

Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces I: the compact case

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

Gaussian processes are widely used as machine learning models for learning unknown functions, particularly when data is scarce. Using Bayesian methods, Gaussian processes are able to quantify uncertainty, thereby allowing what is known and what is unknown to be balanced when modeling and making decisions. To maximize statistical efficiency, and ensure their posterior error bars are appropriate and can be trusted, the Gaussian process models we

*Joint first author.

employ in practice should be adapted to their respective settings. Their prior information should correctly encode the structure of the problem at hand.

Many problems, particularly in the physical sciences and engineering, require modeling functions on non-Euclidean spaces. For example, spatial statistics on the surface of earth (Jeong et al., 2017; Guinness and Fuentes, 2016), cosmology (Lang and Schwab, 2015) and medical imaging (Andersson and Sotiropoulos, 2015) often call for Gaussian process models on the sphere, which is a very well-studied setting—see for instance Marinucci and Peccati (2011). Other applications, however, require Gaussian process models on different spaces: for many of these, appropriate statistical tools are only partially developed, particularly if the application demands careful consideration of symmetries. For example, control parameter tuning and parametric policy adaptation in robotics (Jaquier et al., 2020, 2022), latent representation learning in neuroscience (Jensen et al., 2020) or modeling unknown dynamics of physical systems (Borovitskiy et al., 2020; Hutchinson et al., 2021) require models on various *Lie groups* and their *homogeneous spaces*. At present, these applications tend to rely on ad-hoc heuristic or partially understood techniques, rather than on a principled and comprehensive formal treatment.

In this two-part work, we study Gaussian process priors on a large class of non-Euclidean spaces. We focus on *stationary* priors, whose covariance kernels are invariant under symmetries of the space, generalizing the usual Euclidean notion of stationarity as shift-invariance. We work with spaces whose group of symmetries is rich enough, including, in part I, compact *Lie groups* such as the special orthogonal group $\text{SO}(n)$ or the special unitary group $\text{SU}(n)$, and their *homogeneous spaces* such as the sphere \mathbb{S}_n or the Stiefel manifold $\text{V}(k, n)$. In part II (Azangulov et al., 2024), we consider certain *non-compact symmetric spaces* such as the hyperbolic spaces \mathbb{H}^n and spaces of symmetric positive definite matrices $\text{SPD}(n)$. As part of our work, we contribute implementations of these ideas to the GEOMETRICKERNELS¹ library. We now proceed to specify notation and more formally introduce the setting.

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

1.1 Kernels and Gaussian Processes

Let X be a set. We say that a stochastic process f indexed by X is a *Gaussian process* if, for any finite set of inputs $\mathbf{x} \in X^N$, the random vector $f(\mathbf{x})$ is multivariate Gaussian. Every such process is determined by its *mean function* $\mu(\cdot) = \mathbb{E}(f(\cdot))$ and *covariance kernel* $k(\cdot, \cdot') = \text{Cov}(f(\cdot), f(\cdot'))$, so we write $f \sim \text{GP}(\mu, k)$. We assume all Gaussian processes we consider are real-valued, unless stated otherwise or apparent from context.

Assume $y_i = f(x_i) + \varepsilon_i$ with $\varepsilon \sim \text{N}(\mathbf{0}, \Sigma)$, and let \mathbf{x}, \mathbf{y} be the data. If $f \sim \text{GP}(0, k)$ is the prior Gaussian process, then the posterior $f \mid \mathbf{y}$ is also a Gaussian process, whose mean

1. Available at [HTTPS://GEOMETRIC-KERNELS.GITHUB.IO](https://geometric-kernels.github.io). A prototypical software implementation, as well as our experiments, are can be found at [HTTPS://GITHUB.COM/IMBIK/LIESTATIONARYKERNELS](https://github.com/imbirik/liestationarykernels).

and covariance kernel are

$$\mathbb{E}(f \mid \mathbf{y}) = \mathbf{K}_{(\cdot)x}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \boldsymbol{\Sigma})^{-1}\mathbf{y} \quad \text{Cov}(f \mid \mathbf{y}) = \mathbf{K}_{(\cdot,\cdot')} - \mathbf{K}_{(\cdot)x}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \boldsymbol{\Sigma})^{-1}\mathbf{K}_{\mathbf{x}(\cdot')} \quad (1)$$

where (\cdot) and (\cdot') denote arbitrary sets of evaluation locations. We will also use another way of thinking about the posterior: using *pathwise conditioning* (Chiles and Delfiner, 2009; Wilson et al., 2020, 2021). From this viewpoint, the posterior can be written

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

where equality is taken to be in distribution.

In order to perform Bayesian learning with Gaussian processes, a prior Gaussian process, or a parametric family thereof, must be chosen. In the Euclidean setting, priors given by *stationary kernels* are a widely-used choice, which we now discuss.

1.2 Stationary Kernels and Gaussian Processes in Euclidean Spaces

A kernel k over \mathbb{R}^n is called *stationary* if it is translation invariant in the sense that

$$k(\mathbf{x} + \mathbf{c}, \mathbf{x}' + \mathbf{c}) = k(\mathbf{x}, \mathbf{x}') \quad (3)$$

for all $\mathbf{c} \in \mathbb{R}^n$. In this case, there is a function $\mathbb{k} : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{k}(\mathbf{x} - \mathbf{x}'). \quad (4)$$

A Gaussian process, similarly, is called stationary if its mean function is translation invariant—hence, constant, usually assumed zero—and its covariance kernel is stationary.

Stationary kernels are widely used in applications. This includes the widely-popular family of *Matérn* kernels advocated by Stein (1999), which uses three interpretable parameters: amplitude σ^2 which controls variability, length scale κ which controls spatial decay, and smoothness ν which controls differentiability of the process. The limiting case as $\nu \rightarrow \infty$ of the Matérn family is arguably the single most popular kernel used, the squared exponential (heat, Gaussian, RBF, diffusion) kernel which induces an infinitely differentiable Gaussian process prior. We will study these kernels in more general settings: as a glimpse of what we will be able to carry out, we illustrate Gaussian process regression with a Matérn kernel on a real projective space, using the techniques of Section 4, in Figure 1.

Stationary kernels have important properties. Bochner’s Theorem says that continuous stationary kernels are in bijective correspondence with finite nonnegative measures via the inverse Fourier transform, namely

$$\mathbb{k}(\mathbf{x} - \mathbf{x}') = \int_{\mathbb{R}^n} e^{2\pi i \langle \boldsymbol{\psi}, \mathbf{x} - \mathbf{x}' \rangle} d\Psi(\boldsymbol{\psi}) = \int_{\mathbb{R}^n} \cos(2\pi \langle \boldsymbol{\psi}, \mathbf{x} - \mathbf{x}' \rangle) d\Psi(\boldsymbol{\psi}) \quad (5)$$

where Ψ is the spectral measure of \mathbb{k} , and the second equality follows from \mathbb{k} being real-valued. This both gives a way of defining new kernels—by specifying their respective spectral measures—and calculating them numerically. In particular, defining $\boldsymbol{\phi}$ to be a vector of complex exponentials, and, following Rahimi and Recht (2008), discretizing the integral by using Monte Carlo, we obtain the approximation

$$\mathbb{k}(\mathbf{x} - \mathbf{x}') \approx \boldsymbol{\phi}(\mathbf{x})^\top \boldsymbol{\phi}(\mathbf{x}'). \quad (6)$$




Figure 1: We illustrate a regression problem on a real projective plane \mathbb{RP}_2 , with a Matérn-3/2 kernel in the sense of Section 4. Here and below we use the Boy’s surface immersion of \mathbb{RP}_2 into \mathbb{R}^3 (Bryant, 1988), which is non-isometric. Regression is applied to model the function $y(x) = \sum_{j=1}^5 \cos(d_{\mathbb{RP}_2}(x, \xi_j))$ where ξ_j are uniform samples on \mathbb{RP}_2 . We show the resulting posterior mean and standard deviation, and one posterior sample.

Letting ϕ_s be the components of ϕ , this further leads to the expression

$$f(\cdot) \approx \sum_{s=1}^S w_s \phi_s(\cdot) \quad w_\ell \sim \mathcal{N}(0, 1) \quad (7)$$

for approximating the Gaussian process itself. This enables one to efficiently draw approximate samples from the prior, and transform them into approximate posterior samples via (2) by plugging in the approximate prior in place of $f(\cdot)$ and $f(\mathbf{x})$.

The *feature map* ϕ is not uniquely determined: one choice is to use sines and cosines

$$\phi_{2s-1}(\mathbf{x}) = \sigma \sqrt{\frac{2}{S}} \cos(2\pi \langle \psi_s, \mathbf{x} \rangle) \quad \phi_{2s}(\mathbf{x}) = \sigma \sqrt{\frac{2}{S}} \sin(2\pi \langle \psi_s, \mathbf{x} \rangle) \quad \psi_s \sim \frac{1}{\sigma^2} \Psi \quad (8)$$

where $s \in 1, \dots, S/2$ for S even, and $\sigma^2 = \Psi(\mathbb{R}^n)$. Observe that, here, sines and cosines come in pairs: each sample from the spectral measure corresponds to exactly two functions used in the approximation. An alternative choice is to use only cosines, but add a random phase

$$\phi_s(\mathbf{x}) = \sigma \sqrt{\frac{2}{S}} \cos(2\pi \langle \psi_s, \mathbf{x} \rangle + \beta_s) \quad \psi_s \sim \frac{1}{\sigma^2} \Psi \quad \beta_s \sim \mathcal{U}(0, 2\pi) \quad (9)$$

for $s \in 1, \dots, S$ and $\sigma^2 = \Psi(\mathbb{R}^n)$, where \mathcal{U} is the uniform distribution. An important advantage of ϕ in (8) and (9) over the complex exponentials we began with is their real-valuedness. General analogs of Bochner’s theorem and Fourier feature expansions similar to (8) and (9) will be the cornerstone of our considerations in the geometric settings we study.

1.3 Previous Work and Contributions

Recently, authors in the statistics and machine learning communities have begun exploring generalizations of Gaussian process models to non-Euclidean settings—and, in particular to Riemannian manifolds, of which the spaces studied in this work are special cases.

The first difficulty here occurs in defining a positive semi-definite covariance kernel to begin with. For instance, Feragen et al. (2015) have shown that, for a Riemannian manifold

X with geodesic distance d , if one generalizes the widely used squared exponential kernel naïvely by substituting the Euclidean distance with d , then the obtained expression

$$\sigma^2 \exp\left(-\frac{d(x, x')^2}{2\kappa^2}\right) \quad (10)$$

will be positive semi-definite for all $\kappa > 0$ if and only if X is isometric to a Euclidean space. In some cases, there exist even stronger results of this type. For example, Da Costa et al. (2023) show that the kernel in (10) is not positive semi-definite for all values of κ —rather than for some—if the space X of interest belongs to the class of *symmetric spaces*,² thereby solving the open problem posed by Feragen and Hauberg (2016) in the case of symmetric spaces. This distance-based approach is therefore a technical disaster in need of alternatives.

One alternative, pursued by Lindgren et al. (2011), Said et al. (2013), Coveney et al. (2020), and Borovitskiy et al. (2020) and mostly restricted to compact manifolds, involves defining Riemannian Gaussian processes using stochastic partial differential equations. A key result in this setting is that the kernels of certain such processes admit *manifold Fourier feature* expansions

$$k(x, x') = \sigma^2 \sum_{j=0}^{\infty} \Psi(\lambda_j) f_j(x) f_j(x') \quad (11)$$

where (λ_j, f_j) are eigenpairs of the (negated) Laplace–Beltrami operator $-\Delta$, and $\Psi(\lambda_j)$ can be viewed as the generalized spectral measure—now supported on the nonnegative integers.

The upside of this approach is that, in principle, it allows one to compute covariance kernels without needing to solve stochastic partial differential equations. The downside is that computing the eigenpairs analytically or numerically is for many manifolds a non-trivial problem in its own right. One can observe that (11) coincides with various explicit expansions known for spheres and tori, derived from Bochner-type symmetry-based considerations, rather than from differential equations. This suggests that to build a deeper understanding, one should also study these phenomena from the perspective of stationarity under group action. This will be the starting point for our work. For instance, if viewed through an appropriate lens, in part I of this work we present techniques to bypass or solve the problem of finding Laplace–Beltrami eigenpairs for a large class of compact Riemannian manifolds.

If we set $S(\lambda) = e^{-\lambda t}$ in (11), the obtained expression coincides with the *heat kernel* (also known as the *diffusion kernel*), the fundamental solution of the heat equation, which generalizes the squared exponential kernel in a sound manner, without incurring the problems of geodesic-based approaches such as (10) (Grigoryan, 2009). This kernel has been studied by many authors in different specific settings (Lafferty et al., 2005; Kondor, 2008; Zhao and Song, 2018), and is a landmark target of development for our work. Niu et al. (2019) and Ye et al. (2020) suggest a general approach to approximate heat kernel’s values by simulating random walks, leveraging connections between the heat equation and Brownian motion. While this approach can work, it is also a computationally-expensive method-of-last-resort, since it can converge slowly and is not guaranteed to give a positive semi-definite kernel for finite samples. We develop different techniques that simultaneously apply to more general

2. The class of symmetric spaces is a subclass of homogeneous spaces that includes, for example, spheres and hyperbolic spaces.

kernels, including analogs of Matérn kernels, and obtain approximations which are guaranteed positive semi-definite.

We begin with an abstract description of stationary Gaussian processes on a Riemannian manifold X acted on by a suitably well-behaved Lie group G , given by Yaglom (1961) at the dawn of stochastic process theory. This description relies on certain far-reaching generalizations of harmonic analysis built upon the notions of *representation theory*. We develop this abstract description into a fully-accessible computational tool suitable for day-to-day statistical practice. We do this by making Yaglom’s abstract expressions constructive: specifically, we bring together and synthesize various theory and techniques to develop computational methods for statistical learning based on representation-theoretic expressions. This includes the following classes of approximations.

1. Algorithms for pointwise kernel evaluation, with automatic differentiation support.
2. Algorithms for efficiently drawing sample paths, both from priors and posteriors.

These algorithms approximate stationary kernels by linear combinations of functions which depend only on the given space, and admit closed-form expressions arising from algebraic considerations. The coefficients in these linear combinations depend on the particular kernel, but for the aforementioned heat and Matérn kernels can be computed exactly. Our sampling routines rely on the same quantities, and additionally require only that one can sample uniformly on the space of interest.

Together, these algorithmic primitives provide practitioners with all of the tools needed to define Gaussian processes for spatial models (Cressie, 1992) and decision-making algorithms such as Bayesian optimization (Frazier, 2018), along with primitives for constructing vector-valued processes (Hutchinson et al., 2021) and other extensions, and studying their theoretical properties (Rosa et al., 2023). The techniques developed provide a blueprint for how one might derive other non-Euclidean kernels from symmetry-based considerations, including potentially on discrete spaces (Borovitskiy et al., 2021, 2023) or more general algebraic structures where representation-theoretic considerations apply (Fukumizu et al., 2008). This is a purely theoretical paper: for applications, we refer the reader to the aforementioned works.

In total, our work makes it possible to use Gaussian processes on a much wider collection of manifolds admitting analytic descriptions than was previously practically available.

2. Stationary Gaussian Processes under Group Action

Let X be a set, and let G be a group acting on X with group action $\triangleright : G \times X \rightarrow X$. A kernel is called *stationary*³ if it satisfies

$$k(g \triangleright x, g \triangleright x') = k(x, x') \tag{12}$$

for all $g \in G$ and all $x, x' \in X$. Unless otherwise stated, we assume that k is real valued and continuous. From this, one can easily see that if $f \sim \text{GP}(0, k)$ has a stationary kernel, then the processes $f(g \triangleright \cdot)$ and $f(\cdot)$ have the same distribution. In this case, we say that f is stationary. Without loss of generality, we work with processes with zero mean. We

3. Note that the term *homogeneous* is also used in the literature, notably by Yaglom (1961).

now introduce the key results of Yaglom (1961) for different classes of spaces—these are, effectively, far-reaching generalizations of the Bochner’s theorem.

2.1 Stationary Kernels on Lie Groups

A *Lie group* (G, \cdot) is a group which is also a smooth manifold. We focus on the case where G acts on itself from both the left and the right, and defer the case where only one of these actions is present to Section 2.2. This falls into the general notion of stationarity introduced earlier by considering the product group $G \times G$ acting on the set G by $(g_1, g_2) \triangleright x = g_1 \cdot x \cdot g_2^{-1}$. In that case, for $(g_1, g_2) \in G \times G$, we have

$$k((g_1, g_2) \triangleright x, (g_1, g_2) \triangleright x') = k(g_1 \cdot x \cdot g_2^{-1}, g_1 \cdot x' \cdot g_2^{-1}) = k(x, x'). \quad (13)$$

Kernels—and general functions—that satisfy this property are called *bi-invariant*. Letting $e \in G$ be the identity element, we can define a single-argument function $\mathbb{k} : G \rightarrow \mathbb{R}$ such that

$$k(x_1, x_2) = k(x_2^{-1} \cdot x_1, e) = \mathbb{k}(x_2^{-1} \cdot x_1) = \mathbb{k}(x_1 \cdot x_2^{-1}) \quad (14)$$

which directly generalizes the one-argument notion of the Euclidean case.

We will need notions and results from the *representation theory* of compact Lie groups. The idea of representation theory is to study the structure of a group G by studying how it can be instantiated as the group of transformations of some vector space, which is often more concrete and easier to understand. More formally, a *representation* $\pi : G \rightarrow \text{GL}(V)$ is a homomorphism from G to the general linear group, consisting of bijective linear transformations of a (complex) vector space V . A representation π is called *unitary* if its range consists of unitary matrices. Representations that do not have any invariant subspaces except $\{0\}$ and V , and therefore cannot be decomposed into a direct sum of smaller representations, are called *irreducible*.

We work with smooth representations of compact Lie groups. In the setting at hand, there are at most a countable number of irreducible unitary representations—let Λ be an index set in bijective correspondence with the set of such representations. We defer the question of how to actually obtain such a set to subsequent sections. Let $\lambda \in \Lambda$ be an index. One can show that every smooth irreducible unitary representation

$$\pi^{(\lambda)} : G \rightarrow \text{GL}(V_\lambda) \quad (15)$$

is finite-dimensional, in the sense that the vector space V_λ is finite-dimensional. Let $d_\lambda = \dim V_\lambda$. If we introduce a G -invariant inner product $\langle \cdot, \cdot \rangle_{V_\lambda}$ and an orthonormal basis $e_1, \dots, e_{d_\lambda}$ in V_λ , we can define the *matrix coefficients* of the representation $\pi^{(\lambda)}$ by

$$\pi_{jk}^{(\lambda)} : G \rightarrow \mathbb{C} \quad \pi_{jk}^{(\lambda)}(g) = \langle \pi^{(\lambda)}(g)e_j, e_k \rangle_{V_\lambda} \quad (16)$$

where $1 \leq j, k \leq d_\lambda$. Let μ_G be the unique probabilistic Haar measure on G , which by compactness always exists. By the Peter–Weyl Theorem—see Folland (1995) as well as Weil (1940)—the matrix coefficients form an orthogonal basis in the space $L^2(G, \mu_G)$ with $\|\pi_{jk}^{(\lambda)}\|_{L^2(G, \mu_G)} = 1/\sqrt{d_\lambda}$. Thus, $f \in L^2(G, \mu_G)$ may be expressed in this basis in the form of an infinite series using the functions $\sqrt{d_\lambda}\pi_{jk}^{(\lambda)}(\cdot)$. This generalizes Fourier series to the case

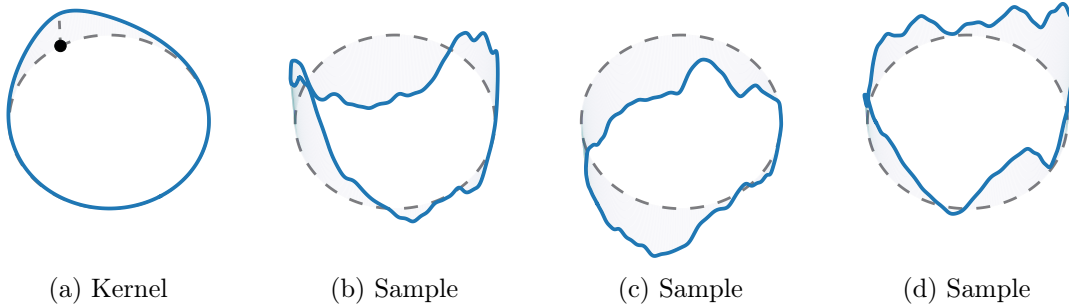


Figure 2: Illustration of the Matérn-3/2 kernel of Section 4, and three random samples from its respective prior Gaussian process, on the circle group \mathbb{S}_1 , which acts on itself by rotation.

of a compact topological group G . Define the *character* $\chi^{(\lambda)} : G \rightarrow \mathbb{C}$ corresponding to the representation $\pi^{(\lambda)}$ by the expression

$$\chi^{(\lambda)}(g) = \text{tr } \pi^{(\lambda)}(g) = \sum_{j=1}^{d_\lambda} \pi_{jj}^{(\lambda)}(g). \quad (17)$$

Note that the character does not depend on a choice of a basis $e_1, \dots, e_{d_\lambda}$ —a property inherited from the trace operator tr . Having introduced the necessary notions, we are now ready to present the representation-theoretic description of stationary Gaussian processes on compact Lie groups.

