Manifold learning: an informal introduction

Felix Dietrich*

February 3, 2019

1 Introduction

This note provides an informal introduction into the the concepts used in data-driven manifold learning. The presentation is not meant to be rigorous or complete, but should enable the reader to understand the core concepts of the field, including their relations to each other. We present them by using examples that illustrate the ideas. The main questions addressed here are the following:

1. What is a manifold?
2. What is a function on a manifold?
3. What is a measurement of a point on a manifold?
4. What is an embedding?
5. What is (data-driven) manifold learning?
6. What is the difference between linear and nonlinear manifold learning?

2 Manifolds

A **manifold** is a collection of (abstract) points with additional information about their relations. One way to think about manifolds is to consider them as objects in the real world: a piece of paper, a cell, the earth, or the water that is flowing in a river. These are all collections of points (atoms, molecules) with additional structure that is particular to them. A (flat) piece of paper is made of atoms, arranged in a particular way so that it can be distinguished from a ring-shaped object, for example. The ring can also be comprised purely of "paper molecules", but would still not be "the same manifold" because it has a different shape–the relation between the points is different!.

An important concept for manifolds is their **dimension**. Intuitively, the dimension of a manifold is its local resemblance of Euclidean space. For a sheet of paper, this notion means that when we zoom into it at any point, the sheet has to look like a two-dimensional plane - independent of its global shape, which might be crumpled up, folded, or stretched out. This means the sheet of paper is a two-dimensional manifold. If we look (not too closely) at a thin strand of

*felix.dietrich@tum.de

hair, it looks one-dimensional at every point, so we can think of it as a one-dimensional manifold. The crust of the earth (not considering its interior) can also be considered two-dimensional around every point - even though globally, the earth is spherical rather than flat. So, if we do not think of its height, the earths crust can be approximated as a two-dimensional manifold.

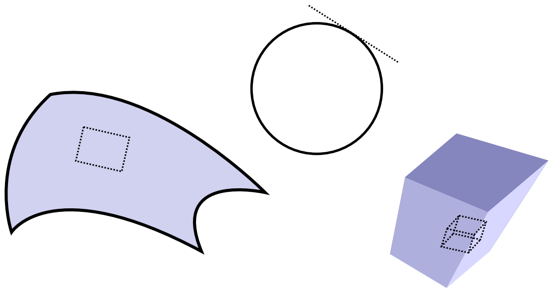

Figure 1: Examples for manifolds of dimension one (circle), two (surface), and three (cube). The dashed insets demonstrate that the manifolds locally are similar to the n-dimensional Euclidean space (represented here as a line, a rectangle, and a small cube). The colors are only to visualize that the sheet and the cube have points in their interior, while the circle does not. Notice that we did not yet choose a parametrization - the manifolds exist "as is", without the need to assign coordinates to each point.

Since manifolds can be real objects, the notation used to describe them is usually rather ambiguous. For example, we can capture a bounded, two-dimensional surface (like a sheet of paper) by denoting it simply as "a two-dimensional manifold with boundary".

2.1 Functions on manifolds

Crucially important for the following concepts is the idea of a (real-valued) **function on a manifold**. The typical notation for such a function is

$$f : M \rightarrow \mathbb{R}, \tag{1}$$
$$p \mapsto f(p). \tag{2}$$

The first line means that the function f maps from points on a (previously defined) manifold M into the real numbers in $\mathbb{R}$. Every point on the manifold is assigned exactly one real number this way. The second line then specifies how a point p in M is mapped to a real number $f(p)$. An important class of functions on manifolds are coordinate functions, which we discuss in the next section. Concrete examples for functions on a manifold are the temperature distribution on the ground of the earth, or the population density over a country. Every point of the country (the manifold!) is assigned a single number: the density of people living there, usually measured in persons per square kilometer.

