a) If $\nu = \infty$, compute $\boldsymbol{\lambda}^{\infty,\kappa} = \boldsymbol{\lambda}^\infty/\kappa$.
 - (b) If $\nu < \infty$, compute $\boldsymbol{\lambda}^{\infty,\gamma} = \boldsymbol{\lambda}^\infty/\gamma$ where $\gamma = \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2 \right)^{-1/2}$ and sample $y \sim \chi_{2\nu}$ from a chi distribution with 2ν degrees of freedom, compute $\boldsymbol{\lambda}^{\nu,\kappa} = \boldsymbol{\lambda}^{\infty,\gamma}/y$.
4. Reject $\boldsymbol{\lambda}^{\nu,\kappa}$ with probability equal to $1 - \prod_{1 \leq i < j \leq d} \tanh \pi |\lambda_i - \lambda_j|$ returning to step 1, otherwise output $\boldsymbol{\lambda}^{\nu,\kappa}$.

Remark 27. *Setting $\gamma = \left(\frac{2\nu}{\kappa^2} \right)^{-1/2}$ instead of $\gamma = \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2 \right)^{-1/2}$ results in an algorithm for sampling the measure $\widehat{\mu}_{\nu,\kappa,\sigma^2}$ corresponding to the shifted Laplacian, instead of $\mu_{\nu,\kappa,\sigma^2}$.*

This completes our development: on $\text{SPD}(d)$, we can now compute heat and Matérn kernels, and perform efficient sampling from the prior.

5.2.4 ALTERNATIVE EXPRESSIONS FOR HEAT KERNELS IN SPECIAL CASES

For $d = 2$, there is a relatively simple alternative formula due to Sawyer (1992), which appeared in Jaquier et al. (2022). Take $\mathbf{X}, \mathbf{Y} \in \text{SPD}(d)$ and denote by $\mathbf{S}_\mathbf{X}, \mathbf{S}_\mathbf{Y}$ their respective Cholesky factors. Let $H_1 \geq H_2$ be the singular values of $\mathbf{S}_\mathbf{X} \mathbf{S}_\mathbf{Y}^{-1}$ and define $\alpha = H_1 - H_2$. Then we have

$$k_{\infty,\kappa,\sigma^2}(\mathbf{X}, \mathbf{Y}) = \frac{\sigma^2}{C_\infty} \exp\left(-\frac{H_1^2 + H_2^2}{2\kappa^2}\right) \int_0^\infty \frac{(2s + \alpha) \exp(-s(s + \alpha)/\kappa^2)}{(\sinh(s) \sinh(s + \alpha))^{1/2}} ds \quad (58)$$




Figure 7: The smallest zonal spherical function $\pi^{(0)}$, which can be non-constant.

where C_∞ is a normalizing constant ensuring that $k_{\infty, \kappa, \sigma^2}(\mathbf{X}, \mathbf{X}) = \sigma^2$. Note that this expression, when implemented numerically using for instance quadrature, is not guaranteed to be positive semi-definite.

The expression (58) and the corresponding hyperbolic space expression (41) look similar: this is not a coincidence. Following Section 5.2.1, we have $\text{SPD}(2) = \mathbb{R}_{>0} \times \text{SSPD}(2)$. If $d = 2$, and only in this case, then $\text{SSPD}(d) = \mathbb{H}_d$ (Helgason, 2000, Introduction, §4), allowing one to apply techniques which are otherwise specific to hyperbolic spaces. Leveraging this connection, one can obtain formulas like (58). One can also derive expressions for $d = 3$, which are unfortunately rather complicated: see Sawyer (1992) for additional details.

5.3 Approximation Error, Computational Costs, and Practical Considerations

Since kernels on non-compact symmetric spaces admit integral representations, computing these kernels numerically generally involves Monte Carlo sampling, as opposed to truncation of infinite series which occurs in the compact case. We now explore these Monte Carlo approximations and some practical considerations which arise from non-compactness.

5.3.1 CORRELATIONS IN THE LARGE LENGTH SCALE LIMIT

The first important practical consideration which occurs for heat and Matérn kernels on non-compact symmetric spaces is that, in the large length scale limit, the correlation between any two given points need not converge to one. This behavior, which differs from both the Euclidean and compact Riemannian settings, is a consequence of the following proposition.

Proposition 28. *Let $X = G/H$ be a symmetric space that satisfies Assumption A. Recall that $\pi^{(\lambda)}$ denote zonal spherical functions on X and $\lambda \in \mathfrak{a}^*$ are vectors. Consider $\mathbb{k} \in L^2(G)$ such that $k(g_1 \cdot H, g_2 \cdot H) = \mathbb{k}(g_2^{-1} \cdot g_1)$ is a stationary kernel on X . Then*

$$\mathbb{k}(g) \leq \underbrace{\mathbb{k}(e)}_{\text{variance of } k} \pi^{(0)}(g) \quad (59)$$

for all $g \in G$. Moreover, there exists $g \in G$ such that $g \notin e \cdot H$ and $\pi^{(0)}(g) < 1$.

Proof. Appendix C. □

That is, under our usual assumptions, all stationary kernels are bounded above by the zonal spherical function $\pi^{(0)}$ which can be strictly less than one at points outside of the

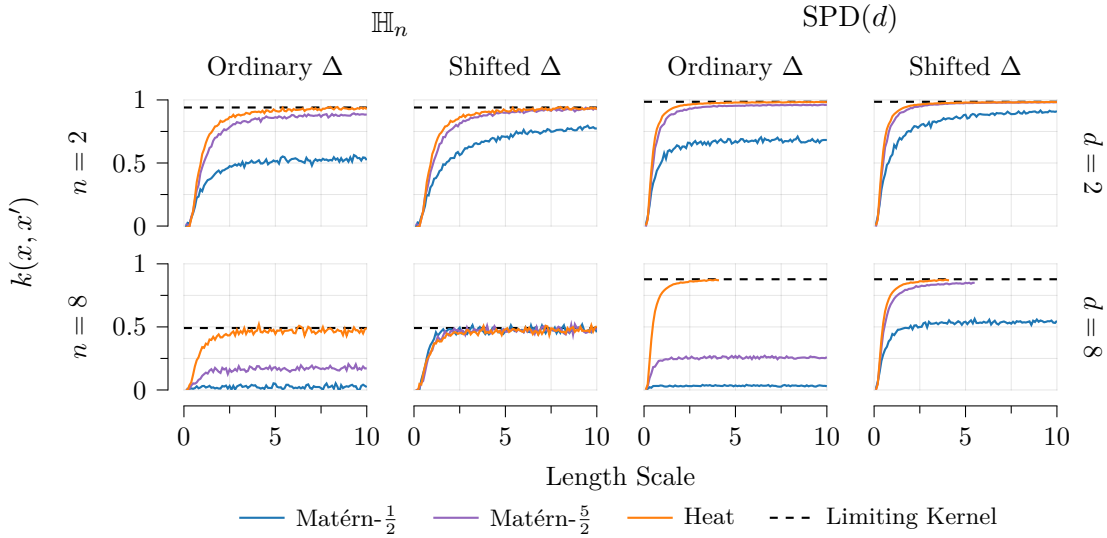


Figure 8: Illustration of how $k(x, x')$ changes as κ increases. To reduce noise stemming from approximation error, we average over x, x' for which $k(x, x')$ is theoretically the same. We see that behavior depends on the smoothness and choice of Laplacian: for the ordinary Laplacian, or ν small enough, the limiting value may be significantly smaller than the value of the *limiting kernel* $k^{(0)}(g_1 \cdot H, g_2 \cdot H) = \pi^{(0)}(g_2^{-1} \cdot g_1)$. In contrast, kernels defined using the shifted Laplacian with ν large enough are guaranteed to reach this limit, albeit with a rate of convergence that depends on the geometry of the space and the kernel’s smoothness, where smoother kernels tend to converge to their respective limiting values faster.

identity orbit $e \cdot H$. In the Euclidean or compact Riemannian cases, the analog of $\pi^{(0)}$ would be the constant function 1, and this phenomenon does not occur. Figure 7 shows $\pi^{(0)}$ for hyperbolic space and the space of symmetric positive-definite matrices.

Overall, the behavior of heat and Matérn kernels—whether one works with the shifted Laplacian, or not—is similar to that of Euclidean or compact Riemannian kernels, except for the aforementioned large length scale behavior. For Matérn kernels specifically, this behavior also depends on the choice of Laplacian. Recall that the density $p_{\nu, \kappa}(\lambda)$ of the measure $\mu_{\nu, \kappa, \sigma^2}$ with respect to the measure $|c(\lambda)|^{-2} d\lambda$ is

$$p_{\nu, \kappa}(\lambda) \propto \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu - n/2} \quad (60)$$

where $\nu < \infty$. If we had $\|\rho\|^2 = 0$ and $c(\lambda) = 1$ as in the Euclidean case, then $\mu_{\nu, \kappa, \sigma^2}$ would be the t -distribution, which converges to Dirac delta function as $\kappa \rightarrow \infty$. The respective kernel, when evaluated at $g_1 \cdot H, g_2 \cdot H$, would then converge to $\pi^{(0)}(g_2^{-1} \cdot g_1)$ —the closest thing one can get to a constant in our setting. This is indeed what happens when the shifted Laplacian is used, granted that ν is large enough for the singularity of $p_{\nu, \kappa}$ to be stronger than the zero of $c(\lambda)^{-2}$. However, this Euclidean-like behavior never happens when the ordinary Laplace–Beltrami operator is used. This claim is formalized in the following proposition.

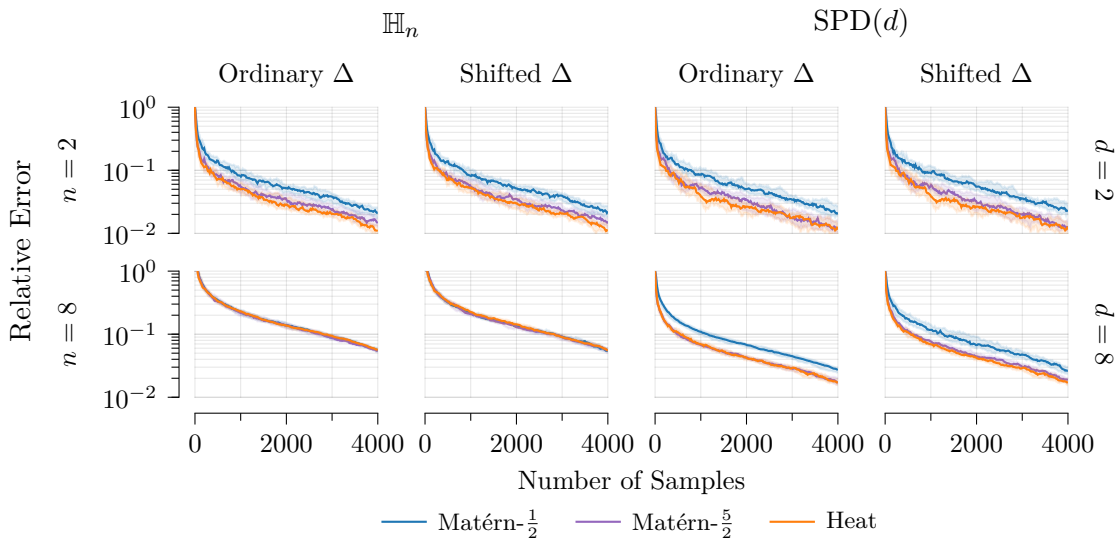


Figure 9: Approximation error for the random Fourier feature approximation (62), as a function of the number of Monte Carlo samples. Specifically, we plot relative error, which is equivalent to selecting a normalization for each curve so that its value at the origin equals one. We display the median, along with 25% and 75% quantiles, calculated via 20 different random seeds. We see that error decreases at the rate expected given the Monte Carlo nature of the approximation.

