

# The GeometricKernels Package: Heat and Matérn Kernels for Geometric Learning on Manifolds, Meshes, and Graphs

**Peter Mostowsky**<sup>†</sup>

*Independent*

PMOSTOWSKY@GMAIL.COM

**Vincent Dutordoir**

*University of Cambridge*

VD309@CAM.AC.UK

**Iskander Azangulov**

*University of Oxford*

ISKANDER.AZANGULOV@SPC.OX.AC.UK

**Noémie Jaquier**

*KTH Royal Institute of Technology*

JAQUIER@KTH.SE

**Michael John Hutchinson**

*University of Oxford*

MICHAEL.HUTCHINSON@UNIV.OX.AC.UK

**Aditya Ravuri**

*University of Cambridge*

AR847@CAM.AC.UK

**Leonel Rozo**

*Italian Institute of Artificial Intelligence for Industry*

LEONEL.ROZO@AI4I.IT

**Alexander Terenin**<sup>†</sup>

*Cornell University*

AVT28@CORNELL.EDU

**Viacheslav Borovitskiy**<sup>†</sup>

*ETH Zürich and University of Edinburgh*

VIACHESLAV.BOROVITSKIY@GMAIL.COM

**Editor:** Zeyi Wen

## Abstract

Kernels are a fundamental technical primitive in machine learning. In recent years, kernel-based methods such as Gaussian processes are becoming increasingly important in applications where quantifying uncertainty is of key interest. In settings that involve structured data defined on graphs, meshes, manifolds, or other related spaces, defining kernels with good uncertainty-quantification behavior, and computing their value numerically, is less straightforward than in the Euclidean setting. To address this difficulty, we present `GEOMETRICKERNELS`, a Python software package which implements the geometric analogs of classical Euclidean squared exponential—also known as heat—and Matérn kernels, which are widely-used in settings where uncertainty is of key interest. As a byproduct, we obtain the ability to compute Fourier-feature-type expansions, which are widely used in their own right, on a wide set of geometric spaces. Our implementation supports automatic differentiation in every major current framework simultaneously via a backend-agnostic design. In this companion paper to the package and its documentation, we outline the capabilities of the package and present an illustrated example of its interface. We also include a brief overview of the theory the package is built upon and provide some historic context in the appendix.

---

<sup>†</sup>Corresponding authors.

Website: [HTTPS://GEOMETRIC-KERNELS.GITHUB.IO](https://geometric-kernels.github.io).

Code available at: [HTTPS://GITHUB.COM/GEOMETRIC-KERNELS/GEOMETRICKERNELS](https://github.com/geometric-kernels/geometrickernels).

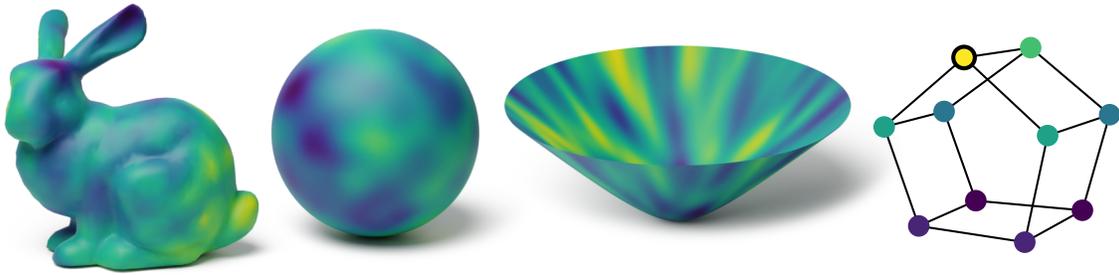


Figure 1: Left: GeometricKernels-computed samples from a set of Gaussian processes, with covariance given by the heat kernel, on three spaces: the Stanford bunny mesh, the unit sphere  $\mathbb{S}_2$ , and the hyperbolic space  $\mathbb{H}_2$ . Right: a GeometricKernels-computed Matérn kernel  $k(\bullet, \cdot)$  on a graph.

## 1. Introduction

Kernel methods, and methods built on them such as Gaussian processes, are a fundamental part of the modern machine learning toolkit. The growth of machine learning has led to models being adopted in *geometric* settings that involve structured data defined over spaces such as graphs, meshes, and their smooth analogs such as Riemannian manifolds. An increasing need for kernel methods in such settings—particularly those with principled and intuitive handling of uncertainty, such as symmetry-based kernels—has led to a recent growth in research on *geometric kernels* and their applications in fields such as robotics and neuroscience: examples are given in Appendix B.

A key challenge in geometric settings is that kernels, to be well-defined, need to be positive semi-definite functions. In geometric settings—in contrast to the classical Euclidean case—it is generally unclear how to define such functions directly using closed-form analytic expressions. Straightforward generalizations, such as the distance-based squared exponential kernel, fail to be positive semi-definite in most cases (Schoenberg, 1938; Feragen et al., 2015; Da Costa et al., 2025). As a result, many kernels are implicitly-defined—so, their numerical computation tends to be more involved than of Euclidean kernels. Our key motivation is to develop software that handles this difficulty well.

We present GEOMETRICKERNELS, a software package designed to make geometric kernels as easy to work with in practice as possible. We provide implementations of the geometric analogs of classical Euclidean squared exponential and Matérn kernel classes, which can be used within kernel learning pipelines in a largely plug-and-play manner. We support graphs, meshes, and a collection of mathematically tractable manifolds. Our implementation is designed to seamlessly support GPUs and automatic differentiation using all major frameworks—PYTORCH, JAX, and TENSORFLOW—simultaneously. We also support NUMPY to facilitate debugging and other use cases.