Theorem 1. *A Gaussian process $f \sim \text{GP}(0, k)$ on a compact Lie group G is stationary with respect to left-and-right action of G on itself if and only if k is of form*

$$k(g_1, g_2) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \text{Re } \chi^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (18)$$

where $a^{(\lambda)} \geq 0$ satisfy $\sum_{\lambda \in \Lambda} d_\lambda a^{(\lambda)} < \infty$. Moreover, for all λ , $\text{Re } \chi^{(\lambda)}$ is positive-definite.

Proof. This follows mostly by specializing Theorem 2 of Yaglom (1961), which describes complex-valued stationary processes on general topological groups, to the case where f is real-valued and Gaussian. A detailed proof is given in Appendix E. \square

This is a powerful result: it says that to compute a general stationary kernel on a Lie group numerically, it suffices to compute the coefficients $a^{(\lambda)}$, and the characters $\chi^{(\lambda)}$, up to some order of truncation. To carry this out, we will need some way to represent the index set Λ of irreducible unitary representations numerically. We describe methods for doing so and the methods for efficient approximate sampling from such processes in Section 3. For the circle group $G = \mathbb{S}_1 = \text{SO}(2) = \text{U}(1)$, an illustration of such a kernel, and samples of a corresponding process, computed using the techniques of Section 3, is given in Figure 2. Before proceeding to these, we study the situation on other spaces of interest.

2.2 Stationary Kernels on Homogeneous Spaces of Compact Lie Groups

Let G be a compact Lie group, and let H be a closed, not necessarily normal subgroup of G . We say that $X = G/H$ is a *homogeneous space*.⁴ X consists of cosets of the form $g \cdot H$ —the equivalence classes defined in the natural manner, and the action of G on X is defined by $g_1 \triangleright (g_2 \cdot H) = (g_1 \cdot g_2) \cdot H$. For simplicity, where unambiguous we will often drop the $(\cdot) \cdot H$ component from notation, and write elements of G/H as $x \in G/H$.

Consider complex-valued functions $f : G/H \rightarrow \mathbb{C}$ defined on the homogeneous space. It is not hard to see that such functions are in bijective correspondence with functions $\tilde{f} : G \rightarrow \mathbb{C}$ which are constant on all cosets, namely $\tilde{f}(g) = \tilde{f}(g \cdot h)$ for all $h \in H$. Moreover, if we define the natural projection $\phi : G \rightarrow G/H$ by $\phi : g \mapsto g \cdot H$, then every function $f \in L^2(G/H)$ can be written as $\tilde{f} \circ \phi \in L^2(G)$. This shows homogeneous spaces inherit square-integrability from the Lie group G of their automorphisms (Folland, 1995).

We can therefore expand a function $f \in L^2(G/H)$ using the $L^2(G)$ -orthonormal bases introduced previously. In particular, every such function can be written as an infinite weighted sum of matrix coefficients $\pi_{jk}^{(\lambda)}(\cdot)$. On the other hand, not every function constructed this way will be constant on the cosets, as all functions $f \in L^2(G/H)$ must be. To get around this, we revisit the bases chosen in the vector spaces V_λ , which were previously used to define the matrix coefficients. For an $f \in L^2(G/H)$, let $\tilde{f}(g) = (f \circ \phi)(g)$, $\tilde{f} \in L^2(G)$. Write

$$\tilde{f}(g) = \sum_{\lambda \in \Lambda} d_\lambda \sum_{j,k=1}^{d_\lambda} \langle \tilde{f}, \pi_{jk}^{(\lambda)} \rangle_{L^2(G)} \pi_{jk}^{(\lambda)}(g) \quad (19)$$

$$= \sum_{\lambda \in \Lambda} d_\lambda \int_G \tilde{f}(u) \underbrace{\left(\sum_{j,k=1}^{d_\lambda} \pi_{jk}^{(\lambda)}(g) \overline{\pi_{jk}^{(\lambda)}(u)} \right)}_{\chi^{(\lambda)}(u^{-1} \cdot g)} d\mu_G(u) = \sum_{\lambda \in \Lambda} d_\lambda (\tilde{f} * \chi^{(\lambda)})(g). \quad (20)$$

where $*$ denotes convolution, and the calculation follows by interchanging sums and integrals. We therefore see that the convolution with a character $\chi^{(\lambda)}$ corresponds to orthogonal projection onto the subspace $\text{span } \pi_{jk}^{(\lambda)} \subseteq L^2(G)$. This correspondence—between projections and convolution against characters—will reappear in different guises throughout this work. It is easy to see that since \tilde{f} is by definition constant on cosets, its convolution against a character is also constant on cosets, thus $\tilde{f} * \chi^{(\lambda)} \in L^2(G/H)$. Therefore to find an orthogonal (generalized Fourier) basis on $L^2(G/H)$, it suffices, for every λ , to find an orthogonal basis of the vector subspace of $\text{span } \pi_{jk}^{(\lambda)}$ that consists of functions constant on cosets.

Consider a function $f \in \text{span } \pi_{jk}^{(\lambda)}$. If we write the condition $f(g \cdot h) = f(g)$ for all $g \in G, h \in H$ in the basis $\pi_{jk}^{(\lambda)}(\cdot)$, where f is expressed as $f(\cdot) = \sum_{j,k=1}^{d_\lambda} c_{jk} \pi_{jk}^{(\lambda)}(\cdot)$, we get

$$\sum_{j,k=1}^{d_\lambda} \left(\sum_{l=1}^{d_\lambda} c_{jl} \pi_{kl}^{(\lambda)}(h) \right) \pi_{jk}^{(\lambda)}(g) = \sum_{j,k=1}^{d_\lambda} c_{jk} \pi_{jk}^{(\lambda)}(g). \quad (21)$$

4. In geometry, a homogeneous space is a topological space X upon which some group G of automorphisms acts continuously and transitively (Folland, 1995, Section 2.6). In this case, if we fix any point $x \in X$ and define H to be the subgroup of G consisting of such elements $g \in G$ that $g \triangleright x = x$, it can be proved that X is isomorphic to G/H .

Let \mathbf{C} be a matrix with entries c_{jk} . The above condition can be more compactly written as $\pi^{(\lambda)}(h)\mathbf{C}^\top = \mathbf{C}^\top$ for all $h \in H$. Thus, the rows of \mathbf{C} are invariant vectors of the representation $\pi^{(\lambda)}$ that is restricted onto the subgroup H —note that these can be linearly dependent. Let r_λ be the dimension of the space of invariant vectors of $\pi^{(\lambda)}|_H$, and choose an orthonormal basis $e_1, \dots, e_{d_\lambda}$ of V_λ such that $e_1, \dots, e_{r_\lambda}$ are invariant vectors of $\pi^{(\lambda)}|_H$. Then, the rows of \mathbf{C} are linear combinations of $e_1, \dots, e_{r_\lambda}$ and thus the dimension of the subspace of $\text{span } \pi_{jk}^{(\lambda)}$ that consists of functions constant on cosets is exactly $r_\lambda d_\lambda$.

If $\pi_{jk}^{(\lambda)}$ are defined using the preceding basis of V_λ , a direct computation shows that $\pi_{jk}^{(\lambda)}$, $1 \leq j \leq d_\lambda$, $1 \leq k \leq r_\lambda$ are constant on cosets and thus form the orthogonal basis we were originally looking for (Weil, 1940).

Using these notions, define the *spherical functions* to be the entries of the first r_λ columns of the matrix $\pi_{jk}^{(\lambda)}$ for $1 \leq j \leq d_\lambda$ and $1 \leq k \leq r_\lambda$, and define the *zonal spherical functions* similarly, but with $1 \leq j \leq r_\lambda$ and $1 \leq k \leq r_\lambda$. Note that not all spherical functions are zonal, due to the differing indices, and that zonal spherical functions are constant on double cosets, meaning that

$$\pi_{jk}^{(\lambda)}(h_1 \cdot g \cdot h_2) = \pi_{jk}^{(\lambda)}(g) \quad h_1, h_2 \in H \quad 1 \leq j \leq r_\lambda \quad (22)$$

which can be seen by a direct computation. With these notions at hand, we are now ready to describe stationary Gaussian processes on G/H .

Theorem 2. *A Gaussian process $f \sim \text{GP}(0, k)$ on a compact homogeneous space G/H is stationary with respect to the action of G if and only if k is of form*

$$k(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \Lambda} \sum_{j,k=1}^{r_\lambda} a_{jk}^{(\lambda)} \text{Re } \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (23)$$

where the coefficients $a_{jk}^{(\lambda)} \in \mathbb{R}$ form symmetric positive semi-definite matrices $\mathbf{A}^{(\lambda)}$ of size $r_\lambda \times r_\lambda$ satisfying $\sum_{\lambda} \text{tr } \mathbf{A}^{(\lambda)} < \infty$, and $\pi_{jk}^{(\lambda)}$ are the zonal spherical functions. Moreover, for each individual $\lambda \in \Lambda$, the corresponding sum over j, k is positive semi-definite.

Proof. Similar to before, this follows mostly by specializing Theorem 5 of Yaglom (1961), which describes complex-valued stationary processes to the real-valued Gaussian case. A detailed proof is given in Appendix E. \square

Theorem 2 simplifies for *symmetric* homogeneous spaces, which possess certain additional properties that hold, for instance, for all homogeneous spaces of constant sectional curvature. For such spaces, we study non-compact analogs of this result in part II of this work (Azangulov et al., 2024). For symmetric spaces, $r_\lambda \leq 1$, thus every representation of G corresponds to at most one zonal spherical function $\pi_{11}^{(\lambda)}$. The matrix $\mathbf{A}^{(\lambda)}$, then, simplifies into a scalar similar to the one which appeared in the Lie group setting. This occurs, for instance, on spheres $\mathbb{S}_n = \text{SO}(n+1)/\text{SO}(n)$ and Grassmann manifolds $\text{Gr}(r, n) = \text{O}(n)/(\text{O}(r) \times \text{O}(n-r))$, including projective spaces RP_n as a special case. We illustrate a Matérn kernel on the sphere, and samples of the corresponding process, in Figure 3.

Theorem 2 can also be used to give a description of left-invariant kernels on Lie groups, complementing the preceding results on bi-invariant kernels.

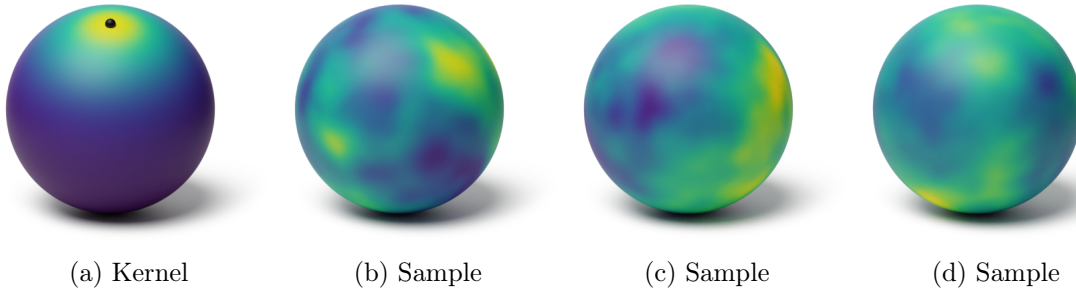


Figure 3: Illustration of the Matérn-3/2 kernel of Section 4, and samples from its corresponding prior Gaussian process, on the sphere $G/H = \text{SO}(3)/\text{SO}(2) = \mathbb{S}_2$.

Corollary 3. *Every left-invariant but not necessarily bi-invariant kernel on a compact Lie group G can be expressed according to (23), with $r_\lambda = d_\lambda$.*

Proof. Letting $H = \{e\}$ be the trivial group, the claim is immediate from Theorem 2. \square

Note that any stationary positive semi-definite function k on a homogeneous space G/H induces a positive semi-definite function

$$\mathbb{k}(g) = k(g \cdot H, e \cdot H) = k(g_1 \cdot H, g_2 \cdot H) \quad g_2^{-1} \cdot g_1 = g \quad (24)$$

over the group G which is constant on all double cosets $H \cdot g \cdot H$, namely

$$\mathbb{k}(h_1 \cdot g \cdot h_2) = \mathbb{k}(g) \quad (25)$$

for $g \in G$ and $h_1, h_2 \in H$. Conversely, any positive semi-definite function on G constant on all double cosets $H \cdot g \cdot H$ induces a positive semi-definite kernel on the homogeneous space G/H .⁵

3. Computational Techniques for Stationary Gaussian Processes

We now develop computational techniques for implementing stationary Gaussian process models from the preceding section. To work with the posterior mean, posterior covariance, and pathwise conditioning expressions of Section 1.1, we need two computational primitives:

- (a) Evaluating the covariance kernel $k(\cdot, \cdot)$ pointwise.
- (b) Efficiently sampling from the prior $f(\cdot)$.⁶

In general, (a) and (b) need to be performed in an automatically-differentiable manner, for instance to optimize model hyperparameters using maximum likelihood, or differentiate the prior density in Hamiltonian Monte Carlo.

5. Note that \mathbb{k} should only represent a left-invariant kernel. In particular, this means that $\mathbb{k}(g_2^{-1} \cdot g_1) = \mathbb{k}(g_1 \cdot g_2^{-1})$ is *not* necessarily true for all $g_1, g_2 \in G$ and that these kernels do not fall under the assumptions of Theorem 1 of the preceding section.

6. 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*.

Assuming the scaling coefficients $a^{(\lambda)}$ or the scaling matrices $\mathbf{A}^{(\lambda)}$ are a priori fixed, Theorems 1 and 2 offer avenues for evaluating $k(\cdot, \cdot)$ and, with some additional work, approximately sampling $f(\cdot)$. This can be achieved by truncating the respective infinite series to the desired accuracy. To do so, one must be able to traverse the index set Λ : we study this in Section 3.1. For pointwise kernel evaluation, one must also compute the characters $\chi^{(\lambda)}$ in the Lie group case, and the zonal spherical functions $\pi_{j,k}^{(\lambda)}$ in the homogeneous space case: we study this in Section 3.2.

For sampling, in Section 3.3 we develop a technique we term *generalized random phase Fourier features* that allows efficient sampling using only the aforementioned quantities, as long as one can also draw random samples from the manifold’s Haar measure. Moreover, it allows one—if needed—to compute the Karhunen–Loève basis in terms of the characters or zonal spherical functions, respectively.

With these quantities at hand, all that remains is to obtain the scaling coefficients $a^{(\lambda)}$ or scaling matrices $\mathbf{A}^{(\lambda)}$, which are kernel-dependent. We study these for specific kernel classes in Section 4.

3.1 Enumerating Representations of Compact Lie Groups

Both for (approximate) pointwise kernel evaluation and for sampling on compact Lie groups and their homogeneous spaces we will need a way to enumerate the set Λ of irreducible unitary representations. This is a central question of the representation theory. We now sketch a practical approach to doing so.

Every compact Lie group G contains a connected abelian Lie subgroup, which is isomorphic to a torus—call the maximal (by inclusion) of these subgroups *maximal tori*. All maximal tori have the same dimension and are conjugate to one another, in the sense that if $T^{(1)}, T^{(2)} \subseteq G$ are two maximal tori, then there exists $g \in G$ such that $T^{(1)} = g^{-1} \cdot T^{(2)} \cdot g$. Moreover, every element of G is conjugate to an element of a given maximal torus—for example, every unitary matrix which is an element $U \in G = \mathbf{U}(d)$ is diagonalizable by a unitary transformation, namely $U = VDV^{-1}$ where D is an element of the torus of diagonal unitary matrices and $V \in \mathbf{U}(d)$.

Representations of G can be therefore be studied by studying their relationship with representations on the torus, which are simpler. Specifically, on an m -dimensional torus \mathbb{T}^m , the standard theory of Fourier analysis enables one to express a suitably regular function $f : \mathbb{T}^m \rightarrow \mathbb{C}$ as an infinite sum

$$f(x) = \sum_{n \in \mathbb{Z}^m} c_n e^{2\pi i \langle n, x \rangle} = \sum_{n \in \Lambda} c_n \pi_n(x) \quad (26)$$

which we have expressed in two ways—analytic, and representation-theoretic. In the expression above $\Lambda = \mathbb{Z}^m$ and $\pi_n : \mathbb{T}^m \rightarrow \mathbb{T} \subseteq \mathbf{GL}(\mathbb{C}^1)$ are the irreducible unitary representations, which on a torus are one-dimensional and can be explicitly written as complex exponentials. This establishes a bijective correspondence between such representations and Fourier basis functions, at least as long as one studies the torus.

Since every compact Lie group contains a maximal torus T , this gives us information about representations on more general compact Lie groups. When restricted to such a torus, a representation of G splits into a sum of one-dimensional irreducible representations,

parametrized by tuples of integers, which are called the *weights* of the representation. One can show these tuples, and the original representation on G , are both uniquely determined by a specific weight called the *highest weight*. By mapping the set of highest weights with a linear transform, one obtains the set of *signatures*.

Theorem 4. *The set of irreducible unitary representations of G is in explicit bijective correspondence with the set of SIGNATURES $\Lambda \subseteq \mathbb{Z}^m$, where m is the dimension of T .*

Proof. Appendix A. □

Henceforth, we use Λ to refer to the set of signatures, which can be calculated explicitly for most Lie groups. We provide explicit expressions for $\text{SO}(n)$ and $\text{SU}(n)$ in Appendix B.

We can use this bijective correspondence to impose an ordering on Λ , and define truncated analogs of the infinite series given by Yaglom (1961). Specifically, for a stationary kernel k on a compact Lie group G , and $L < \infty$, define

$$k^{(L)}(g_1, g_2) = \sum_{\lambda \in \tilde{\Lambda}} a^{(\lambda)} \text{Re } \chi^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (27)$$

where $\tilde{\Lambda} \subseteq \mathbb{Z}^m$ is of size L and consists of signatures with the (potentially up to a constant) largest values of $a^{(\lambda)}$, and we take the real part of each character since k is real-valued. Analogously, on a compact homogeneous space G/H define

$$k^{(L)}(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \tilde{\Lambda}} \sum_{j,k=1}^{r_\lambda} a_{jk}^{(\lambda)} \text{Re } \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (28)$$

where $\tilde{\Lambda} \subseteq \mathbb{Z}^m$ is of size L , and consists of signatures with the (potentially up to a constant) largest values under a suitable matrix ordering, for instance the one given by $\text{tr } \mathbf{A}^{(\lambda)}$.

For the concrete kernels we consider in Section 4, signatures determine the respective Laplace–Beltrami eigenvalues, and the coefficients $a^{(\lambda)}$ grow monotonically with these eigenvalues. Similarly, the matrices $\mathbf{A}^{(\lambda)}$ turn out to be scalar multiples of the identity, where the scalar multiple grows monotonically with the respective eigenvalues. This allows one to obtain a suitable ordering, and therefore implement the truncated series numerically—as long as one can calculate the characters and zonal spherical functions, respectively. For any truncation level, the resulting expressions are stationary and positive semi-definite. We therefore proceed to study how to calculate characters and zonal spherical functions.

3.2 Pointwise Kernel Evaluation

To evaluate a stationary kernel pointwise, we now study how to calculate the characters $\chi^{(\lambda)}$ and zonal spherical functions $\pi_{jk}^{(\lambda)}$ corresponding to a signature $\lambda \in \Lambda$.

3.2.1 CALCULATING THE CHARACTERS OF LIE GROUPS

To calculate the characters of a compact Lie group G , we will use the observations that (i) every element of G is conjugate to an element of a fixed maximal torus T , (ii) characters are conjugation-invariant, and (iii) on a maximal torus, they admit an explicit description via the Weyl character formula.




Figure 4: Values of the heat kernel $k(\bullet, \cdot)$ on \mathbb{T}^2 , \mathbb{RP}_2 , and \mathbb{S}_2 . The torus \mathbb{T}^2 is equipped with the flat metric, so the depicted embedding is not isometric.

Theorem 5. *Every character $\chi^{(\lambda)}$ with signature λ can be expressed as*

$$\chi^{(\lambda)}(g) = \chi^{(\lambda)}(t) = \frac{P_1(t)}{P_2(t)} \quad t = \tilde{g}^{-1} \cdot g \cdot \tilde{g} \quad (29)$$

where $t \in T$ is an element of a maximal torus conjugate to g . For a fixed identification of T with $\mathbb{T}^m \subseteq \mathbb{C}^m$, P_1 and P_2 are explicit polynomials with respect to the variables $z_1, \dots, z_m, \bar{z}_1, \dots, \bar{z}_m$ of \mathbb{T}^m . Moreover, the number of terms in P_1, P_2 does not grow with λ .