Proposition 29. *For all $\nu < \infty$, the spectral measure $\mu_{\nu, \kappa, \sigma^2}$ does not converge weakly to a Dirac measure as $\kappa \rightarrow \infty$. In contrast, if either $\nu = \infty$, or if we instead work with $\hat{\mu}_{\nu, \kappa, \sigma^2}$ and ν is large enough, then the respective measure weakly converges to a Dirac measure.*

Proof. Appendix C. □

The practical effect of this difference is that it further restricts the ability of Matérn kernels to capture large length scale behavior. The reason for this difference is that, in the heat kernel, the term $\|\rho\|^2$ enters the spectral density in a multiplicative manner that can be absorbed into the normalizing constant. This does not occur in Matérn kernels defined using the ordinary Laplacian. If one instead works with Matérn kernels defined using the shifted Laplacian, the term $\|\rho\|^2$ vanishes, giving

$$\hat{p}_{\nu, \kappa}(\lambda) \propto \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu - n/2} \quad (61)$$

which resembles the Euclidean case more closely. This means that, Matérn kernels defined using the shifted Laplacian are better able to represent functions which are more spatially constant, and any remaining limitations in this ability arise from the geometry of the space, that is, from the zonal spherical functions $\pi^{(\lambda)}$, and not from the choice of kernel. This provides a practical reason to work with the shifted rather than ordinary Laplacian.

We conclude by examining large length scale behavior empirically. Figure 8 shows how $k(x, x')$, for a fixed set of points $x \neq x'$, changes as one varies the length scale, for the

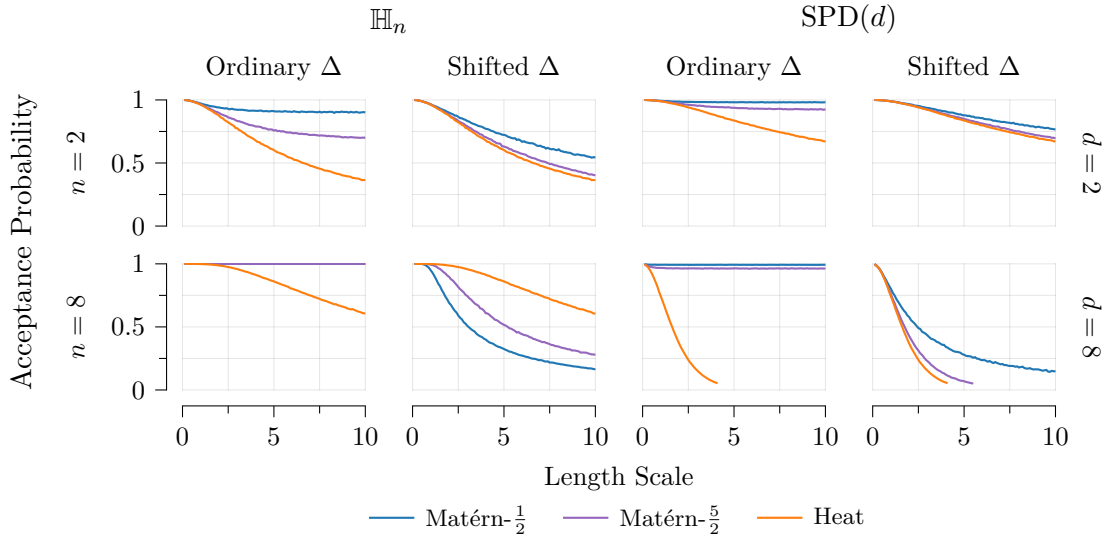


Figure 10: Variability in acceptance probability of the rejection samplers for the respective Matérn and heat kernels’ spectral measures, depending on smoothness, length scale, and dimension of the space. Acceptance probabilities concentrate around one for small length scales, and decrease for larger ones. This rate of decrease is faster in higher dimension.

hyperbolic space and space of symmetric positive definite-matrices examples developed in this work. We see that, if the ordinary Laplacian is used, $k(x, x')$ can converge to a significantly smaller value, compared to if the shifted Laplacian is used—especially for less-smooth kernels. We also see that the rate of convergence depends on the interplay between the geometry of the space, the kernel’s smoothness, and the choice of Laplacian. We will examine this behavior further in the sequel, in the context of rejection sampler efficiency—to do so, we proceed to discuss further approximation considerations. To conclude, note that \mathbb{H}_2 and $\text{SPD}(2)$ behave similarly: this occurs because of the factorization discussed in Section 5.2.4 which connects the two spaces. In contrast, \mathbb{H}_8 and $\text{SPD}(8)$ exhibit a wider range of behavior.

5.3.2 EMPIRICAL APPROXIMATION ERROR AND REJECTION SAMPLER EFFICIENCY

We now examine approximation error for the kernels considered. We focus on the random Fourier feature approximation (15), given by

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{\sigma^2}{L} \sum_{l=1}^L e^{(i\lambda_l + \rho)a(h_l \cdot g_1)} \overline{e^{(i\lambda_l + \rho)a(h_l \cdot g_2)}} \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k \quad h_l \sim \mu_H \quad (62)$$

because, unlike the approximation in (12), it is guaranteed to be symmetric and positive semi-definite for all L , and is thus a more practical choice for many applications. By Theorem 8, the variance of this approximation is uniformly bounded on compact subsets. To get an idea for how accurate this approximation is, we evaluate it empirically on a set of hyperbolic spaces \mathbb{H}_n , and spaces of symmetric positive definite matrices $\text{SPD}(d)$.

Figure 9 shows that approximation error, quantified using relative L^2 -error from an approximate ground-truth kernel computed using 5000 Monte Carlo samples. This quantity is seen to decrease at the usual square-root Monte Carlo rate.

For Matérn and heat kernels, the approximation (62) requires sampling from the respective spectral measure, which for the spaces of interest—except \mathbb{H}_n with odd n —involves rejection sampling. To understand the efficiency of these samplers, we evaluate their acceptance probabilities for \mathbb{H}_n and $\text{SPD}(d)$, and examine how this changes as a function of the dimension of the space, as well as smoothness and length scale parameters ν and κ .

Figure 10 shows results. For sufficiently-small length scales, the proposed rejection sampling scheme is very efficient, with near-unit acceptance probabilities. As the length scale increases, acceptance probabilities drop, with higher-dimensional spaces exhibiting a faster drop than lower-dimensional spaces. This rate is smoothness-dependent: for \mathbb{H}_2 , $\text{SPD}(2)$, and $\text{SPD}(8)$, the rejection sampler for the less-smooth Matérn-1/2 kernel is more efficient compared to that of the smoother heat kernel, while for \mathbb{H}_8 the opposite holds. This indicates that efficiency depends on the interplay between the kernel and the geometry of the space.

While this may at first seem like a problem, observe by comparing Figures 8 and 10, that regions where acceptance probabilities are low are also regions where the kernel has converged to its large-length-scale limit. In those cases, one therefore can simply approximate k by its limit, and avoid Monte Carlo error as an extra benefit. We conclude, empirically, that the proposed rejection samplers are inefficient in the same situations where they are unneeded.

6. Conclusion

In this paper, we developed computationally tractable Gaussian process models on a natural class of non-compact Riemannian manifolds where this is possible, namely symmetric spaces. We derived specific kernel expressions and sampling algorithms for two concrete classes of spaces which are important in applications: hyperbolic spaces and the manifolds of symmetric positive definite matrices. Our work brings together a number of results from probability and statistics, differential geometry, harmonic analysis, and representation theory. In total, we obtain kernel matrix approximations which are guaranteed to be positive semi-definite, making them simple and reliable to use. We hope our techniques enable Gaussian processes to be applied in a wider set of potential applications with geometric structure.

Acknowledgements

The authors are grateful to Prof. Mikhail Lifshits for providing useful feedback. IA and AS were supported by RSF grant №21-11-00047. AT was supported by Cornell University, jointly via the Center for Data Science for Enterprise and Society, the College of Engineering, and the Ann S. Bowers College of Computing and Information Science. VB was supported by an ETH Zürich Postdoctoral Fellowship. We further acknowledge support from Huawei Research and Development.

Appendix A. Additional Details on Non-compact Symmetric Spaces

Here, provide further details on the geometric and algebraic structure of symmetric spaces. We mostly follow Helgason (1962, 1979, 1994). We begin with the arguably more intuitive geometric picture.

Definition 30. *Let M be a connected manifold with Riemannian metric g . For a point $p \in M$ let $T_p M$ be a tangent space at p . Let s_p be the GEODESIC SYMMETRY defined on the image U of a sufficiently small ball $B(0_p, \varepsilon) \subset T_p M$ under the exponential map evaluated at p —or, more precisely, by*

$$s_p \exp(v) = \exp(-v) \quad v \in B(0_p, \varepsilon). \quad (63)$$

The space M is called a LOCALLY SYMMETRIC SPACE if, for all points $p \in M$, the map s_p is an isometry. It is called a (globally) SYMMETRIC SPACE if this isometry can be extended to an isometry of the whole space M .

Roughly speaking, this means that a Riemannian manifold (M, g) is a symmetric space if all point reflections are isometries of the manifold. This holds for such spaces as spheres, Grassmannians, Euclidean, hyperbolic spaces, and many others.

From the definition of a symmetric space, it immediately follows that M is geodesically complete, since every geodesic can be continued by reflection about its endpoints. Geodesic completeness implies, by the Hopf–Rinow Theorem, that for every $p, q \in M$, we can find a geodesic $\gamma : [0, 1] \rightarrow M$ that connects p and q , namely $\gamma(0) = p$ and $\gamma(1) = q$ (Helgason, 1979, Theorem 10.4). Reflecting points around the midpoint of a geodesic, $\gamma(1/2)$, sends p to q and q to p . Thus, the group of isometries acts transitively on M .

Let G be the isometry group of M . By the Myers–Steenrod Theorem (Myers and Steenrod, 1939) G can naturally be given the structure of a Lie group. Fix an arbitrary point $p \in M$ on the manifold. One can show that p generates a Lie subgroup H of isometries that leaves p fixed in the sense that $g \triangleright p = p$ if and only if $g \in H$: call H the *isotropy group*. This allows one to identify M as a homogeneous space G/H .

We can use this representation to study symmetric spaces in an algebraic, Lie-theoretic way. The reflection map s_p defines an involution σ on G , namely

$$\sigma(g) = s_p \circ g \circ s_p \quad g \in G. \quad (64)$$

The main technical tool for studying Lie groups is the notion of a *Lie algebra*, which allows one to linearize various questions that concern Lie groups.

Definition 31. *A LIE ALGEBRA \mathfrak{g} is a vector space equipped with a bilinear operation $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$, called the LIE BRACKET, which is a bilinear form that is ALTERNATING, namely $[X, X] = 0$, and satisfies the JACOBI IDENTITY*

$$[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0. \quad (65)$$

Every Lie group gives rise to a Lie algebra, which is constructed as follows.

Definition 32. Let G be a Lie group. Define the LIE ALGEBRA OF A LIE GROUP to be $\mathfrak{g} = T_e G$, the tangent space of G at the identity element e , equipped with the Lie bracket constructed as follows. One can show that each element of \mathfrak{g} is in bijective correspondence with left-invariant vector fields on G , namely vector fields X such that $(DL_g)X_h = X_{g \cdot h}$, where $L_g : G \rightarrow G$ is the left multiplication map $L_g(h) = g \cdot h$. Every vector field X defines a derivation, namely $X(Y)(g)$, which is a directional derivative of Y along the tangent vector X_g . With this notation, the Lie bracket of a pair of vector fields is defined as $[X, Y] = X(Y) - Y(X)$, which carries over to the Lie algebra by the aforementioned bijective correspondence.

In the most important case of a linear Lie group, that is, a closed subgroup of $GL(n)$, the tangent space is naturally identified with a subspace of the vector space of all matrices $M(n, \mathbb{R})$. Here, the Lie bracket is simply the additive commutator $[\mathbf{X}, \mathbf{Y}] = \mathbf{X}\mathbf{Y} - \mathbf{Y}\mathbf{X}$.

Let \mathfrak{g} and \mathfrak{h} be the Lie algebras of G and H , respectively, where G and H are as above. Since σ is an involution, meaning σ^2 is the identity map, the differential of σ at the identity, $\theta = D_e \sigma$ is also an involution, so it is a linear map with eigenvalues $\{+1, -1\}$. The eigenspace corresponding to $+1$ is exactly \mathfrak{h} . Let \mathfrak{p} be the eigenspace corresponding to -1 , which can be identified with the tangent space $T_p M$ (Helgason, 1979, Theorem 3.3(iii)). We call the pair $(\mathfrak{h}, \mathfrak{p})$ a *Cartan pair*.