## 2. The GeometricKernels Package

The GeometricKernels Python package implements heat and Matérn kernels for a broad class of geometric spaces. To aid understanding, we review these kernels in Appendix A. In addition, Appendix B includes a brief historical overview of these and related kernel classes. We now discuss this package’s technical capabilities, some of which are illustrated in Figure 1, as well as its design.

### 2.1 Classes for Geometric Spaces

Supported spaces are represented as classes in the `spaces` subpackage. Current support includes:

1. Compact Riemannian manifolds: the unit circle  $\mathbb{S}_1$  via the class `Circle`; unit hyperspheres  $\mathbb{S}_n$ ,  $n \geq 2$ , via the class `Hypersphere`; special orthogonal Lie groups  $\text{SO}(n)$ ,  $n \geq 3$ , via the class `SpecialOrthogonal`; special unitary Lie groups  $\text{SU}(n)$ ,  $n \geq 2$  via the class `SpecialUnitary`.
  2. Non-compact Riemannian manifolds: the hyperbolic spaces  $\mathbb{H}_n$ ,  $n \geq 2$ , via the class `Hyperbolic`; the manifold  $\text{SPD}(n)$ ,  $n \geq 2$ , of symmetric positive definite matrices  $n \times n$  endowed with the affine-invariant Riemannian metric via the class `SymmetricPositiveDefiniteMatrices`.
  3. Meshes, viewed as approximations of smooth two-dimensional surfaces: via the class `Mesh`.
  4. Node sets of general undirected graphs with non-negative weights: via the class `Graph`.
  5. Node sets of specialized graphs which are used for representing combinatorial domains: via, for instance, the `HypercubeGraph` class.
  6. Edge sets of unweighted undirected graphs, or more generally edge sets of simplicial 2-complexes: via the `GraphEdges` class.
  7. Products of discrete spectrum spaces: among the preceding spaces, this includes everything except non-compact Riemannian manifolds, and is accessed via `ProductDiscreteSpectrumSpace`.
- A brief discussion of additional spaces one might incorporate in future work is given in Appendix A.

## 2.2 Classes for Geometric Kernels

Although multiple kernel classes technically exist in the library, the class `MaternGeometricKernel` in the `kernels` subpackage universally covers most use cases, without any need to understand the library internals or the mathematics behind the kernels. This `MaternGeometricKernel` dispatches on the provided space and automatically chooses the internal implementation of the kernel. Another useful kernel class is `ProductGeometricKernel` which takes a number of geometric kernels  $k_i : X_i \times X_i \rightarrow \mathbb{R}$  for  $1 \leq i \leq m$ —where  $X_i$  can be any of the supported spaces, without restrictions—and constructs the product kernel  $k : X \times X \rightarrow \mathbb{R}$  on  $X = X_1 \times \dots \times X_m$  defined by

$$k((x_1, \dots, x_m), (x'_1, \dots, x'_m)) = k_1(x_1, x'_1) \cdot \dots \cdot k_m(x_m, x'_m) \quad (1)$$

This is more computationally efficient than constructing a `MaternGeometricKernel` instance corresponding to a `ProductDiscreteSpectrumSpace` and provides additional flexibility, as each of the factors  $k_i$  can have its own independent set of hyperparameters, providing automatic-relevance-determination-like (Rasmussen and Williams, 2006) behavior.

## 2.3 Approximate Finite-dimensional Feature Maps

Beyond kernels, the library offers the *approximate finite-dimensional feature maps* corresponding to all the kernels provided. These are functions  $\phi : X \rightarrow \mathbb{R}^\ell$  such that the kernel  $k$  satisfies

$$k(x, x') \approx \phi(x)^T \phi(x') \quad (2)$$

where  $\ell$  is a hyperparameter that controls approximation quality. Concrete feature map implementations can be found in the `feature_maps` subpackage. For most use cases, an appropriate feature map can be obtained via the function `default_feature_map` from the `kernels` subpackage.

One can use approximate finite-dimensional feature maps for many purposes. They are particularly helpful, for instance, to efficiently—that is, without incurring cubic costs—obtain an approximate sample from a Gaussian process  $f \sim \text{GP}(0, k)$  via

$$f(x) = \zeta^T \phi(x) \quad \zeta \sim N(\mathbf{0}, \mathbf{I}). \quad (3)$$

We provide convenience utilities for efficiently sampling Gaussian process priors in the `sampling` subpackage. Throughout, care is taken to ensure that computation of the map  $x \rightarrow \phi(x)$  is done in a fully-automatically-differentiable manner. For more on feature maps, see the appendices.

## 2.4 Multi-backend Design

GeometricKernels is designed to support multiple numerical computation and automatic-differentiation backends, including `PYTORCH`, `JAX`, `TENSORFLOW`, all with GPU support, as well as a basic `NUMPY` backend for lighter-weight use that does not require automatic-differentiation or GPUs. This is implemented via multiple dispatch, using function and operator overloading, which is provided by the library `LAB` (Bruinsma, 2019). From the user’s viewpoint, one simply needs to include an appropriate import statement, for example

```
import geometrickernels.torch
```

to enable PyTorch support. Alternatively, one can use

```
import geometrickernels.jax
```

to enable JAX support. GeometricKernels preserves the array types it is given: as long as all inputs to its functions are JAX-arrays, the output will also be a JAX-array, and similarly for other backends. All of our computations support batching in the same way as each backend. Each backend’s code works independently: as a result, in principle, all backends can be imported simultaneously—though, we do not do so unless the user explicitly requests it, to avoid bringing in unnecessary dependencies.

