



HAL
open science

Geomstats: A Python Package for Riemannian Geometry in Machine Learning

Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou,
Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti,
et al.

► **To cite this version:**

Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, et al.. Geomstats: A Python Package for Riemannian Geometry in Machine Learning. Journal of Machine Learning Research, 2020, 21 (223), pp.1-9. hal-02536154v1

HAL Id: hal-02536154

<https://inria.hal.science/hal-02536154v1>

Submitted on 8 Apr 2020 (v1), last revised 21 Dec 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Geomstats: A Python Package for Riemannian Geometry in Machine Learning

Nina Miolane
Alice Le Brigant
Johan Mathe
Benjamin Hou
Nicolas Guigui
Yann Thanwerdas
Stefan Heyder
Olivier Peltre
Niklas Koep
Hadi Zaatiti
Hatem Hajri
Yann Cabanes
Thomas Gerald
Paul Chauchat
Christian Shewmake
Bernhard Kainz
Claire Donnat
Susan Holmes
Xavier Pennec

NMIOLANE@STANFORD.EDU
 ALICE.LE-BRIGANT@UNIV-PARIS1.FR
 JOHAN@FROGLABS.AI
 BENJAMIN.HOU11@IMPERIAL.AC.UK
 NICOLAS.GUIGUI@INRIA.FR
 YANN.THANWERDAS@INRIA.FR
 STEFAN.HEYDER@TU-ILMENAU.DE
 OPELTRE@GMAIL.COM
 NIKLAS.KOEP@GMAIL.COM
 HADI.ZAATITI@IRT-SYSTEMX.FR
 HATEM.HAJRI@IRT-SYSTEMX.FR
 YANN.CABANES@GMAIL.COM
 THOMAS.GERALD@LIP6.FR
 PCHAUCHAT@GMAIL.COM
 CSHEWMAKE2@GMAIL.COM
 B.KAINZ@IMPERIAL.AC.UK
 CDONNAT@STANFORD.EDU
 SUSAN@STAT.STANFORD.EDU
 XAVIER.PENNEC@INRIA.FR

Editor: Francis Bach, David Blei and Bernhard Schölkopf

Abstract

We introduce GEOMSTATS, an open-source Python toolbox for computations and statistics on nonlinear manifolds, such as hyperbolic spaces, spaces of symmetric positive definite matrices, Lie groups of transformations, and many more. We provide object-oriented and extensively unit-tested implementations. Among others, manifolds come equipped with families of Riemannian metrics, with associated exponential and logarithmic maps, geodesics and parallel transport. Statistics and learning algorithms provide methods for estimation, clustering and dimension reduction on manifolds. All associated operations are vectorized for batch computation and provide support for different execution backends, namely NumPy, PyTorch and TensorFlow, enabling GPU acceleration. This paper presents the package, compares it with related libraries and provides relevant code examples. We show that GEOMSTATS provides reliable building blocks to foster research in differential geometry and statistics, and to democratize the use of Riemannian geometry in machine learning applications. The source code is freely available under the MIT license at geomstats.ai.

1. Introduction

Data on manifolds naturally arise in different fields. **Hyperspheres** model directional data in molecular and protein biology (Kent and Hamelryck, 2005), and are used through Principal Nested Spheres to model some aspects of 3D shapes (Jung et al., 2012; Hong et al., 2016). Computations on **hyperbolic spaces** arise for impedance density estimation (Huckemann et al., 2010), geometric network comparison (Asta and Shalizi, 2014) or the analysis of reflection coefficients extracted from a radar signal (Chevallier et al., 2015). **Symmetric positive definite (SPD) matrices** characterize data from diffusion tensor imaging (DTI) (Yuan et al., 2012) and functional magnetic resonance