Recall that a Lie algebra is called *semi-simple* if it contains no non-zero abelian ideals.¹³ Equivalently, a Lie algebra is semi-simple if it can be decomposed into a direct sum of *simple*¹⁴ Lie algebras. A Lie group is called *semi-simple* if its Lie algebra is *semi-simple*. The principal source of examples of semi-simple Lie groups are the simple Lie groups, which are defined as connected non-abelian groups containing no connected normal subgroups—note that such groups can have a discrete normal subgroup, as $SL(n)$ does, for example.

Our spaces of interest have the form G/H where G is a non-compact *semi-simple* Lie group with *finite center* and H is a maximal, by inclusion, compact subgroup of G . Most semi-simple Lie groups which occur in practice satisfy the latter assumption. Let \mathfrak{g} and \mathfrak{h} be Lie algebras corresponding to G and to H . Recall that we use e to refer to the identity element of G . Define the *adjoint action* of $X \in \mathfrak{g}$ by

$$\text{ad}(X) : \mathfrak{g} \rightarrow \mathfrak{g}, \quad \text{ad}(X)(Y) = [X, Y]. \quad (66)$$

where $[\cdot, \cdot]$ denotes the Lie bracket of \mathfrak{g} . Define the *Killing form* B of \mathfrak{g} to be a symmetric bilinear form on $\mathfrak{g} \times \mathfrak{g}$ defined by

$$B(X, Y) = \text{tr}(\text{ad}(X) \circ \text{ad}(Y)) \quad (67)$$

where \circ denotes composition of functions. One can show that a non-zero Lie algebra is semi-simple if and only if its Killing form is non-degenerate.

The subalgebra \mathfrak{p} defined previously can be viewed as the orthogonal complement of \mathfrak{h} with respect to B . The differential at e of involution σ , for $X \in \mathfrak{h}$ and $Y \in \mathfrak{p}$, can be written

$$d_e \sigma(X + Y) = X - Y. \quad (68)$$

13. A Lie algebra \mathfrak{g} is called *abelian* if $[X, Y] = 0$ for all X, Y in \mathfrak{g} . A subspace \mathfrak{i} of a Lie algebra \mathfrak{g} is called an *ideal* if it is stable under the Lie bracket action of \mathfrak{g} : that is, for any $X \in \mathfrak{g}$ and $Y \in \mathfrak{i}$ one has $[X, Y] \in \mathfrak{i}$. This means that \mathfrak{i} is a subalgebra.

14. A non-abelian Lie algebra is called *simple* if it does not contain any non-trivial ideals.

Within the subalgebra \mathfrak{p} , let \mathfrak{a} be a maximal abelian subalgebra, and let \mathfrak{a}^* be its dual vector space. For a linear functional $\alpha \in \mathfrak{a}^*$ we can define the *root space*

$$\mathfrak{g}_\alpha = \{X \in \mathfrak{g} : [X, Y] = \alpha(Y)X, \forall Y \in \mathfrak{a}\}. \quad (69)$$

This is a vector space: let m_α be its dimension. If $m_\alpha > 0$ and $\alpha \neq 0$, then α is called a *restricted root*. Let Σ be the set of restricted roots. One can show that the number of restricted roots is finite. A detailed description of restricted roots and their multiplicities for irreducible symmetric spaces can be found in Araki (1962).

From the restricted roots, one can choose a set of *positive roots* Σ^+ defined by the following conditions.

1. For each restricted root α , exactly one of α and $-\alpha$ is contained in Σ^+ .
2. For any $\alpha, \beta \in \Sigma^+$, either $\alpha + \beta \in \Sigma^+$ or $\alpha + \beta$ is not a restricted root.

With these notions, we can define Lie subalgebra \mathfrak{n} of \mathfrak{g} by

$$\mathfrak{n} = \bigoplus_{\alpha \in \Sigma^+} \mathfrak{g}_\alpha. \quad (70)$$

Putting these components together, we arrive at the *Iwasawa decomposition* of the Lie algebra \mathfrak{g} , which is

$$\mathfrak{g} = \mathfrak{n} \oplus \mathfrak{a} \oplus \mathfrak{h}. \quad (71)$$

This decomposition can be lifted from the Lie algebra to the Lie group.

Theorem 33. *Let G and H be defined as above and let $\mathfrak{g} = \mathfrak{n} \oplus \mathfrak{a} \oplus \mathfrak{h}$ be an Iwasawa decomposition of the Lie algebra \mathfrak{g} of G . Then there are subgroups A and N of G , whose respective Lie algebras are \mathfrak{a} and \mathfrak{n} , such that the multiplication map $N \times A \times H \rightarrow G$ given by $(n, a, h) \mapsto n \cdot a \cdot h$ is diffeomorphism. Moreover, the subgroups A and N are simply connected.*

For $g \in G$, let $N(g) \in N$, $A(g) \in A$, and $H(g) \in H$ be the elements for which $g = N(g) \cdot a(g) \cdot H(g)$. We now illustrate a few examples.

Example 34. *For $G = \text{SL}(n)$ let N be the subgroup of upper-triangular matrices with 1 on the diagonal, A be the subgroup of diagonal matrices with non-negative entries, and $H = \text{SO}(n)$. Then the Iwasawa decomposition becomes the well-known RQ-decomposition.*

Example 35. *Let $G = \text{SO}_0(1, n)$ be the identity component of the Lorentz group—this group is introduced in Section 5.1. Its Lie algebra \mathfrak{g} takes the form*

$$\mathfrak{g} = \{\mathbf{X} \in \text{M}(n+1, \mathbb{R}) : \forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^{n+1} \langle \mathbf{X}\mathbf{u}, \mathbf{v} \rangle_M + \langle \mathbf{u}, \mathbf{X}\mathbf{v} \rangle_M = 0\} \quad (72)$$

$$= \{\mathbf{X} \in \text{M}(n+1, \mathbb{R}) : \mathbf{X}^\top \mathbf{B} + \mathbf{B}\mathbf{X} = 0\} \quad (73)$$

$$= \left\{ \begin{pmatrix} 0 & \mathbf{b}^\top \\ \mathbf{b} & \mathbf{D} \end{pmatrix} : \mathbf{b} \in \mathbb{R}^n, \mathbf{D}^\top = -\mathbf{D} \right\} \quad (74)$$

where $\mathbf{B} = \text{diag}(-1, 1, \dots, 1)$ is the Gram matrix of $\langle \cdot, \cdot \rangle_M$ defined in (28) and, as before, $\text{M}(n+1, \mathbb{R})$ denotes the vector space of $(n+1) \times (n+1)$ matrices.

We now present a computational recipe for obtaining the Iwasawa decomposition, adapted from Sawyer (2016). Let \mathbf{Q} be the following orthogonal matrix:

$$\mathbf{Q} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \mathbf{0} \\ \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{n-1} \end{pmatrix}. \quad (75)$$

If $\mathbf{Q}^\top \mathbf{M} \mathbf{Q} = \mathbf{N} \mathbf{A} \mathbf{H}$ is the Iwasawa decomposition inside $\mathrm{SL}(n+1)$, then the Iwasawa decomposition of $\mathbf{M} \in G$ is

$$\mathbf{M} = \mathbf{Q} \mathbf{N} \mathbf{Q}^\top \cdot \mathbf{Q} \mathbf{A} \mathbf{Q}^\top \cdot \mathbf{Q} \mathbf{H} \mathbf{Q}^\top = N(\mathbf{M}) \cdot A(\mathbf{M}) \cdot H(\mathbf{M}). \quad (76)$$

The resulting matrices can be described explicitly. Firstly, $A(\mathbf{M})$ has the following form, for some $t > 0$:

$$\mathbf{A}_t = \exp(a(t)) = \begin{pmatrix} \cosh t & \sinh t & \mathbf{0} \\ \sinh t & \cosh t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{n-1} \end{pmatrix} \quad a(t) = \begin{pmatrix} 0 & t & \mathbf{0} \\ t & 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0}_{n-1} \end{pmatrix} \in \mathfrak{a} \quad (77)$$

where $\mathbf{0}_{n-1}$ is the $(n-1) \times (n-1)$ matrix of zeros. For a matrix $\mathbf{X} \in \mathfrak{g}$ partitioned as

$$\mathbf{X} = \begin{pmatrix} 0 & a & \mathbf{b}^\top \\ a & 0 & -\mathbf{c}^\top \\ \mathbf{b} & \mathbf{c} & \mathbf{D} \end{pmatrix} \quad \mathbf{b}, \mathbf{c} \in \mathbb{R}^{n-1} \quad \mathbf{D}^\top = -\mathbf{D} \quad (78)$$

the commutator $[\mathbf{X}, a(t)]$ takes the form

$$[\mathbf{X}, a(t)] = t \cdot \begin{pmatrix} 0 & 0 & \mathbf{c}^\top \\ 0 & 0 & -\mathbf{b}^\top \\ \mathbf{c} & \mathbf{b} & \mathbf{0}_{n-1} \end{pmatrix}. \quad (79)$$

Equating this to $\alpha(a(t))\mathbf{X}$ with $\alpha(a(t)) = \lambda t$ for some λ , one gets from the definition of the root spaces (69) that $\lambda = \pm 1$. Choosing $\alpha(a(t)) = t$ as the positive root results in the elements of \mathfrak{n} taking the form

$$\begin{pmatrix} 0 & 0 & \mathbf{b}^\top \\ 0 & 0 & -\mathbf{b}^\top \\ \mathbf{b} & \mathbf{b} & \mathbf{0}_{n-1} \end{pmatrix} \quad \mathbf{b} \in \mathbb{R}^{n-1} \quad (80)$$

thus the matrix $N(\mathbf{M}) = \exp(\mathbf{X})$ is of the form

$$\mathbf{N}_{\mathbf{b}} = \begin{pmatrix} 1 + \Delta & \Delta & \mathbf{b}^\top \\ -\Delta & 1 - \Delta & -\mathbf{b}^\top \\ \mathbf{b} & \mathbf{b} & \mathbf{I}_{n-1} \end{pmatrix} \quad \mathbf{b} \in \mathbb{R}^{n-1} \quad \Delta = \frac{\|\mathbf{b}\|^2}{2}. \quad (81)$$

The abelian Lie algebra \mathfrak{a} endowed with the inner product given by the restriction of the Killing form $B|_{\mathfrak{a} \times \mathfrak{a}}$ can be identified with the Euclidean space \mathbb{R}^l for some l . We say that $\dim \mathfrak{a} = l$ is the *rank* of the symmetric space G/H .

A vector $X \in \mathfrak{a}$ is called *regular* if $\alpha(X) \neq 0$ for all $\alpha \in \Sigma$. Since Σ is a finite set, the connected components of the set of regular elements are convex cones, which are called *Weyl chambers*. By construction, the sign of $\alpha(X)$ is constant on each Weyl chamber. Let \mathfrak{a}_+ be the *positive Weyl chamber*, which is the unique Weyl chamber such that all elements of \mathfrak{a}_+ take positive values on Σ^+ . Define *Weyl group* W to be the finite subgroup of the group of bijective linear transforms $\text{GL}(\mathfrak{a})$, which is generated by reflections relative to the hyperplanes $L_\alpha = \{\alpha(X) = 0 | X \in \mathfrak{a}\}$, where $\alpha \in \Sigma$. This group acts by permutations on the set of Weyl chambers.

Using the Killing form B , we can identify \mathfrak{a} and \mathfrak{a}^* : for $\lambda \in \mathfrak{a}^*$ we associate $X_\lambda \in \mathfrak{a}$ such that $B(X_\lambda, X) = \lambda(X)$ for all $X \in \mathfrak{a}$. Thus, the space \mathfrak{a}^* is naturally endowed with the inner product $\langle \lambda, \mu \rangle = B(X_\lambda, X_\mu)$. Finally, the Weyl chamber and the Weyl group can be lifted to \mathfrak{a}^* by writing

$$\mathfrak{a}_+^* = \{\lambda \in \mathfrak{a}^* : X_\lambda \in \mathfrak{a}_+\}. \quad (82)$$

Similarly, the Weyl group W is generated by reflections relative to the hyperplanes $L_\alpha^* = \{\alpha(X_\lambda) = 0 : \lambda \in \mathfrak{a}^*\}$, where $\alpha \in \Sigma$, which are defined on the dual space of \mathfrak{a} .