## 2.5 Integration with Other Packages

We support integration with Gaussian process packages `GPYTORCH` (Gardner et al., 2018), `GPJAX` (Pinder and Dodd, 2022), and `GPFLOW` (Matthews et al., 2017), which are currently the most popular packages in their ecosystems. This integration is provided via the respective kernel wrapper classes `GPYtorchGeometricKernel`, `GPJaxGeometricKernel`, `GPflowGeometricKernel` in the `frontends` subpackage. We also support standalone use without a frontend, for instance for kernel PCA.

## 3. An Illustrative Example

We now walk through a simple example, which is mirrored on our documentation. Our goal is to compute the kernel matrix for a Matérn kernel on the sphere. First, let us import the package using

```
import numpy
import geometric_kernels
from geometric_kernels.spaces import Hypersphere
from geometric_kernels.kernels import MaternGeometricKernel
```

where, since we have not imported a backend, the package will use the `NUMPY` backend by default. To create a kernel one uses

```
hypersphere = Hypersphere(dim=2)
kernel = MaternGeometricKernel(hypersphere)
```

which first creates the space—here, a simple two-dimensional unit sphere  $\mathbb{S}_2$ —and then, the kernel over the space. To ensure maximum support for different backends, which might handle trainable parameters differently, we adopt a JAX-inspired design where kernel hyperparameters are stored within a Python dictionary. This is initialized, and then set, using

```
params = kernel.init_params()
params["nu"] = np.array([5/2])
params["lengthscale"] = np.array([1.])
```

which gives us a parameter object, here called `params`. We now assemble a kernel matrix using three points on the sphere, parameterized by their coordinates in  $\mathbb{R}^3$ . For this, we evaluate the kernel via

```

xs = np.array([[0., 0., 1.],
               [0., 1., 0.],
               [1., 0., 0.]])
Kxx = kernel.K(params, xs)

```

which returns a matrix `Kxx`, with whatever array type was supplied within the input variables, namely `params`, `x1`, and `x2`. We can now use this kernel matrix for whatever we want—for instance, kernel ridge regression—but for illustrative purposes here simply print it. The obtained values are

```

print(np.around(Kxx, 2))

# [[1.  0.36 0.36]
#  [0.36 1.  0.36]
#  [0.36 0.36 1.  ]]

```

which, since our inputs are Numpy arrays, produces a Numpy array. An illustration of how our code works with different backends is shown in Appendix C.

## 4. Conclusion

In this work, we introduced the GeometricKernels Python package. We documented the kinds of kernels, spaces and other primitives that it supports, and described the key elements of its design philosophy, including seamless GPU and automatic differentiation support within a number of different backends. We concluded with a simple example illustrating these capabilities of the package.

## Acknowledgments

We are grateful to Wessel Bruinsma for his prompt responses and invaluable assistance with the LAB library (Bruinsma, 2019), which is the foundation of the backend-independence of our package. Furthermore, we express our gratitude to Thomas Pinder for contributing the first version of the geometric kernel wrapper for his library GPJax (Pinder and Dodd, 2022). While at St. Petersburg State University, PM was supported by the RSF grant N°21-11-00047. IA was supported by the UK Engineering and Physical Sciences Research Council, grant number EP/T517811/1. NJ is supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. AR is supported by the Accelerate Programme for Scientific Discovery. 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.

## References

- M. Alain, S. Takao, B. Paige, and M. P. Deisenroth. Gaussian Processes on Cellular Complexes. In *International Conference on Machine Learning*, 2024. Cited on page 13.
- I. Azangulov, A. Smolensky, A. Terenin, and V. Borovitskiy. Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces I: the compact case. *arXiv:2208.14960*, 2022. Cited on pages 11–13.
- 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. *arXiv:2301.13088*, 2023. Cited on pages 11, 13.

- D. Bolin, A. B. Simas, and J. Wallin. Gaussian Whittle–Matérn fields on metric graphs. *Bernoulli*, 2024. Cited on page 13.
- K. M. Borgwardt, E. Ghisu, F. Llinares-López, L. O’Bray, and B. Rieck. Graph Kernels: State-of-the-Art and Future Challenges. *Foundations and Trends in Machine Learning*, 2020. Cited on page 13.
- K. M. Borgwardt, H.-P. Kriegel, S. V. N. Vishwanathan, and N. N. Schraudolph. Graph kernels for disease outcome prediction from protein-protein interaction networks. In *Biocomputing*, 2007. Cited on page 13.
- K. M. Borgwardt, C. S. Ong, S. Schönauer, S. V. N. Vishwanathan, A. J. Smola, and H.-P. Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 2005. Cited on page 13.
- V. Borovitskiy, I. Azangulov, A. Terenin, P. Mostowsky, M. Deisenroth, and N. Durrande. Matérn Gaussian processes on graphs. In *Artificial Intelligence and Statistics*, 2021. Cited on pages 11, 13.
- 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 13.
- V. Borovitskiy, A. Terenin, P. Mostowsky, and M. Deisenroth. Matérn Gaussian Processes on Riemannian Manifolds. In *Advances in Neural Information Processing Systems*, pages 12426–12437, 2020. Cited on pages 11–13.
- W. Bruinsma. LAB: A Generic Interface for Linear Algebra Backends. [HTTPS://GITHUB.COM/WESSELB/LAB](https://github.com/wesselb/lab), 2019. Cited on pages 4, 5.
- S. Coveney, C. Cantwell, and C. Roney. Atrial conduction velocity mapping: Clinical tools, algorithms and approaches for understanding the arrhythmogenic substrate. *Medical and Biological Engineering and Computing*, 2022. Cited on page 13.
- N. Da Costa, C. Mostajeran, J.-P. Ortega, and S. Said. Invariant kernels on Riemannian symmetric spaces: a harmonic-analytic approach. *SIAM Journal on Mathematics of Data Science*, 2025. Cited on page 2.
- C. Doumont, V. Picheny, V. Borovitskiy, and H. Moss. Omnipresent Yet Overlooked: Heat Kernels in Combinatorial Bayesian Optimization. In *Advances in Neural Information Processing Systems*, 2025. Cited on page 13.
- D. B. Dunson, H.-T. Wu, and N. Wu. Graph based Gaussian processes on restricted domains. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2022. Cited on page 13.
- V. Dutordoir, N. Durrande, and J. Hensman. Sparse Gaussian processes with spherical harmonic features. In *International Conference on Machine Learning*, 2020. Cited on page 13.
- A. Feragen, F. Lauze, and S. Hauberg. Geodesic exponential kernels: when curvature and linearity conflict. In *Conference on Computer Vision and Pattern Recognition*, 2015. Cited on page 2.
- B. Fichera, V. Borovitskiy, A. Krause, and A. Billard. Implicit Manifold Gaussian Process Regression. *Advances in Neural Information Processing Systems*, 2023. Cited on page 13.
- P. I. Frazier. Bayesian Optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*. INFORMS, 2018. Cited on page 10.