imaging (fMRI) (Sporns et al., 2005). Covariance matrices, which are also SPD matrices, appear in many fields like automatic speech recognition (ASR) systems (Shinohara et al., 2010), image and video descriptors (Harandi et al., 2014), or air traffic complexity representation (Le Brigant and Puechmorel, 2018). The **Lie groups of transformations** $SO(3)$ and $SE(3)$ appear naturally when dealing with articulated objects like the human spine (Arsigny, 2006; Boisvert et al., 2006), or the pose of a camera (Kendall and Cipolla, 2017; Hou et al., 2018). **Stiefel manifolds** are used to process video action data or to analyze two vector-cardiograms in (Chakraborty and Vemuri, 2019). **Grassmannians** appear in computer vision to perform video-based face recognition and shape recognition (Chellappa, 2008). Statistics on **landmark spaces** are used in anthropology and many applied fields to describe biological shapes (Richtsmeier and Cheverud, 1992). More generally, Kendall shape spaces gave rise to a very important literature in shape statistics (Dryden and Mardia, 1998). A variety of applications use infinite-dimensional Riemannian geometry tools to study the shapes of **discretized curves**, sampled e.g. from closed two or three-dimensional curves defining the contours of organs in computational anatomy (Younes, 2012). We also find open curves embedded in any of the previously mentioned manifolds that describe the temporal evolution of physical phenomena, for example in a Lie group (Celledoni et al., 2015) or in the hyperbolic plane (Le Brigant, 2017).

Yet, the adoption of differential geometry computations has been inhibited by the lack of a reference implementation. Code sequences are often custom-tailored for specific problems and are not easily reused. Some Python packages do exist, but focus on optimization: `Pymanopt` (Townsend et al., 2016), `Geoopt` (Bécigneul and Ganea, 2018; Kochurov, 2019), and `McTorch` (Meghwanshi et al., 2018). Others are dedicated to a single manifold: `PyRiemann` on SPD matrices (Barachant, 2015), `PyQuaternion` on 3D rotations (Wynn, 2014), `PyGeometry` on spheres, toruses, 3D rotations and translations (Censi, 2012). Lastly, others lack unit-tests and continuous integration: `Theano-Geometry` (Kühnel and Sommer, 2017). There is a need for an open-source low-level implementation of differential geometry, and associated learning algorithms for manifold-valued data.

We present `GEOMSTATS`, an open-source Python package of computations and statistics on non-linear manifolds. `GEOMSTATS` has three main objectives: (i) foster research in differential geometry and geometric statistics by providing low-level code to get intuition or test a theorem and a platform to share algorithms; (ii) democratize the use of geometric statistics by implementing user-friendly geometric learning algorithms using Scikit-Learn API; and (iii) provide educational support to learn “hands-on” differential geometry and geometric statistics, through its examples and visualizations.

2. Implementation overview

The package `geomstats` is organized into two main modules: `geometry` and `learning`. **The module `geometry` implements concepts in Riemannian geometry** with an object-oriented approach. Manifolds mentioned in the introduction are available as classes that inherit from the base class `Manifold`. The base class `RiemannianMetric` provides methods such as the geodesic distance between two points, the exponential and logarithm maps at a base point, etc. As examples, `HyperbolicMetric`, `StiefelCanonicalMetric`, the Lie groups’ `InvariantMetrics`, or the curves’ `L2Metric` and `SRVMetric` (Srivastava et al., 2011) inherit from `RiemannianMetric`. Going beyond Riemannian geometry, the class `Connection` implements affine connections. The implementation uses automatic differentiation with `autograd` to allow computations on manifolds when closed-form formulae do not exist. The API of the module `geometry` uses terms from differential geometry, to foster contributions of researchers from this field. **The module `learning` implements statistics and learning algorithms** for data on manifolds. The code is object-oriented and classes inherit from Scikit-Learn base classes and mixin: `BaseEstimator`, `ClassifierMixin`, `RegressorMixin`, etc. This module provides classes such as `FréchetMean`, `KMeans`, `TangentPCA`, for implementations of Fréchet mean estimators (Fréchet, 1948), K -means and principal component analysis (PCA)

designed for manifold data. The API of the module `learning` follows Scikit-Learn’s API, being therefore user-friendly to machine learning researchers and engineers.

The code follows international standards for readability and ease of collaboration, is vectorized for batch computations, undergoes unit-testing with continuous integration using Travis, relies on TensorFlow and PyTorch backends for GPU acceleration, and is partially ported to R. The package comes with a `visualization` module to provide intuitions on differential geometry, see Figure 1. The GitHub repository at github.com/geomstats/geomstats offers a convenient way to ask for help or request features by raising issues. The website geomstats.github.io provides documentation for users and important guidelines for those wishing to contribute to the project.




