

Pymanopt: A Python Toolbox for Optimization on Manifolds using Automatic Differentiation

James Townsend

University College London, London, UK

JAMES.TOWNSEND.14@UCL.AC.UK

Niklas Koep

RWTH Aachen University, Germany

NIKLAS.KOEP@RWTH-AACHEN.DE

Sebastian Weichwald

Max Planck Institute for Intelligent Systems, Tübingen, Germany

SWEICHWALD@TUE.MPG.DE

Editor: Antti Honkela

Abstract

Optimization on manifolds is a class of methods for optimization of an objective function, subject to constraints which are smooth, in the sense that the set of points which satisfy the constraints admits the structure of a differentiable manifold. While many optimization problems are of the described form, technicalities of differential geometry and the laborious calculation of derivatives pose a significant barrier for experimenting with these methods.

We introduce PYMANOPT (available at pymanopt.github.io), a toolbox for optimization on manifolds, implemented in Python, that—similarly to the Manopt¹ Matlab toolbox—implements several manifold geometries and optimization algorithms. Moreover, we lower the barriers to users further by using automated differentiation² for calculating derivative information, saving users time and saving them from potential calculation and implementation errors.

Keywords: Riemannian optimization, non-convex optimization, manifold optimization, projection matrices, symmetric matrices, rotation matrices, positive definite matrices

1. Introduction

Optimization on manifolds, or Riemannian optimization, is a method for solving problems of the form

$$\min_{x \in \mathcal{M}} f(x)$$

where $f: \mathcal{M} \rightarrow \mathbb{R}$ is a (cost) function and the search space \mathcal{M} is smooth, in the sense that it admits the structure of a differentiable manifold. Although the definition of differentiable manifold is technical and abstract, many familiar sets satisfy this definition and are therefore compatible with the methods of optimization on manifolds. Examples include the *sphere* (the set of points with unit Euclidean norm) in \mathbb{R}^n , the set of *positive definite matrices*, the set of *orthogonal matrices* as well as the set of p -dimensional subspaces of \mathbb{R}^n with $p < n$, also known as the *Grassmann* manifold.

To perform optimization, the function f needs to be defined for points on the manifold \mathcal{M} . Elements of \mathcal{M} are often represented by elements of \mathbb{R}^n or $\mathbb{R}^{m \times n}$, and f is often well defined on some or all of this “ambient” Euclidean space. If f is also differentiable, it makes sense for an optimization algorithm to use the derivatives of f and adapt them to the manifold setting in order to iteratively refine solutions based on curvature information. This is one of the key aspects of Manopt (Boumal et al., 2014), which allows the user to pass a function’s gradient and Hessian to state of the art

1. Manopt is available at manopt.org and was introduced in Boumal et al. (2014).

2. We use the term *automated differentiation* to refer to the automatic calculation of derivatives, whether using the method commonly known as *automatic differentiation*, as implemented by Autograd (Maclaurin et al., 2015) and TensorFlow (Abadi et al., 2015), or *symbolic differentiation* as implemented by Theano (Al-Rfou et al., 2016).

solvers which exploit this information to optimize over the manifold \mathcal{M} . However, working out and implementing gradients and higher order derivatives is a laborious and error prone task, particularly when the objective function acts on matrices or higher rank tensors. Manopt’s state of the art Riemannian Trust Regions solver, described in Absil et al. (2007), requires second order directional derivatives (or a numerical approximation thereof), which are particularly challenging to work out for the average user, and more error prone and tedious even for an experienced mathematician.

It is these difficulties which we seek to address with this toolbox. PYMANOPT supports a variety of modern Python libraries for automated differentiation of cost functions acting on vectors, matrices or higher rank tensors. Combining optimization on manifolds and automated differentiation enables a convenient workflow for rapid prototyping that was previously unavailable to practitioners. All that is required of the user is to instantiate a manifold, define a cost function, and choose one of PYMANOPT’s solvers. This means that the Riemannian Trust Regions solver in PYMANOPT is just as easy to use as one of the derivative-free or first order methods.