Appendix B. Spherical Fourier Transform

Here we introduce the notion of Fourier transform appropriate for non-compact symmetric spaces. We use notation and definitions from Appendix A. Additionally, we define $L_H^p(G, \mu_G)$ to be the subspace of $L^p(G, \mu_G)$ which consists of H -bi-invariant functions, meaning all $f \in L^p(G, \mu_G)$ such that $f(h_1 \cdot g \cdot h_2) = f(g)$.

Result 36. *The SPHERICAL FOURIER TRANSFORM of $f \in L_H^1(G, \mu_G)$ is given by*

$$\widehat{f}(\lambda) = \int_G f(g) \overline{\pi^{(\lambda)}(g)} d\mu_G(g) \quad (83)$$

for $\lambda \in \mathfrak{a}^*$, where the function \widehat{f} is uniquely determined by its values on the positive Weyl chamber $\mathfrak{a}_+^* \subseteq \mathfrak{a}^*$. Moreover, every $f \in L_H^1(G, \mu_G) \cap L_H^2(G, \mu_G)$ can be recovered from its spherical Fourier transform using the inversion formula

$$f(g) = \int_{\mathfrak{a}_+^*} \widehat{f}(\lambda) \pi^{(\lambda)}(g) |c(\lambda)|^{-2} d\lambda = \frac{1}{|W|} \int_{\mathfrak{a}^*} \widehat{f}(\lambda) \pi^{(\lambda)}(g) |c(\lambda)|^{-2} d\lambda \quad (84)$$

where integration is taken with respect to the Lebesgue measure $d\lambda$ on \mathfrak{a}^* , $c(\lambda)$ is a space-dependent function called HARISH-CHANDRA'S c -FUNCTION, and $|W|$ denotes the order of the Weyl group. Finally, the map $f \mapsto \widehat{f}$ is an isometry between the spaces $L_H^2(G, \mu_G)$ and $L^2(\mathfrak{a}_+^*, |c(\lambda)|^{-2} d\lambda)$.

Proof. Helgason (1994, Chapter III, Theorem 1.5). □

Crucially, the zonal spherical functions $\pi^{(\lambda)}$ are eigenfunctions of the Laplace–Beltrami operator on G (Gangolli, 1968). Thanks to this, and to the fact that the spherical Fourier transform is built from these zonal spherical functions, we get the following.

Result 37. *The solution of the heat equation on a symmetric space G/H is given by*

$$\mathcal{P}(t, g_1 \cdot H, g_2 \cdot H) \propto \int_{\mathfrak{a}^*} e^{-t(\|\lambda\|^2 + \|\rho\|^2)} \cdot \pi^{(\lambda)}(g_2^{-1}g_1) \cdot |c(\lambda)|^{-2} d\lambda \quad (85)$$

where $\rho = \frac{1}{2} \sum_{\alpha \in \Sigma^+} m_\alpha \alpha$ is the half-sum of all positive roots, integration is taken with respect to the Lebesgue measure on \mathfrak{a}^* , and $c(\lambda)$ is Harish-Chandra's c -function.

Proof. The full proof can be found in Gangolli (1968). A sketch of the argument is as follows: let Δ denote the Laplace-Beltrami operator on G/H . One can show that $\Delta \pi^{(\lambda)} = -(\|\lambda\|^2 + \|\rho\|^2) \pi^{(\lambda)}$. Then, similarly to the Euclidean case, one can apply the spherical Fourier transform to the sides of the heat equation, solve the resulting algebraic equation, and take the inverse spherical Fourier transform to obtain the solution. \square

Appendix C. Proofs

Proposition 6. *The standard deviation of the estimator (12) is uniformly bounded by $\frac{\sigma^2}{\sqrt{L}}$.*

Proof. Since the estimator is an average of IID random variables, it is enough to estimate the variance of one term. Write

$$\text{Var}(\sigma^2 e^{(i\lambda+\rho)^\top a(h \cdot g)}) \leq \mathbb{E} \left| \sigma^2 e^{(i\lambda+\rho)^\top a(h \cdot g)} \right|^2 = \sigma^2 \int_{\mathfrak{a}^*} \int_H \left| e^{(i\lambda+\rho)^\top a(h \cdot g)} \right|^2 d\mu_H(h) d\mu_k(\lambda). \quad (86)$$

Denote the identity element of the group G by e and apply Result 7 with $g_1 = g_2 = g$ to get

$$\pi^{(\lambda)}(e) = \int_H e^{(i\lambda+\rho)^\top a(h \cdot g)} \overline{e^{(i\lambda+\rho)^\top a(h \cdot g)}} d\mu_H(h) = \int_H \left| e^{(i\lambda+\rho)^\top a(h \cdot g)} \right|^2 d\mu_H(h). \quad (87)$$

On the other hand, by definition, we have

$$\pi^{(\lambda)}(e) = \int_H e^{(i\lambda+\rho)^\top a(h)} d\mu_H(h). \quad (88)$$

Since $h = e \cdot e \cdot h$, by the uniqueness of Iwasawa decomposition, we have $A(h) = e$ and therefore $a(h) = 0 \in \mathfrak{a}$, implying $\pi^{(\lambda)}(e) = 1$. Hence

$$\text{Var}(\sigma^2 e^{(i\lambda+\rho)^\top a(h \cdot g)}) \leq \sigma^2 \int_{\mathfrak{a}^*} 1 d\mu_k(\lambda) = \sigma^4. \quad (89)$$

Since the variance of an IID sum is equal to the sum of variances of its components, the claim follows. \square

Theorem 8. *For any compact subset $U \subset G/H$ there is a constant $C_U > 0$ such that, for any $g_1 \cdot H, g_2 \cdot H \in U$, the right-hand side of (15) is an unbiased estimator of $k(g_1 \cdot H, g_2 \cdot H)$, whose standard deviation is bounded by $C_U \frac{\sigma^2}{\sqrt{L}}$. The constant C_U , however, does not admit an upper bound: for any fixed L , the standard deviation on the full space G/H is unbounded.*

Proof. Recall that the estimator under consideration is

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{\sigma^2}{L} \sum_{l=1}^L e^{(i\lambda_l + \rho)^\top a(h_l \cdot g_1)} \overline{e^{(i\lambda_l + \rho)^\top a(h_l \cdot g_2)}} \quad \lambda_l \sim \frac{1}{\sigma^2} \mu_k \quad h_l \sim \mu_H. \quad (90)$$

Since it is an average of IID variables, it is enough to estimate variance of only one term. To this end, write

$$\text{Var} \left(\sigma^2 e^{(i\lambda + \rho)^\top a(h \cdot g_1)} \overline{e^{(i\lambda + \rho)^\top a(h \cdot g_2)}} \right) \leq \mathbb{E} \left| \sigma^2 e^{(i\lambda + \rho)^\top a(h \cdot g_1)} \overline{e^{(i\lambda + \rho)^\top a(h \cdot g_2)}} \right|^2 \quad (91)$$

$$= \sigma^2 \int_{\mathfrak{a}^*} \int_H |e^{(i\lambda + \rho)^\top a(h \cdot g_1)} e^{(-i\lambda + \rho)^\top a(h \cdot g_2)}|^2 d\mu_H(h) d\mu_k(\lambda) \quad (92)$$

$$= \sigma^2 \int_{\mathfrak{a}^*} \int_H e^{2\rho^\top a(h \cdot g_1)} e^{2\rho^\top a(h \cdot g_2)} d\mu_H(h) d\mu_k(\lambda) \quad (93)$$

$$= \sigma^4 \int_H e^{2\rho^\top a(h \cdot g_1)} e^{2\rho^\top a(h \cdot g_2)} d\mu_H(h) \quad (94)$$

$$\leq \sigma^4 \left(\int_H e^{4\rho^\top a(h \cdot g_1)} d\mu_H(h) \right)^{\frac{1}{2}} \left(\int_H e^{4\rho^\top a(h \cdot g_2)} d\mu_H(h) \right)^{\frac{1}{2}}. \quad (95)$$

Recall that for $\lambda \in \mathfrak{a}^* = \mathbb{R}^n$ we have

$$\pi^{(\lambda)}(g) = \int_H e^{(i\lambda + \rho)^\top a(h \cdot g)} d\mu_H(h). \quad (96)$$

Let us formally extend the definition of $\pi^{(\lambda)}$ to $\lambda \in \mathbb{C}^n$. Then (95) directly implies

$$\text{Var} \left(\sigma^2 e^{(i\lambda + \rho)^\top a(h \cdot g_1)} \overline{e^{(i\lambda + \rho)^\top a(h \cdot g_2)}} \right) \leq \sigma^4 \left(\pi^{(-3i\rho)}(g_1) \right)^{1/2} \left(\pi^{(-3i\rho)}(g_2) \right)^{1/2}. \quad (97)$$

By Helgason (2000, Chapter IV, §2.1), the functions $\pi^{(\lambda)}$ are well-defined and continuous for all $\lambda \in \mathbb{C}^n$: more precisely, these functions are *defined* as continuous solutions of certain functional equation, after which Helgason (2000, Ch. IV, §4.2, Th. 4.3) proves that for $\lambda \in \mathbb{C}^n$ they have the form (96). Therefore, on any compact $U \subseteq G$ the function $\pi^{(-3i\rho)}$ is bounded. Denoting $C_U = \sigma^4 \max_{g \in U} \pi^{-3i\rho}(g)$, we get

$$\text{Var} \left(\sigma^2 e^{(i\lambda + \rho)^\top a(h \cdot g_1)} \overline{e^{(i\lambda + \rho)^\top a(h \cdot g_2)}} \right) \leq \sigma^4 C_U \quad g_1, g_2 \in U \subseteq G. \quad (98)$$

This proves the first part of the claim.

To prove the second part of the claim we show that the estimator of the variance, i.e. in the case of $g_1 = g_2 = g \in G$, is globally unbounded. By the proof of Proposition 6, we have

$$1 = \pi^\lambda(e) = \int_H e^{2\rho^\top a(h \cdot g)} d\mu_H(h) = \mathbb{E} e^{2\rho^\top a(h \cdot g)}. \quad (99)$$

Thus

$$\text{Var} \left(\sigma^2 e^{(i\lambda + \rho)^\top a(h \cdot g)} \overline{e^{(i\lambda + \rho)^\top a(h \cdot g)}} \right) = \text{Var} \left(\sigma^2 e^{2\rho^\top a(h \cdot g)} \right) \quad (100)$$

$$= \mathbb{E} \left| \sigma^2 e^{2\rho^\top a(h \cdot g)} \right|^2 - \left| \mathbb{E} \sigma^2 e^{2\rho^\top a(h \cdot g)} \right|^2 \quad (101)$$

$$= \sigma^4 \left(\pi^{(-3i\rho)}(g) - 1 \right). \quad (102)$$

By Helgason (2000, Chapter IV, Theorem 8.1), the necessary and sufficient condition for $\pi^{(\lambda)} : G \rightarrow \mathbb{C}$ to be bounded, in the sense $\sup_{g \in G} |\pi^{(\lambda)}(g)| < \infty$, is

$$\lambda \in \mathfrak{a}^* + i \operatorname{Conv}(\{w\rho : w \in W\}) \quad (103)$$

where W is the Weyl group equipped with its natural action on \mathfrak{a}^* —see Appendix A for the respective definition—and Conv denotes the convex hull of a set. Observe that $-3i\rho \notin \mathfrak{a}^* + i \operatorname{Conv}(\{w\rho : w \in W\})$ because $|w\rho| = |\rho|$ and thus $3\rho \notin \operatorname{Conv}(\{w\rho : w \in W\})$. Hence, the right-hand side of (102) is unbounded, which completes the proof. \square

Theorem 13. *For all $\nu, \kappa, \sigma^2 > 0$ and all $g_2 \in G$, the Matérn kernels $k_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$ and $\widehat{k}_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$ lie in the Sobolev space $H^{\nu+n/2}$ if $\nu < \infty$, and in all Sobolev spaces H^α , $\alpha > 0$, if $\nu = \infty$. Thus, they are continuous and possess continuous derivatives of all integer orders strictly less than ν .*

Proof. The characterization of Sobolev spaces on a non-compact symmetric space given by Strichartz (1983, Section 6)—see also Große and Schneider (2013) to clarify equivalence with the classical definition—implies that $f \in H^{\nu+n/2}$ if and only if

$$\int_{\mathfrak{a}^*} \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{\nu+n/2} |\widehat{f}(\lambda)|^2 |c(\lambda)|^{-2} d\lambda < \infty \quad (104)$$

where \widehat{f} is as in (83). Denote $f(g) = k_{\nu, \kappa, \sigma^2}(g \cdot H, e \cdot H)$ and $h(g) = \widehat{k}_{\nu, \kappa, \sigma^2}(g \cdot H, e \cdot H)$. If we assume $\nu < \infty$, then

$$\widehat{f}(\lambda) \propto \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} \quad \widehat{h}(\lambda) \propto \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu-n/2}. \quad (105)$$

Direct computation shows

$$\left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} c(\nu, \kappa, \rho, n) \leq |\widehat{f}(\lambda)| \leq C(\nu, \kappa, \rho, n) \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} \quad (106)$$

$$\left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} \widetilde{c}(\nu, \kappa, \rho, n) \leq |\widehat{h}(\lambda)| \leq \widetilde{C}(\nu, \kappa, \rho, n) \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2}. \quad (107)$$

Observing that, for any $\nu > 0$, we have

$$\int_{\mathfrak{a}^*} \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda < \infty \quad (108)$$

which is shown in the proof of Proposition 29, implies that $f, h \in H^{\nu+n/2}$. The functions $k_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$, $\widehat{k}_{\nu, \kappa, \sigma^2}(\cdot, g_2 \cdot H)$ for $g_2 \neq e$ are simple shifts of the functions f and h , thus the same property holds for them as well. Finally, in the case of $\nu = \infty$, the Fourier transform is upper bounded by $C(\alpha, \kappa, \rho, n) \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\alpha}$ for all $\alpha > 0$, which proves that the respective kernels lie in the Sobolev spaces H^α for all $\alpha > 0$.