Figure 1: Left: Minimization of a scalar field on the sphere \mathbb{S}^2 using Riemannian gradient descent. Middle: Regular geodesic grid on the hyperbolic space \mathbb{H}^2 in Poincaré disk representation. Right: Geodesic on the Lie group $SE(3)$ for the canonical left-invariant metric. These and more examples are available at geomstats.ai.

3. Comparison and interaction with existing packages

This section compares the package `geomstats` with related Python implementations on differential geometry and learning. Table 1 compares the geometric operations and Table 2 compares the engineering infrastructures.

The library `TheanoGeometry` (Kühnel and Sommer, 2017) is the most closely related to `GEOMSTATS` and provides nonlinear statistics and stochastic equations on Riemannian manifolds. The differential geometric tensors are computed with automatic differentiation. However, this library does not provide statistical learning algorithms and lacks engineering maintenance. Then, there are several packages that focus on optimization on Riemannian manifolds. `Pymanopt` (Townsend et al., 2016) computes gradients and Hessian-vector products on Riemannian manifolds with automatic differentiation, and provides the following solvers: steepest descent, conjugate gradients, the Nelder-Mead algorithm, particle swarm optimization and the Riemannian trust regions. `Geoopt` (Kochurov, 2019) focuses on stochastic adaptive optimization on Riemannian manifolds, for machine learning problems. The library provides stochastic solvers, stochastic gradient descent and Adam, as well as the following samplers: Stochastic Gradient Langevin Dynamics, Hamiltonian Monte-Carlo, Stochastic Gradient Hamiltonian Monte-Carlo. Lastly, `McTorch` (Meghwanshi et al., 2018) provides optimization on Riemannian manifold for deep learning by adding a “Manifold” parameter to PyTorch’s network layers and optimizers. The library provides the following solvers: stochastic gradient descent, AdaGrad and conjugate gradients. As these libraries focus on optimization, they substitute potentially computationally expensive operations by practical proxies, for example, by replacing exponential maps by so-called retractions. However, they are less modular than `GEOMSTATS` in terms

	Manifolds	Geometry
Pymanopt	Euclidean manifold, symmetric matrices, sphere, complex circle, $SO(n)$, Stiefel, Grassmannian, oblique manifold, $SPD(n)$, elliptope, fixed-rank PSD matrices	Exponential and logarithmic maps, retraction, vector transport, <code>egrad2rgrad</code> , <code>ehess2rhess</code> , inner product, distance, norm
Geoopt	Euclidean manifold, sphere, Stiefel, Poincaré ball	Same as Pymanopt
McTorch	Stiefel, $SPD(n)$	Same as Pymanopt
TheanoGeometry	Sphere, ellipsoid, $SPD(n)$, Landmarks, $GL(n)$, $SO(n)$, $SE(n)$	Inner product, exponential and logarithmic maps, parallel transport, Christoffel symbols, Riemann, Ricci and scalar curvature, geodesics, Fréchet mean
Geomstats	Euclidean manifold, Minkowski and hyperbolic space, sphere, $SO(n)$, $SE(n)$, $GL(n)$, Stiefel, Grassmannian, $SPD(n)$, discretized curves, Landmarks	Levi-Civita connection, parallel transport, exponential and logarithmic maps, inner product, distance, norm, geodesics, invariant metrics

Table 1: Comparison of libraries in terms of geometric operations

	Backends	Continuous integration (CI) and coverage
Pymanopt	Autograd, PyTorch, TensorFlow, Theano	CI, coverage 85%
Geoopt	PyTorch	Not disclosed
McTorch	PyTorch	CI, coverage 84%
TheanoGeometry	Theano	No CI, no unit tests
Geomstats	NumPy, PyTorch, TensorFlow	CI, coverage 89% (NumPy), 45% (TensorFlow), 47% (PyTorch)

Table 2: Comparison of code infrastructure

of the Riemannian geometry. For example, each manifold comes with a single Riemannian metric, in contrast to GEOMSTATS where families of Riemannian metrics are implemented. Furthermore, they do not provide statistical learning algorithms.