Proof. This is a reformulation of Fegan (1991, Theorem 9.9)—see Appendix C. □

These expressions are ratios of polynomials, and can therefore be evaluated and automatically differentiated. To find t , one employs techniques that depend on the group. In many cases, such as $\text{SU}(n)$ or $\text{SO}(2n+1)$, this reduces to finding the eigenvalues of g . Additionally, for many groups, the polynomials P_1 and P_2 can often be rewritten as certain determinants of order $\dim T$, or short sums thereof. In cases where $P_2 = 0$, the ratio should be understood as a limit when the denominator approaches zero, for instance when g has repeated eigenvalues. This is not an issue, since one can show that the denominator is always a divisor of the numerator, so the ratio $\frac{P_1}{P_2}$ is itself a polynomial. Further details on these expressions can be found in Appendix C, and explicit expressions for the Lie groups $\text{SU}(n)$ and $\text{SO}(n)$ can be found in Appendix D.

3.2.2 CALCULATING THE ZONAL SPHERICAL FUNCTIONS OF HOMOGENEOUS SPACES

The class of homogeneous spaces is very general, and thus a unified treatment similar to the Lie group setting is more difficult. Consequently, the most efficient ways to compute zonal spherical functions tend to involve space-specific considerations, even in symmetric spaces which contain at most one zonal spherical function per signature. Nonetheless, we present a general technique to express zonal spherical function on G/H in terms of the characters on G , which we term *generalized periodic summation* by analogy with the Euclidean case. Before presenting this, however, we consider homogeneous spaces where case-specific descriptions of spherical functions are available, at least to some extent.

- $\mathbb{S}_n = \text{SO}(n+1)/\text{SO}(n)$: *the n -dimensional sphere*. This is a symmetric homogeneous space. The spherical functions on this space are called *spherical harmonics*, and are very well-studied. In particular, zonal spherical harmonics can be explicitly written in terms of the *Gegenbauer polynomials*.
- $V(k, n) = \text{SO}(n)/\text{SO}(n-k)$: *Stiefel manifolds*. This is the space of all k -tuples of orthonormal vectors in \mathbb{R}^n . The spherical functions are called *Stiefel harmonics*, and are studied by Gelbart (1974).
- $\text{RP}_n = \text{SO}(n)/\text{O}(n-1)$: *real projective space*. This is the space of all one-dimensional linear subspaces of \mathbb{R}^n , and is a symmetric space. One can obtain the zonal spherical functions by leveraging connections between this space and the sphere, or by rewriting it as $\text{P}(n) = \text{O}(n)/(\text{O}(1) \times \text{O}(n-1))$, which is as a special case of a Grassmannian.
- $\text{Gr}(k, n) = \text{O}(n)/(\text{O}(k) \times \text{O}(n-k))$: *Grassmannians*. This is the space of all k -dimensional linear subspaces of \mathbb{R}^n , and is a symmetric homogeneous space. Spherical functions for Grassmannians are described by Davis (1999a,b).

In cases where space-specific descriptions of spherical functions are available, efficient algorithmic techniques can often be constructed from them. We illustrate the values of kernels defined in Section 4, computed by means of such techniques, in Figure 4. Unfortunately, many descriptions involve heavy mathematical machinery, and it can be highly non-trivial to translate them into a practical algorithm. In cases where applications merit it, we view the formulation of such algorithms as a promising avenue for future work, and now return to the general case.

If the coefficients matrices $\mathbf{A}^{(\lambda)} = a^{(\lambda)}\mathbf{I}$ are multiples of the identity matrix, as will be the case for the heat and Matérn kernels studied in Section 4, one can compute kernels on G/H in terms of characters of G . The next claim shows how.

Theorem 6. *Let $f \sim \text{GP}(0, k)$, where k is stationary and satisfies $\mathbf{A}^{(\lambda)} = a^{(\lambda)}\mathbf{I}$. Then*

$$k(g_1 \cdot H, g_2 \cdot H) = \int_H k_G(g_1 \cdot h, g_2) d\mu_H(h) \quad f(g \cdot H) = \int_H f_G(g \cdot h) d\mu_H(h) \quad (30)$$

where μ_H is the Haar measure on H , and $f_G \sim \text{GP}(0, k_G)$ with

$$k_G(g_1, g_2) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \text{Re } \chi^{(\lambda)}(g_2^{-1} \cdot g_1). \quad (31)$$

Proof. Appendix E. □

This generalizes the *periodic summation* idea in the Euclidean setting, which says that a kernel on the torus can be obtained by starting with a kernel on \mathbb{R}^d , and summing over multiples of 2π (Borovitskiy et al., 2020). One can therefore compute such kernels on a homogeneous space G/H by computing their analogs on the underlying Lie group G using the techniques of Section 3.2.1, and averaging over uniform random samples from H , obtaining

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{1}{S} \sum_{s=1}^S k_G(g_1 \cdot h_s, g_2) \quad h_s \sim \mu_H. \quad (32)$$




Figure 5: Samples from a Gaussian process with heat kernel covariance, on \mathbb{T}^2 , \mathbb{RP}_2 , and \mathbb{S}_2 .

Note that this procedure—in the finite-sample case—does *not* necessarily preserve positive semi-definiteness (or even symmetry) of the kernel. To guarantee a well-defined kernel at the cost of additional computational complexity, it can therefore be preferable to average over both arguments of the kernel, obtaining

$$k(g_1 \cdot H, g_2 \cdot H) \approx \frac{1}{S^2} \sum_{s=1}^S \sum_{s'=1}^S k_G(g_1 \cdot h_s, g_2 \cdot h_{s'}) \quad h_s \sim \mu_H \quad (33)$$

where we emphasize that a total of S samples are computed from μ_H , which are summed over the set of all pairs of samples. The resulting expression is guaranteed to be symmetric and positive semi-definite. We now proceed to consider sampling.

3.3 Efficient Sampling

Since a sample from a posterior Gaussian process can be obtained by transforming prior samples into posterior samples using (2) we now study the problem of sampling from prior Gaussian process f . Given the set of signatures Λ of Section 3.1, we can form a prior sample from f by summing a sequence of processes $f^{(\lambda)} \sim \text{GP}(0, k^{(\lambda)})$ where $k^{(\lambda)} = a^{(\lambda)} \text{Re} \chi^{(\lambda)}(g_2^{-1} \cdot g_1)$ or its homogeneous space analog. We now introduce techniques for efficiently sampling from such processes, which further allow one to compute a Karhunen–Loève basis as a byproduct.

3.3.1 GENERALIZED RANDOM PHASE FOURIER FEATURES

The key idea for constructing samples comes from the following observation.

Proposition 7. *Let \mathcal{S} be a measurable space, let μ be probability measure on \mathcal{S} , and let ϕ_j with $j = 1, \dots, N$ be an $L^2(\mathcal{S}, \mu)$ -orthonormal system. Let $K(x, x') = \sum_{j=1}^N \phi_j(x) \overline{\phi_j(x')}$. Then for all $x, x' \in \mathcal{S}$ we have*

$$K(x, x') = \int_{\mathcal{S}} K(x, u) \overline{K(x', u)} d\mu(u). \quad (34)$$

Proof. This follows by direct computation, given for completeness in Appendix E. □

We can therefore Monte Carlo approximate the function K as

$$K(x, x') \approx \frac{1}{S} \sum_{s=1}^S K(x, u_s) \overline{K(x', u_s)} \quad u_s \sim \mu. \quad (35)$$

If $f \sim \text{GP}(0, K)$ then we have the approximation

$$f(x) \approx \frac{1}{\sqrt{S}} \sum_{s=1}^S w_s K(x, u_s) \quad u_s \sim \mu \quad w_s \sim \text{N}(0, 1). \quad (36)$$

Thus, starting from a covariance kernel of the form $K(x, x') = \sum_{j=1}^N \phi_j(x) \overline{\phi_j(x')}$ for an orthonormal system ϕ_j , we have constructed a family of random functions $K(\cdot, u_l)$ which enable us to approximately express the Gaussian process $\text{GP}(0, K)$ in a form which is directly amenable to efficient sampling.

If $\mathcal{S} = \mathbb{T}^d$, and ϕ_j are sinusoids from the sine-cosine random Fourier features, then $K(x, u_l)$ are exactly cosines with a random phase shift. We therefore say that u_l is the *random phase* and K is the *phase function*, and call this construction *generalized random phase Fourier features* for general measure spaces, equipped only with a probability measure μ . We illustrate samples from such processes, computed additionally using the techniques of Section 4, in Figure 5. Note that the process (36) may fail to be stationary: we examine this further in Section 3.3.4. We now consider spaces of interest.

3.3.2 PHASE FUNCTIONS AND EFFICIENT SAMPLING FOR LIE GROUPS

For a compact Lie group G , there are usually efficient methods for sampling from the Haar measure on G : see Mezzadri (2007). Thus, generalized random phase Fourier features enable us to efficiently sample from approximate priors using only the characters. If $\chi^{(\lambda)}$ is real-valued, namely $\text{Re } \chi^{(\lambda)} = \chi^{(\lambda)}$, then we let $K^{(\lambda)}(g_1, g_2) = d_\lambda \chi^{(\lambda)}(g_2^{-1} \cdot g_1)$, corresponding to $\phi_\lambda = \sqrt{d_\lambda} \pi_{jk}^{(\lambda)}$. Otherwise, note that the map $g \mapsto (\pi^{(\lambda)}(g))^*$, which sends a group element g to the conjugate transpose of the matrix defined by the matrix coefficients $\pi_{ij}^{(\lambda)}(g)$, is an irreducible unitary representation. It is thus equal to $\pi^{(\lambda')}$ for some $\lambda' \in \Lambda$, $\lambda' \neq \lambda$ and whose character $\chi^{(\lambda')}$ obviously equals $\overline{\chi^{(\lambda)}}$. In this case, we associate the aggregate of $\sqrt{d_\lambda} \pi_{jk}^{(\lambda)}$ and $\sqrt{d_{\lambda'}} \pi_{jk}^{(\lambda')}$ with ϕ_l , and further associate $d_\lambda \chi^{(\lambda)} + d_{\lambda'} \chi^{(\lambda')} = 2d_\lambda \text{Re } \chi^{(\lambda)}$ with $K^{(\lambda)}$. With this, we can compute each of the terms in (18) using only the characters, and the resulting approximation is always real-valued. We therefore obtain

$$f(g) \approx \sum_{\lambda \in \tilde{\Lambda}} \frac{1}{\sqrt{S}} \sum_{s=1}^S w_s^{(\lambda)} K^{(\lambda)}(g, u_s) \quad u_s \sim \mu_G \quad w_s^{(\lambda)} \sim \text{N}\left(0, \frac{a^{(\lambda)}}{d_\lambda}\right). \quad (37)$$

Here, $\tilde{\Lambda}$ are the truncated signatures defined previously, μ_G is the probabilistic Haar measure on the group G . $\tilde{\Lambda}$ should be chosen so as *not* to contain λ' together with any λ except in the case $\lambda = \lambda'$. The size of $\tilde{\Lambda}$ and value of L control approximation quality. One can interpret the function $g \mapsto [K^{(\lambda)}(g, u_s)]_{s=1, \dots, S, \lambda \in \tilde{\Lambda}}$ as an approximate finite-dimensional feature map: one can therefore also use this construction for other purposes, beyond just sampling, where feature maps are helpful. This can include, for instance, speeding up training using Bayesian-linear-model-based techniques.

3.3.3 PHASE FUNCTIONS AND EFFICIENT SAMPLING FOR HOMOGENEOUS SPACES

If G/H is a homogeneous space, the key question for defining generalized random phase Fourier features is how to correctly choose the function K in Proposition 7. A natural idea is to associate each of zonal spherical function $\pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1)$ with a separate function $\tilde{K}(g_1 \cdot H, g_2 \cdot H)$: however, this is not the right thing to do, as we obtain

$$\tilde{K}(g_1 \cdot H, g_2 \cdot H) = \sum_{j=1}^{d_\lambda} \pi_{jk_1}^{(\lambda)}(g_1) \overline{\pi_{jk_2}^{(\lambda)}(g_2)} \quad (38)$$

for $1 \leq k_1, k_2 \leq r_\lambda$, which is not the form required by Proposition 7 unless $k_1 = k_2$. Moreover, \tilde{K} and the induced approximation f can be complex-valued. A more delicate approach is therefore needed. Let $a_{k_1 k_2}^{(\lambda)}$ be a set of coefficients such that they form a symmetric positive semi-definite matrix whose precise form is to be determined later. Define

$$K(g_1 \cdot H, g_2 \cdot H) = \sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda)} \pi_{k_1 k_2}^{(\lambda)}(g_2^{-1} \cdot g_1) = \sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda)} \sum_{j=1}^{d_\lambda} \pi_{jk_1}^{(\lambda)}(g_1) \overline{\pi_{jk_2}^{(\lambda)}(g_2)} \quad (39)$$

and observe the following.

Proposition 8. *The function K defined by (39) satisfies*

$$K(x, x') = \int_{G/H} K(x, u) \overline{K(x', u)} d\mu_{G/H}(u) \quad (40)$$

if and only if the matrix $\mathbf{A} = d_\lambda^{-1} (a_{k_1 k_2}^{(\lambda)})$ is idempotent, namely $\mathbf{A}^2 = \mathbf{A}$.

Proof. Appendix E. □

It follows that if the scaled coefficient matrix \mathbf{A} whose entries are $d_\lambda^{-1} a_{k_1 k_2}^{(\lambda)}$ is an orthogonal projection, then generalized random phase Fourier features sampling is possible. By positive semi-definiteness, \mathbf{A} has an eigendecomposition

$$\mathbf{A} = \mathbf{U}_{(\lambda)} \text{diag}(\alpha_1^{(\lambda)}, \dots, \alpha_{r_\lambda}^{(\lambda)}) \mathbf{U}_{(\lambda)}^\top \quad \mathbf{U}_{(\lambda)} = (\mathbf{u}_1^{(\lambda)}, \dots, \mathbf{u}_{r_\lambda}^{(\lambda)}) \quad (41)$$

where $\alpha_k^{(\lambda)}$ are non-negative eigenvalues, and the columns of $\mathbf{U}_{(\lambda)}$ are an orthonormal basis of eigenvectors of \mathbf{A} . Reinterpreting this matrix product as a sum, we can write

$$\mathbf{A} = \sum_{k=1}^{r_\lambda} \alpha_k^{(\lambda)} \underbrace{\mathbf{u}_k^{(\lambda)} \mathbf{u}_k^{(\lambda)\top}}_{\mathbf{U}_{(\lambda)}^{(k)}} \quad (42)$$

where $\mathbf{U}_{(\lambda)}^{(k)}$ are orthogonal projection matrices. Now, define

$$\tilde{K}_k^{(\lambda)}(g_1 \cdot H, g_2 \cdot H) = \sum_{k_1, k_2=1}^{r_\lambda} d_\lambda \mathbf{U}_{\lambda, k_1 k_2}^{(k)} \pi_{k_1 k_2}^{(\lambda)}(g_2^{-1} \cdot g_1). \quad (43)$$

Recall that the aforementioned map $g \mapsto (\pi^{(\lambda)}(g))^*$ defines an irreducible unitary representation and thus equal to $\pi^{(\lambda')}$ for some $\lambda' \in \Lambda$. If λ is such that $\lambda = \lambda'$, then $\pi_{jk}^{(\lambda)} = \overline{\pi_{kj}^{(\lambda)}}$ which, by virtue of $\mathbf{U}_{(\lambda)}^{(k)}$ being symmetric, implies $\widetilde{K}_k^{(\lambda)}(x, x') \in \mathbb{R}$. We have in this case

$$\sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda)} \operatorname{Re} \pi_{k_1 k_2}^{(\lambda)}(g_2^{-1} \cdot g_1) = \sum_{k=1}^{r_\lambda} \frac{\alpha_k^{(\lambda)}}{d_\lambda} \widetilde{K}_k^{(\lambda)}(g_1 \cdot H, g_2 \cdot H) \quad (44)$$

where $\widetilde{K}_k^{(\lambda)}$ are real-valued phase functions which for $\lambda = \lambda'$ we rename to $K_k^{(\lambda)}$.

If $\lambda \neq \lambda'$, to ensure the resulting samples are real-valued, write

$$\sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda)} \operatorname{Re} \pi_{k_1 k_2}^{(\lambda)}(g_2^{-1} \cdot g_1) + \sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda')} \operatorname{Re} \pi_{k_1 k_2}^{(\lambda')}(g_2^{-1} \cdot g_1) \quad (45)$$

$$= \sum_{k_1, k_2=1}^{r_\lambda} a_{k_1 k_2}^{(\lambda)} \pi_{k_1 k_2}^{(\lambda)}(g_2^{-1} \cdot g_1) + a_{k_1 k_2}^{(\lambda')} \pi_{k_1 k_2}^{(\lambda')}(g_2^{-1} \cdot g_1) = \sum_{k=1}^{r_\lambda} \frac{\alpha_k^{(\lambda)}}{d_\lambda} 2 \operatorname{Re} \widetilde{K}_k^{(\lambda)}(g_1 \cdot H, g_2 \cdot H) \quad (46)$$

where we used the fact that $a_{k_1 k_2}^{(\lambda)} = a_{k_2 k_1}^{(\lambda')} \in \mathbb{R}$. From the preceding argument it is clear that $2 \operatorname{Re} \widetilde{K}_k^{(\lambda)}$ are real-valued phase functions which for $\lambda \neq \lambda'$ we rename to $K_k^{(\lambda)}$.

Since terms with different λ values in the generalized random phase Fourier feature approximation are independent, it suffices to sample from each $K_k^{(\lambda)}$ by (36), multiply by $\sqrt{\alpha_k^{(\lambda)}/d_\lambda}$, and reassemble the expression by summing terms. This gives the approximation

$$f(g \cdot H) \approx \sum_{\lambda \in \widetilde{\Lambda}} \sum_{k=1}^{r_\lambda} \frac{1}{\sqrt{S}} \sum_{s=1}^S w_s^{(\lambda)} K_k^{(\lambda)}(g \cdot H, u_s) \quad w_s^{(\lambda)} \sim \mathcal{N}\left(0, \frac{\alpha_k^{(\lambda)}}{d_\lambda}\right) \quad u_s \sim \mu_{G/H} \quad (47)$$

where $\widetilde{\Lambda}$ are the truncated signatures, and $\mu_{G/H}$ is the probabilistic Haar measure on G/H . $\widetilde{\Lambda}$ should be chosen so as *not* to contain λ' together with any λ unless $\lambda = \lambda'$, where this notation was previously defined in the context of the matrix-coefficient conjugate-transpose map $g \mapsto (\pi^{(\lambda)}(g))^*$. The size of $\widetilde{\Lambda}$ and the value of S control the quality of the approximation. If G/H is a symmetric space, then $r_\lambda \leq 1$, and the above formulas simplify considerably. Mirroring the Lie group setting, the function $x \mapsto [K_k^{(\lambda)}(x, u_s)]_{s=1, \dots, S, k=1, \dots, r_\lambda, \lambda \in \widetilde{\Lambda}}$ gives an approximate finite-dimensional feature map for the kernel.

3.3.4 CALCULATING A KARHUNEN–LOÈVE BASIS

It is not hard to see that matrix coefficients or general spherical functions provide a Karhunen–Loève basis for the Gaussian process, whose covariance is given by the respective character or a combination of zonal spherical functions. However, the former are available on a much smaller class of spaces compared to the latter. This was one of the main reasons we explored how to sample from a stationary Gaussian process using only the former. We now ask and answer a complementary question: given only the characters or zonal spherical functions, is it possible to obtain a Karhunen–Loève basis?

Our strategy will be to find a non-orthonormal basis, and Gram–Schmidt orthogonalize it to obtain a Karhunen–Loève basis. This generalizes the approach of Dai and Xu (2013, Section 1.3) from spheres to more general spaces. We begin from an abstract formulation. Consider an orthonormal system ϕ_j in the space $L^2(\mathcal{S}, \mu)$, where (\mathcal{S}, μ) is a measure space and μ is a probability measure. Define, as before, the function $K(x, x') = \sum_{j=1}^N \phi_j(x) \overline{\phi_j(x')}$. Consider a set of N points x_j in \mathcal{S} and define the matrices

$$\mathbf{M}_k = \begin{pmatrix} \phi_1(x_1) & \cdots & \phi_1(x_k) \\ \vdots & \ddots & \vdots \\ \phi_k(x_1) & \cdots & \phi_k(x_k) \end{pmatrix}. \quad (48)$$

Lemma 9. *There exist $\{x_1, \dots, x_N\} \subseteq \mathcal{S}$ such that $\det \mathbf{M}_k \neq 0$ for all $1 \leq k \leq N$.*

Proof. Appendix E. □

Definition 10. *A collection $\{x_j\}_{j=1}^N \subseteq \mathcal{S}$ is a FUNDAMENTAL SET if $\det \mathbf{M}_N \neq 0$.*

The existence of such sets, which are non-unique, is guaranteed by Lemma 9.