- T. Gao, S. Z. Kovalsky, and I. Daubechies. Gaussian process landmarking on manifolds. *SIAM Journal on Mathematics of Data Science*, 2019. Cited on page 13.
- J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. In *Advances in Neural Information Processing Systems*, 2018. Cited on page 4.
- R. Garnett. *Bayesian Optimization*. Cambridge University Press, 2023. Cited on page 10.
- T. Gärtner. *Kernels for Structured Data*. World Scientific, 2008. Cited on page 13.
- T. Gärtner, J. W. Lloyd, and P. A. Flach. Kernels and distances for structured data. *Machine Learning*, 2004. Cited on page 13.
- S. Ghosh, N. Das, T. Gonçalves, P. Quaresma, and M. Kundu. The journey of graph kernels through two decades. *Computer Science Review*, 2018. Cited on page 13.
- A. Grigoryan. *Heat kernel and analysis on manifolds*. American Mathematical Society, 2009. Cited on page 13.
- J. He, G. Ma, J. Kang, and Y. Yang. Scalable Bayesian inference for heat kernel Gaussian processes on manifolds. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2025. Cited on page 13.
- M. Hutchinson, A. Terenin, V. Borovitskiy, S. Takao, Y. Teh, and M. Deisenroth. Vector-valued Gaussian Processes on Riemannian Manifolds via Gauge Independent Projected Kernels. *Advances in Neural Information Processing Systems*, 2021. Cited on page 13.
- 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 page 13.
- N. Jaquier, L. Rozo, M. González-Duque, V. Borovitskiy, and T. Asfour. Bringing motion taxonomies to continuous domains via GPLVM on hyperbolic manifolds. In *International Conference on Machine Learning*, 2024. Cited on page 13.
- B. Jie, D. Zhang, W. Gao, Q. Wang, C.-Y. Wee, and D. Shen. Integration of network topological and connectivity properties for neuroimaging classification. *IEEE Transactions on Biomedical Engineering*, 2013. Cited on page 13.
- H. Kashima, K. Tsuda, and A. Inokuchi. Marginalized kernels between labeled graphs. In *International Conference on Machine Learning*, 2003. Cited on page 13.
- R. I. Kondor and J. Lafferty. Diffusion kernels on graphs and other discrete structures. In *International Conference on Machine Learning*, 2002. Cited on page 13.
- R. Kondor and H. Pan. The Multiscale Laplacian Graph Kernel. *Advances in Neural Information Processing Systems*, 2016. Cited on page 13.
- N. M. Kriege, F. D. Johansson, and C. Morris. A survey on graph kernels. *Applied Network Science*, 2020. Cited on page 13.
- J. Lafferty, G. Lebanon, and T. Jaakkola. Diffusion kernels on statistical manifolds. *Journal of Machine Learning Research*, 2005. Cited on page 13.

- C. Leslie, E. Eskin, and W. S. Noble. The spectrum kernel: A string kernel for SVM protein classification. In *Biocomputing*, 2002. Cited on page 13.
- D. Li, W. Tang, and S. Banerjee. Inference for Gaussian Processes with Matérn Covariogram on Compact Riemannian Manifolds. *Journal of Machine Learning Research*, 2023. Cited on page 13.
- F. Lindgren, D. Bolin, and H. Rue. The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running. *Spatial Statistics*, 2022. Cited on page 11.
- 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 11.
- H. Lodhi, C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. Watkins. Text classification using string kernels. *Journal of Machine Learning Research*, 2002. Cited on page 13.
- A. G. d. G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 2017. Cited on page 4.
- M. Niu, Z. Dai, P. Cheung, and Y. Wang. Intrinsic Gaussian process on unknown manifolds with probabilistic metrics. *Journal of Machine Learning Research*, 2023. Cited on page 13.
- R. L. Peach, M. Vaino-Carl, N. Grossman, M. David, E. Mallas, D. Sharp, P. A. Malhotra, P. Vanderghenst, and A. Gosztolai. Implicit Gaussian process representation of vector fields over arbitrary latent manifolds. In *International Conference on Learning Representations*, 2024. Cited on page 13.
- T. Pinder and D. Dodd. GPJax: A Gaussian Process Framework in JAX. *Journal of Open Source Software*, 2022. Cited on pages 4, 5.
- E. Porcu, M. Bevilacqua, R. Schaback, and C. J. Oates. The Matérn model: A journey through statistics, numerical analysis and machine learning. *Statistical Science*, 2024. Cited on page 13.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2008. Cited on page 12.
- C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. Cited on pages 3, 10.
- G. Riutort-Mayol, P.-C. Bürkner, M. R. Andersen, A. Solin, and A. Vehtari. Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming. *Statistics and Computing*, 2023. Cited on page 13.
- D. Robert-Nicoud, A. Krause, and V. Borovitskiy. Intrinsic Gaussian Vector Fields on Manifolds. In *Artificial Intelligence and Statistics*, 2024. Cited on page 13.
- 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 13.
- B. Ru, X. Wan, X. Dong, and M. Osborne. Interpretable Neural Architecture Search via Bayesian Optimisation with Weisfeiler–Lehman Kernels. In *International Conference on Learning Representations*, 2020. Cited on page 13.