The optimization libraries are complementary to GEOMSTATS and interact easily with it. GEOMSTATS provides low-level implementations of Riemannian geometry that can be used to define optimization costs. In turn, an optimization library can provide an efficient solver to use within the implementation of GEOMSTATS’ learning algorithms. An example of such interactions, between Pymanopt and GEOMSTATS can be found in GEOMSTATS’ `examples` folder.

4. Usage: Examples of Learning on Riemannian manifolds

The folder `examples` provides sample codes to get started with GEOMSTATS. More involved GEOMSTATS applications can be found at: github.com/geomstats/applications. This section illustrates the use of GEOMSTATS to generalize learning algorithms to Riemannian manifolds. The following code snippet illustrates the use of K -means on the hypersphere.

The following code snippet shows the use of tangent PCA on the 3D rotations.

All geometric computations are performed behind the scenes. The user only needs a high-level understanding of Riemannian geometry. Each algorithm can be used with any of the manifolds and metric implemented in the package.

Acknowledgments

This work is partially supported by the National Science Foundation, grant NSF DMS RTG 1501767, the Inria-Stanford associated team GeomStats, and the European Research Council (ERC) under

the European Union’s Horizon 2020 research and innovation program (grant agreement G-Statistics No. 786854).

References

- Vincent Arsigny. *Processing Data in Lie Groups: An Algebraic Approach. Application to Non-Linear Registration and Diffusion Tensor MRI*. PhD thesis, École polytechnique, 11 2006.
- Dena Asta and Cosma Rohilla Shalizi. Geometric Network Comparison. *Journal of Machine Learning Research*, 2014. ISSN 1939-1374. doi: 10.1109/PES.2006.1709566. URL <http://arxiv.org/abs/1411.1350>.
- Alexandre Barachant. PyRiemann: Python package for covariance matrices manipulation and Biosignal classification with application in Brain Computer interface, 2015. URL <https://github.com/alexandrebarachant/pyRiemann>.
- Gary Bécigneul and Octavian-Eugen Ganea. Riemannian Adaptive Optimization Methods. In *Proc. of ICLR 2019*, pages 1–16, 2018. URL <http://arxiv.org/abs/1810.00760>.
- Jonathan Boisvert, Xavier Pennec, Hubert Labelle, Farida Cheriet, and Nicholas Ayache. Principal spine shape deformation modes using riemannian geometry and articulated models. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 4069 LNCS:346–355, 2006. ISSN 16113349. doi: 10.1007/11789239{_}36.
- Elena Celledoni, Markus Eslietzichler, and Alexander Schmeding. Shape analysis on Lie groups with applications in computer animation. *arXiv preprint arXiv:1506.00783*, 2015.
- Andrea Censi. PyGeometry: Library for handling various differentiable manifolds., 2012. URL <https://github.com/AndreaCensi/geometry>.
- Rudrasis Chakraborty and Baba C Vemuri. Statistics on the (compact) Stiefel manifold : Theory and Applications. *Annals of Statistics*, 47(1):415–438, 2019.
- Pavan Turaga ; Ashok Veeraraghavan ; Rama Chellappa. Statistical analysis on Stiefel and Grassmann Manifolds with applications in Computer Vision Pavan Turaga , Ashok Veeraraghavan and Rama Chellappa Center for Automation Research University of Maryland , College Park. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- Emmanuel Chevallier, Frédéric Barbaresco, and Jesus Angulo. Probability density estimation on the hyperbolic space applied to radar processing. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9389:753–761, 2015. ISSN 16113349.
- I L Dryden and K V Mardia. *Statistical shape analysis*. John Wiley & Sons, New York, 1998.
- Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l’institut Henri Poincaré*, 10(4):215–310, 1948.
- Mehrtash Tafazzoli Harandi, Richard I Hartley, Brian C Lovell, and Conrad Sanderson. Sparse Coding on Symmetric Positive Definite Manifolds using Bregman Divergences. *CoRR*, abs/1409.0, 2014. URL <http://arxiv.org/abs/1409.0083>.
- Junpyo Hong, Jared Vicory, Jörn Schulz, Martin Styner, J S Marron, and StephenM Pizer. Non-Euclidean Classification of Medically Imaged Objects via s-reps. *Med Image Anal*, 31:37–45, 2016.
- Benjamin Hou, Nina Miolane, Bishesh Khanal, Mathew C. H. Lee, Amir Alansary, Steven McDonagh, Jo V. Hajnal, Daniel Rueckert, Ben Glocker, and Bernhard Kainz. Computing CNN loss and gradients for pose estimation with Riemannian geometry. In *MICCAI*, volume 11070 LNCS, 2018. ISBN 9783030009274.
- Stephan Huckemann, Peter Kim, Ja Yong Koo, and Axel Munk. Möbius deconvolution on the hyperbolic plane with application to impedance density estimation. *Annals of Statistics*, 38(4):2465–2498, 2010. ISSN 00905364. doi: 10.1214/09-AOS783.
- Sungkyu Jung, Ian L. Dryden, and J. S. Marron. Analysis of principal nested spheres. *Biometrika*, 99(3):551–568, 2012. ISSN 00063444. doi: 10.1093/biomet/ass022.
- Alex Kendall and Roberto Cipolla. Geometric Loss Functions for Camera Pose Regression with Deep Learning. In *CVPR*, 2017. doi: 10.1109/CVPR.2017.694.
- J.T. JT Kent and Thomas Hamelryck. Using the Fisher-Bingham distribution in stochastic models for protein structure. *Quantitative Biology, Shape Analysis, and Wavelets*, 24(1):57–60, 2005. URL <http://www.amsta.leeds.ac.uk/statistics/workshop/lasr2005/Proceedings/kent.pdf>.
- Maxim Kochurov. Geopt: Riemannian Adaptive Optimization Methods with pytorch optim, 2019.
- Line Kühnel and Stefan Sommer. Computational Anatomy in Theano. *CoRR*, 2017.
- Alice Le Brigant. Computing distances and geodesics between manifold-valued curves in the SRV framework. *Journal of Geometric Mechanics (forthcoming)*, *arXiv preprint arXiv:1601.02358*, 2017.
- Alice Le Brigant and Stéphane Puechmorel. Optimal Riemannian quantization with an application to air traffic analysis. *arXiv preprint arXiv:1806.07605*, 2018.
- Mayank Meghwanshi, Pratik Jawanpuria, Anoop Kunchukuttan, Hiroyuki Kasai, and Bamdev Mishra. McTorch, a manifold optimization library for deep learning, 2018. URL <http://arxiv.org/abs/1810.01811>.