Proposition 11. *The notion of a fundamental set does not depend on a specific orthonormal system ϕ_j but rather only on the space $\text{span } \phi_j$ —that is, on K .*

Proof. We have

$$\mathbf{M}_N^* \mathbf{M}_N = \begin{pmatrix} K(x_1, x_1) & \cdots & K(x_N, x_1) \\ \vdots & \ddots & \vdots \\ K(x_1, x_N) & \cdots & K(x_N, x_N) \end{pmatrix}. \quad (49)$$

We have $\det \mathbf{M}_N \neq 0$ if and only if $\det \mathbf{M}_N^* \mathbf{M}_N > 0$. The latter depends on K only, and hence only on $\text{span } \phi_j$. □

Theorem 12. *If x_j is a fundamental set with respect to the system ϕ_j , then*

$$\text{span } \phi_j = \text{span } K(\cdot, x_j) \quad (50)$$

Proof. Appendix E. □

Notice that (49) allows one to check whether a set of points is a fundamental set without using the basis ϕ_j . Thus, one can find a fundamental set by uniformly sampling a set of points $x_1, \dots, x_N \in G$, and checking whether the determinant $\det \mathbf{M}_N^* \mathbf{M}_N$ vanishes. This is generally a rare event, so in practice it typically suffices to sample once.

The basis $K(x_j, \cdot)$ is not necessarily orthonormal, but pairwise inner products of basis functions can nonetheless be easily calculated using its special form, by writing

$$\langle K(\cdot, x_j), K(\cdot, x_k) \rangle_{L^2(\mathcal{S}, \mu)} = \sum_{n=1}^N \overline{\phi_n(x_j)} \phi_n(x_k) = K(x_k, x_j) = \overline{K(x_j, x_k)}. \quad (51)$$

Here, we have used (34) and $K(x, x') = \overline{K(x', x)}$. One therefore obtains a Karhunen–Loève basis by Gram–Schmidt orthogonalizing $K(\cdot, x_j)$ at $\mathcal{O}(N^3)$ cost.

4. Computational Techniques for Heat and Matérn Kernels

In the preceding sections, we obtained Lie-theoretic and representation-theoretic expressions for stationary kernels and Gaussian processes under group action. In these expressions, only the scaling factors $a^{(\lambda)}$ and $\mathbf{A}^{(\lambda)}$, respectively, depend on the specific kernel. We therefore study how to compute them for the *Matérn* and *heat* kernel classes: our approach will be to define heat kernels first, and extend to the Matérn class using connections between them. To proceed, we need to turn our spaces into Riemannian manifolds by defining a metric tensor, for which we make the canonical choice.

Assumption I. *X is equipped with the metric tensor induced by its Killing form.*

If $X = G$ is a Lie group, this metric is automatically bi-invariant under the action of G on itself from both sides. If $X = G/H$ is a homogeneous space, this metric is the quotient metric obtained from G , and is invariant under the canonical action of G on G/H .

4.1 Squared Exponential Kernels as Solutions of the Heat Equation

Before defining the general Riemannian heat kernel, consider first its Euclidean analog. In this setting, it is widely known as the squared exponential, or Gaussian, or RBF kernel. We denote it by $k_{\infty, \kappa, \sigma^2}$. It has two positive numerical parameters κ and σ^2 , representing its amplitude (variance) and length scale, respectively. For $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ this kernel is

$$k_{\infty, \kappa, \sigma^2}(\mathbf{x}, \mathbf{x}') = \sigma^2 e^{-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\kappa^2}}. \quad (52)$$

Observe that if we let $\kappa = \sqrt{2t}$ and $\sigma^2 = \frac{1}{(4\pi t)^{n/2}}$ then, defining

$$\mathcal{P}(t, \mathbf{x}, \mathbf{x}') = k_{\infty, \sqrt{2t}, (4\pi t)^{-n/2}}(\mathbf{x}, \mathbf{x}') \quad (53)$$

we can express this kernel as the fundamental solution of the heat equation

$$\frac{\partial \mathcal{P}}{\partial t}(t, \mathbf{x}, \mathbf{x}') = \Delta_{\mathbf{x}} \mathcal{P}(t, \mathbf{x}, \mathbf{x}') \quad \mathcal{P}(0, \mathbf{x}, \mathbf{x}') = \delta(\mathbf{x} - \mathbf{x}') \quad (54)$$

where $\Delta_{\mathbf{x}}$ is the Laplacian acting on the \mathbf{x} -argument, and the initial condition $\delta(\mathbf{x} - \mathbf{x}') \in \mathcal{S}'(X)$ is the Dirac delta function, understood as a vector in the Schwartz space $\mathcal{S}'(X)$ of tempered distributions. This gives the general way of defining general analogs of squared exponential kernels, which we call heat kernels, and will be our starting point.

The heat equation generalizes to a Riemannian manifold X under mild regularity conditions, so long as Δ is replaced with the Laplace–Beltrami operator. The Riemannian heat kernel \mathcal{P} can now be defined as the fundamental solution of the Riemannian heat equation. We also define its squared-exponential-like reparametrization by

$$k_{\infty, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\kappa}} \mathcal{P}(\kappa^2/2, x, x') \quad (55)$$

where C_{κ} is a constant chosen depending on how one wishes to normalize the kernel's amplitude. For spaces where $\mathcal{P}(\kappa^2/2, x, x)$ does not depend on x , such as homogeneous spaces, a natural choice is $C_{\kappa} = \mathcal{P}(\kappa^2/2, x, x)$. For other spaces, if X is compact, one can

choose $C_\kappa = \int_X \mathcal{P}(\kappa^2/2, x, x) \text{dvol}(x)$. By general theory (Grigoryan, 2009), these functions are positive definite kernels. At this stage, however, it is not clear whether or not they are stationary, and if so, what the corresponding scaling coefficients $a^{(\lambda)}$ or scaling matrices $\mathbf{A}^{(\lambda)}$ are. To understand this, we specialize to the settings at hand.

4.1.1 COMPACT LIE GROUPS

To understand how to express solutions of the heat equation in a more explicit manner, we begin by connecting the representation-theoretic notions introduced previously with spectral properties of the Laplace–Beltrami operator Δ .

Result 13. *For every λ , the matrix coefficients $\pi_{jl}^{(\lambda)}$ are eigenfunctions of $-\Delta$ corresponding to the same eigenvalue $\alpha_\lambda \geq 0$. Moreover, the eigenvalues α_λ can be written⁷*

$$\alpha_\lambda = \|\rho\|_B^2 - \|w + \rho\|_B^2 \quad (56)$$

where w is the highest weight of the representation with signature λ , $\|g\|_B^2 = -B(g, g)$, B is the Killing form of G , and ρ is the half-sum of positive roots of G —see Appendix A.

Proof. The first part of the claim is well-known: we sketch a simple but instructive proof in Appendix E. The second part is given in Fegan (1991, Theorem 10.6). \square

For technical details on these quantities, see Appendix A. With this connection established, one can calculate an explicit expression for the heat kernel on G in the sense of the heat equation using the representation-theoretic notions at hand.

Proposition 14. *The heat kernel is stationary, and is given by*

$$k_{\infty, \sqrt{2t}, (4\pi t)^{-n/2}}(g_1, g_2) = \mathcal{P}(t, g_1, g_2) = \sum_{\lambda \in \Lambda} e^{-\alpha_\lambda t} d_\lambda \chi_\lambda(g_2^{-1} \cdot g_1). \quad (57)$$

Proof. Write out both sides of the heat equation in the basis $\{\sqrt{d_\lambda} \pi_{jk}^{(\lambda)}\}_{\lambda, j, k}$. \square

Therefore, to evaluate heat kernels pointwise, and to sample from the corresponding Gaussian processes, it suffices to apply the machinery of Section 3. In particular, note that $\lambda \mapsto e^{-\alpha_\lambda t}$ is monotone—we will see this holds for all concrete kernels that we consider. Therefore, to choose the set of truncated signatures $\tilde{\Lambda}$, we select the signatures corresponding to the smallest eigenvalues, and apply Result 13.

4.1.2 HOMOGENEOUS SPACES OF COMPACT LIE GROUPS

We again begin by connecting representation-theoretic notions with spectral theory of Δ .

Result 15. *The spherical functions $\pi_{jk}^{(\lambda)}$ on a compact homogeneous space G/H are the eigenfunctions of $-\Delta$. Moreover, the eigenvalue corresponding to every spherical function $\pi_{jk}^{(\lambda)}$, which by definition is also a matrix coefficient of G , is given in (56).*

Proof. Appendix E. \square

7. This expression is sometimes called *Freudenthal’s formula for the eigenvalues*.

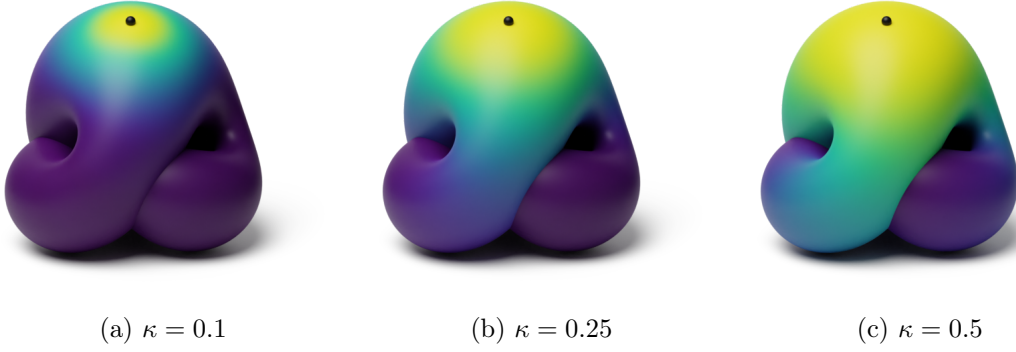


Figure 6: Values of heat kernels on a projective plane \mathbb{RP}_2 , with various length scales.

We can use this to compute explicit expressions for the heat kernel on G/H .

Proposition 16. *The heat kernel is stationary, and is given by*

$$k_{\infty, \sqrt{2t}, (4\pi t)^{-n/2}}(g_1 \cdot H, g_2 \cdot H) = \mathcal{P}(t, g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \Lambda} e^{-\alpha_\lambda t} d_\lambda \sum_{j,k=1}^{r_\lambda} \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1). \quad (58)$$

Proof. Write out both sides of the heat equation in the basis of spherical functions. \square

In homogeneous spaces, too, it therefore suffices to apply the machinery of Section 3. We illustrate heat kernels computed in this manner in Figure 6.

4.2 Matérn Kernels

We now use the preceding notions to introduce Matérn kernels. On \mathbb{R}^n , these constitute (Rasmussen and Williams, 2006; Stein, 1999) a widely used three parameter family

$$k_{\nu, \kappa, \sigma^2}(\mathbf{x}, \mathbf{x}') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{\kappa} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{\kappa} \right) \quad (59)$$

with parameters $\sigma^2 > 0$, $\kappa > 0$, and $\nu > 0$ being the variance, length scale, and smoothness, respectively. Here, K_ν is the modified Bessel function of the second kind (Gradshteyn and Ryzhik, 2014). We are interested in defining them in settings beyond \mathbb{R}^n .

We approach the problem by connecting Matérn kernels with squared exponential kernels, and leveraging the preceding results. Still on \mathbb{R}^n , as $\nu \rightarrow \infty$, we have

$$\lim_{\nu \rightarrow \infty} k_{\nu, \kappa, \sigma^2}(\mathbf{x}, \mathbf{x}') = k_{\infty, \kappa, \sigma^2}(\mathbf{x}, \mathbf{x}') \quad (60)$$

but this is not the only connection between these families of kernels. One can show

$$k_{\nu, \kappa, \sigma^2}(\mathbf{x}, \mathbf{x}') = \frac{(2\nu)^\nu}{\Gamma(\nu) \kappa^{2\nu}} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2} u} k_{\infty, \sqrt{2u}, \sigma^2}(\mathbf{x}, \mathbf{x}') du \quad (61)$$

$$= \sigma^2 \frac{(2\nu)^\nu (4\pi)^{n/2}}{\Gamma(\nu) \kappa^{2\nu}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2} u} \mathcal{P}(u, \mathbf{x}, \mathbf{x}') du. \quad (62)$$

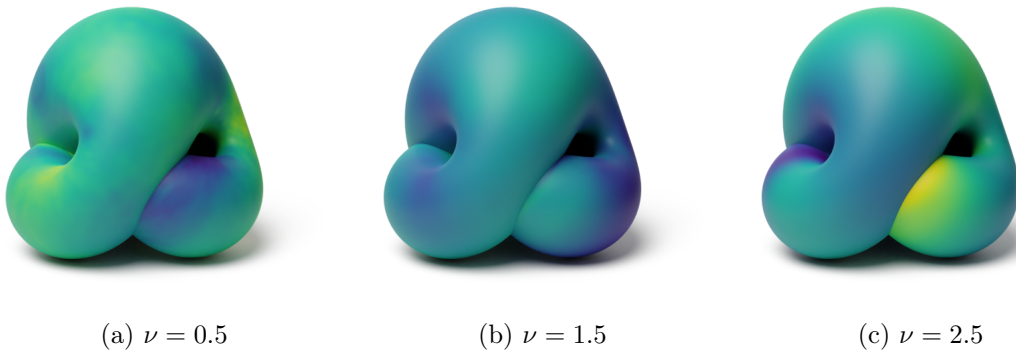


Figure 7: Samples from Matérn Gaussian processes on \mathbb{RP}_2 , with varying smoothness.

This can be understood by observing that the spectral measure of the Matérn kernel, which is a T distribution, can be written as a gamma mixture of Gaussians, which are the spectral measure of the squared exponential kernel. (61) and (62) are the kernel analog of this: Appendix E gives a proof. More generally, we adopt this expression—after dropping inessential constants and renormalizing—as the *definition* of a Matérn kernel.

Definition 17. Let X be a n -dimensional compact Lie group or a homogeneous space of a compact Lie group, and let $\mathcal{P}(u, x, x')$ be the heat kernel. Define the MATÉRN KERNEL by

$$k_{\nu, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\nu, \kappa}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \mathcal{P}(u, x, x') du \quad (63)$$

where $C_{\nu, \kappa}$ is chosen so that $k_{\nu, \kappa, \sigma^2}(x, x) = \sigma^2$.

It is easy to check that such kernels are positive (semi)definite whenever the respective squared exponential kernels are positive (semi)definite, regardless of the underlying space X —see Appendix E. The main issue at this stage is that working with Matérn kernels ostensibly requires one to evaluate an integral. We therefore study the expression for the spaces at hand to understand simplifications, given as follows.

Proposition 18. The Matérn kernel on a compact Lie group G is given by

$$k_{\nu, \kappa, \sigma^2}(g_1, g_2) = \frac{\sigma^2}{C'_{\nu, \kappa}} \sum_{\lambda \in \Lambda} \left(\frac{2\nu}{\kappa^2} + \alpha_\lambda \right)^{-\nu-n/2} d_\lambda \chi_\lambda(g_2^{-1} \cdot g_1) \quad (64)$$

where $C'_{\nu, \kappa}$ is a normalizing constant which ensures that $k_{\nu, \kappa, \sigma^2}(g, g) = \sigma^2$.

Proof. Appendix E. □

Proposition 19. The Matérn kernel on a homogeneous space G/H is given by

$$k_{\nu, \kappa, \sigma^2}(g_1 \cdot H, g_2 \cdot H) = \frac{\sigma^2}{C'_{\nu, \kappa}} \sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} \left(\frac{2\nu}{\kappa^2} + a_\lambda \right)^{-\nu-n/2} d_\lambda \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1). \quad (65)$$

where $C'_{\nu, \kappa}$ is a normalizing constant which ensures that $k_{\nu, \kappa, \sigma^2}(x, x) = \sigma^2$.

Proof. Appendix E. □

Note that the constants $C'_{\nu,\kappa}$ in (64) and (65) can differ, owing to the different settings. All these kernels can be computed using the techniques of Section 3. One can similarly derive formulas for random sampling from Matérn priors. We illustrate such samples in Figure 7. We conclude by verifying that these kernels satisfy the same differentiability properties as the usual Euclidean Matérn class.

Theorem 20. *For all $\nu > 0$, Matérn kernels on compact Lie groups and their homogeneous spaces are positive definite and lie in the Sobolev space $H^{\nu+n/2}$. Thus, they are continuous and possess continuous derivatives of all orders strictly less than ν .*

Proof. Appendix E. □

4.3 Approximation Error and Computational Costs

To compute the preceding kernels or efficiently sample the respective Gaussian processes, one must truncate infinite series, and potentially draw Monte Carlo samples from the respective probability measures. These result in approximation error, which we analyze both analytically and empirically in this section.

4.3.1 ASYMPTOTICS

Truncation error on general manifolds. We begin with understanding truncation error, since except for very specific cases where the series can be summed analytically,⁸ this approximation error is always present. Truncation error is governed by the rate of convergence of the respective series, which depends in turn on the growth rate of the Laplace–Beltrami eigenvalues α_λ , the representation degrees d_λ , as well as r_λ for homogeneous spaces, and on the kernel’s parameters. Throughout this section we work on a compact homogeneous space X . However, to get the basic asymptotics, we initially ignore the homogeneous space structure of X , and treat it as a general compact Riemannian manifold: then, k_{ν,κ,σ^2} can be obtained via the general manifold Fourier feature expansion. Let

$$k_{\nu,\kappa,\sigma^2}^{(J)}(x, x') = \frac{\sigma^2}{C_{\nu,\kappa}} \sum_{j=0}^J \Psi_{\nu,\kappa}(\lambda_j) f_j(x) f_j(x') \quad \Psi_{\nu,\kappa}(\lambda) = \begin{cases} \left(\frac{2\nu}{\kappa^2} + \lambda\right)^{-\nu-n/2} & \nu < \infty \\ e^{-\frac{\kappa^2}{2}\lambda} & \nu = \infty \end{cases} \quad (66)$$

where $(-\lambda_j, f_j)$ are Laplace–Beltrami eigenpairs, defined such that $\{f_j\}_{j=0}^\infty$ is an orthonormal basis of $L^2(X)$ and $n = \dim(X)$. Note that $k_{\nu,\kappa,\sigma^2}^{(\infty)} = k_{\nu,\kappa,\sigma^2}$ where the right-hand side is the Matérn or heat kernel, as discussed in Sections 4.1 and 4.2. For simplicity, we set $\sigma^2 = C_{\nu,\kappa}$ and omit the corresponding index. Since the presented kernel expressions for Lie groups and homogeneous spaces can be viewed as simplifications of this general expression, which group together and combine certain eigenfunctions into characters and zonal spherical functions, respectively, one can understand truncation error by analyzing the respective tail of (66).

8. The only such case known to the authors is the circle $\mathbb{S}_1 = \text{SO}(2)$, see Example 8 in Borovitskiy et al. (2020), which adapts the argument in the supplementary material of Guinness and Fuentes (2016).

Specifically, we focus on the L^2 -truncation error $\|k_{\nu,\kappa} - k_{\nu,\kappa}^{(J)}\|_{L^2(X \times X)}$. To bound this, we apply Weyl's Law (Chavel, 1984), which for a general compact Riemannian manifold of dimension n says that there is a constant $C > 0$ such that

$$C^{-1}j^{2/n} \leq \alpha_j \leq Cj^{2/n}. \quad (67)$$

We obtain the following.

Proposition 21. *For $J \in \mathbb{N}$ sufficiently large, and \tilde{C} depending only on ν , κ , and C , we have*

$$\|k_{\nu,\kappa} - k_{\nu,\kappa}^{(J)}\|_{L^2(X \times X)} \leq \tilde{C} \begin{cases} J^{-2\nu/n-1/2} & \nu < \infty, \\ J^{1/2-1/n} e^{-\frac{\kappa^2}{2C} J^{2/n}} & \nu = \infty. \end{cases} \quad (68)$$

Proof. Appendix E. □

This gives a bound on the truncation error, which is seen to decrease in a manner that depends on the kernel's smoothness as quantified by parameter ν , and on the manifold's geometry through the growth rate of λ_j —in this simple case, determined by the dimension n and constant C only. It is therefore clear that manifold Fourier feature truncation error decreases polynomially for Matérn kernels and exponentially for the heat kernel.

Truncation error on Lie groups. On a Lie group $X = G$, one can think of the character-based kernel expression as a simplification of the general manifold expression, where certain terms corresponding to repeated eigenvalues are combined. Recall the expression

$$k_{\nu,\kappa}^{(L)}(g_1, g_2) = \sum_{\lambda \in \tilde{\Lambda}} a^{(\lambda)} \operatorname{Re} \chi^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (69)$$

of Section 3, where $\tilde{\Lambda}$ is the truncated set of L irreducible representations, constructed according to Section 3.2.1. Note that each summand in this expression corresponds to d_λ^2 eigenfunctions in the manifold Fourier feature expansion. This means that this series converges faster than manifold Fourier features do, with the precise speedup determined by the behavior of d_λ for the group of interest.