- I. J. Schoenberg. Metric spaces and positive definite functions. *Transactions of the American Mathematical Society*, 1938. Cited on page 2.
- B. Schölkopf and A. J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press, 2002. Cited on page 10.
- N. Shervashidze, P. Schweitzer, E. J. Van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. Weisfeiler–Lehman graph kernels. *Journal of Machine Learning Research*, 2011. Cited on page 13.
- A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing*, 2020. Cited on page 13.
- S. J. Swamidass, J. Chen, J. Bruand, P. Phung, L. Ralaivola, and P. Baldi. Kernels for small molecules and the prediction of mutagenicity, toxicity and anti-cancer activity. *Bioinformatics*, 2005. Cited on page 13.
- S. Takerkart, G. Auzias, B. Thirion, and L. Ralaivola. Graph-Based Inter-Subject Pattern Analysis of fMRI Data. *PLoS One*, 2014. Cited on page 13.
- G. Thomas, P. Flach, and W. Stefan. On Graph Kernels: Hardness Results and Efficient Alternatives. In *Learning Theory and Kernel Machines*. Springer, 2003. Cited on page 13.
- S. V. N. Vishwanathan, N. N. Schraudolph, R. Kondor, and K. M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 2010. Cited on page 13.
- P. Whittle. Stochastic processes in several dimensions. *Bulletin of the International Statistical Institute*, 1963. Cited on page 11.
- 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 10.
- 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 10.
- K. Wyrwal, A. Krause, and V. Borovitskiy. Residual Deep Gaussian Processes on Manifolds. In *The Thirteenth International Conference on Learning Representations*, 2025. Cited on page 13.
- K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*, 2018. Cited on page 13.
- P. Yanardag and S. V. N. Vishwanathan. Deep Graph Kernels. In *Knowledge Discovery and Data Mining*, 2015. Cited on page 13.
- M. Yang, V. Borovitskiy, and E. Isufi. Hodge-Compositional Edge Gaussian Processes. In *Artificial Intelligence and Statistics*, 2024. Cited on page 13.

## Appendix A. Heat and Matérn Kernels on Non-Euclidean Spaces

Here we introduce the necessary technical notions, beginning with a brief overview of kernels in general, Gaussian processes, and the specific kernel classes of interest.

### A.1 Brief Review of Kernels

Let  $X$  be a set. A *kernel* is defined to be a positive semi-definite function  $k : X \times X \rightarrow \mathbb{R}$ . By *positive semi-definite function*, we mean that for any  $N$  and any collection of data points  $\mathbf{x} = (x_1, \dots, x_N)$ , the *kernel matrix*

$$\mathbf{K}_{\mathbf{x}\mathbf{x}} = \begin{bmatrix} k(x_i, x_j) : & i = 1, \dots, N \\ & j = 1, \dots, N \end{bmatrix} \quad (4)$$

must be positive semi-definite. We will be interested in kernels whose domain  $X$  is a manifold, mesh, or graph rather than a Euclidean space  $\mathbb{R}^n$ .

Kernels are a fundamental building block of many machine learning algorithms (Schölkopf and Smola, 2002). Examples include non-linear classification and regression algorithms, such as for instance kernel ridge regression, as well as kernel-based distance functions such as the maximum mean discrepancy and dimensionality reduction techniques like kernel PCA. For most of these algorithms, the key computational primitive involves evaluating the kernel function  $k(x, x')$  pointwise for some set of data points  $x, x' \in X$  chosen by the algorithm. In addition, certain algorithms of key interest, such as Gaussian process models which quantify uncertainty, often require further computational primitives: to understand what is needed, we introduce them now.

### A.2 Gaussian Processes

One of the most important classes of kernel-based models today is the class of *Gaussian process* models (Rasmussen and Williams, 2006), which are used for Bayesian learning and decision-making algorithms such as Bayesian optimization (Frazier, 2018; Garnett, 2023) and various forms of active learning. We introduce these now.

In the simplest case, one considers the non-linear regression model  $y_i = f(x_i) + \varepsilon_i$ , where  $\varepsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$  is the likelihood noise variance, and  $f$  is the unknown regression function to be learned. Gaussian processes are a class of prior distributions on the space of functions  $f : X \rightarrow \mathbb{R}$  which are uniquely determined by their mean and covariance kernel, written  $f \sim \text{GP}(m, k)$ . Such priors are conjugate: letting  $m = 0$  without loss of generality, the resulting posterior distribution  $f \mid \mathbf{y}$  is another Gaussian process, with mean and covariance

$$\mu_{f \mid \mathbf{y}}(\cdot) = \mathbf{K}_{(\cdot, \mathbf{x})}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \Sigma)^{-1} \mathbf{y} \quad k_{f \mid \mathbf{y}}(\cdot, \cdot') = \mathbf{K}_{(\cdot, \mathbf{x})}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \Sigma)^{-1} \mathbf{K}_{(\mathbf{x}, \cdot')}. \quad (5)$$