- Joan T Richtsmeier and James M Cheverud. Advances in Anthropological Morphometrics. *Annu. Rev. Anthropology*, 21(5):283–305, 1992.
- Yusuke Shinohara, Takashi Masuko, and Masami Akamine. Covariance clustering on Riemannian manifolds for acoustic model compression. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4326–4329, 3 2010.
- Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: A structural description of the human brain. *PLoS Computational Biology*, 1(4):0245–0251, 2005. ISSN 15537358. doi: 10.1371/journal.pcbi.0010042.
- Anuj Srivastava, Eric Klassen, Shantanu H Joshi, and Ian H Jermyn. Shape analysis of elastic curves in euclidean spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(7):1415–1428, 2011.
- James Townsend, Niklas Koep, and Sebastian Weichwald. Pymanopt: A Python Toolbox for Manifold Optimization using Automatic Differentiation. *arXiv preprint arXiv:1603.03236*, 17:1–4, 2016. ISSN 15337928. URL <http://arxiv.org/abs/1603.03236>.
- Kieran Wynn. PyQuaternions: A fully featured, pythonic library for representing and using quaternions, 2014. URL <https://github.com/KieranWynn/pyquaternion>.
- Laurent Younes. Spaces and manifolds of shapes in computer vision: An overview. *Image and Vision Computing*, 30(6-7):389–397, 2012.
- Ying Yuan, Hongtu Zhu, Weili Lin, and J S Marron. Local polynomial regression for symmetric positive definite matrices. *Journal of the Royal Statistical Society Series B*, 74(4):697–719, 2012. URL <https://econpapers.repec.org/RePEc:bla:jorssb:v:74:y:2012:i:4:p:697-719>.