2. The Potential of Optimization on Manifolds and Pymanopt Use Cases

Much of the theory of how to adapt Euclidean optimization algorithms to (matrix) manifolds can be found in Smith (1994); Edelman et al. (1998); Absil et al. (2008). The approach of optimization on manifolds is superior to performing free (Euclidean) optimization and projecting the parameters back onto the search space after each iteration (as in the projected gradient descent method), and has been shown to outperform standard algorithms for a number of problems.

Hosseini and Sra (2015) demonstrate this advantage for a well-known problem in machine learning, namely inferring the maximum likelihood parameters of a mixture of Gaussian (MoG) model. Their alternative to the traditional expectation maximization (EM) algorithm uses optimization over a product manifold of positive definite (covariance) matrices. Rather than optimizing the likelihood function directly, they optimize a reparameterized version which shares the same local optima. The proposed method, which is on par with EM and shows less variability in running times, is a striking example why we think a toolbox like PYMANOPT, which allows the user to readily experiment with and solve problems involving optimization on manifolds, can accelerate and pave the way for improved machine learning algorithms.³

Further successful applications of optimization on manifolds include matrix completion tasks (Vandereycken, 2013; Boumal and Absil, 2015), robust PCA (Podosinnikova et al., 2014), dimension reduction for independent component analysis (ICA) (Theis et al., 2009), kernel ICA (Shen et al., 2007) and similarity learning (Shalit et al., 2012).

Many more applications to machine learning and other fields exist. While a full survey on the usefulness of these methods is well beyond the scope of this manuscript, we highlight that at the time of writing, a search for the term “manifold optimization” on the IEEE Xplore Digital Library lists 1065 results; the Manopt toolbox itself is referenced in 90 papers indexed by Google Scholar.

3. Implementation

Our toolbox is written in Python and uses NumPy and SciPy for computation and linear algebra operations. Currently PYMANOPT is compatible with cost functions defined using Autograd (Maclaurin et al., 2015), Theano (Al-Rfou et al., 2016) or TensorFlow (Abadi et al., 2015). PYMANOPT itself and all the required software is open source, with no dependence on proprietary software.

To calculate derivatives, Theano uses symbolic differentiation, combined with rule-based optimizations, while both Autograd and TensorFlow use reverse-mode automatic differentiation. For a discussion of the distinctions between the two approaches and an overview of automatic differentiation in the context of machine learning, we refer the reader to Baydin et al. (2015).

3. A quick example implementation for inferring MoG parameters is available at pymanopt.github.io/MoG.html.

Much of the structure of PYMANOPT is based on that of the Manopt Matlab toolbox. For this early release, we have implemented all of the solvers and a number of the manifolds found in Manopt, and plan to implement more, based on the needs of users. The codebase is structured in a modular way and thoroughly commented to make extension to further solvers, manifolds, or backends for automated differentiation as straightforward as possible. Both a user and developer documentation are available. The GitHub repository at github.com/pymanopt/pymanopt offers a convenient way to ask for help or request features by raising an issue, and contains guidelines for those wishing to contribute to the project.

4. Usage: A Simple Instructive Example

All automated differentiation in PYMANOPT is performed behind the scenes so that the amount of setup code required by the user is minimal. Usually only the following steps are required:

- (a) Instantiation of a manifold \mathcal{M}
- (b) Definition of a cost function $f: \mathcal{M} \rightarrow \mathbb{R}$
- (c) Instantiation of a PYMANOPT solver

We briefly demonstrate the ease of use with a simple example. Consider the problem of finding an $n \times n$ positive semi-definite (PSD) matrix S of rank $k < n$ that best approximates a given $n \times n$ (symmetric) matrix A , where closeness between A and its low-rank PSD approximation S is measured by the following loss function

$$L_\delta(S, A) \triangleq \sum_{i=1}^n \sum_{j=1}^n H_\delta(s_{i,j} - a_{i,j})$$

for some $\delta > 0$ and $H_\delta(x) \triangleq \sqrt{x^2 + \delta^2} - \delta$ the pseudo-Huber loss function. This loss function is robust against outliers as $H_\delta(x)$ approximates $|x| - \delta$ for large values of x while being approximately quadratic for small values of x (Huber, 1964).