Since symmetric spaces, as well as all homogeneous spaces with bi-invariant Riemannian metric, are spaces of bounded geometry (Eldering, 2012), we can use Hebey (2000, Theorem 3.4) to get the claimed continuous differentiability properties from the appropriate analogs of the classical Sobolev embedding theorems. \square

Lemma 15. *We have $A(\mathbf{UM}) = A(\mathbf{M})$ for all $\mathbf{M} \in \text{SO}_0(1, n)$ and all $\mathbf{U} \in H \subseteq \text{SO}_0(1, n)$ of the form*

$$\mathbf{U} = \begin{pmatrix} 1 & 0 & \\ 0 & 1 & \\ & & \tilde{\mathbf{U}} \end{pmatrix} \quad \tilde{\mathbf{U}} \in \text{SO}(n-1). \quad (31)$$

Proof. We use the Iwasawa decomposition of Example 35. With this, every $\mathbf{M} \in \text{SO}_0(1, n)$ can be decomposed into the product

$$\mathbf{M} = \mathbf{N}_{\mathbf{b}} \mathbf{A}_t \mathbf{H} = \begin{pmatrix} 1 + \Delta & \Delta & \mathbf{b}^\top \\ -\Delta & 1 - \Delta & -\mathbf{b}^\top \\ \mathbf{b} & \mathbf{b} & \mathbf{I}_{n-1} \end{pmatrix} \begin{pmatrix} \cosh(t) & \sinh(t) & \\ \sinh(t) & \cosh(t) & \\ & & \mathbf{I}_{n-1} \end{pmatrix} \mathbf{H} \quad (109)$$

where $\mathbf{b} \in \mathbb{R}^{n-1}$, $\Delta = \frac{\|\mathbf{b}\|^2}{2}$, $t \in \mathbb{R}$, $\mathbf{H} \in H$.

Observe that

$$\begin{pmatrix} 1 & & \\ & 1 & \\ & & \tilde{\mathbf{U}} \end{pmatrix} \begin{pmatrix} 1 + \Delta & \Delta & \mathbf{b}^\top \\ -\Delta & 1 - \Delta & -\mathbf{b}^\top \\ \mathbf{b} & \mathbf{b} & \mathbf{I}_{n-1} \end{pmatrix} = \begin{pmatrix} 1 + \Delta & \Delta & \mathbf{b}^\top \\ -\Delta & 1 - \Delta & -\mathbf{b}^\top \\ \tilde{\mathbf{U}}\mathbf{b} & \tilde{\mathbf{U}}\mathbf{b} & \tilde{\mathbf{U}} \end{pmatrix} \quad (110)$$

$$= \begin{pmatrix} 1 + \Delta & \Delta & (\tilde{\mathbf{U}}\mathbf{b})^\top \\ -\Delta & 1 - \Delta & -(\tilde{\mathbf{U}}\mathbf{b})^\top \\ \tilde{\mathbf{U}}\mathbf{b} & \tilde{\mathbf{U}}\mathbf{b} & \mathbf{I}_{n-1} \end{pmatrix} \begin{pmatrix} 1 & & \\ & 1 & \\ & & \tilde{\mathbf{U}} \end{pmatrix}. \quad (111)$$

Also note that \mathbf{U} commutes with \mathbf{A}_t , i.e. $\mathbf{UA}_t = \mathbf{A}_t\mathbf{U}$. Using these two properties, we infer that $\mathbf{UM} = \mathbf{N}_{\tilde{\mathbf{U}}\mathbf{b}} \cdot \mathbf{A}_t \cdot \mathbf{UH}$. Hence, by the uniqueness of the Iwasawa decomposition, we have

$$N(\mathbf{UM}) = \mathbf{N}_{\tilde{\mathbf{U}}\mathbf{b}} \quad A(\mathbf{UM}) = \mathbf{A}_t \quad H(\mathbf{UM}) = \mathbf{UH}. \quad (112)$$

In particular, $A(\mathbf{UM}) = \mathbf{A}_t = A(\mathbf{M})$, which proves the claim. \square

Proposition 16. *For any $m \in \mathbb{N}$, and any coefficients $\alpha_j \geq 0$, the density*

$$p_{\infty, \kappa}(\lambda) \propto \sum_{j=0}^m \alpha_j \lambda^j e^{-\frac{\kappa^2}{2} \lambda^2} \quad \lambda \in \mathbb{R}_{>0} \quad (36)$$

is a mixture of scaled chi-distributed random variables. Specifically, we have

$$p_{\infty, \kappa}(\lambda) = \sum_{j=0}^m c_j p_{y_j}(\lambda) \quad c_j = \frac{\alpha_j / \beta_j}{\sum_{j=0}^m \alpha_j / \beta_j} \quad \beta_j = 2^{\frac{1-j}{2}} \Gamma\left(\frac{j+1}{2}\right)^{-1} \kappa \quad (37)$$

where $\Gamma(\cdot)$ is the gamma function, and $p_{y_j}(\lambda)$ is the density of the random variable $y_j = x_j / \kappa$ where $x_j \sim \chi_{j+1}$ is a chi-distributed random variable with $j+1$ degrees of freedom.

Proof. The densities p_{x_j} of x_j and p_{y_j} of y_j are given by

$$p_{x_j}(\xi) = 2^{\frac{1-j}{2}} \Gamma\left(\frac{j+1}{2}\right)^{-1} \xi^j e^{-\xi^2/2} \quad p_{y_j}(\psi) = p_{x_j}(\kappa\psi)\kappa \quad (113)$$

where the second expression follows from the change of variables formula for densities. Write

$$p_{\infty,\kappa}(\lambda) \propto \sum_{j=0}^m \alpha_j \lambda^j e^{-\frac{\kappa^2}{2}\lambda^2} = \sum_{j=0}^m \alpha_j \beta_j^{-1} \underbrace{\beta_j \lambda^j e^{-\frac{\kappa^2}{2}\lambda^2}}_{p_{y_j}(\lambda)} \propto \sum_{j=0}^m c_j p_{y_j}(\lambda). \quad (114)$$

The right-hand side is an actual density, which corresponds to the mixture of p_{y_j} with coefficients c_j , which proves the claim. \square

Proposition 17. *For any $m \in \mathbb{N}$, $m \leq n-1$, and any coefficients $\alpha_j \geq 0$, the density*

$$p_{\nu,\kappa}(\lambda) \propto \sum_{j=0}^m \alpha_j \lambda^j \left(\frac{2\nu}{\kappa^2} + \lambda^2 + \left(\frac{n-1}{2} \right)^2 \right)^{-\nu - \frac{n}{2}} \quad \lambda \in \mathbb{R}_{>0} \quad (39)$$

is a mixture of transformed beta-prime distributed random variables. Specifically, we have

$$p_{\nu,\kappa}(\lambda) = \sum_{j=0}^m c_j p_{y_j}(\lambda) \quad c_j = \frac{\alpha_j/\beta_j}{\sum_{j=0}^m \alpha_j/\beta_j} \quad \beta_j = 2B\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)^{-1} \gamma^{\nu + \frac{n-j-1}{2}} \quad (40)$$

where B is the beta function, $\gamma = \frac{2\nu}{\kappa^2} + \left(\frac{n-1}{2}\right)^2$, and p_{x_j} is the density of the random variable $y_j = \sqrt{\gamma x_j}$ where $x_j \sim \text{Beta}'\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)$ is beta-prime distributed.

Proof. We have $y_j = \Phi(x_j)$ with $\Phi : [0, \infty) \rightarrow [0, \infty)$ a bijective map given by $\Phi(\xi) = \sqrt{\gamma\xi}$. Then, using the change of variable formula for densities, write

$$p_{y_j}(\psi) = p_{x_j}(\psi^2/\gamma) 2\psi/\gamma \quad p_{x_j}(\xi) = \frac{\xi^{\frac{j-1}{2}} (1+\xi)^{-\nu - \frac{n}{2}}}{B\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)}. \quad (115)$$

Then

$$p_{y_j}(\psi) = \frac{(\psi^2/\gamma)^{\frac{j-1}{2}} (1 + \psi^2/\gamma)^{-\nu - \frac{n}{2}}}{B\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)} 2\psi/\gamma = \frac{2\gamma^{\nu + \frac{n-j-1}{2}} \psi^j (\gamma + \psi^2)^{-\nu - \frac{n}{2}}}{B\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)} \quad (116)$$

$$= \underbrace{2B\left(\frac{j+1}{2}, \nu + \frac{n-j-1}{2}\right)^{-1}}_{\beta_j} \gamma^{\nu + \frac{n-j-1}{2}} \psi^j (\gamma + \psi^2)^{-\nu - \frac{n}{2}}. \quad (117)$$

It follows that

$$p_{\nu,\kappa}(\lambda) \propto \sum_{j=0}^m a_j \lambda^j (\gamma + \lambda^2)^{-\nu - n/2} = \sum_{j=0}^m a_j \beta_j^{-1} \underbrace{\beta_j \lambda^j (\gamma + \lambda^2)^{-\nu - n/2}}_{p_{y_j}(\lambda)} \propto \sum_{j=0}^m c_j p_{y_j}(\lambda). \quad (118)$$

The right-hand side is an actual density, which corresponds to the mixture of p_{y_j} with coefficients c_j , which proves the claim. \square

Proposition 20. *The group $G = \mathrm{GL}(d)$ acts on $\mathrm{SPD}(d)$ transitively by conjugation as*

$$\triangleright : \mathrm{GL}(d) \times \mathrm{SPD}(d) \rightarrow \mathrm{SPD}(d) \quad \mathbf{M} \triangleright \mathbf{S} = \mathbf{M}\mathbf{S}\mathbf{M}^\top \quad (46)$$

where $H = \mathrm{O}(d) \subseteq G$ is the isotropy subgroup corresponding to the element $\mathbf{I}_d \in \mathrm{SPD}(d)$. $\mathrm{SPD}(d)$ can thus be regarded as the homogeneous space $\mathrm{GL}(d)/\mathrm{O}(d)$. Moreover, the map

$$s : \mathrm{GL}(d) \rightarrow \mathrm{GL}(d) \quad s(\mathbf{M}) = \mathbf{M}^{-\top} \quad (47)$$

is an involution, namely $s(s(\mathbf{M})) = \mathbf{M}$, and $s(\mathbf{M}) = \mathbf{M}$ if and only if $\mathbf{M} \in \mathrm{O}(d)$. Using this representation, $\mathrm{SPD}(d)$ satisfies all aspects of the definition of a symmetric space, except for connectedness of the group G .