To get an estimate of d_λ , one can apply the results of Augarten (2020), obtaining that the number of irreducible representations of dimension at most N is at most $(r!)^3 N^{r/R}$, where r is the rank of G and R is the number of positive roots of G , which grows quadratically with r . This reveals the Lie group's *rank* to be the primary factor influencing the degree of speedup: the larger it is, the bigger d_λ will be. In general, however, understanding how geometric factors affect the precise interplay between the growth of a_λ and d_λ is non-trivial, and we therefore examine it empirically in the sequel.

Truncation error on homogeneous spaces. On a homogeneous space $X = G/H$ the situation is similar: recall that the zonal spherical functions play the role of characters, and

$$k_{\nu,\kappa}^{(L)}(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \tilde{\Lambda}} a^{(\lambda)} \sum_{j=1}^{r_\lambda} \operatorname{Re} \pi_{jj}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (70)$$

where $\tilde{\Lambda}$ is as above and $\pi_{jj}^{(\lambda)}$ are zonal spherical functions. Each inner sum of r_λ terms corresponds to $r_\lambda d_\lambda$ eigenfunctions, which up to a constant coincide with spherical functions. Here, the rate of speedup is determined by the behavior of d_λ and r_λ for the particular homogeneous space under study.

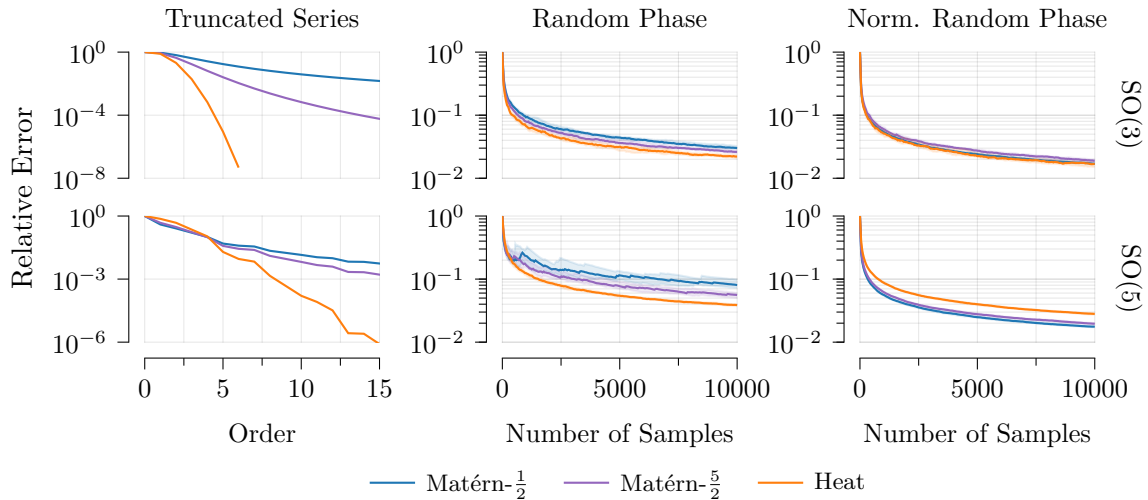


Figure 8: Here we compute the approximation error for the truncated character-based series expression, and random phase approximation, as a function of the truncation order, and number of Monte Carlo samples, respectively. For the random phase variants, we plot the median, as well as 25% and 75% quantiles, computed using 20 different random seeds. We examine relative error, which is equivalent to normalizing each curve so that it equals one at the origin. We see that, for the character-based expansion, approximation error decreases rapidly with the truncation order, especially for the smoother heat kernel. Note that the rate also depends on the length scale. Random-phase approximations converge at a slower rate. Normalized random phase approximations outperform their non-normalized variants.

Other asymptotics. The generalized periodic summation of Section 3.2.2 and generalized random phase Fourier feature estimators of Sections 3.2 and 3.3 are Monte Carlo estimators. Therefore, they converge at the usual \sqrt{S} rate, in the Monte Carlo sense, where S is the number of samples. The constants are governed by the interplay between the geometry of the space under consideration and properties of the kernel. To understand how these settings perform and get a more quantitative view of truncation, we now turn to empirical evaluation.

4.3.2 EMPIRICAL EVALUATION

We now evaluate the error in all of the proposed kernel approximations—based on truncation, periodic summation, and generalized random phase Fourier features—empirically. For this, we selected a total of two Lie groups, namely the special orthogonal groups $SO(3)$ and $SO(5)$, and three homogeneous spaces, namely the Stiefel manifolds $V(1, 5)$, $V(2, 5)$, and $V(3, 5)$. We consider three kernels of varying smoothness, namely a Matérn- $\frac{1}{2}$ kernel corresponding to a non-differentiable Gaussian process, a Matérn- $\frac{5}{2}$ kernel corresponding to a twice-differentiable Gaussian process, and heat kernel corresponding to an infinitely-differentiable Gaussian process. We select length scales that ensure the processes are neither approximately spatially constant, nor resemble white noise. All experiments—where applicable due to stochasticity—were repeated 20 times to assess variability.

In our first experiment, we evaluate the truncation error of the character-based kernel expression, and the approximation error of the generalized random phase approximation, on $\text{SO}(3)$ and $\text{SO}(5)$. For this, we first precompute the kernel $k_{\nu,\kappa}^{(L_{\max})}$ using $L_{\max} = 50$ total characters, to create a suitable set of ground-truth values for the kernel. We assume σ^2 is chosen in such a way that $k_{\nu,\kappa}^{(L_{\max})}(x, x) = 1$. We then vary the truncation order L , namely the number of terms used in the series in (69), and compare the resulting kernel $k_{\nu,\kappa}^{(L)}$ to the precomputed kernel. We similarly vary the number of Monte Carlo samples S used for the random phase approximation given by (37) with number of characters fixed to L_{\max} . We denote the covariance function of the resulting approximation by $k_{\nu,\kappa}^{(L,S)}$. We also consider an alternative approximation $\tilde{k}_{\nu,\kappa}^{(L,S)}(x, x')$, which we call the *normalized random phase approximation*, which forces $\tilde{k}_{\nu,\kappa}^{(L,S)}(x, x) = 1$. This approximation is defined by

$$\tilde{k}_{\nu,\kappa}^{(L,S)}(x, x') = \frac{k_{\nu,\kappa}^{(L,S)}(x, x')}{\sqrt{k_{\nu,\kappa}^{(L,S)}(x, x)k_{\nu,\kappa}^{(L,S)}(x', x')}}. \quad (71)$$

In all cases, we examine the L^2 -error $\|k_{\nu,\kappa}^{L_{\max}} - \cdot\|_{L^2(X \times X)}$ where (\cdot) is either $k_{\nu,\kappa}^{(L)}$, $k_{\nu,\kappa}^{(L,S)}$ or $\tilde{k}_{\nu,\kappa}^{(L,S)}$. In the former case, this error can be computed analytically in terms of $\Psi_{\nu,\kappa}(a^{(\lambda)})$. In the latter cases, we draw 50 uniformly random points $x_1, \dots, x_{50} \in X$ and approximate

$$\left\|k_{\nu,\kappa}^{(L_{\max})} - k_{\nu,\kappa}^{(L,S)}\right\|_{L^2(X \times X)}^2 \approx \frac{1}{50^2} \sum_{i=1}^{50} \sum_{j=1}^{50} \left|k_{\nu,\kappa}^{(L_{\max})}(x_i, x_j) - k_{\nu,\kappa}^{(L,S)}(x_i, x_j)\right|^2 \quad (72)$$

and similarly for $\tilde{k}_{\nu,\kappa}^{(L,S)}$ instead of $k_{\nu,\kappa}^{(L,S)}$.

Results are given in Figure 8. We see that, mirroring the theoretical picture given in Section 4.3.1, the truncation error decay rate varies according to the smoothness of the kernel. For the heat kernel, we observe superexponential decay in truncation error. For the two Matérn kernels, we observe decay that is slightly slower than exponential on the scales considered. This indicates that the extra d_λ -factor of the character-based formula produces a significant increase in accuracy on these manifolds, reinforcing the theoretical picture presented previously.

Next, we examine the random phase approximation, again on $\text{SO}(3)$ and $\text{SO}(5)$, whose results are also given in Figure 8. Here, we see that error decreases at the expected Monte Carlo rate. Compared to the previous setting, we also see less variability in behavior among kernels with different levels of smoothness. In all cases, the normalized approximation $\tilde{k}_{\nu,\kappa}^{L,S}$ performs better: when applicable, we thus recommend using it in practice.

For our next experiment, we evaluate the performance of periodic summation on the Stiefel manifolds $V(1, 5)$, $V(2, 5)$, and $V(3, 5)$. Note that, of these, $V(1, 5) = \mathbb{S}_4$ is the four-dimensional sphere. Similarly to the preceding experiment, we compute a ground-truth kernel: on $V(1, 5)$ we do so using the zonal-spherical-function-based available for the sphere⁹, whereas on $V(2, 5)$ and $V(3, 5)$ where such an expansion is not available we instead compute a ground-truth kernel using 150 Monte Carlo samples. We then vary the number of Monte Carlo

9. We use the first 10 eigenspaces on sphere. Note that from Gelbart (1974), these correspond to eigenspaces on $\text{SO}(5)$ with signature $(k, 0)$ —we use this to improve the kernel approximation.




Figure 9: Here we compute the approximation error for the periodic-summation-based kernel computation technique. We plot the median, as well as 25% and 75% quantiles, computed using 20 different random seeds. We examine relative error, which is equivalent to normalizing each curve so that it equals one at the origin. We see that error decreases at a modest \sqrt{S} rate, as one would expect given the Monte Carlo nature of the approximation.

samples used for periodic summation, and examine the L^2 -error between the approximate kernel and the ground truth, again computed as in (72), but this time using 20 random points instead of 50.

Results are given in Figure 9. Here, we see that on $V(1, 5)$, approximation accuracy varies slightly according to smoothness, and decays slightly faster for the heat kernel compared to Matérn kernels. On $V(2, 5)$ and $V(3, 5)$, no such differences are visible, and all techniques perform comparably—though, it is possible this similarity is caused by the comparatively-imprecise ground truth used in these cases.

In summary, we see that empirically, truncation error tends to decay very fast on the examples considered: a few dozen characters or zonal spherical functions is usually more than enough to accurately approximate the limiting kernel. The Monte-Carlo-based approximations considered—generalized random phase Fourier features and generalized periodic summation—decay at the usual Monte Carlo rate. For these, a reasonable number of samples may be on the order of hundreds or thousands to accurately represent the limiting kernel. Finally, note that even if the error is large, all of the kernel approximations considered here are positive semi-definite, and therefore define perfectly valid and usable Gaussian processes in their own right, irrespective of the truncation level chosen.

4.3.3 COMPUTATIONAL COMPLEXITY

We conclude by summarizing the computational costs of the three proposed kernel approximations: (i) truncated series expansions based on characters and zonal spherical functions, (ii) generalized periodic summation, and (iii) generalized random phase Fourier features.

Given a set of truncated signatures $\tilde{\Lambda}$ of size L , our proposed truncated series expansions are linear combinations of either characters or zonal spherical functions, depending on the setting. The computational cost of evaluating the truncated series, given that characters or zonal spherical functions are known and cost $\mathcal{O}(1)$ to evaluate, is therefore linear: this gives a complexity of $\mathcal{O}(L)$. In cases where these respective functions are known, they are usually either polynomials or ratios of such, with powers increasing as $\mathcal{O}(L)$ grows. Thus,

the $\mathcal{O}(1)$ assumption is arguably only appropriate in cases where these expansions are very efficient, and can be evaluated in a numerically stable manner using only lower-order terms.¹⁰ Fortunately, in practice this is often the case: in all of the examples we considered, a few dozen terms sufficed to obtain an accurate approximation.

Generalized periodic summation generally requires one to use Monte Carlo to evaluate the respective integral. The computational costs of doing this directly—that is, without using a double sum to improve accuracy—are $\mathcal{O}(LS)$, where S is the number of samples. Recall that, to guarantee positive semi-definiteness, one can modify the approximation by introducing a double sum: this results in $\mathcal{O}(LS^2)$ kernel evaluation cost. This is the most costly algorithm in this work, and shows that having an explicit form for the zonal spherical functions can be significantly beneficial in terms of efficiency.

Generalized random phase Fourier features provide a feature map with LS features in total, which can be sampled in $\mathcal{O}(S)$ time, where S is the number of samples. Evaluating this feature map costs $\mathcal{O}(LS)$ for Lie groups, and $\mathcal{O}(LRS)$ for homogeneous spaces, where $R = \max_{\lambda \in \tilde{\Lambda}} r_\lambda$. The complexity of approximately sampling a Gaussian process prior using this technique is the same.

5. Conclusion

In this work, we derived a set of techniques that make it possible to deploy Gaussian process models on a number of compact Riemannian manifolds which admit analytic descriptions and are important in applications. This was done by combining and synthesizing results from a variety of mathematical fields, including probability, statistics, differential geometry, and representation theory. Our techniques yield approximations which are essentially-closed-form algebraic expressions, and are guaranteed positive semi-definite, making them simple and reliable to use. We hope our ideas enable practitioners to deploy these models in new prospective applications.

Acknowledgments

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.

10. In theory, there exists a character formula which can be calculated in constant time: unfortunately, this formula involves ratios where the denominator can evaluate to zero, rendering the calculation unstable. We therefore use an alternative formula, which leads to an explicit polynomial and is numerically stable, but whose size unfortunately grows with L .

Appendix A. Traversing Representations via Highest Weights and Signatures

Here, primarily following Bröcker and Tom Dieck (2013) and Fegan (1991), we review the relevant notions of representation theory and describe a general approach that allows one to traverse the set of irreducible unitary representations.

Recall that a substantial part of the mathematical structure of a Lie group is encoded within its tangent spaces. Due to the group structure, these are linked rigidly enough that it suffices to study the tangent space at one point, such as the identity element e of the group. This tangent space, equipped with some additional algebraic structure, is called the *Lie algebra* \mathfrak{g} of the group G .¹¹

Following the discussion from the main text, every irreducible representation of a connected compact semi-simple Lie group G produces a representation of its maximal torus $T_m \subseteq G$ by restriction. Using the structure of the torus, this representation splits into a sum of one-dimensional irreducible representations, parametrized by tuples of integers, which are called the *weights* of the representation. It turns out these representations possess even more structure: they are in fact uniquely determined by just one of the tuples, called the *highest weight*.

We proceed to define the above-mentioned weights, along with the related notion of the *roots* of a Lie group in an abstract and general fashion, which produces a configuration of vectors in \mathfrak{t}^* , the space of linear functionals on the Lie algebra \mathfrak{t} of a fixed maximal torus. This combinatorial structure enables us to enumerate all irreducible representations.

A.1 Weights and Roots of Lie groups

Let G be a Lie group, and let \mathfrak{g} be its Lie algebra. Usually G is a linear group, namely it is realized as a closed subgroup of $GL(V)$ for some finite-dimensional complex vector space V , in which case \mathfrak{g} is naturally realized as a subspace of the space of endomorphisms $\text{End}(V)$, defined as the space of all linear operators from V to V . In what follows, we use \bar{x} and \overline{A} , respectively, to denote the conjugate of a complex number, and the closure of a subset of a topological space, with the intended meaning clear by context. We now introduce the notion of a maximal torus.

Result 22. *Every compact Lie group G contains a torus T which is maximal with respect to inclusion. Any other torus is contained in a maximal one. Maximal tori are not unique, but are conjugate to one another.*

Proof. Bröcker and Tom Dieck (2013, Chapter IV). □

We can use this to study representations of G . Let \mathfrak{t} be the Lie algebra of T . Given a representation $\pi : G \rightarrow GL(V)$ of G on a vector space V , consider its restriction $\tilde{\pi} = \pi|_T$ to T . Since T is a torus, this restriction decomposes into a direct sum of one-dimensional

11. A Lie algebra \mathfrak{g} can also be viewed as the set of left-invariant vector fields on G : using this perspective, we define the Lie bracket $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$ as the commutator of vector fields $[x, y] = xy - yx$. More generally, we define an abstract Lie algebra to be a finite-dimensional vector space equipped with an antisymmetric bilinear binary operation satisfying the Jacobi identity $[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$

(complex-valued) irreducible unitary representations of T , given as

$$\tilde{\pi}(x) = \text{diag}(1, \dots, 1, \exp(\theta_1(x)), \dots, \exp(\theta_k(x))) \quad (73)$$

where θ_i are homomorphisms $T \rightarrow \mathbb{T} \subset \mathbb{C}$. We define the *weights* of a representation in terms of these homomorphisms as follows.

Definition 23. Define the **WEIGHTS** of a representation $\pi : G \rightarrow \text{GL}(V)$ to be the real differentials $d_e\theta_i : \mathfrak{t} \rightarrow \mathbb{R}$, where θ_i are defined above. By linearity, we have $d_e\theta_i \in \mathfrak{t}^*$ where \mathfrak{t}^* is the space of linear functionals over the space \mathfrak{t} .

For each $x \in G$, define the *adjoint action* $I_x : G \rightarrow G$ by $I_x(g) = x \cdot g \cdot x^{-1}$, and the *adjoint representation* $\text{Ad} : G \rightarrow \text{GL}(\mathfrak{g})$ by $\text{Ad}(x) = d_e I_x$. If G is a linear group, then the adjoint action is realized by conjugation of matrices in \mathfrak{g} . This allows us to introduce the *roots* of G .

Definition 24. Define the **ROOTS** of a Lie group G to be the non-zero weights of the adjoint representation $\text{Ad} : G \rightarrow \text{GL}(\mathfrak{g})$.

The adjoint representation gives rise to a particular bilinear form on \mathfrak{g} that we will need further, called the *Killing form*. We define it now. Let $\text{ad} : \mathfrak{g} \rightarrow \text{End}(\mathfrak{g})$ be defined by $\text{ad}(X)(Y) = [X, Y]$.

Definition 25. Define the symmetric bilinear **KILLING FORM** on the Lie algebra \mathfrak{g} by $B(X, Y) = \text{tr}(\text{ad}(X) \circ \text{ad}(Y))$ for $X, Y \in \mathfrak{g}$, where \circ denotes composition. We say G is **SEMI-SIMPLE** if B is non-degenerate in the sense that $B(X, Y) = 0$ for all Y implies $X = 0$.

It is not hard to see that the Killing form is invariant under the adjoint action, that is,

$$B(\text{Ad}(g)X, \text{Ad}(g)Y) = B(X, Y). \quad (74)$$

Different choices of maximal tori $T^{(1)}$ and $T^{(2)}$ produce different sets of weights in the respective tangents spaces $\mathfrak{t}^{(1)}$ and $\mathfrak{t}^{(2)}$, but since all maximal tori are conjugate, the corresponding differential $\mathfrak{t}^{(1)} \rightarrow \mathfrak{t}^{(2)}$ is an isometry with respect to the Killing form, sending one set of the weights into another.

One can show that a connected semi-simple Lie group is compact if and only if B is negative definite. Reinterpreted as a tensor field, $-B$ therefore canonically defines a bi-invariant Riemannian structure. Hereinafter by default we assume that the groups we are working with are semi-simple.

A.2 The Highest Weight of a Representation

To enumerate representations we will use certain weights called the *highest weights*. For a given representation its highest weight is the maximal weight with respect to a certain partial order. We proceed to define and formalize these concepts.

Definition 26. For each root α , let $L_\alpha \subset \mathfrak{t}$ be the hyperplane

$$L_\alpha = \{X \in \mathfrak{t} : \alpha(X) = 0\} \quad (75)$$

which, by virtue of α being linear, passes through the origin. Then $\mathfrak{t} \setminus \cup_\alpha L_\alpha$ is by construction a union of disjoint convex polyhedral cones, each of which we call a **WEYL CHAMBER**.

Define $\sigma : \mathfrak{t} \rightarrow \mathfrak{t}^*$ according to $-B(X, Y) = \sigma(X)(Y)$, where B is the Killing form. By non-degeneracy of B this establishes an isomorphism between \mathfrak{t} and \mathfrak{t}^* . The subsets of \mathfrak{t}^* that are images of Weyl chambers under σ will also be called Weyl chambers.

To define the ordering, we fix one arbitrary Weyl chamber $\mathfrak{t}_+ \subseteq \mathfrak{t}$ and call it *the positive chamber*. We also denote $\mathfrak{t}_+^* = \sigma(\mathfrak{t}_+) \subseteq \mathfrak{t}^*$. This defines the order on \mathfrak{t}^* as follows.

Definition 27. *Let $\theta, \theta' \in \mathfrak{t}^*$. We say that $\theta \geq \theta'$ if $\theta - \theta' \in \overline{\mathfrak{t}_+^*}$, which defines a partial order. We say that $\theta \in \mathfrak{t}^*$ is POSITIVE if $\theta(X) > 0$ for any $X \in \mathfrak{t}_+$.*

For purposes of calculating Laplace–Beltrami eigenvalues on compact Lie groups, we now introduce a particular weight ρ .