This can be formulated as an optimization problem on the manifold of PSD matrices:

$$\min_{S \in \mathcal{PSD}_k^n} L_\delta(S, A)$$

where $\mathcal{PSD}_k^n \triangleq \{M \in \mathbb{R}^{n \times n} : M \succeq 0, \text{rank}(M) = k\}$. This task is easily solved using PYMANOPT:

```

from pymanopt.manifolds import PSDFixedRank
import autograd.numpy as np
from pymanopt import Problem
from pymanopt.solvers import TrustRegions

# Let A be a (n x n) matrix to be approximated

# (a) Instantiation of a manifold
# points on the manifold are parameterized as YY^T
# where Y is a matrix of size n x k
manifold = PSDFixedRank(A.shape[0], k)
# (b) Definition of a cost function (here using autograd.numpy)
def cost(Y):
    S = np.dot(Y, Y.T)
    delta = .5
    return np.sum(np.sqrt((S - A)**2 + delta**2) - delta)

```

```

# define the Pymanopt problem
problem = Problem(manifold=manifold, cost=cost)
# (c) Instantiation of a Pymanopt solver
solver = TrustRegions()

# let Pymanopt do the rest
Y = solver.solve(problem)
S = np.dot(Y, Y.T)

```

The examples folder within the PYMANOPT toolbox holds further instructive examples, such as performing inference in mixture of Gaussian models using optimization on manifolds instead of the expectation maximization algorithm. Also see the examples section on pymanopt.github.io.

5. Conclusion

PYMANOPT enables the user to experiment with different state of the art solvers for optimization problems on manifolds, like the Riemannian Trust Regions solver, without any extra effort. Experimenting with different cost functions, for example by changing the pseudo-Huber loss $L_\delta(S, A)$ in the code above to the Frobenius norm $\|S - A\|_F$, a p -norm $\|S - A\|_p$, or some more complex function, requires just a small change in the definition of the cost function. For problems of greater complexity, PYMANOPT offers a significant advantage over toolboxes that require manual differentiation by enabling users to run a series of related experiments without returning to pen and paper each time to work out derivatives. Gradients and Hessians only need to be derived if they are required for other analysis of a problem. We believe that these advantages, coupled with the potential for extending PYMANOPT to large-scale applications using TensorFlow, could lead to significant progress in applications of optimization on manifolds.

Acknowledgments

We would like to thank the developers of the Manopt Matlab toolbox, in particular Nicolas Boumal and Pierre-Antoine Absil, for developing Manopt, and for the generous help and advice they have given. We would also like to thank Heiko Strathmann for his thoughtful advice as well as the anonymous reviewers for their constructive feedback and idea for a more suitable application example.