Proof. Recall that every symmetric positive-definite matrix $\mathbf{S} \in \mathrm{SPD}(d)$ admits a Cholesky decomposition $\mathbf{S} = \mathbf{C}\mathbf{C}^\top$ where \mathbf{C} is a lower-triangular matrix. The map

$$\mathrm{GL}(d) \rightarrow \mathrm{SPD}(d) \quad \mathbf{M} \mapsto \mathbf{M} \triangleright \mathbf{I}_d = \mathbf{M}\mathbf{M}^\top \quad (119)$$

is thus surjective. For \mathbf{S}_1 and $\mathbf{S}_2 \in \mathrm{SPD}(d)$ with Cholesky factors \mathbf{C}_1 and \mathbf{C}_2 , respectively,

$$\mathbf{C}_2\mathbf{C}_1^{-1} \triangleright \mathbf{S}_1 = \mathbf{C}_2\mathbf{C}_2^\top = \mathbf{S}_2 \quad (120)$$

hence this group action is transitive. Note in particular that

$$\mathbf{H} \triangleright \mathbf{I}_d = \mathbf{H}\mathbf{H}^\top = \mathbf{I}_d \quad (121)$$

holds if and only if $\mathbf{H} \in \mathrm{O}(d)$. Thus the isotropy group of the element \mathbf{I}_d is exactly $\mathrm{O}(d)$. This means that $\mathrm{SPD}(d)$ may be regarded as the quotient $\mathrm{GL}(\mathbb{R}^d)/\mathrm{O}(d)$, which is a smooth homogeneous space. The identity $s(s(\mathbf{M})) = s$ is immediate. The condition $s(\mathbf{M}) = \mathbf{M}$ means $\mathbf{M}^{-\top} = \mathbf{M}$ and is equivalent to $\mathbf{M}\mathbf{M}^\top = \mathbf{I}_d$: hence it is true if and only if $\mathbf{M} \in \mathrm{O}(d)$. \square

Proposition 21. *We have $\mathrm{SSPD}(d) \cong \mathrm{SL}(d)/\mathrm{SO}(d)$. Moreover, this space is symmetric and satisfies Assumption A.*

Proof. First, we restrict the action \triangleright defined by (46) and the involution s defined by (47) from the group $\mathrm{GL}(d)$ to its subgroup $\mathrm{SL}(d)$ consisting of unit-determinant matrices. Applying the proof of Proposition 20 *mutatis mutandis* allows us to regard $\mathrm{SSPD}(d)$ as the smooth homogeneous space $\mathrm{SL}(d)/\mathrm{SO}(d)$. Using the fact that $\mathrm{SL}(d)$, in contrast to $\mathrm{GL}(d)$, is connected, we conclude that $\mathrm{SSPD}(d)$ is a symmetric space. We now verify Assumption A.

1. By standard theory, the group $\mathrm{SL}(d)$ for $d > 1$ is a simple Lie group, and hence semi-simple (Helgason, 1979, Ch. III, §8).
2. The center of $\mathrm{SL}(d)$ consists of the scalar matrices corresponding to d th roots of unity in the base field (Rotman, 1995, Theorem 8.9), which is a finite set.
3. $\mathrm{SO}(d)$ is a maximal compact subgroup of $\mathrm{SL}(d)$: every compact subgroup K of $\mathrm{GL}(d)$ preserves some inner product—to find one, take an arbitrary inner product and average it over K —while the orthogonal groups of all inner products in d -dimensional space are isomorphic, and a Lie group cannot have a proper Lie subgroup with the same dimension and number of connected components.

The claim follows. \square

Proposition 22. *Let $\mathbb{R}_{>0}$ and $\text{SSPD}(d)$ be equipped with their respective canonical metrics, and let $\text{SPD}(d)$ be equipped with the affine-invariant metric. Then*

$$\mathcal{T} : \mathbb{R}_{>0} \times \text{SSPD}(d) \rightarrow \text{SPD}(d) \quad \mathcal{T}(a, \mathbf{S}) = a^{\sqrt{d}} \mathbf{S} \quad (48)$$

is a Riemannian isometry.

Proof. From the definition of \mathcal{T} , it is clear that this is a diffeomorphism. We will show that this map preserves geodesic distance, which by virtue of the Myers–Steenrod theorem (Myers and Steenrod, 1939) implies that \mathcal{T} is a Riemannian isometry.

Recall that the canonical metric on $\text{SSPD}(d)$ is defined by pushforward of the bilinear form on the tangent space at \mathbf{I}_d that itself is induced by the Killing form. The geodesic distance induced by this metric is known to correspond to the geodesic distance $d_{\text{SSPD}(d)}$ coinciding with the right-hand side of (44). Indeed, the identification of the metric on $\text{SPD}(d)$ (Bhatia, 2009, Lemma 6.1.1) allows one to show that the unique geodesic between two unit-determinant matrices lies inside $\text{SSPD}(d)$ (Bhatia, 2009, Theorem 6.1.6).

Since the Frobenius norm of a symmetric matrix is the l^2 -norm of the vector of its eigenvalues, both $d_{\text{SPD}(d)}$ and $d_{\text{SSPD}(d)}$, for a pair of appropriate matrices \mathbf{S}_1 and \mathbf{S}_2 , are equal to

$$\left(\sum_{j=1}^d \log^2 \lambda_j \right)^{1/2} \quad (122)$$

where $\lambda_1, \lambda_2, \dots, \lambda_d$ are eigenvalues of $\mathbf{S}_1^{-1/2} \mathbf{S}_2 \mathbf{S}_1^{-1/2}$.

The canonical metric on $\mathbb{R}_{>0}$ corresponds to the distance $d_{\mathbb{R}_{>0}}(a_1, a_2) = |\log(v_1 v_2^{-1})| = |\log(v_1) - \log(v_2)|$, which is the pullback of the Euclidean distance through the logarithm.

From this, the remaining argument is straightforward. Take $\mathbf{S}_1, \mathbf{S}_2 \in \text{SSPD}(d)$ and $a_1, a_2 \in \mathbb{R}_{>0}$. Denote the eigenvalues of $\mathbf{S}_1^{-1/2} \mathbf{S}_2 \mathbf{S}_1^{-1/2}$ by $\lambda_1, \lambda_2, \dots, \lambda_d$. Note that in this case $\sum_{j=1}^d \log \lambda_j = 0$, since the determinant is equal to 1, hence

$$\text{dist}_{\text{SPD}(d)}(v_1 \mathbf{S}_1, v_2 \mathbf{S}_2)^2 = \|v_1 v_2^{-1} \mathbf{S}_1^{-1/2} \mathbf{S}_2 \mathbf{S}_1^{-1/2}\|_F^2 \quad (123)$$

$$= \sum_{j=1}^d \log(v_1 v_2^{-1} \lambda_j)^2 \quad (124)$$

$$= \sum_{j=1}^d \log(\lambda_j)^2 + 2 \log\left(\frac{v_1}{v_2}\right) \sum_{j=1}^d \log \lambda_j + \log\left(v_1^{\sqrt{d}} v_2^{-\sqrt{d}}\right)^2 \quad (125)$$

$$= d_{\text{SSPD}(d)}(\mathbf{S}_1, \mathbf{S}_2)^2 + d_{\mathbb{R}_{>0}}(v_1^{\sqrt{d}}, v_2^{\sqrt{d}})^2. \quad (126)$$

The last equation is exactly the square of the product metric between the pairs (a_1, \mathbf{S}_1) and (a_2, \mathbf{S}_2) , which completes the proof. \square

Proposition 25- ε . Consider a random vector \mathbf{x} on \mathbb{R}^d with density $p_{\mathbf{x}}(\boldsymbol{\xi}) \propto e^{-\gamma \|\boldsymbol{\xi}\|^2/2} f(\boldsymbol{\xi})$ where $\gamma > 0$ and $f : \mathbb{R}^d \rightarrow [0, \infty)$ satisfies $f(a\boldsymbol{\xi}) = a^k f(\boldsymbol{\xi})$ for some $k \in \mathbb{N}$ and all $a > 0$. Let

y be a random variable independent from \mathbf{x} , distributed according to the $\chi_{2\nu}$ distribution, namely $p_y(\psi) \propto \psi^{2\nu-1} e^{-\psi^2/2}$ and $p_{\mathbf{x},y}(\boldsymbol{\xi}, \psi) = p_{\mathbf{x}}(\boldsymbol{\xi})p_y(\psi)$. Then the random vector $\mathbf{z} = \mathbf{x}/y$ has the following density:

$$p_{\mathbf{z}}(\boldsymbol{\zeta}) \propto f(\boldsymbol{\zeta})(\gamma^{-1} + \|\boldsymbol{\zeta}\|^2)^{-\alpha} \quad \alpha = \nu + (d+k)/2. \quad (127)$$

Proof. Let us compute the joint density $p_{\mathbf{z},y}$. We have $(\mathbf{z}, y) = \boldsymbol{\Phi}(\mathbf{x}, y)$, where the function $\boldsymbol{\Phi} : \mathbb{R}^d \times \mathbb{R}_{>0} \rightarrow \mathbb{R}^d \times \mathbb{R}_{>0}$ is given by $\boldsymbol{\Phi}(\boldsymbol{\xi}, \psi) = (\boldsymbol{\xi}/\psi, \psi)$. Since $\boldsymbol{\Phi}$ is bijective, we can use the change of variables formula for densities to write

$$p_{\mathbf{z},y}(\boldsymbol{\zeta}, \psi) = p_{\mathbf{x},y}(\psi\boldsymbol{\zeta}, \psi) \det(\mathbf{J}_{\boldsymbol{\Phi}^{-1}}(\boldsymbol{\zeta}, \psi)) \quad (128)$$

where $\mathbf{J}_{\boldsymbol{\Phi}^{-1}}$ is the Jacobian matrix of the map $\boldsymbol{\Phi}^{-1} : \mathbb{R}^d \times \mathbb{R}_{>0}$ given by $\boldsymbol{\Phi}^{-1}(\boldsymbol{\zeta}, \psi) = (\psi\boldsymbol{\zeta}, \psi)$. It is easy to see that

$$\mathbf{J}_{\boldsymbol{\Phi}^{-1}}(\boldsymbol{\zeta}, \psi) = \begin{pmatrix} \psi & 0 & \dots & 0 & \zeta_1 \\ 0 & \psi & \dots & 0 & \zeta_2 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & \psi & \zeta_d \\ 0 & 0 & \dots & 0 & 1 \end{pmatrix} \quad \det(\mathbf{J}_{\boldsymbol{\Phi}^{-1}}(\boldsymbol{\zeta}, \psi)) = \psi^d. \quad (129)$$

Thus

$$p_{\mathbf{z},y}(\boldsymbol{\zeta}, \psi) = p_{\mathbf{x},y}(\psi\boldsymbol{\zeta}, \psi)\psi^d = p_{\mathbf{x}}(\psi\boldsymbol{\zeta})p_y(\psi)\psi^d \propto e^{-\gamma\|\psi\boldsymbol{\zeta}\|^2/2}\psi^k f(\boldsymbol{\zeta})\psi^{2\nu-1}e^{-\psi^2/2}\psi^d \quad (130)$$

$$\propto \psi^{d+k+2\nu-1}e^{-\psi^2/2(\gamma\|\boldsymbol{\zeta}\|^2+1)}f(\boldsymbol{\zeta}). \quad (131)$$

Now, we can compute the marginal distribution of \mathbf{z} , up to a constant, by writing

$$p_{\mathbf{z}}(\boldsymbol{\zeta}) = \int_{\mathbb{R}_{>0}} p_{\mathbf{z},\psi}(\boldsymbol{\zeta}, \psi) d\psi \propto f(\boldsymbol{\zeta}) \int_{\mathbb{R}_{>0}} e^{-\psi^2(\gamma\|\boldsymbol{\zeta}\|^2+1)/2}\psi^{d+k+2\nu-1} d\psi \quad (132)$$

$$\propto f(\boldsymbol{\zeta})(\gamma\|\boldsymbol{\zeta}\|^2 + 1)^{-(d+k+2\nu)/2} \propto f(\boldsymbol{\zeta})(\gamma^{-1} + \|\boldsymbol{\zeta}\|^2)^{-(d+k+2\nu)/2} \quad (133)$$

where the integral was computed using Gradshteyn and Ryzhik (2014, Section 3.326, Item 2). This finishes the proof of the auxiliary statement. \square

Proposition 25. *Consider the density defined by*

$$p_{\nu,\kappa}(\boldsymbol{\lambda}) \propto \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2 \right)^{-\nu-n/2} \prod_{1 \leq i < j \leq d} \pi |\lambda_i - \lambda_j| \quad \boldsymbol{\lambda} \in \mathbb{R}^d. \quad (57)$$

Let $\boldsymbol{\lambda}^{\infty,\gamma} \sim p_{\infty,\gamma}$, where $\gamma = \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2 \right)^{-1/2}$ and the right-hand side is defined in Corollary 24, and let $y \sim \chi_{2\nu}$ be a chi-distributed random variable with $2\nu \in \mathbb{R}_{>0}$ degrees of freedom. If $\boldsymbol{\lambda}^{\nu,\kappa} = \boldsymbol{\lambda}^{\infty,\gamma}/y$, then $\boldsymbol{\lambda}^{\nu,\kappa} \sim p_{\nu,\kappa}$.