One can also express posterior random functions using *pathwise conditioning* (Wilson et al., 2020, 2021) as transformations of prior random functions, given by

$$(f \mid \mathbf{y})(\cdot) = f(\cdot) + \mathbf{K}_{(\cdot, \mathbf{x})}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \Sigma)^{-1}(\mathbf{y} - f(\mathbf{x}) - \varepsilon). \quad (6)$$

To use these expressions, we require two computational primitives: (1) mirroring general kernel methods, the ability to compute  $k(x, x')$ , and (2) the ability to sample the prior  $f(\cdot)$  at a potentially-large set of evaluation points jointly, in a computationally-efficient manner with respect to the number of evaluation points. Of these, (2) generally requires computational primitives beyond pointwise kernel evaluation: a standard approach involves using an *approximate feature map*—that is, a function  $\phi : X \rightarrow \mathbb{R}^l$  that  $k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathbb{R}^l}$ .

### A.3 Heat Kernels

Principled universal kernels that adapt to the geometry of the input domain are of key importance in the tasks where well-calibrated uncertainty is needed. In the Euclidean case, the most popular

such classes are the squared exponential (RBF, Gaussian, diffusion, heat) and Matérn kernel. It turns out—see, for instance, Borovitskiy et al. (2020, 2021) and Azangulov et al. (2022, 2023)—these classes can be generalized to a number of geometric settings in principled way. These generalizations can be derived in many different ways: in what follows, we describe them in, arguably, the simplest way possible, along the lines of Azangulov et al. (2022, 2023).

We begin by introducing the *heat kernel*  $k_{\infty, \kappa, \sigma^2}$ , which is parameterized by a pair of parameters: the *length scale*  $\kappa > 0$  and *amplitude* (also called *outputscale*)  $\sigma^2$ . This kernel is the infinite-smoothness limit of the more general class of Matérn kernels  $k_{\nu, \kappa, \sigma^2}$  which we define in terms of the heat kernel in the sequel. Let  $X$  be a Riemannian manifold, mesh, or graph. Define

$$k_{\infty, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\infty, \kappa}} \mathcal{P}(\kappa^2/2, x, x') \quad (7)$$

where  $\mathcal{P}$  is the solution of the differential equation

$$\frac{\partial \mathcal{P}}{\partial t}(t, x, x') = \Delta_x \mathcal{P}(t, x, x') \quad \mathcal{P}(0, x, x') = \delta(x, x'), \quad (8)$$

where  $\Delta_x$  is a suitable notion of the Laplace operator operator on the space of interest, acting on the  $x$  variable, and  $\delta(x, x')$  is the Dirac delta function, defined appropriately. The constant  $C_{\infty, \kappa}$  is defined to ensure  $\int_X k_{\infty, \kappa, \sigma^2}(x, x) dx = \sigma^2$ . On all of the classes of spaces of interest, one can show that Equation (8) admits a unique solution within a suitable class of functions. For our purposes, the key computational question will be how to work with this kernel numerically: we will return to this question momentarily, but first define the full class of Matérn kernels.

#### A.4 Matérn Kernels

Heat kernels are infinitely-differentiable functions: in many applications, such as Bayesian optimization, one would rather work with finitely-differentiable functions because in many cases this has been shown to lead to improved performance. To do so, we introduce the class of Matérn kernels, defined for  $\nu > 0$  using the expression

$$k_{\nu, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\nu, \kappa}} \int_0^\infty u^{\nu-1+n/2} e^{-\frac{2\nu}{\kappa^2}u} \mathcal{P}(u, x, x') du \quad (9)$$

where  $C_{\nu, \kappa}$  is the constant that ensures  $\int_X k_{\nu, \kappa, \sigma^2}(x, x) dx = \sigma^2$ , and  $n$  is a space-dependent constant which is equal to  $\dim X$  for manifolds, 0 for graphs, and 2 for meshes approximating smooth surfaces. For the spaces of interest, one can show these integrals are well-defined whenever  $\mathcal{P}$  is defined.

In the Euclidean setting, Matérn kernels are usually defined as stationary kernels whose spectral measure is a  $t$ -distribution, which is a gamma mixture of normals. Our definition is equivalent to this one in the Euclidean case, but moves the gamma-mixture-of-normals representation into the kernel, rather than spectral, domain—and, as a result, is much more general, as it requires no structure on the space of interest, other than the ability to define a heat kernel. Heat kernels, in turn, are of fundamental interest in mathematics, and a substantial amount of research has gone into how to define them on a large class of spaces. This makes our definition particularly simple and convenient.

One can derive the same class of kernels from other viewpoints, as well. In particular, it is easy to show these kernels coincide with ones arising from the stochastic partial differential equation view of Gaussian processes of Whittle (1963) and Lindgren et al. (2011, 2022), by examining their spectral representations like that of Borovitskiy et al. (2020, 2021) and applying the argument in Azangulov et al. (2022, Proposition 18). One can similarly obtain this expression starting from the perspective of reproducing kernels on Sobolev spaces (Borovitskiy et al., 2020, Appendix D), which, as a byproduct, implies the same differentiability properties for these kernels as for their Euclidean counterparts.

From a computational viewpoint, working with Matérn kernels turns out to be similar to working with the heat kernel: as occurs in all cases currently considered in the library, the extra integral can be passed into whatever computational scheme one uses for  $\mathcal{P}$ , and need not be evaluated numerically. To understand this better, we now review the class of computational techniques we work with.