References

- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. URL <http://tensorflow.org>.
- P.-A. Absil, C.G. Baker, and K.A. Gallivan. Trust-Region Methods on Riemannian Manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, 2007.
- P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13298-3.
- R. Al-Rfou, G. Alain, A. Almahairi, C. Angermueller, D. Bahdanau, N. Ballas, F. Bastien, J. Bayer, A. Belikov, A. Belopolsky, Y. Bengio, A. Bergeron, J. Bergstra, V. Bisson, J. Blecher Snyder, N. Bouchard, N. Boulanger-Lewandowski, X. Bouthillier, A. de Brébisson, O. Breuleux, P.-L. Carrier, K. Cho, J. Chorowski, P. Christiano, T. Cooijmans, M.-A. Côté, M. Côté, A. Courville, Y.N. Dauphin, O. Delalleau, J. Demouth, G. Desjardins, S. Dieleman, L. Dinh, M. Ducoffe, V. Dumoulin, S. Ebrahimi Kahou, D. Erhan, Z. Fan, O. Firat, M. Germain, X. Glorot, I. Goodfellow, M. Graham, C. Gulcehre, P. Hamel, I. Harlouchet, J.-P. Heng, B. Hidasi, S. Honari, A. Jain, S. Jean, K. Jia, M. Korobov, V. Kulkarni, A. Lamb, P. Lamblin, E. Larsen, C. Laurent, S. Lee, S. Lefrancois, S. Lemieux, N. Léonard, Z. Lin, J. A. Livezey, C. Lorenz, J. Lowin, Q. Ma, P.-A. Manzagol, O. Mastropietro, R.T. McGibbon, R. Memisevic, B. van Merriënboer, V. Michalski, M. Mirza, A. Orlandi, C. Pal, R. Pascanu, M. Pezeshki, C. Raffel, D. Renshaw, M. Rocklin, A. Romero, M. Roth, P. Sadowski, J. Salvatier, F. Savard, J. Schlüter, J. Schulman, G. Schwartz, I.V. Serban, D. Serdyuk, S. Shabanian, É. Simon, S. Spiekermann, S.R. Subramanyam, J. Synowski, J. Tanguay, G. van Tulder, J. Turian, S. Urban, P. Vincent, F. Visin, H. de Vries, D. Warde-Farley, D.J. Webb, M. Willson, K. Xu, L. Xue, L. Yao, S. Zhang, and Y. Zhang. Theano: A Python framework for fast computation of mathematical expressions. *arXiv preprint arXiv:1605.02688*, 2016. URL <http://deeplearning.net/software/theano>.
- A.G. Baydin, B.A. Pearlmutter, A.A. Radul, and J.M. Siskind. Automatic differentiation in machine learning: a survey. *arXiv preprint arXiv:1502.05767*, 2015.
- N. Boumal and P.-A. Absil. Low-rank matrix completion via preconditioned optimization on the Grassmann manifold. *Linear Algebra and its Applications*, 475:200–239, 2015. doi: 10.1016/j.laa.2015.02.027.
- N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab Toolbox for Optimization on Manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014. URL <http://manopt.org>.
- A. Edelman, T.A. Arias, and S.T. Smith. The Geometry of Algorithms with Orthogonality Constraints. *SIAM J. Matrix Anal. & Appl.*, 20(2):303–353, 1998.
- R. Hosseini and S. Sra. Matrix Manifold Optimization for Gaussian Mixtures. In *Advances in Neural Information Processing Systems*, pages 910–918, 2015.
- P.J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- D. Maclaurin, D. Duvenaud, M. Johnson, and R.P. Adams. Autograd: Reverse-mode differentiation of native Python, 2015. URL <http://github.com/HIPS/autograd>.
- A. Podosinnikova, S. Setzer, and M. Hein. Robust PCA: Optimization of the Robust Reconstruction Error over the Stiefel Manifold. In *36th German Conference on Pattern Recognition (GCPR)*, 2014.
- U. Shalit, D. Weinshall, and G. Chechik. Online Learning in the Embedded Manifold of Low-rank Matrices. *Journal of Machine Learning Research*, 13(1):429–458, 2012.
- H. Shen, S. Jegelka, and A. Gretton. Fast Kernel ICA using an Approximate Newton Method. In *International Conference on Artificial Intelligence and Statistics*, pages 476–483, 2007.
- S.T. Smith. Optimization techniques on Riemannian manifolds. *Fields institute communications*, 3(3):113–135, 1994.
- F.J. Theis, T.P. Cason, and P.-A. Absil. Soft dimension reduction for ICA by joint diagonalization on the Stiefel manifold. In *Independent Component Analysis and Signal Separation*, pages 354–361. Springer, 2009.
- B. Vandereycken. Low-Rank Matrix Completion by Riemannian Optimization. *SIAM J. Optim.*, 23(2):1214–1236, 2013.