Riemannian Metric estimation and the problem of geometric recovery

Dominique C. Perrault-Joncas
Department of Statistics
University of Washington
Seattle, WA 98195
dcpj@stat.washington.edu

Marina Meilă
Department of Statistics
University of Washington
Seattle, WA 98195
mmp@stat.washington.edu

In recent years, manifold learning has become a popular tool for performing nonlinear dimensionality reduction. This has led to the development of a flurry of algorithms of varying degree of complexity that aim to recover manifold geometry using either local or global features of the data. Building on the Laplacian Eigenmap framework, we unify all these algorithms through the Riemannian Metric g. The Riemannian Metric allows us to compute geometric quantities (such as angle, length, or volume) on the original manifold for any coordinate system. This geometric faithfulness, which is not guaranteed for most algorithms, allows us to define geometric measurements that are indepent of the algorithm used, and hence seamlessly move from one algorithm to the next. A more significant consequence of this geometric faithfulness is that it allows for regressions, predictions, and other statistical anlyses to be conducted directly on the low-dimensional representation of the data. If the geometry of the data were not fully and accurately preserved, these analyses would, in general, not be correct.

Generally speaking, a *manifold learning* or *manifold embedding* algorithm is a method of non-linear dimension reduction. Hence, its input is a set of points $\mathcal{D} = \{p_1, \ldots p_N\}$ in $\mathbb{R}^n$, where n is typically high. The points $\mathcal{D}$ are assumed to be sampled from a manifold $\mathcal{M} \subset \mathbb{R}^n$ which is equipped with a *Riemannian metric* g. These points are mapped into vectors $\{f(p_1), \ldots f(p_N)\} \subset \mathbb{R}^m$, with $m \ll n$, where $f : \mathcal{M} \rightarrow \mathbb{R}^m$ is assumed to be an embedding of $\mathcal{M}$ into $\mathbb{R}^m$.

Defining an optimal embedding map f that retains as many features as possible from the orginal data $\mathcal{D}$ while embedding the data in the space of lowest possible dimensions is the goal of manifold learning algorithms. Unfortunately, these algorithms suffer important shortcomings, which have been documented in [3, 4]. Two of the more significant criticisms are that 1) the algorithms fail to actually recover the geometry of the manifold in many instances and 2) there is no coherent framework in which the multitude of existing algorithms can easily be compared and selected for a given application.

Much effort has been invested into finding manifold learning algorithms that "recover the geometry". In this work, we translate this vague notion in mathematical terms, equating it with the concepts of *isometry*, *isometric embedding*, and *pushforward metric*. Not surprisingly, the majority of manifold learning algorithms output embeddings that are not isometric.

In contrast to previous work in manifold learning, we take the following point of view: if recovering the geometry is important, then instead of finding a single embedding algorithm that has this property, we will provide for a way to augment any reasonable embedding algorithm with a Riemannian metric represented in the algorithm's own coordinates. This addresses the two main criticisms of manifold learning algorithms referred to above, and offers a convenient framework for moving between embeddings.

To acheive this, we set out to augment the coordinate representation f produced by any manifold learning algorithm with the information necessary for reconstructing the geometry of the data. This information is embodied in the Riemannian metric. Expressing the metric g defined on $\mathcal{M}$ on $\mathcal{N} = f(\mathcal{M})$ is formally known as obtaining the pushforward of the metric g under map f.

Pushforward Metric Let f be and embedding between Riemannian manifolds $(\mathcal{M}, g)$ and $(\mathcal{N}, h)$, then the pushforward $\varphi^* g_p$ of the metric g_p along $f^{-1} = \varphi$ is given by

$$\langle u, v \rangle_{\varphi^* g_p} \quad = \quad \left\langle df_p^\dagger u, df_p^\dagger v \right\rangle_{g_p} \ ,$$

for $u, v \in T_{f(p)}\mathcal{N}$ and where $df_p^\dagger$ denotes the Jacobian of f^{-1}.

In effect, recovering the pushforward of g under f redefines the inner product on the tangent spaces $T_p f(\mathcal{M})$ for every point $p \in f(\mathcal{M})$, so that the traditional notions of Euclidean geometry of angles, length, and volume on $f(\mathcal{M})$ carry over from $\mathcal{M}$.