Proof. Applying Proposition 25- ε , taking $p_{\mathbf{x}}$ equal to $p_{\infty, \gamma}$, where

$$p_{\infty, \gamma}(\boldsymbol{\lambda}) \propto \exp\left(-\left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2\right)^{-1} (\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2)/2\right) \prod_{1 \leq i < j \leq d} \pi|\lambda_i - \lambda_j| \quad (134)$$

$$\propto \exp\left(-\underbrace{\left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\rho}\|^2\right)^{-1}}_{\gamma^2} \|\boldsymbol{\lambda}\|^2/2\right) \underbrace{\prod_{1 \leq i < j \leq d} \pi|\lambda_i - \lambda_j|}_{f(\boldsymbol{\lambda})}. \quad (135)$$

We have $f(a\boldsymbol{\lambda}) = a^{\frac{d(d-1)}{2}} f(\boldsymbol{\lambda})$. Then, Proposition 25- ε shows that for a chi-distributed random variable $y \sim \chi_{2\nu}$, the random vector $\boldsymbol{\lambda}^{\nu, \kappa} = \boldsymbol{\lambda}^{\infty, \gamma}/y$ has density

$$p_{\boldsymbol{\lambda}^{\nu, \kappa}}(\boldsymbol{\lambda}) \propto f(\boldsymbol{\lambda})(\gamma^{-2} + |\boldsymbol{\lambda}|^2)^{-(d + \frac{d(d-1)}{2} + 2\nu)/2} \quad (136)$$

$$= \left(\frac{2\nu}{\kappa^2} + \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\rho}\|^2\right)^{-(\nu + \frac{d(d+1)}{4})} \prod_{1 \leq i < j \leq d} \pi|z_i - z_j|. \quad (137)$$

From this, and noting that the dimension of $\text{SPD}(d)$ is $n = \frac{d(d+1)}{2}$, the claim follows. \square

Proposition 28. *Let $X = G/H$ be a symmetric space that satisfies Assumption A. Recall that $\pi^{(\lambda)}$ denote zonal spherical functions on X and $\lambda \in \mathfrak{a}^*$ are vectors. Consider $\mathbb{k} \in L^2(G)$ such that $k(g_1 \cdot H, g_2 \cdot H) = \mathbb{k}(g_2^{-1} \cdot g_1)$ is a stationary kernel on X . Then*

$$\mathbb{k}(g) \leq \underbrace{\mathbb{k}(e)}_{\text{variance of } k} \pi^{(0)}(g) \quad (59)$$

for all $g \in G$. Moreover, there exists $g \in G$ such that $g \notin e \cdot H$ and $\pi^{(0)}(g) < 1$.

Proof. Recall the definition of zonal spherical functions $\pi^{(\lambda)}$, namely

$$\pi^{(\lambda)}(g) = \int_H e^{(i\lambda + \rho)^\top a(h \cdot g)} d\mu_H(h). \quad (138)$$

Write

$$|\pi^{(\lambda)}(g)| \leq \int_H |e^{(i\lambda + \rho)^\top a(h \cdot g)}| d\mu_H(h) = \int_H e^{\rho^\top a(h \cdot g)} d\mu_H(h) = \pi^{(0)}(g). \quad (139)$$

Let μ_k be the spectral measure of the kernel k . Then

$$|\mathbb{k}(g)| = |k(g \cdot H, e)| = \left| \int_{\mathfrak{a}^*} \pi^{(\lambda)}(g) d\mu_k(\lambda) \right| \leq \int_{\mathfrak{a}^*} |\pi^{(\lambda)}(g)| d\mu_k(\lambda) = \mathbb{k}(e) \pi^{(0)}(g) \quad (140)$$

because $\pi^{(\lambda)}(e) = 1$, and therefore

$$\mathbb{k}(e) = \int_{\mathfrak{a}^*} \pi^{(\lambda)}(e) d\mu_k(\lambda) = \int_{\mathfrak{a}^*} 1 d\mu_k(\lambda). \quad (141)$$

Finally, note that $\pi^{(0)} \neq 1$, since $\Delta \pi^{(0)} = -\|\rho\|^2 \pi^{(0)}$ for $\rho \neq 0$, while $\Delta 1 = 0$. What is more, $\pi^{(0)}(g) \leq \pi^{(0)}(e) = 1$ since $g_1, g_2 \mapsto \pi^{(0)}(g_2^{-1} \cdot g_1)$ is a positive semi-definite function. \square

Proposition 29. *For all $\nu < \infty$, the spectral measure $\mu_{\nu, \kappa, \sigma^2}$ does not converge weakly to a Dirac measure as $\kappa \rightarrow \infty$. In contrast, if either $\nu = \infty$, or if we instead work with $\widehat{\mu}_{\nu, \kappa, \sigma^2}$ and ν is large enough, then the respective measure weakly converges to a Dirac measure.*

Proof. Recall that

$$\mu_{\nu, \kappa, \sigma^2}(\lambda) \propto \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \quad (142)$$

$$\mu_{\infty, \kappa, \sigma^2}(\lambda) \propto e^{-\frac{\kappa^2}{2}(\|\lambda\|^2 + \|\rho\|^2)} |c(\lambda)|^{-2} d\lambda. \quad (143)$$

The proof is based on analysis of Harish-Chandra's c -function. This function is meromorphic and has (Helgason, 1994, p. 109) the following form

$$c(\lambda) \propto \prod_{\alpha \in \Sigma^+} \frac{2^{-i\lambda^\top \alpha / \|\alpha\|^2} \Gamma(i\lambda^\top \alpha / \|\alpha\|^2)}{\Gamma\left(\frac{1}{2}\left(\frac{1}{2}m_\alpha + 1 + i\lambda^\top \alpha / \|\alpha\|^2\right)\right) \Gamma\left(\frac{1}{2}\left(\frac{1}{2}m_\alpha + m_{2\alpha} + i\lambda^\top \alpha / \|\alpha\|^2\right)\right)} \quad (144)$$

where the product is taken over the the set of positive roots Σ^+ and m_α denotes the multiplicity of the root α , see Appendix A for a brief introduction into these notions.¹⁵

By Strichartz (1983, Lemma 6.2), there exists a constant $c_0 > 0$ such that for all $s > 1$

$$\int_{\|\lambda\| \leq s} |c(\lambda)|^{-2} d\lambda \leq c_0 s^n. \quad (145)$$

With these expressions, we are ready to prove the proposition. Without loss of generality, we assume $\sigma^2 = 1$. We prove the first part of statement by showing that, as $\kappa \rightarrow \infty$, the measures $\mu_{\nu, \kappa, 1}$ converge to the measure

$$\mu_{\nu, \infty}(\lambda) \propto \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \quad (146)$$

rather than to a Dirac delta function. We start by showing that $\mu_{\nu, \infty}$ is a finite measure.

Applying (145), we bound

$$\int_{s \leq \|\lambda\| \leq 2s} 1 d\mu_{\nu, \infty} \propto \int_{s \leq \|\lambda\| \leq 2s} \left(\|\lambda\|^2 + \|\rho\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \quad (147)$$

$$\leq \left(s^2 + \|\rho\|^2 \right)^{-\nu-n/2} \int_{\|\lambda\| \leq 2s} |c(\lambda)|^{-2} d\lambda \lesssim s^{-n-2\nu} s^n = s^{-2\nu} \quad (148)$$

where \lesssim denotes inequality up to a constant. Using this, we can write

$$\int_{\mathfrak{a}^*} 1 d\mu_{\nu, \infty} = \int_{\|\lambda\| \leq 1} 1 d\mu_{\nu, \infty} + \sum_{k=0}^{\infty} \int_{2^k \leq \|\lambda\| < 2^{k+1}} 1 d\mu_{\nu, \infty} \lesssim c_0 \|\rho\|^{-2\nu-n} + \sum_{k=0}^{\infty} 2^{-2\nu k} < \infty. \quad (149)$$

This proves that the measure $\mu_{\nu, \infty}$ is finite.

¹⁵. Our notion of positive roots coincides with the notion of *indivisible* positive roots in Helgason (1994).

When $\kappa \rightarrow \infty$, the densities of $\mu_{\nu, \kappa, 1}$ with respect to the Lebesgue measure $d\lambda$ converge to the density of $\mu_{\nu, \infty}$ pointwise. Thus, by Scheffé's lemma, recalling that $d\mu_{\nu, \infty}$ is finite, the measures $\mu_{\nu, \kappa, 1}$ converge in distribution to $\mu_{\nu, \infty}$. This proves the first part of the statement.

Now, we prove the second part of the statement. For this, we will show that if either $\nu = \infty$, or if we instead work with $\widehat{\mu}_{\nu, \kappa, \sigma^2}$ and ν is large enough, the mass of the measures concentrates at $\lambda = 0$ as $\kappa \rightarrow \infty$. We prove this only for $\widehat{\mu}_{\nu, \kappa, \sigma^2}$, $\nu < \infty$, as the case of $\nu = \infty$ is similar, but simpler. As before, we also assume $\sigma^2 = 1$ without any loss of generality. First, write

$$\int_{\|\lambda\| \geq 1} \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \leq \int_{\|\lambda\| \geq 1} \|\lambda\|^{-2\nu-n} |c(\lambda)|^{-2} d\lambda \quad (150)$$

$$= \sum_{k=0}^{\infty} \int_{2^k \leq \|\lambda\| < 2^{k+1}} \|\lambda\|^{-2\nu-n} |c(\lambda)|^{-2} d\lambda \quad (151)$$

$$\leq \sum_{k=0}^{\infty} 2^{-2\nu k - nk} 2^{(k+1)n} \leq c_1 < \infty. \quad (152)$$

Next, because of (145), we have for any $\varepsilon > 0$ that

$$\int_{\|\lambda\| \geq \varepsilon} \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \leq c_1 + \int_{1 > \|\lambda\| \geq \varepsilon} \varepsilon^{-2\nu-n} |c(\lambda)|^{-2} d\lambda \quad (153)$$

$$\leq c_1 + c_0 \varepsilon^{-2\nu-n} < \infty. \quad (154)$$

Let us focus on the case $\|\lambda\| < \varepsilon$. The formula for $c(\lambda)$ gives us that, in this case, $|c(\lambda)|^{-2} \approx c_3 \prod_{\alpha \in \Sigma^+} |\lambda^\top \alpha|^2$, since $|\Gamma(xi)|^{-2} = \frac{x \sinh(\pi x)}{\pi} = x^2 + \mathcal{O}(x^4)$. So, we have

$$\|\lambda\|^{-2\nu-n} |c(\lambda)|^{-2} \approx c_3 \|\lambda\|^{-2\nu-n} \prod_{\alpha \in \Sigma^+} |\lambda^\top \alpha|^2 \quad (155)$$

and therefore, if $-2\nu - n + 2|\Sigma^+| + \dim \mathfrak{a}^* < 0$, the integral

$$\int_{\|\lambda\| < \varepsilon} \|\lambda\|^{-2\nu-n} \prod_{\alpha} |\lambda^\top \alpha|^2 d\lambda \quad (156)$$

does not converge. This condition must hold for large enough ν . We can even be more specific and give an explicit sufficient condition for this: note that since $m_\alpha \in \mathbb{Z}_{>0}$, we have $\sum_{\alpha \in \Sigma^+} m_\alpha \geq \sum_{\alpha \in \Sigma^+} 1 = |\Sigma^+|$. Also, by Strichartz (1983, p. 409), we have $n = \sum m_\alpha + \dim \mathfrak{a}^*$. Hence $-2\nu - n + 2|\Sigma^+| + \dim \mathfrak{a}^* \leq -2\nu - n + |\Sigma^+| + \sum_{\alpha \in \Sigma^+} m_\alpha + \dim \mathfrak{a}^* = -2\nu + |\Sigma^+|$. Thus $\nu > |\Sigma^+|/2$ is a sufficient—but not always necessary—condition for the integral to diverge. Because of the above, for any fixed $\varepsilon > 0$, as $\kappa \rightarrow \infty$, we have

$$\int_{\|\lambda\| \geq \varepsilon} \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \leq C_\varepsilon \quad \int_{\|\lambda\| < \varepsilon} \left(\frac{2\nu}{\kappa^2} + \|\lambda\|^2 \right)^{-\nu-n/2} |c(\lambda)|^{-2} d\lambda \rightarrow \infty. \quad (157)$$