## A.5 Fourier Features

The primary class of computational techniques we implement are based on suitable forms of Fourier analysis to the spaces of interest. For instance, for *discrete spectrum spaces* like compact Riemannian manifolds, graphs and meshes, these are based off of the Laplacian eigenfunctions—namely, functions  $\phi_n$  which satisfy  $\Delta_x \phi_n = -\lambda_n \phi$ . Using these functions, one can solve the heat equation analytically: after some algebra, obtaining formulas such as

$$k_{\infty, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\infty, \kappa}} \sum_{n=0}^{\infty} \Phi_{\infty, \kappa}(\lambda_n) f_n(x) f_n(x') \quad (10)$$

where  $\Phi_{\infty, \kappa}(\lambda) = e^{-\frac{\kappa^2}{2}\lambda}$  in the heat kernel case. By passing the integral through the sum, one can show that Matérn kernels have the same form, with a certain function  $\Phi_{\nu, \kappa}$  with an explicit analytic form in place of the function  $\Phi_{\infty, \kappa}$ . This expression defines a Mercer-type expansion for the kernel, also known as *manifold Fourier features*.

In our work, we implement a large class of expansions of the kind described above—where the precise expansion used, and precise form of  $f_n$ , are computed on a case-by-case basis for various manifolds. For meshes and graphs, computing their analogues—graph Laplacian and mesh Laplacian eigenvectors—is a standard and well-studied problem in numerical linear algebra, with broad use in computer graphics and other areas.

For many spaces, especially for compact Riemannian manifolds, we encounter repeated eigenvalues  $\lambda_n$ . In such cases, we can sometimes compute the sums  $G_l(x, x') = \sum_{s=1}^{d_l} f_{l,s}(x) f_{l,s}(x')$ , where  $\{f_{l,s}\}_{s=1}^{d_l}$  are some sets of eigenfunctions corresponding to the same eigenvalue  $\lambda_l$ , without actually computing the individual eigenfunctions  $f_{l,s}$ . This gives a simplified and faster-converging expansion:<sup>1</sup>

$$k_{\nu, \kappa, \sigma^2}(x, x') = \frac{\sigma^2}{C_{\nu, \kappa}} \sum_{l=0}^{\infty} \Phi_{\nu, \kappa}(\lambda_l) G_l(x, x'). \quad (11)$$

For example, the classical addition theorem on the hyperspheres  $\mathbb{S}_n$  identifies the  $G_l$ , where  $\{f_{l,s}\}_{s=1}^{d_l}$  are sets of eigenfunctions spanning an entire eigenspace, with Gegenbauer polynomials (see, e.g., Borovitskiy et al. (2020, Appendix B)). Another situation like this happens on compact Lie groups. Here, each eigenfunction  $f_n$  can be shown to arise from a unitary irreducible representation of the group. The functions  $G_l$  associated to the sets of eigenfunctions corresponding to the same unitary irreducible representation turn out to be proportional to the character of said representation, something that can be computed exactly and efficiently using abstract algebraic considerations (Azangulov et al., 2022). Importantly, in this case,  $G_l$  need not correspond to an entire eigenspace.

GeometricKernels takes advantage of this: we introduce the concept of a *level*, which groups together multiple functions  $f_n$  whose associated eigenvalues  $\lambda_n$  are the same, but does not necessarily include all eigenfunctions of this kind. More specifically, levels correspond to the indices  $l$  and the corresponding sets of eigenfunctions  $\{f_{l,s}\}_{s=1}^{d_l}$ , as in (11) above, depending on the context.

For non-compact Riemannian manifolds—which, in the terminology of our package, are not discrete spectrum spaces—the machinery is more akin to the traditional random Fourier features (Rahimi and Recht, 2008): in this case, the kernels are represented by intractable integrals that can be approximated by Monte Carlo techniques. Importantly, these approximations—and, indeed, all others available the library—are guaranteed to be positive semi-definite.

---

1. We do not require  $\{f_{l,s}\}_{s=1}^{d_l}$  to span the whole eigenspace corresponding to a certain eigenvalue in general.

## A.6 Future Work

There are a number of geometric settings which admit kernels that can be computed using methods similar to the ones we consider. This includes spaces of graphs (Borovitskiy et al., 2023) and other combinatorial spaces (Doumont et al., 2025), vector fields on manifolds (Hutchinson et al., 2021; Robert-Nicoud et al., 2024), metric graphs (Bolin et al., 2024), general cellular complexes (Alain et al., 2024), and compact homogeneous spaces (Azangulov et al., 2022). In principle, one could consider kernels for implicit geometry (Fichera et al., 2023; Peach et al., 2024; Dunson et al., 2022; Niu et al., 2023; He et al., 2025), though supporting this in a backend-independent manner would be considerably more involved. Incorporating these new settings could increase the package’s generality and support new use cases.

## Appendix B. Brief Historical Overview of Kernels on Structured Domains

The development of kernels targeting geometric data began mostly in the the 2000s, during a time when kernel methods such as support vector machines were particularly prominent in machine learning. The types of data considered include strings (Lodhi et al., 2002; Leslie et al., 2002), points on statistical manifolds (Lafferty et al., 2005), or even data represented as terms in a formal logic (Gärtner et al., 2004). The most popular setting within this broad area (Ghosh et al., 2018), by a large margin, are kernels *between* graphs (Thomas et al., 2003; Kashima et al., 2003) and kernels *on* graphs—that is, kernels between graph nodes (Kondor and Lafferty, 2002). The interest towards the former extended into the 2010s (Vishwanathan et al., 2010; Yanardag and Vishwanathan, 2015; Kondor and Pan, 2016) and led to creation of such important tools as Weisfeiler–Lehman kernels (Shervashidze et al., 2011) closely connected with graph neural networks (Xu et al., 2018).