To define $\varphi^* g_p$ for each point in $\mathcal{D}$, we take advantage of the fact that the Laplace-Beltrami operator $\Delta_\mathcal{M}$ can be estimated from $\mathcal{D}$ [2]. The key point here is that the operator $\Delta_\mathcal{M}$ is *coordinate free*: once estimated, it can be applied to any function on $\mathcal{M}$, including coordinates. This applies to the embedding map f, which is obviously defined on $\mathcal{M}$. From the dependency of the Laplace-Beltrami operator on g:

$$\Delta_\mathcal{M} f = \frac{1}{\sqrt{\det(g)}} \frac{\partial}{\partial x^l} \left(\sqrt{\det(g)} g^{lk} \frac{\partial}{\partial x^k} f \right) \ . \tag{1}$$

recovering $\varphi^* g_p$ follows from the following proposition:

Proposition 0.1

$$\varphi^* g^{ij} \ = \ \frac{1}{2} \Delta_\mathcal{M} \left(f_i - f_i(p)\right) \left(f_j - f_j(p)\right) |_{f_i = f_i(p), f_j = f_j(p)} \tag{2}$$

Proof To find $\varphi^* g(p)$, the pushforward of g under f at point p, we apply (1) to the coordinate products $\frac{1}{2}\left(f_i - f_i(p)\right)\left(f_j - f_j(p)\right)$ for $i, j = 1, \ldots, d$. By direct evaluation of the r.h.s. using (1), we obtain

$$g^{lk} \frac{\partial}{\partial x^l} \left(f_i - f_i(p)\right) \times \frac{\partial}{\partial x^k} \left(f_j - f_j(p)\right) |_{f_i = f_i(p), f_j = f_j(p)} \ = \ \varphi^* g^{ij}.$$

Figure 1 provides a demonstration of our method. Figure 1a displays a dome, on which we sample 2000 points uniformly, and trace a geodesic between fixed points on the manifold $\mathcal{M}$. Figure 1b displays an embedding of this sample using the eigenmap [1], and traces the same geodesic, computed by taking into account the Riemannian metric. The theoretical length of geodesic is $\pi/2$, with an estimated geodesic of 1.56 on the original data, and 1.62 on the embedding. Thus, the embedding's relative error is 3.5%, despite the fact that the it is drawn on a much smaller scale and has an entirely different shape than the original.

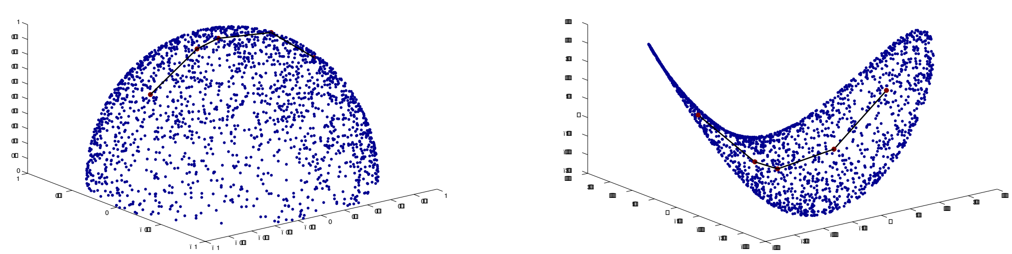

a) b)

Figure 1: a) Points sampled uniformly on dome with geodesic in black with red endpoints. b) Eigenmap of the dome along with geodesic in black with red endpoints.

References

[1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2002.

[2] M. Belkin, J. Sun, and Y. Wang. Constructing laplace operator from point clouds in rd. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 1031–1040, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.

[3] Y. Goldberg, A. Zakai, D. Kushnir, and Y. Ritov. Manifold Learning: The Price of Normalization. *Journal of Machine Learning Research*, 9:1909–1939, AUG 2008.

[4] T. Wittman. MANIfold learning matla demo. http://www.math.umn.edu/ wittman/mani/, retrived June 2010.