Since $\widehat{\mu}_{\nu, \kappa, 1}$ is a probability measure, namely $\widehat{\mu}_{\nu, \kappa, 1}(\mathfrak{a}^*) = 1$, we have

$$\widehat{\mu}_{\nu, \kappa, 1}(\{\lambda \in \mathfrak{a}^* : \|\lambda\| \geq \varepsilon\}) \rightarrow 0 \quad \widehat{\mu}_{\nu, \kappa, 1}(\{\lambda \in \mathfrak{a}^* : \|\lambda\| < \varepsilon\}) \rightarrow 1 \quad (158)$$

which proves that $\widehat{\mu}_{\nu, \kappa, 1} \rightarrow \delta$ as $\kappa \rightarrow \infty$. \square

References

- J.-P. Anker, E. Papageorgiou, and H.-W. Zhang. Asymptotic behavior of solutions to the heat equation on noncompact symmetric spaces. *Journal of Functional Analysis*, 2023. Cited on page 13.
- S. Araki. On root systems and an infinitesimal classification of irreducible symmetric spaces. *Journal of Mathematics, Osaka City University*, 1962. Cited on page 33.
- M. Arnaudon, F. Barbaresco, and L. Yang. Riemannian Medians and Means With Applications to Radar Signal Processing. *IEEE Journal of Selected Topics in Signal Processing*, 2013. Cited on page 21.
- A. Arvanitogeorgos. *An introduction to Lie groups and the geometry of homogeneous spaces*. American Mathematical Society, 2003. Cited on pages 5, 6, 12.
- I. Azangulov, A. Smolensky, A. Terenin, and V. Borovitskiy. Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces I: the compact case. *Journal of Machine Learning Research*, 2024. Cited on pages 1, 4, 8, 11, 12, 14.
- R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2009. Cited on pages 21, 42.
- V. Borovitskiy, I. Azangulov, A. Terenin, P. Mostowsky, M. Deisenroth, and N. Durrande. Matérn Gaussian processes on graphs. In *International Conference on Artificial Intelligence and Statistics*, 2021. Cited on page 2.
- V. Borovitskiy, M. R. Karimi, V. R. Somnath, and A. Krause. Isotropic Gaussian Processes on Finite Spaces of Graphs. In *Artificial Intelligence and Statistics*, 2023. Cited on page 2.
- V. Borovitskiy, A. Terenin, P. Mostowsky, and M. Deisenroth. Matérn Gaussian Processes on Riemannian Manifolds. In *Advances in Neural Information Processing Systems*, 2020. Cited on page 2.
- J.-P. Chiles and P. Delfiner. *Geostatistics: Modeling Spatial Uncertainty*. Wiley, 2009. Cited on page 3.
- S. Coveney, C. Corrado, C. H. Roney, D. O’Hare, S. E. Williams, M. D. O’Neill, S. A. Niederer, R. H. Clayton, J. E. Oakley, and R. D. Wilkinson. Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity. *Philosophical Transactions of the Royal Society A*, 2020. Cited on page 2.
- N. Da Costa, C. Mostajeran, J.-P. Ortega, and S. Said. Invariant kernels on Riemannian symmetric spaces: a harmonic-analytic approach. *arXiv:2310.19270*, 2023. Cited on page 2.
- G. Dong and G. Kuang. Target Recognition in SAR Images via Classification on Riemannian Manifolds. *IEEE Geoscience and Remote Sensing Letters*, 2015. Cited on page 21.

- I. L. Dryden, A. Koloydenko, and D. Zhou. Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *Annals of Applied Statistics*, 2009. Cited on page 21.
- J. Eldering. Persistence of noncompact normally hyperbolic invariant manifolds in bounded geometry. *Comptes Rendus Mathématique*, 2012. Cited on page 38.
- R. Gangolli. Asymptotic behaviour of spectra of compact quotients of certain symmetric spaces. *Acta Mathematica*, 1968. Cited on pages 14, 35, 36.
- I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, 7th edition, 2014. Cited on page 43.
- A. Grigoryan. Estimates of heat kernels on Riemannian manifolds. In *Instructional Conference on Spectral Theory and Geometry*, 1998. Cited on page 20.
- A. Grigoryan. *Heat kernel and analysis on manifolds*. American Mathematical Society, 2009. Cited on pages 12, 23.
- N. Große and C. Schneider. Sobolev spaces on Riemannian manifolds with bounded geometry: General coordinates and traces. *Mathematische Nachrichten*, 2013. Cited on page 38.
- Harish-Chandra. Representations of a semisimple Lie group on a Banach space. I. *Transactions of the American Mathematical Society*, 1953. Cited on page 8.
- E. Hebey. *Nonlinear Analysis on Manifolds: Sobolev Spaces and Inequalities: Sobolev Spaces and Inequalities*. American Mathematical Society, 2000. Cited on page 38.
- S. Helgason. *Differential Geometry and Symmetric Spaces*. Academic Press, 1962. Cited on pages 2, 8, 31.
- S. Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces*. Elsevier, 1979. Cited on pages 31, 32, 41.
- S. Helgason. *Geometric Analysis on Symmetric Spaces*. American Mathematical Society, 1994. Cited on pages 11, 14, 31, 35, 45.
- S. Helgason. *Groups and Geometric Analysis: Integral Geometry, Invariant Differential Operators, and Spherical Functions*. American Mathematical Society, 2000. Cited on pages 10, 26, 37, 38.
- S. Helgason. Huygens' principle for wave equations on symmetric spaces. *Journal of Functional Analysis*, 1992. Cited on page 14.
- M. Hutchinson, A. Terenin, V. Borovitskiy, S. Takao, Y. W. Teh, and M. Deisenroth. Vector-valued Gaussian Processes on Riemannian Manifolds via Gauge Independent Projected Kernels. *Advances in Neural Information Processing Systems*, 2021. Cited on page 2.

- N. Jaquier, V. Borovitskiy, A. Smolensky, A. Terenin, T. Asfour, and L. Rozo. Geometry-aware Bayesian Optimization in Robotics using Riemannian Matérn Kernels. In *Conference on Robot Learning, 2022*. Cited on pages 2, 25.
- N. Jaquier, L. Rozo, S. Calinon, and M. Bürger. Bayesian optimization meets Riemannian manifolds in robot learning. In *Conference on Robot Learning, 2020*. Cited on pages 2, 20, 21.
- N. Jaquier, L. Rozo, M. González-Duque, V. Borovitskiy, and T. Asfour. Bringing motion taxonomies to continuous domains via GPLVM on hyperbolic manifolds. In *International Conference on Machine Learning, 2024*. Cited on page 2.
- S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel Methods on the Riemannian Manifold of Symmetric Positive Definite Matrices. In *Computer Vision and Pattern Recognition, 2013*. Cited on page 21.
- K. Jensen, T.-C. Kao, M. Tripodi, and G. Hennequin. Manifold GPLVMs for discovering non-Euclidean latent structure in neural data. In *Advances in Neural Information Processing Systems, 2020*. Cited on page 2.
- R. Kondor. *Group theoretical methods in machine learning*. PhD thesis, Columbia University, 2008. Cited on page 3.
- J. Lester. Orthochronous subgroups of $O(p, q)$. *Linear and Multilinear Algebra, 1993*. Cited on page 17.
- Z. Lin. Riemannian Geometry of Symmetric Positive Definite Matrices via Cholesky Decomposition. *SIAM Journal on Matrix Analysis and Applications, 2019*. Cited on page 21.
- F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 2011*. Cited on page 2.
- I. Livshits. Use of Shifted Laplacian Operators for Solving Indefinite Helmholtz Equations. *Numerical Mathematics: Theory, Methods and Applications, 2015*. Cited on page 14.
- D. G. Maher. Wrapping Brownian Motion and Heat Kernels II: Symmetric Spaces. *Journal of Lie Theory, 2014*. Cited on page 14.
- M. L. Mehta. *Random Matrices*. Academic Press, 1991. Cited on page 24.
- F. Mezzadri. How to generate random matrices from the classical compact groups. *Notices of the American Mathematical Society, 2007*. Cited on page 10.
- S. B. Myers and N. E. Steenrod. The Group of Isometries of a Riemannian Manifold. *Annals of Mathematics, 1939*. Cited on pages 31, 42.
- M. A. Naimark. *Normed Rings*. P. Noordhoff, 1959. Cited on page 6.

- W. Neiswanger, K. A. Wang, and S. Ermon. Bayesian Algorithm Execution: Estimating Computable Properties of Black-box Functions Using Mutual Information. In *International Conference on Machine Learning*, 2021. Cited on page 7.
- A. Nejatbakhsh, I. Garon, and A. Williams. Estimating noise correlations across continuous conditions with Wishart processes. In *Advances in Neural Information Processing Systems*, 2024. Cited on page 2.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2008. Cited on page 10.
- D. J. Robinson. *An Introduction to Abstract Algebra*. de Gruyter, 2003. Cited on page 3.
- P. Rosa, V. Borovitskiy, A. Terenin, and J. Rousseau. Posterior Contraction Rates for Matérn Gaussian Processes on Riemannian Manifolds. In *Advances in Neural Information Processing Systems*, 2023. Cited on page 2.
- J. J. Rotman. *An Introduction to the Theory of Groups*. Springer, 1995. Cited on page 41.
- P. Sawyer. Computing the Iwasawa decomposition of the classical Lie groups of noncompact type using the QR decomposition. *Linear Algebra and its Applications*, 2016. Cited on page 34.
- P. Sawyer. The heat equation on the spaces of positive definite matrices. *Canadian Journal of Mathematics*, 1992. Cited on pages 25, 26.
- I. E. Segal and J. von Neumann. A Theorem on Unitary Representations of Semisimple Lie Groups. *Annals of Mathematics*, 1950. Cited on page 5.
- R. S. Strichartz. Improved Sobolev inequalities. *Transactions of the American Mathematical Society*, 1983. Cited on pages 38, 45, 46.
- L. Sun. An Uncertainty Principle on Hyperbolic Space. *Proceedings of the American Mathematical Society*, 1994. Cited on page 18.
- A. Terenin, D. R. Burt, A. Artemev, S. Flaxman, M. van der Wilk, C. E. Rasmussen, and H. Ge. Numerically Stable Sparse Gaussian Processes via Minimum Separation using Cover Trees. *Journal of Machine Learning Research*, 2024. Cited on page 7.
- A. Terras. *Harmonic Analysis on Symmetric Spaces—Higher Rank Spaces, Positive Definite Matrix Space and Generalizations*. Springer, 2016. Cited on pages 4, 21, 23.
- J. L. Vázquez. Asymptotic behaviour for the heat equation in hyperbolic space. *Communications in Analysis and Geometry*, 2022. Cited on page 13.
- N. J. Wildberger. Hypergroups, symmetric spaces, and wrapping maps. In *Probability Measures on Groups and Related Structures*. World Scientific, 1995. Cited on page 13.

- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International Conference on Machine Learning*, 2020. Cited on pages 3, 7.
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise Conditioning of Gaussian Processes. *Journal of Machine Learning Research*, 2021. Cited on pages 3, 7.
- J. T. Wilson, F. Hutter, and M. Deisenroth. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2018. Cited on page 7.
- J. A. Wolf. *Spaces of constant curvature*. American Mathematical Society, 2011. Cited on pages 16, 17.
- A. M. Yaglom. Second-order homogeneous random fields. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Contributions to Probability Theory*. University of California Press, 1961. Cited on pages 2, 4–6, 8.