The aforementioned graph kernels have been used in many applications. For example, they have been used to predict mutagenicity, toxicity and anti-cancer activity for small molecules (Swamidass et al., 2005), to predict protein function (Borgwardt et al., 2005), and disease outcome from protein-protein interaction networks (Borgwardt et al., 2007). Other structured kernels have been used in neuroimaging applications (Jie et al., 2013), analyzing fMRI data (Takerkart et al., 2014), and neural architecture search (Ru et al., 2020). Many other applications can be found in the review by Kriege et al. (2020) and books by Borgwardt et al. (2020) and Gärtner (2008).

Kernels, while constrained to be positive semi-definite, are otherwise a very flexible model class. Most kernel design research, historically, has focused on maximizing predictive performance in specialized settings, and not on effective uncertainty quantification—a key aspect in Gaussian process regression and Bayesian optimization. Today, this has changed: due to difficulties with producing meaningful uncertainty using other techniques such as neural networks, designing kernels which produce good uncertainty behavior in practice is a prominent research question.

One line of work, which focuses on good uncertainty quantification (Borovitskiy et al., 2020, 2021; Azangulov et al., 2022, 2023; Yang et al., 2024; Wyrwal et al., 2025), involves generalizing kernel classes with known good uncertainty behavior in the Euclidean setting, such as the squared exponential (also known as heat, RBF, Gaussian, diffusion) and Matérn kernel classes, to geometric settings. This class is becoming popular in modern geometric Gaussian process applications (Coveney et al., 2022; Jaquier et al., 2022, 2024) and comes with an appealing set of theoretical guarantees (Li et al., 2023; Rosa et al., 2023).<sup>2</sup> Thus we focus on this class of kernels for our package. Note that, within it, heat kernels have been classically and extensively studied within the mathematics community (Grigoryan, 2009), with their variants on certain spaces studied in machine learning, for instance, by Kondor and Lafferty (2002), Lafferty et al. (2005), and Gao et al. (2019).

2. Notably, the manifold Fourier features and their close relatives associated to Euclidean boundary value problems find impressive applications as general inference techniques (Dutordoir et al., 2020; Solin and Särkkä, 2020; Riutort-Mayol et al., 2023). Porcu et al. (2024) review diverse applications of Matérn kernels throughout various disciplines.

## Appendix C. Multi-backend example

Here, we illustrate an example of how the example of Section 3 can be adapted to work using all backends. Note that no actual GEOMETRICKERNELS code changes, except for an import statement.

```
import numpy as np
import torch
import geometric_kernels.torch
from geometric_kernels.spaces import Hypersphere
from geometric_kernels.kernels import
↳ MaternGeometricKernel

hypersphere = Hypersphere(dim=2)
kernel = MaternGeometricKernel(hypersphere)

params = kernel.init_params()
params["nu"] = torch.tensor([5/2])
params["lengthscale"] = torch.tensor([1.])

xs = torch.tensor([[0., 0., 1.],
                   [0., 1., 0.],
                   [1., 0., 0.]])
Kxx = kernel.K(params, xs)

print(np.around(Kxx.detach().cpu().numpy(), 2))

# [[1.  0.36 0.36]
# [0.36 1.  0.36]
# [0.36 0.36 1.  ]]
```

(a) PyTorch

```
import numpy as np
import jax.numpy as jnp
import geometric_kernels.jax
from geometric_kernels.spaces import Hypersphere
from geometric_kernels.kernels import
↳ MaternGeometricKernel

hypersphere = Hypersphere(dim=2)
kernel = MaternGeometricKernel(hypersphere)

params = kernel.init_params()
params["nu"] = jnp.array([5/2])
params["lengthscale"] = jnp.array([1.])

xs = jnp.array([[0., 0., 1.],
                [0., 1., 0.],
                [1., 0., 0.]])
Kxx = kernel.K(params, xs)

print(np.around(np.asarray(Kxx), 2))

# [[1.  0.36 0.36]
# [0.36 1.  0.36]
# [0.36 0.36 1.  ]]
```

(b) JAX

```
import numpy as np
import tensorflow as tf
import geometric_kernels.tensorflow
from geometric_kernels.spaces import Hypersphere
from geometric_kernels.kernels import
↳ MaternGeometricKernel

hypersphere = Hypersphere(dim=2)
kernel = MaternGeometricKernel(hypersphere)

params = kernel.init_params()
params["nu"] = tf.convert_to_tensor([5/2])
params["lengthscale"] = tf.convert_to_tensor([1.])

xs = tf.convert_to_tensor([[0., 0., 1.],
                           [0., 1., 0.],
                           [1., 0., 0.]])
Kxx = kernel.K(params, xs)

print(np.around(Kxx.numpy(), 2))

# [[1.  0.36 0.36]
# [0.36 1.  0.36]
# [0.36 0.36 1.  ]]
```

(c) TensorFlow

```
import numpy as np
# no backend other than numpy is needed
import geometric_kernels
from geometric_kernels.spaces import Hypersphere
from geometric_kernels.kernels import
↳ MaternGeometricKernel

hypersphere = Hypersphere(dim=2)
kernel = MaternGeometricKernel(hypersphere)

params = kernel.init_params()
params["nu"] = np.array([5/2])
params["lengthscale"] = np.array([1.])

xs = np.array([[0., 0., 1.],
               [0., 1., 0.],
               [1., 0., 0.]])
Kxx = kernel.K(params, xs)

print(np.around(kernel.K(params, xs), 2))

# [[1.  0.36 0.36]
# [0.36 1.  0.36]
# [0.36 0.36 1.  ]]
```

(d) NumPy