Definition 28. *Define the HALF-SUM OF THE ROOTS to be $\rho = \frac{1}{2} \sum r$, where the sum is taken over all positive roots.*

Before proceeding to the main result on highest weights, we describe a certain algebraic structure generated by weights. Let $\exp : \mathfrak{g} \rightarrow G$ be the Lie-theoretic exponential map.

Definition 29. *Define the WEIGHT LATTICE¹² to be*

$$P = \{c_1\alpha_1 + \dots + c_k\alpha_k : k \in \mathbb{N}, c_j \in \mathbb{Z}, \text{ and } \alpha_j \text{ are weights}\}, \quad (76)$$

which is spanned by the weights of all representations.

With these notions in place, we can state the key structural result which we will use to enumerate the irreducible representations of G .

Result 30. *For a given choice of a positive chamber \mathfrak{t}_+ , every representation π of G has a unique weight w which is largest with respect to the ordering on \mathfrak{t}^* induced by \mathfrak{t}_+ . This weight w is always an element of $\overline{\mathfrak{t}_+^*}$, and we call it the HIGHEST WEIGHT. Moreover, every element of $P \cap \overline{\mathfrak{t}_+^*}$ is the highest weight of a unique irreducible representation.*

Proof. See Bröcker and Tom Dieck (2013, Definition 2.7, Chapter VI) or Fegan (1991, Theorems 8.19 and 8.21). \square

This gives us a way to explicitly enumerate the irreducible representations of a given Lie group G : (1) choose a maximal torus $T \subseteq G$, (2) choose a positive chamber $\mathfrak{t}_+ \subseteq \mathfrak{t}$, where \mathfrak{t} is the Lie algebra of T , and (3) calculate the intersection of the weight lattice P with the closure of the dual positive chamber $\overline{\mathfrak{t}_+^*}$. Every weight in this intersection maps uniquely onto an irreducible representation.

This establishes what actually needs to be done to calculate the index set Λ of all irreducible representations. However, there is a more convenient way to carry this process out numerically. If we identify our maximal torus T with $\mathbb{T}^m \subseteq \mathbb{C}^m$, it turns out we can express the action of the highest weight of a representation simply by raising an element of \mathbb{T}^m to the power of some integer multi-index. This idea leads to the notion of a *signature*, formulated as follows.

¹². We define a *lattice* to be a discrete additive subgroup of \mathbb{R}^m .

Result 31. *Let w be the highest weight of a representation π . Suppose the chosen maximal torus $T \subseteq G$ is isomorphic to $\mathbb{T}^m \subseteq \mathbb{C}^m$ via the isomorphism of form*

$$t \in T \rightarrow (\varepsilon_1(t), \dots, \varepsilon_m(t)) \in \mathbb{T}^m. \quad (77)$$

Then for every $t \in T$, and every $X \in \mathfrak{t}$ satisfying $t = \exp(X)$, w can be expressed as

$$w(X) = \varepsilon_1(t)^{p_1} \cdot \dots \cdot \varepsilon_m(t)^{p_m}. \quad (78)$$

We call $(p_1, \dots, p_m) \in \mathbb{Z}^m$ the SIGNATURE of the representation, which is unique up to a choice of a maximal torus T and its isomorphism with the standard torus \mathbb{T}^m .

Proof. See Naimark and Stern (1982, Chapter VII, Section 2). □

Thus, we can use signatures to identify highest weights with tuples of integers. Not all tuples will correspond to highest weights: various restrictions on the integers forming a signature vary according to G . These are best handled on a case by case basis: see Appendix B for examples.

Appendix B. Calculating Signatures and Highest Weights for $\mathrm{SL}(n)$, $\mathrm{SU}(n)$, and $\mathrm{SO}(n)$

To illustrate the preceding ideas via example, we now calculate the above-introduced notions for three widely-occurring Lie groups. We defer the remaining task of determining which irreducible representations corresponds to which highest weights to subsequent sections.

Example 32. *If $G = \mathrm{SL}(n)$, then*

$$\mathfrak{g} = \{X \in M(n, \mathbb{R}) : \mathrm{tr}(X) = 0\}. \quad (79)$$

The most natural choice of a maximal torus T is the subgroup of all diagonal matrices, then $\mathfrak{t} \subset \mathfrak{g}$ is naturally realized as the subspace of diagonal matrices. The Killing form on \mathfrak{g} is realized as $B(X, Y) = 2n \mathrm{tr}(XY)$, and its restriction to \mathfrak{t} is then

$$B(\mathrm{diag}(x_1, \dots, x_n), \mathrm{diag}(y_1, \dots, y_n)) = 2n(x_1y_1 + \dots + x_ny_n). \quad (80)$$

The space \mathfrak{t}^ , equipped with the symmetric bilinear form, is naturally isomorphic to the $n - 1$ -dimensional Euclidean space. Denote $e_i = (0, \dots, 1, \dots, 0)$ where 1 is in the i th component. A standard realization of roots and weights is as follows. Consider the subspace E of \mathbb{R}^d orthogonal to $e_1 + \dots + e_n$. The roots are precisely the vectors $\alpha \in E$ having integer coordinates and $|\alpha|^2 = 2$. In other words, they are of the form $\alpha = e_i - e_j$ for $i \neq j$. The root lattice can be described as $Q = \mathbb{Z}^n \cap E$.*

The half-sum of the positive roots then equals

$$\rho = \left(\frac{n-1}{2}, \frac{n-3}{2}, \dots, \frac{3-n}{2}, \frac{1-n}{2} \right). \quad (81)$$

The weight lattice P can be described as the dual of Q . Namely, one can find such $\varpi_i \in E$, $i = 1, \dots, n - 1$, that $(\varpi_i, e_j - e_{j+1}) = \delta_{ij}$, then $Q = \sum_i \mathbb{Z}\varpi_i$.

A natural choice for a positive chamber $\overline{\mathfrak{t}}_+^*$ is then the positive cone spanned by ϖ_i , namely the set of their linear combinations with non-negative coefficients. It can also be very conveniently described as

$$\overline{\mathfrak{t}}_+^* = \{(x_1, \dots, x_n) \in E \subset \mathbb{R}^n : x_1 \geq x_2 \geq \dots \geq x_n\}. \quad (82)$$

The action of the Weyl group is by all possible permutations of coordinates, thus $W \cong S_n$, the symmetric group on n points.

Example 33. If $G = \mathrm{SU}(n)$, then its tangent space is

$$\mathfrak{g} = \{X \in M(n, \mathbb{C}) : X^* = -X, \mathrm{tr}(X) = 0\}. \quad (83)$$

The most natural choice of a maximal torus is the subgroup of diagonal matrices, that is,

$$T = \left\{ \mathrm{diag}(e^{i\theta_1}, \dots, e^{i\theta_n}) : \theta_1 + \dots + \theta_n = 0 \right\}. \quad (84)$$

Its tangent space is then

$$\mathfrak{t} = \{ \mathrm{diag}(i\theta_1, \dots, i\theta_n) : \theta_1 + \dots + \theta_n = 0 \}. \quad (85)$$

The Killing form on \mathfrak{t} can be calculated as

$$B(i \mathrm{diag}(\theta_1, \dots, \theta_n), i \mathrm{diag}(\eta_1, \dots, \eta_n)) = -2n(\theta_1\eta_1 + \dots + \theta_n\eta_n). \quad (86)$$

Under the identification $\mathrm{diag}(i\theta_1, \dots, i\theta_n) \mapsto (\theta_1, \dots, \theta_n)$ of \mathfrak{t} with \mathbb{R}^n the description of roots and weights coincides with that of $\mathrm{SL}(n)$.

Example 34. Let $G = \mathrm{SO}(n)$, its tangent is $\mathfrak{g} = \{X \in M(n, \mathbb{R}) : A^\top = -A\}$.

The description of roots and weights for $G = \mathrm{SO}(n)$ depends on the oddity of n . For $n = 2k$ or $2k + 1$ the maximal tori T can be chosen as

$$T = \left\{ \left(\begin{array}{ccc} \cos \theta_1 & \sin \theta_1 & \\ -\sin \theta_1 & \cos \theta_1 & \\ & & \ddots \\ & & \cos \theta_k & \sin \theta_k \\ & & -\sin \theta_k & \cos \theta_k \end{array} \right) \right\} \quad (87)$$

$$T = \left\{ \left(\begin{array}{ccc} \cos \theta_1 & \sin \theta_1 & \\ -\sin \theta_1 & \cos \theta_1 & \\ & & \ddots \\ & & \cos \theta_k & \sin \theta_k \\ & & -\sin \theta_k & \cos \theta_k \\ & & & & 1 \end{array} \right) \right\}. \quad (88)$$

Their tangent spaces are then

$$\mathfrak{t} = \left\{ \left(\begin{array}{ccc} 0 & \theta_1 & \\ -\theta_1 & 0 & \\ & & \ddots \\ & & 0 & \theta_k \\ & & -\theta_k & 0 \end{array} \right) \right\} \quad \mathfrak{t} = \left\{ \left(\begin{array}{ccc} 0 & \theta_1 & \\ -\theta_1 & 0 & \\ & & \ddots \\ & & 0 & \theta_k \\ & & -\theta_k & 0 \\ & & & & 0 \end{array} \right) \right\} \quad (89)$$

and the formulas for the Killing form are

$$B\left(\text{bdiag}\left(\begin{pmatrix} 0 & \theta_1 \\ -\theta_1 & 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 & \theta_k \\ -\theta_k & 0 \end{pmatrix}\right), \text{bdiag}\left(\begin{pmatrix} 0 & \eta_1 \\ -\eta_1 & 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 & \eta_k \\ -\eta_k & 0 \end{pmatrix}\right)\right) = (4 - 4k)(\theta_1\eta_1 + \dots + \theta_k\eta_k) \quad (90)$$

$$B\left(\text{bdiag}\left(\begin{pmatrix} 0 & \theta_1 \\ -\theta_1 & 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 & \theta_k \\ -\theta_k & 0 \end{pmatrix}\right), \text{bdiag}\left(\begin{pmatrix} 0 & \eta_1 \\ -\eta_1 & 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 & \eta_k \\ -\eta_k & 0 \end{pmatrix}\right), 0\right) = (2 - 4k)(\theta_1\eta_1 + \dots + \theta_k\eta_k). \quad (91)$$

In both cases, \mathfrak{t}^* is naturally identified with \mathbb{R}^k . The roots are, then,

$$e_i \pm e_j, \quad i \neq j \quad \text{for } n = 2k, \quad (92)$$

$$\pm e_i \quad \text{and } e_i \pm e_j, \quad i \neq j \quad \text{for } n = 2k + 1. \quad (93)$$

The half-sum of the positive roots equals

$$\rho = (n - 1, n - 2, \dots, 2, 1, 0) \quad \text{for } n = 2k, \quad (94)$$

$$\rho = \left(\frac{n}{2} - 1, \frac{n}{2} - 2, \dots, \frac{3}{2}, \frac{1}{2}\right) \quad \text{for } n = 2k + 1. \quad (95)$$

The weight lattice Q in both cases can be described as \mathbb{Z}^k . The action of the Weyl group is by permutations of coordinates and changing of signs.¹³

The standard choice of the positive chamber is

$$\overline{\mathfrak{t}_+^*} = \{(x_1, \dots, x_k) \in Q : x_i \geq x_{i+1}, x_{k-1} \geq |x_k|\} \quad \text{for } n = 2k, \quad (96)$$

$$\overline{\mathfrak{t}_+^*} = \{(x_1, \dots, x_k) \in Q : x_i \geq x_{i+1} \geq 0\} \quad \text{for } n = 2k + 1. \quad (97)$$

Example 35. For $G = \text{SO}(n)$ and the choices of T and $\overline{\mathfrak{t}_+^*}$ as above, the possible signatures (p_1, \dots, p_k) are

$$p_1 \geq p_2 \geq \dots \geq p_{k-1} \geq |p_k| \quad \text{for } n = 2k, \quad (98)$$

$$p_1 \geq p_2 \geq \dots \geq p_k \geq 0 \quad \text{for } n = 2k + 1. \quad (99)$$

Example 36. For $G = \text{SU}(n)$ the possible signatures (p_1, \dots, p_n) are those with

$$p_1 \geq p_2 \geq \dots \geq p_n. \quad (100)$$

Appendix C. The Weyl Character Formula

To describe the Weyl character formula, we will need the principal group of symmetries of a root system, its Weyl group. Again, denote by T a fixed maximal torus and by \mathfrak{t} its Lie algebra.

Definition 37. Define the WEYL GROUP W as the subgroup of $\text{O}(\mathfrak{t})$ generated by all reflections in L_α of Definition 26. W acts transitively on the set of all Weyl chambers by permutations.

13. In group-theoretic terms, $W \cong S_k \times (\mathbb{Z}/2\mathbb{Z})^k$, the semidirect product of S_k acting on the k -th power of the cyclic group of order 2 by permutations of components.

To proceed, we introduce two ways of expressing certain alternating sums.

Result 38 (Weyl Denominator Formula). *Let $X \in \mathfrak{t}$, and define*

$$j(\exp(X)) = \sum_{\sigma \in W} \operatorname{sgn}(\sigma) \exp(\sigma(\rho)(X)) \quad (101)$$

where W is the Weyl group, ρ is half the sum of positive roots of Definition 28, and $\operatorname{sgn}(\sigma)$ is the determinant of σ as a linear operator.¹⁴ Then

$$j(\exp(X)) = \prod_{\alpha > 0} (\exp(\frac{i}{2}\alpha(X)) - \exp(\frac{-i}{2}\alpha(X))) \quad (102)$$

where the product is taken over all positive roots.

Proof. Fegan (1991), Theorem 9.2. □

We now introduce the Weyl character formula, which expresses characters as the ratio of two alternating sums depending on the chosen tangent vector.

Result 39 (Weyl Character Formula). *For a highest weight w , define*

$$j_w(\exp(X)) = \sum_{\sigma \in W} \operatorname{sgn}(\sigma) \exp(\sigma(w + \rho)(X)). \quad (103)$$

Then for $t \in T$, the character χ_w can be expressed as

$$\chi_w(t) = \frac{j_w(t)}{j(t)}. \quad (104)$$

Proof. Fegan (1991), Theorem 9.9. □

Ostensibly, this gives us a way to calculate character values on the maximal torus T . However, since characters are conjugation-invariant, the Weyl character formula can be used to calculate character values everywhere on G . We state this formally as follows.

Corollary 40. *Let $g \in G$, and let $t \in T$ be such that there is an $h \in G$ such that $g = h \cdot t \cdot h^{-1}$, where t and h always exist by Result 22. Then*

$$\chi_w(g) = \chi_w(h^{-1} \cdot t \cdot h) = \chi_w(t). \quad (105)$$

Proof. This follows because the character is invariant under conjugation:

$$\chi_w(h^{-1} \cdot t \cdot h) = \operatorname{tr} \pi_w(h^{-1}) \pi_w(t) \pi_w(h) = \operatorname{tr} \pi_w(h^{-1}) \pi_w(h) \pi_w(t) = \chi_w(t). \quad (106)$$

□

14. Note that $\operatorname{sgn}(\sigma)$ is equal to $(-1)^m$ when σ is a composition of m reflections.

This establishes how to calculate the character χ_w for a given highest weight w . We now study how to express this using signatures. If $t = (\varepsilon_1, \dots, \varepsilon_n)$ under the correspondence between the maximal torus and the complex exponentials, we can re-express the above formula as the ratio of two polynomials in $\varepsilon_1, \dots, \varepsilon_n$, where the denominator divides the numerator, and the ratio itself is also a polynomial.

In many cases, this expression can be simplified greatly. As an example, for $G = \text{SU}(2)$, the Weyl character formula simply says that for a unique representation of dimension d —with signature (d) —the character value equals

$$\chi_d(g) = \frac{\lambda^d - \lambda^{-d}}{\lambda - \lambda^{-1}} \quad (107)$$

where λ, λ^{-1} are the eigenvalues of g . Since both eigenvalues can be expressed as

$$\frac{1}{2} \left(\text{tr}(g) \pm \sqrt{\text{tr}(g)^2 - 4} \right) \quad (108)$$

one can show that

$$\chi_d(g) = U_{d-1} \left(\frac{1}{2} \text{tr}(g) \right) \quad (109)$$

where U_k is the k th Chebyshev polynomial of the second kind.

In general, this expression can either be calculated symbolically from the Weyl character formula, or obtained from the Kostant multiplicity formula, which expresses the coefficients as certain combinatorial quantities related to the structure of the root system of G . Note in particular that from this expression, one can deduce the Weyl dimension formula

$$\dim V_w = \chi_w(e) = \frac{\prod_{\alpha > 0} (w + \rho, \alpha)}{\prod_{\alpha > 0} (\rho, \alpha)} \quad (110)$$

which is used for calculating normalizing constants d_λ . In most cases, the alternating sums appearing in the above cases can be interpreted with the help of Weyl denominator formula as certain matrix determinants.

Appendix D. Calculating the Characters for $\text{SU}(n)$, and $\text{SO}(n)$

In more practical terms, for $\text{SO}(n)$ and $\text{SU}(n)$ Weyl character formula provides an explicit expression for $\chi(g)$ as a ratio of two Vandermonde-type determinants depending only on the eigenvalues of g , possibly together with an ordering.

Example 41. Let $G = \text{SO}(n)$ and let (p_1, \dots, p_n) be the signature of a representation with the highest weight ϖ . Then the corresponding character χ_ϖ can be calculated as follows. If $n = 2k + 1$, denote $q_i = p_i + k - i + \frac{1}{2}$ and set $(\gamma_1, \dots, \gamma_k, 1, \gamma_k^{-1}, \dots, \gamma_1^{-1})$ to be the eigenvalues of an element $g \in G$. Then

$$\chi_\varpi(g) = \frac{\xi_1(q_1, \dots, q_k)}{\xi_1(k - \frac{1}{2}, \dots, \frac{3}{2}, \frac{1}{2})}, \quad \text{where} \quad \xi_1(d_1, \dots, d_k) = \det \left(\gamma_i^{d_j} - \gamma_i^{-d_j} \right)_{i,j=1}^k. \quad (111)$$

Weyl's dimension formula specializes to

$$d_\varpi = \frac{2^k}{(2k-1)! \dots 3! 1!} \cdot q_1 \dots q_k \cdot \prod_{1 \leq i < j \leq k} (q_i^2 - q_j^2). \quad (112)$$

If $n = 2k$, denote $q_i = p_i + k - i$ for $i = 1, \dots, k - 1$ and $q_k = |p_k|$. Let the quantity $(\gamma_1, \dots, \gamma_k, \gamma_k^{-1}, \dots, \gamma_1^{-1})$ be the eigenvalues of $g \in G$ ordered such that

$$\begin{pmatrix} \cos \theta_1 & \sin \theta_1 & & & & & & & & & \\ -\sin \theta_1 & \cos \theta_1 & & & & & & & & & \\ & & \ddots & & & & & & & & \\ & & & \cos \theta_k & \sin \theta_k & & & & & & \\ & & & -\sin \theta_k & \cos \theta_k & & & & & & \\ & & & & & & & & & & 1 \end{pmatrix}, \quad \text{where } \gamma_j = \cos \theta_j + i \sin \theta_j \quad (113)$$

is conjugated in $SO(n)$ to g (contrary to the case $n = 2k + 1$, here each possible collection of eigenvalues determine two conjugacy classes). Introduce

$$\xi_0(d_1, \dots, d_k) = \det \left(\gamma_i^{d_j} + \gamma_i^{-d_j} \right)_{i,j=1}^k, \quad \xi_1(d_1, \dots, d_k) = \det \left(\gamma_i^{d_j} - \gamma_i^{-d_j} \right)_{i,j=1}^k. \quad (114)$$

Then

$$\chi_\varpi(g) = \frac{\xi_0(q_1, \dots, q_k)}{\xi_0(m-1, m-2, \dots, 1, 0)}, \quad \text{if } p_k = q_k = 0, \quad (115)$$

$$\chi_\varpi(g) = \frac{\xi_0(q_1, \dots, q_k) + (\text{sgn } p_m) \xi_1(q_1, \dots, q_k)}{2\xi_0(m-1, m-2, \dots, 1, 0)}, \quad \text{if } p_m \neq 0. \quad (116)$$

The dimension formula turns into

$$d_\varpi = \frac{2^{k-1}}{(2k-2)!(2k-4)! \dots 4! 2!} \cdot \prod_{1 \leq i < j \leq k} (q_i^2 - q_j^2). \quad (117)$$

Example 42. For $G = \text{SU}(n)$ the character of a representation with the signature (p_1, \dots, p_n) can be calculated as follows. Denote $q_i = p_i + n - i$. If $(\gamma_1, \dots, \gamma_n)$ are the eigenvalues of $g \in G$, then

$$\chi(g) = \frac{\xi_1(q_1, \dots, q_k)}{\xi_1(0, 1, \dots, n-1)}, \quad \text{where } \xi_1(d_1, \dots, d_n) = \det \left(\gamma_i^{d_j} \right)_{i,j=1}^n. \quad (118)$$

Appendix E. Proofs

Theorem 1. A Gaussian process $f \sim \text{GP}(0, k)$ on a compact Lie group G is stationary with respect to left-and-right action of G on itself if and only if k is of form

$$k(g_1, g_2) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \text{Re } \chi^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (18)$$

where $a^{(\lambda)} \geq 0$ satisfy $\sum_{\lambda \in \Lambda} d_\lambda a^{(\lambda)} < \infty$. Moreover, for all λ , $\text{Re } \chi^{(\lambda)}$ is positive-definite.

Proof. Theorem 2 of Yaglom (1961) shows that for any stationary Gaussian process $k(g_1, g_2) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \chi^{(\lambda)}(g_2^{-1} \cdot g_1)$ where $a^{(\lambda)} \geq 0$ satisfy $\sum_{\lambda \in \Lambda} d_\lambda a^{(\lambda)} < \infty$. Under our usual assumption that a Gaussian process is real-valued, we have $\text{Re } k(g_1, g_2) = k(g_1, g_2)$. Using this and taking the real parts of both sides of the equation, we get (18).

To prove positive semi-definiteness, write, using (i) the definition of $\chi^{(\lambda)}$, (ii) the homomorphism property of $\pi^{(\lambda)}$, (iii) the fact that $\pi^{(\lambda)}(g)^{-1} = \pi^{(\lambda)}(g)^*$ due to unitarity, where the star denotes the operation of taking the conjugate transpose matrix, and (iv) the standard properties of traces, that

$$\chi^{(\lambda)}(g_2^{-1}g_1) = \text{tr } \pi^{(\lambda)}(g_2^{-1}g_1) = \text{tr } \pi^{(\lambda)}(g_1)\pi^{(\lambda)}(g_2)^* = \sum_{j=1}^{d_\lambda} \sum_{k=1}^{d_\lambda} \pi_{jk}^{(\lambda)}(g_1) \overline{\pi_{jk}^{(\lambda)}(g_2)}. \quad (119)$$

This immediately shows that a character $\chi^{(\lambda)}$, being an inner product of the feature maps $g \rightarrow \{\pi_{jk}^{(\lambda)}(g)\}_{1 \leq j, k \leq d_\lambda}$, is (complex-valued) positive semi-definite. The real part of a complex-valued positive semi-definite function is positive semi-definite, hence $\text{Re } \chi^{(\lambda)}$ is positive semi-definite. \square

Theorem 2. *A Gaussian process $f \sim \text{GP}(0, k)$ on a compact homogeneous space G/H is stationary with respect to the action of G if and only if k is of form*

$$k(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} a_{jk}^{(\lambda)} \text{Re } \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (23)$$

where the coefficients $a_{jk}^{(\lambda)} \in \mathbb{R}$ form symmetric positive semi-definite matrices $\mathbf{A}^{(\lambda)}$ of size $r_\lambda \times r_\lambda$ satisfying $\sum_{\lambda} \text{tr } \mathbf{A}^{(\lambda)} < \infty$, and $\pi_{jk}^{(\lambda)}$ are the zonal spherical functions. Moreover, for each individual $\lambda \in \Lambda$, the corresponding sum over j, k is positive semi-definite.

Proof. Theorem 5 of Yaglom (1961) immediately implies that any stationary real-valued Gaussian process is of form $k(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1)$ where the coefficients $a_{jk}^{(\lambda)} \in \mathbb{C}$ form Hermitian positive semi-definite matrices $\mathbf{A}^{(\lambda)}$ of size $r_\lambda \times r_\lambda$ with complex entries satisfying $\sum_{\lambda} \text{tr } \mathbf{A}^{(\lambda)} < \infty$, and $\pi_{jk}^{(\lambda)}$ are the zonal spherical functions. Furthermore, any covariance of this form, provided it is real-valued, it distribution-wise uniquely defines a stationary Gaussian process. To obtain the claim, we need to understand when k defined in this way is a real-valued function.

The map $g \rightarrow (\pi^{(\lambda)}(g))^*$ defines an irreducible unitary representation, thus is equal to $\pi^{(\lambda')}$ for some $\lambda' \in \Lambda$. No matter if $\lambda' = \lambda$ or not, it is easy to see that $r_{\lambda'} = r_\lambda$ and, what is more, we can define the zonal spherical functions for $\pi^{(\lambda')}$ using the same basis $e_1, \dots, e_{d_\lambda} \in V_{\lambda'} = V_\lambda$, getting

$$\pi_{jk}^{(\lambda')}(g) = \langle \pi^{(\lambda')}(g)e_j, e_k \rangle_{V_\lambda} = \langle (\pi^{(\lambda)}(g))^* e_j, e_k \rangle_{V_\lambda} = \langle e_j, \pi^{(\lambda)}(g)e_k \rangle_{V_\lambda} = \overline{\pi_{kj}^{(\lambda)}(g)}. \quad (120)$$

Using this property, from $\text{Re } k = k$, or equivalently from $k = \bar{k}$ we infer

$$\sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g) = \sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} \overline{a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g)} = \sum_{\lambda \in \Lambda} \sum_{j, k=1}^{r_\lambda} \overline{a_{kj}^{(\lambda')}} \pi_{jk}^{(\lambda)}(g) \quad (121)$$

Because all $\pi_{jk}^{(\lambda)}$ are orthonormal, it follows that $a_{jk}^{(\lambda)} = \overline{a_{kj}^{(\lambda')}}$.

Now, since $k(g_1 \cdot H, g_2 \cdot H) = k(g_2 \cdot H, g_1 \cdot H)$, write

$$\sum_{\lambda \in \Lambda} \sum_{jk=1}^{r_\lambda} a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g) = \sum_{\lambda \in \Lambda} \sum_{jk=1}^{r_\lambda} a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g^{-1}) = \sum_{\lambda \in \Lambda} \sum_{jk=1}^{r_\lambda} a_{jk}^{(\lambda)} \overline{\pi_{kj}^{(\lambda)}(g)} \quad (122)$$

$$= \sum_{\lambda \in \Lambda} \sum_{jk=1}^{r_\lambda} a_{kj}^{(\lambda')} \pi_{jk}^{(\lambda)}(g) \quad (123)$$

This implies $a_{jk}^{(\lambda)} = a_{kj}^{(\lambda')}$ and therefore $a_{jk}^{(\lambda)} \in \mathbb{R}$. Using these observations, write

$$\operatorname{Re} \sum_{\lambda \in \Lambda} \sum_{j,k=1}^{r_\lambda} a_{jk}^{(\lambda)} \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) = \sum_{\lambda \in \Lambda} \sum_{j,k=1}^{r_\lambda} a_{jk}^{(\lambda)} \operatorname{Re} \pi_{jk}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (124)$$

which gives the first part of the claim.

We now show that for each individual $\lambda \in \Lambda$ the corresponding sum over j, k is positive semi-definite. Yaglom (1961, Theorem 5) proves, as a byproduct, that this sum is a covariance function of a Gaussian process given by

$$f_\lambda(gH) = \sum_{j=1}^{d_\lambda} \sum_{k=1}^{r_\lambda} z_{jk}^\lambda \pi_{jk}^{(\lambda)}(g) \quad (125)$$

where z_{jk}^λ are zero-mean Gaussians with covariances given by $\mathbb{E} z_{j_1 k_1}^\lambda z_{j_2 k_2}^\lambda = \delta_{j_1 j_2} a_{k_1 k_2}^{(\lambda)}$. The claim follows. \square

Theorem 6. *Let $f \sim \operatorname{GP}(0, k)$, where k is stationary and satisfies $\mathbf{A}^{(\lambda)} = a^{(\lambda)} \mathbf{I}$. Then*

$$k(g_1 \cdot H, g_2 \cdot H) = \int_H k_G(g_1 \cdot h, g_2) d\mu_H(h) \quad f(g \cdot H) = \int_H f_G(g \cdot h) d\mu_H(h) \quad (30)$$

where μ_H is the Haar measure on H , and $f_G \sim \operatorname{GP}(0, k_G)$ with

$$k_G(g_1, g_2) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \operatorname{Re} \chi^{(\lambda)}(g_2^{-1} \cdot g_1). \quad (31)$$

Proof. Recall that under a suitable choice of basis in the representation space V_λ , spherical functions $\pi_{jk}^{(\lambda)}$ are matrix coefficients with $1 \leq j \leq d_\lambda$ and $1 \leq k \leq r_\lambda$.

Consider the integrals of form

$$I_{jk}^{(\lambda)}(g) = \int_H \pi_{jk}^{(\lambda)}(g \cdot h) d\mu_H(h). \quad (126)$$

Since $\pi_{jk}^{(\lambda)}$ corresponding to $1 \leq j \leq d_\lambda$ and $1 \leq k \leq r_\lambda$ are spherical functions and these are constant on cosets, we have $\pi_{jk}^{(\lambda)}(g \cdot h) = \pi_{jk}^{(\lambda)}(g)$ and thus $I_{jk}^{(\lambda)}(g) = \pi_{jk}^{(\lambda)}(g)$. Furthermore, since spherical functions span the subspace of $\operatorname{span} \pi_{jk}^{(\lambda)}$ that consists of all functions constant

on cosets, and since the remaining matrix coefficients belong to its orthogonal complement, we have $I_{jk}^{(\lambda)}(g) = 0$ for all remaining indices j, k , proving that

$$\int_H \chi^{(\lambda)}(g \cdot h) d\mu_H(h) = \sum_{j=1}^{d_\lambda} \int_H \pi_{jj}^{(\lambda)}(g \cdot h) d\mu_H(h) = \sum_{j=1}^{r_\lambda} \pi_{jj}^{(\lambda)}(g \cdot h). \quad (127)$$

The same holds for real parts.

With this, we can write

$$k(g_1 \cdot H, g_2 \cdot H) = \sum_{\lambda \in \Lambda} a^{(\lambda)} \sum_{j=1}^{r_\lambda} \operatorname{Re} \pi_{jj}^{(\lambda)}(g_2^{-1} \cdot g_1) \quad (128)$$

$$= \sum_{\lambda \in \Lambda} a^{(\lambda)} \int_H \operatorname{Re} \chi^{(\lambda)}(g_2^{-1} \cdot g_1 \cdot h) d\mu_H(h) \quad (129)$$

$$= \int_H \sum_{\lambda \in \Lambda} a^{(\lambda)} \operatorname{Re} \chi^{(\lambda)}(g_2^{-1} \cdot g_1 \cdot h) d\mu_H(h) \quad (130)$$

$$= \int_H k_G(g_1 \cdot h, g_2) d\mu_H(h). \quad (131)$$

Define $f'(gH) = \int_H f_G(g \cdot h) d\mu_H(h)$ where $f_G \sim \text{GP}(0, k_G)$. Now, $\mathbb{E} f'(gH) = 0$, while writing

$$\operatorname{Cov}(f'(g_1H), f'(g_2H)) = \int_H \int_H \operatorname{Cov}(f_G(g_1 \cdot h_1), f_G(g_2 \cdot h_2)) d\mu_H(h_1) d\mu_H(h_2) \quad (132)$$

$$= \int_H \int_H k_G(g_1 \cdot h_1, g_2 \cdot h_2) d\mu_H(h_1) d\mu_H(h_2) \quad (133)$$

$$= \int_H \int_H k_G(g_1 \cdot h_1 \cdot h_2^{-1}, g_2) d\mu_H(h_1) d\mu_H(h_2) \quad (134)$$

$$= \int_H k_G(g_1 \cdot h, g_2) d\mu_H(h) = k(g_1 \cdot H, g_2 \cdot H) \quad (135)$$

proves the claim. \square

Proposition 8. *The function K defined by (39) satisfies*

$$K(x, x') = \int_{G/H} K(x, u) \overline{K(x', u)} d\mu_{G/H}(u) \quad (40)$$

if and only if the matrix $\mathbf{A} = d_\lambda^{-1}(a_{k_1 k_2}^{(\lambda)})$ is idempotent, namely $\mathbf{A}^2 = \mathbf{A}$.

Proof. Expanding the right-hand side of (40) we get

$$\sum_{\substack{k_1, k_2=1 \\ k_3, k_4=1}}^{r_\lambda} a_{k_1, k_2}^{(\lambda)} a_{k_3, k_4}^{(\lambda)} \sum_{j_1, j_2=1}^{d_\lambda} \overline{\pi_{j_1, k_1}^{(\lambda)}(x) \pi_{j_2, k_3}^{(\lambda)}(x')} \int_{G/H} \overline{\pi_{j_1, k_2}^{(\lambda)}(u) \pi_{j_2, k_4}^{(\lambda)}(u)} d\mu_{G/H}(u) \quad (136)$$

$$= d_\lambda \sum_{k_1, k_3=1}^{r_\lambda} \underbrace{\left(\sum_{k=1}^{r_\lambda} \frac{a_{k_1, k}^{(\lambda)}}{d_\lambda} \frac{a_{k_3, k}^{(\lambda)}}{d_\lambda} \right)}_{\text{denote by } b_{k_1, k_3}} \sum_j^{d_\lambda} \overline{\pi_{j, k_1}^{(\lambda)}(x) \pi_{j, k_3}^{(\lambda)}(x')}. \quad (137)$$

Comparing the last expression to the right-hand side of (39) we see that the necessary and sufficient condition for them to match is that $d_\lambda b_{k_1, k_2} = a_{k_1, k_2}$ for all $1 \leq k_1, k_2 \leq r_\lambda$. If let \mathbf{B} be the matrix consisting of b_{k_1, k_2} , then by (137) we have $\mathbf{B} = \mathbf{A}\mathbf{A}^\top$, along with the preceding condition, which in matrix form is $\mathbf{B} = \mathbf{A}$. Hence, the necessary and sufficient condition for (40) to hold is $\mathbf{A} = \mathbf{A}\mathbf{A}^\top$ and since \mathbf{A} is symmetric by assumption, namely $\mathbf{A} = \mathbf{A}^\top$, the condition simplifies to $\mathbf{A}\mathbf{A}^\top = \mathbf{A}^2 = \mathbf{A}$. \square

Lemma 9. *There exist $\{x_1, \dots, x_N\} \subseteq \mathcal{S}$ such that $\det \mathbf{M}_k \neq 0$ for all $1 \leq k \leq N$.*

Proof. Induction. For $k = 1$, the claim is immediate since $\phi_1 \neq 0$. Now, assume $\det \mathbf{M}_k \neq 0$ for all k in the range $k = 1, \dots, n-1$. Split the matrix \mathbf{M}_n into blocks as follows:

$$\mathbf{M}_n = \begin{pmatrix} \mathbf{M}_{n-1} & \phi_*(x_n) \\ \phi_n(x_*) & \phi_n(x_n) \end{pmatrix} \quad (138)$$

and use Schur complement formula for its determinant to obtain

$$\det \mathbf{M}_n = \det \mathbf{M}_{n-1} \left(\phi_n(x_n) - \sum_{j=1}^{n-1} \alpha_j \phi_j(x_n) \right) \quad (139)$$

where the coefficients α_j are defined as

$$(\alpha_1, \dots, \alpha_{n-1}) = (\phi_n(x_1), \dots, \phi_n(x_{n-1})) \mathbf{M}_{n-1}^{-1} \quad (140)$$

and in particular do not depend on x_n . If the statement $\det \mathbf{M}_n = 0$ for all $x_n \in \mathcal{S}$ were true, it would imply $\phi_n(x_n) - \sum_{j=1}^{n-1} \alpha_j \phi_j(x_n) = 0$ for all $x_n \in \mathcal{S}$, which would contradict linear independence of ϕ_j . The claim follows. \square

Theorem 12. *If x_j is a fundamental set with respect to the system ϕ_j , then*

$$\text{span } \phi_j = \text{span } K(\cdot, x_j) \quad (50)$$

Proof. We have

$$K(x, x_k) = \sum_{j=1}^N \overline{\phi_j(x_k)} \phi_j(x). \quad (141)$$

Notice that

$$\begin{pmatrix} K(x, x_1) \\ \vdots \\ K(x, x_N) \end{pmatrix} = \mathbf{M}_N^* \begin{pmatrix} \phi_1(x) \\ \vdots \\ \phi_N(x) \end{pmatrix}. \quad (142)$$

Hence, by non-degeneracy of \mathbf{M}_N , it follows that $K(\cdot, x_j)$ is a basis of $\text{span } \phi_j$. \square

Result 13. *For every λ , the matrix coefficients $\pi_{jl}^{(\lambda)}$ are eigenfunctions of $-\Delta$ corresponding to the same eigenvalue $\alpha_\lambda \geq 0$. Moreover, the eigenvalues α_λ can be written¹⁵*

$$\alpha_\lambda = \|\rho\|_B^2 - \|w + \rho\|_B^2 \quad (56)$$

where w is the highest weight of the representation with signature λ , $\|g\|_B^2 = -B(g, g)$, B is the Killing form of G , and ρ is the half-sum of positive roots of G —see Appendix A.

Proof. Since G is equipped with the metric induced by the Killing form, the left shift and right shift maps

$$L_g : G \rightarrow G \quad L_g u = g \cdot u \quad R_g : G \rightarrow G \quad R_g u = u \cdot g^{-1} \quad (143)$$

are Riemannian isometries of G . Let Δ be the Laplace–Beltrami operator on G , and recall that this operator commutes with Riemannian isometries. Thus, Δ commutes with L_g and R_g for all $g \in G$. Let $\phi_{jk} = \Delta \pi_{jk}^{(\lambda)}$. Then $(\Delta L_g \pi_{jk}^{(\lambda)})(u) = (L_g \phi_{jk})(u) = \phi_{jk}(g \cdot u)$ and, on the other hand,

$$(\Delta L_g \pi_{jk}^{(\lambda)})(u) = \sum_{l=1}^{d_\lambda} \pi_{jl}^{(\lambda)}(g) (\Delta \pi_{lk}^{(\lambda)})(u) = \sum_{l=1}^{d_\lambda} \pi_{jl}^{(\lambda)}(g) \phi_{lk}(u). \quad (144)$$

Hence $\phi_{jk}(gu) = \sum_{l=1}^{d_\lambda} \pi_{jl}^{(\lambda)}(g) \phi_{lk}(u)$. Denoting $\phi(g) = (\phi_{jk}(g))$ we get $\phi(g \cdot u) = \pi^{(\lambda)}(g) \phi(u)$. Examining the action for the right shifts in the same manner, we get $\phi(u \cdot g^{-1}) = \phi(u) \pi^{(\lambda)}(g^{-1})$. Because of this, $\pi^{(\lambda)}(g) \phi(e) = \phi(g) = \phi(e) \pi^{(\lambda)}(g)$ for all $g \in G$, which, by Schur’s Lemma, implies $\phi(e) = \alpha_\lambda I$ and $\phi(g) = \alpha_\lambda \pi^{(\lambda)}(g)$. This proves that every $\pi_{jk}^{(\lambda)}$ is an eigenfunction of Δ corresponding to the same eigenvalue α_λ . By general theory of the Laplacian we have $\alpha_\lambda \leq 0$.

This shows that $\{\sqrt{d_\lambda} \pi_{jk}^{(\lambda)}\}_{\lambda, j, k}$ is an orthonormal basis of eigenfunctions of Δ in $L^2(G)$, where the Haar measure on G coincides with the Riemannian volume measure on G . \square

Result 15. *The spherical functions $\pi_{jk}^{(\lambda)}$ on a compact homogeneous space G/H are the eigenfunctions of $-\Delta$. Moreover, the eigenvalue corresponding to every spherical function $\pi_{jk}^{(\lambda)}$, which by definition is also a matrix coefficient of G , is given in (56).*

Proof. The metric structure of a homogeneous space $\mathcal{M} = G/H$ with G a compact Lie group is inherited from the group G and thus the Laplacian $\Delta_{G/H}$ on G/H satisfies

$$(\Delta_{G/H} f)(g \cdot H) = (\Delta_G \tilde{f})(g) \quad (145)$$

¹⁵. This expression is sometimes called *Freudenthal’s formula for the eigenvalues*.

where Δ_G is the Laplacian on G , $\tilde{f}(g) = f(\phi(g))$ with $\phi : G \rightarrow G/H$ and $\phi(g) = g \cdot H$.

Consider the spherical functions $\pi_{jk}^{(\lambda)}$ for $1 \leq j \leq d_\lambda$, $1 \leq k \leq r_\lambda$. The functions $\pi_{jk}^{(\lambda)}$, as functions on G , are matrix coefficients of irreducible unitary representations of G under a suitable choice of basis in the representation space V_λ . Because of this, the argument from the previous section implies that they are eigenfunctions of $\Delta_{G/H}$ corresponding to the eigenvalues α_λ . Thus $\{\sqrt{d_\lambda} \pi_{jk}^{(\lambda)}\}_{\lambda,j,k}$ is an orthonormal basis of eigenfunctions of $\Delta_{G/H}$ in the space $L^2(G/H)$. Note that, again, the invariant measure on G/H coincides with the Riemannian volume measure on G/H . \square

Result 43. *Euclidean Matérn kernels can be written in integral form as*

$$k_{\nu,\kappa,\sigma^2}(\mathbf{x}, \mathbf{x}') = \frac{(2\nu)^\nu}{\Gamma(\nu)\kappa^{2\nu}} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2}u} k_{\infty,\sqrt{2u},\sigma^2}(\mathbf{x}, \mathbf{x}') du \quad (146)$$

$$= \sigma^2 \frac{(2\nu)^\nu (4\pi)^{n/2}}{\Gamma(\nu)\kappa^{2\nu}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \mathcal{P}(u, \mathbf{x}, \mathbf{x}') du. \quad (147)$$

Proof. This is a proof from Jaquier et al. (2022), which we provide for completeness. By Gradshteyn and Ryzhik (2014, Section 3.326, Item 2) we have

$$\int_0^\infty u^m e^{-au} du = \Gamma(m+1) a^{-m-1}. \quad (148)$$

Substituting $m = \nu + n/2 - 1$ and $a = 2\nu/\kappa^2 + 4\pi^2\lambda^2$ into this equation and then performing a simple rearrangement of terms, we get

$$\left(\frac{2\nu}{\kappa^2} + 4\pi^2\lambda^2\right)^{-\nu-\frac{n}{2}} = \Gamma(\nu + n/2)^{-1} \int_0^\infty u^{\nu+\frac{n}{2}-1} e^{-\frac{2\nu}{\kappa^2}u} e^{-4\pi^2\lambda^2u} du, \quad (149)$$

$$= (4\pi)^{-\frac{n}{2}} \Gamma(\nu + n/2)^{-1} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2}u} (4\pi u)^{\frac{n}{2}} e^{-4\pi^2\lambda^2u} du, \quad (150)$$

$$= \sigma^{-2} (4\pi)^{-\frac{n}{2}} \Gamma(\nu + n/2)^{-1} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2}u} S_{\infty,\sqrt{2u},\sigma^2}(\lambda) du \quad (151)$$

where

$$S_{\infty,\kappa,\sigma^2}(\lambda) = \sigma^2 (2\pi\kappa^2)^{n/2} e^{-2\pi^2\kappa^2\lambda^2} \quad (152)$$

is the spectral density of the squared exponential kernel, namely

$$k_{\infty,\kappa,\sigma^2}(\mathbf{x}, \mathbf{x}') = \int_{\mathbb{R}^n} S_{\infty,\kappa,\sigma^2}(\|\boldsymbol{\xi}\|) e^{2\pi i \langle \mathbf{x} - \mathbf{x}', \boldsymbol{\xi} \rangle} d\boldsymbol{\xi}. \quad (153)$$

Now, using the spectral representation of the Matérn kernel together with (151), we write

$$k_{\nu,\kappa,\sigma^2}(\mathbf{x}, \mathbf{x}') = \frac{\sigma^2}{C_\nu} \int_{\mathbb{R}^n} \left(\frac{2\nu}{\kappa^2} + 4\pi^2\|\boldsymbol{\xi}\|^2\right)^{-\nu-\frac{n}{2}} e^{2\pi i \langle \mathbf{x} - \mathbf{x}', \boldsymbol{\xi} \rangle} d\boldsymbol{\xi} \quad (154)$$

$$= \frac{1}{C_\nu (4\pi)^{\frac{n}{2}} \Gamma(\nu + n/2)} \int_{\mathbb{R}^n} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2}u} S_{\infty,\sqrt{2u},\sigma^2}(\|\boldsymbol{\xi}\|) du e^{2\pi i \langle \mathbf{x} - \mathbf{x}', \boldsymbol{\xi} \rangle} d\boldsymbol{\xi} = \dots \quad (155)$$

By changing the order of integration, rearranging terms and using formula (153) we get

$$\dots = \frac{1}{C_\nu(4\pi)^{\frac{n}{2}}\Gamma(\nu+n/2)} \int_0^\infty u^{\nu-1} e^{-\frac{2\nu}{\kappa^2}u} k_{\infty, \sqrt{2u}, \sigma^2}(\mathbf{x}, \mathbf{x}') du. \quad (156)$$

Using the fact that $C_\nu = \frac{\Gamma(\nu)\kappa^{2\nu}}{2^n\pi^{n/2}\Gamma(\nu+n/2)(2\nu)^\nu}$ and noticing that

$$\frac{1}{C_\nu(4\pi)^{\frac{n}{2}}\Gamma(\nu+n/2)} \frac{2^n\pi^{n/2}\Gamma(\nu+n/2)(2\nu)^\nu}{\Gamma(\nu)\kappa^{2\nu}} = \frac{(2\nu)^\nu}{\Gamma(\nu)\kappa^{2\nu}} \quad (157)$$

proves the claim. \square

Result 44. *Euclidean Matérn kernels, defined in the integral sense, are positive semi-definite.*

Proof. This is a proof from Jaquier et al. (2022), which we provide for completeness. Assume $k_{\nu, \kappa, \sigma^2}$ is defined by Definition 17 and that the corresponding heat kernel $k_{\infty, \kappa, \sigma^2}$ is positive (semi)definite. Take some locations $x_1, \dots, x_m = \mathbf{x}$ and consider the matrix $\mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu$ with elements $k_{\nu, \kappa, \sigma^2}(x_i, x_j)$. In order to prove that $k_{\nu, \kappa, \sigma^2}$ is positive (semi)definite we need to show that $\mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu$ is a positive (semi)definite matrix for an arbitrary choice of m and \mathbf{x} . This means that we need to show that $\mathbf{y}^\top \mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu \mathbf{y}^\top > 0$ for all nonzero vectors $\mathbf{y} \in \mathbb{R}^m$.

To prove that $\mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu$ is positive definite, consider the matrices $\mathbf{K}_{\mathbf{x}\mathbf{x}}^{\infty, \kappa}$ with elements $k_{\infty, \kappa, \sigma^2}(x_i, x_j)$. Then, extending equation (63) to matrices, we have

$$\mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu = C \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} K_{\mathbf{x}\mathbf{x}}^{\infty, \sqrt{2u}} du. \quad (158)$$

Because of this, we obtain

$$\mathbf{y}^\top \mathbf{K}_{\mathbf{x}\mathbf{x}}^\nu \mathbf{y} = C \int_0^\infty \underbrace{u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u}}_* \underbrace{\mathbf{y}^\top K_{\mathbf{x}\mathbf{x}}^{\infty, \sqrt{2u}} \mathbf{y}}_{**} du \quad (159)$$

where the factor $*$ of the integrand is obviously positive and the factor $**$ of the integrand is positive (non-negative) because $\mathbf{K}_{\mathbf{x}\mathbf{x}}^{\infty, \sqrt{2u}}$ is positive (semi)definite by assumption, thus the integral is positive (non-negative). Thus $k_{\nu, \kappa, \sigma^2}$ is positive (semi)definite. \square

Proposition 18. *The Matérn kernel on a compact Lie group G is given by*

$$k_{\nu, \kappa, \sigma^2}(g_1, g_2) = \frac{\sigma^2}{C'_{\nu, \kappa}} \sum_{\lambda \in \Lambda} \left(\frac{2\nu}{\kappa^2} + \alpha_\lambda \right)^{-\nu-n/2} d_\lambda \chi_\lambda(g_2^{-1} \cdot g_1) \quad (64)$$

where $C'_{\nu, \kappa}$ is a normalizing constant which ensures that $k_{\nu, \kappa, \sigma^2}(g, g) = \sigma^2$.

Proof. Substituting the expression (57) for heat kernels on a compact Lie group G into (62), we get, for $x, y \in G$,

$$k_{\nu, \kappa, \sigma^2}(g_1 \cdot H, g_2 \cdot H) = \frac{\sigma^2}{C_{\nu, \kappa}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \mathcal{P}(u, g_1 \cdot H, g_2 \cdot H) du \quad (160)$$

$$= \frac{\sigma^2}{C_{\nu, \kappa}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \sum_{\lambda \in \Lambda} e^{-\alpha_\lambda u} d_\lambda \chi_\lambda(g_2^{-1} \cdot g_1) du. \quad (161)$$

Rearranging the terms we get

$$k_{\nu,\kappa,\sigma^2}(g_1 \cdot H, g_2 \cdot H) = \frac{\sigma^2}{C_{\nu,\kappa}} \sum_{\lambda \in \Lambda} \underbrace{\int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} e^{-\alpha_\lambda u} du}_{\Phi_{\kappa,\nu+n/2}(a_\lambda)} d\lambda \chi_\lambda(g_2^{-1} \cdot g_1) \quad (162)$$

where $\Phi_{\kappa,\nu+n/2}(a_\lambda)$ is defined as the inner integral. By Gradshteyn and Ryzhik (2014, Section 3.326, Item 2), we have

$$\int_0^\infty u^n e^{-au} du = \Gamma(n+1) a^{-n-1}. \quad (163)$$

Hence, $\Phi_{\kappa,\nu+n/2}(a_\lambda) = \Gamma(\nu+n/2) \left(\frac{2\nu}{\kappa^2} + a_\lambda\right)^{-\nu-n/2}$. We obtain

$$k_{\nu,\kappa,\sigma^2}(g_1 \cdot H, g_2 \cdot H) = \frac{\sigma^2}{C'_{\nu,\kappa}} \sum_{\lambda \in \Lambda} \left(\frac{2\nu}{\kappa^2} + a_\lambda\right)^{-\nu-n/2} d\lambda \chi_\lambda(g_2^{-1} \cdot g_1) \quad (164)$$

where $C'_{\nu,\kappa}$ is a normalizing constant ensuring that $k_{\nu,\kappa,\sigma^2}(x, x) = \sigma^2$. \square

Theorem 20. *For all $\nu > 0$, Matérn kernels on compact Lie groups and their homogeneous spaces are positive definite and lie in the Sobolev space $H^{\nu+n/2}$. Thus, they are continuous and possess continuous derivatives of all orders strictly less than ν .*

Proof. See Appendix E. First, note that

$$\min(1, 2\nu/\kappa^2)(1 + a_\lambda) \leq \frac{2\nu}{\kappa^2} + a_\lambda \leq \max(1, 2\nu/\kappa^2)(1 + a_\lambda). \quad (165)$$

Using this, by change of metric expressions it suffices to study k_{ν,κ,σ^2} only in the special case when $2\nu/\kappa^2 = 1$.

Recall that $\{\sqrt{d_\lambda} \pi_{jk}^{(\lambda)}\}_{\lambda,j,k}$ is an orthonormal basis of eigenfunctions of the Laplace–Beltrami operator. Then under our assumption that $2\nu/\kappa^2 = 1$, by De Vito et al. (2020, Proposition 2), we see that k_{ν,κ,σ^2} is the reproducing kernel of $H^{\nu+n/2}$, and therefore an element of this space. Therefore, positive definiteness and continuous differentiability follow from one of the Sobolev embedding theorems—see De Vito et al. (2020, Remark 4). The claim follows. \square

Proposition 21. *For $J \in \mathbb{N}$ sufficiently large, and \tilde{C} depending only on ν , κ , and C , we have*

$$\|k_{\nu,\kappa} - k_{\nu,\kappa}^{(J)}\|_{L^2(X \times X)} \leq \tilde{C} \begin{cases} J^{-2\nu/n-1/2} & \nu < \infty, \\ J^{1/2-1/n} e^{-\frac{\kappa^2}{2C} J^{2/n}} & \nu = \infty. \end{cases} \quad (68)$$

Proof. To begin, write

$$\|k_{\nu,\kappa} - k_{\nu,\kappa}^{(J)}\|_{L^2(X \times X)}^2 = \left\| \sum_{j=J+1}^\infty S_{\nu,\kappa}(\lambda_j) f_j(\cdot) f_j(\cdot) \right\|_{L^2(X \times X)}^2 = \sum_{j=J+1}^\infty S_{\nu,\kappa}(\lambda_j)^2 \quad (166)$$

where we have used the fact that f_j are orthonormal in $L^2(X)$. Since the function $S_{\nu,\kappa}(\cdot)$ is monotonically decreasing, Weyl's law given by (67) implies

$$\sum_{j=J+1}^{\infty} S_{\nu,\kappa}(\lambda_j)^2 \leq \sum_{j=J+1}^{\infty} S_{\nu,\kappa}(C^{-1}j^{2/n})^2. \quad (167)$$

First, assume $\nu < \infty$. Using elementary algebra and the Maclaurin–Cauchy test, we have

$$\sum_{j=J+1}^{\infty} S_{\nu,\kappa}(C^{-1}j^{2/n})^2 = \sum_{j=J+1}^{\infty} \left(\frac{2\nu}{\kappa^2} + C^{-1}j^{2/n} \right)^{-2\nu-n} \leq C^{2\nu+n} \sum_{j=J+1}^{\infty} j^{-4\nu/n-2} \quad (168)$$

$$\leq C^{2\nu+n} \int_J^{\infty} u^{-4\nu/n-2} du \quad (169)$$

$$= \underbrace{\frac{C^{2\nu+n}}{4\nu/n+1}}_{\tilde{C}^2} J^{-4\nu/n-1}. \quad (170)$$

This proves the claim for $\nu < \infty$. Now, assume $\nu = \infty$. First, notice the identity

$$\int_J^{\infty} e^{-au^b} du = a^{-1/b} b^{-1} \int_{aJ^b}^{\infty} v^{\frac{1-b}{b}} e^{-v} dv \quad (171)$$

which holds by change of variables: $au^b = v$.

Using the Maclaurin–Cauchy test and the variable change above with $a = \frac{\kappa^2}{C}$ and $b = 2/n$, write

$$\sum_{j=J+1}^{\infty} S_{\nu,\kappa}(C^{-1}j^{2/n})^2 = \sum_{j=J+1}^{\infty} e^{-\frac{\kappa^2}{C}j^{2/n}} \leq \int_J^{\infty} e^{-\frac{\kappa^2}{C}u^{2/n}} du \quad (172)$$

$$= \left(\frac{\kappa^2}{C} \right)^{-n/2} \frac{n}{2} \int_{\frac{\kappa^2}{C}J^{2/n}}^{\infty} v^{\frac{n-2}{2}} e^{-v} dv \quad (173)$$

$$= \left(\frac{\kappa^2}{C} \right)^{-n/2} \frac{n}{2} \Gamma\left(\frac{n}{2}, \frac{\kappa^2}{C} J^{2/n} \right) \quad (174)$$

where $\Gamma(\cdot, \cdot)$ is the incomplete Gamma function. Olver et al. (2023, (8.11)) immediately implies

$$\frac{\Gamma(x, y)}{y^{x-1} e^{-y}} \rightarrow 1 \quad y \rightarrow \infty. \quad (175)$$

Hence, for large enough J , we have

$$\sum_{j=J+1}^{\infty} S_{\infty,\kappa}(C^{-1}j^{2/n})^2 \leq \left(\frac{\kappa^2}{C} \right)^{-n/2} \frac{n}{2} \Gamma\left(\frac{n}{2}, \frac{\kappa^2}{C} J^{2/n} \right) \leq C' \left(\frac{\kappa^2 J^{2/n}}{C} \right)^{n/2-1} e^{-\frac{\kappa^2 J^{2/n}}{C}} \quad (176)$$

$$= \tilde{C}^2 J^{1-2/n} e^{-\frac{\kappa^2 J^{2/n}}{C}}. \quad (177)$$

The claim follows. \square

References

- J. L. Andersson and S. N. Sotiropoulos. Non-parametric representation and prediction of single-and multi-shell diffusion-weighted MRI data using Gaussian processes. *Neuroimage*, 2015. Cited on page 2.
- T. Augarten. Representation growth of the classical Lie algebras. *Communications in Algebra*, 2020. Cited on page 26.
- I. Azangulov, A. Smolensky, A. Terenin, and V. Borovitskiy. Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces II: non-compact symmetric spaces. *Journal of Machine Learning Research*, 2024. Cited on pages 2, 10.
- 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 6.
- 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 6.
- 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 pages 2, 5, 15, 25.
- T. Bröcker and T. Tom Dieck. *Representations of compact Lie groups*. Springer, 2013. Cited on pages 31, 33.
- R. L. Bryant. Surfaces in Conformal Geometry. In *Proceedings of Symposia in Pure Mathematics*. 1988. Cited on page 4.
- I. Chavel. *Eigenvalues in Riemannian geometry*. Academic Press, 1984. Cited on page 26.
- 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 5.
- N. Cressie. *Statistics for Spatial Data*. Wiley, 1992. Cited on page 6.
- 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 5.
- F. Dai and Y. Xu. *Approximation theory and harmonic analysis on spheres and balls*. Springer, 2013. Cited on page 20.

- A. W. Davis. Spherical functions on the Grassmann manifold and generalized Jacobi polynomials – Part 1. *Linear Algebra and its Applications*, 1999. Cited on page 15.
- A. W. Davis. Spherical functions on the Grassmann manifold and generalized Jacobi polynomials – Part 2. *Linear Algebra and its Applications*, 1999. Cited on page 15.
- E. De Vito, N. Mücke, and L. Rosasco. Reproducing kernel Hilbert spaces on manifolds: Sobolev and diffusion spaces. *Analysis and Applications*, 2020. Cited on page 47.
- H. D. Fegan. *Introduction to compact Lie groups*. World Scientific, 1991. Cited on pages 14, 22, 31, 33, 37.
- A. Feragen and S. Hauberg. Open Problem: Kernel methods on manifolds and metric spaces. What is the probability of a positive definite geodesic exponential kernel? In *Conference on Learning Theory*, 2016. Cited on page 5.
- A. Feragen, F. Lauze, and S. Hauberg. Geodesic exponential kernels: when curvature and linearity conflict. In *Computer Vision and Pattern Recognition*, 2015. Cited on page 4.
- G. B. Folland. *A course in abstract harmonic analysis*. CRC Press, 1995. Cited on pages 7, 9.
- P. I. Frazier. Bayesian Optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*. Institute for Operations Research and the Management Sciences, 2018. Cited on page 6.
- K. Fukumizu, A. Gretton, B. Schölkopf, and B. K. Sriperumbudur. Characteristic kernels on groups and semigroups. In *Advances in Neural Information Processing Systems*, 2008. Cited on page 6.
- S. S. Gelbart. A theory of Stiefel harmonics. *Transactions of the American Mathematical Society*, 1974. Cited on pages 15, 28.
- I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, 7th edition, 2014. Cited on pages 23, 45, 47.
- A. Grigoryan. *Heat kernel and analysis on manifolds*. American Mathematical Society, 2009. Cited on pages 5, 22.
- J. Guinness and M. Fuentes. Isotropic covariance functions on spheres: Some properties and modeling considerations. *Journal of Multivariate Analysis*, 2016. Cited on pages 2, 25.
- 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 pages 2, 6.
- 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, 45, 46.

- 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 page 2.
- 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.
- J. Jeong, M. Jun, and M. G. Genton. Spherical process models for global spatial statistics. *Statistical Science*, 2017. Cited on page 2.
- R. Kondor. *Group theoretical methods in machine learning*. PhD thesis, Columbia University, 2008. Cited on page 5.
- J. Lafferty, G. Lebanon, and T. Jaakkola. Diffusion kernels on statistical manifolds. *Journal of Machine Learning Research*, 2005. Cited on page 5.
- A. Lang and C. Schwab. Isotropic Gaussian random fields on the sphere: regularity, fast simulation and stochastic partial differential equations. *Annals of Applied Probability*, 2015. Cited on page 2.
- 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 5.
- D. Marinucci and G. Peccati. *Random fields on the sphere: representation, limit theorems and cosmological applications*. Cambridge University Press, 2011. Cited on page 2.
- F. Mezzadri. How to generate random matrices from the classical compact groups. *Notices of the American Mathematical Society*, 2007. Cited on page 17.
- M. A. Naimark and A. I. Stern. *Theory of Group Representations*. Springer, 1982. Cited on page 34.
- M. Niu, P. Cheung, L. Lin, Z. Dai, N. Lawrence, and D. Dunson. Intrinsic Gaussian processes on complex constrained domains. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2019. Cited on page 5.
- F. W. J. Olver, A. B. O. Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. V. S. B. R. Mille and, H. S. Cohl, and M. A. McClain. NIST Digital Library of Mathematical Functions, 2023. URL: [HTTPS://DLMF.NIST.GOV/](https://dlmf.nist.gov/). Cited on page 48.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2008. Cited on page 3.
- C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. Cited on page 23.

- 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 6.
- S. Said, P.-O. Amblard, and J. H. Manton. Stationary Random Fields Arising From Second-Order Partial Differential Equations on Compact Lie Groups. *IEEE Transactions on Information Theory*, 2013. Cited on page 5.
- M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, 1999. Cited on pages 3, 23.
- A. Weil. *L'intégration dans les groupes topologiques et ses applications*. Hermann, 1940. Cited on pages 7, 10.
- 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 page 3.
- 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 page 3.
- 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 6–8, 10, 13, 39–41.
- K. Ye, M. Niu, and P. Cheung. Heat kernel and intrinsic Gaussian processes on manifolds. *Biometrika*, 2020. Cited on page 5.
- C. Zhao and J. S. Song. Exact heat kernel on a hypersphere and its applications in kernel SVM. *Frontiers in Applied Mathematics and Statistics*, 2018. Cited on page 5.