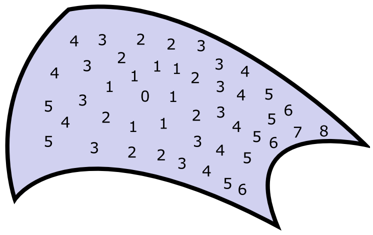

Figure 2: A given real-valued function $f : M \to \mathbb{R}$ assigns all points on the manifold a real number. For this example, only the numbers from zero to eight are shown. Some of the numbers $f(p)$ are shown at their positions p.

2.2 Embedding manifolds in Euclidean space

An **embedding** of a manifold into another one can be taken literally: given a manifold, embedding it into another one is possible if the other manifold can "contain" the given one. Mathematically, an embedding is a function that assigns every point of a given manifold another point in the embedding manifold in such a way that the function is invertible on its image and has a smooth derivative. In these notes, we only consider embeddings of a manifold into Euclidean space, and not into arbitrary other manifolds. A first example for an embedding of a manifold is shown in Fig. 3, where a two-dimensional manifold is embedded into three-dimensional Euclidean space (depicted as the three coordinate axis). For this manifold, an embedding into two-dimensional Euclidean space would also be possible. A good example for a slightly more complicated embedding is the typical illustration of a sphere embedded into three-dimensional Euclidean space: it is more complicated, because even though the sphere is a two-dimensional manifold (locally, it resembles a two-dimensional plane!), it cannot be embedded into two-dimensional Euclidean space as a whole.

Given a particular embedding of a manifold M into Euclidean space, points on M all have coordinate values given through their position in the embedding space.

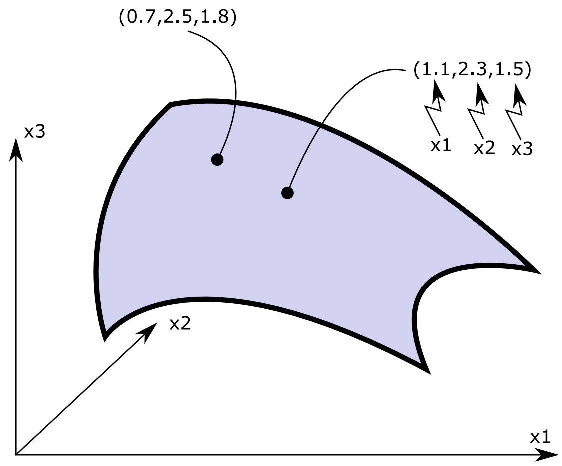

Figure 3: An embedding of a manifold in three-dimensional Euclidean space. Two points on the manifold are shown with their coordinates.

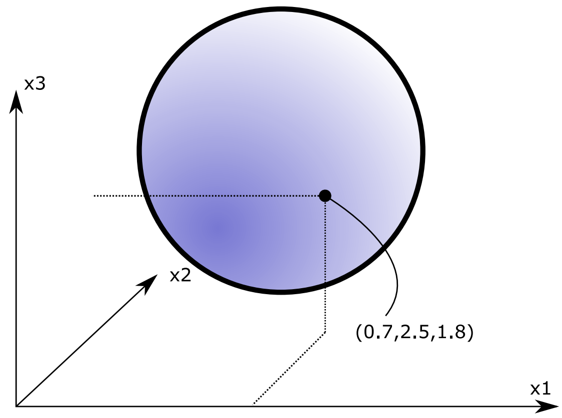

Figure 4: An embedding of a sphere in three-dimensional Euclidean space. A point on the sphere is shown with its coordinates given through the embedding space.

3 Learning manifolds from data

The term **learning** in this context refers to an approximate description of an object that is as accurate as possible, given finite information about the object. **Learning a manifold** then means to understand as much as possible about the underlying manifold from the given amount of information. In these notes, **information** will always refer to a collection of measurements for a finite number of points on a manifold. A good understanding of a given manifold is possible if we can find an embedding into two- or three-dimensional Euclidean space, such that it is possible to visualize the entire manifold.

There are many different ways that manifolds can be "learned". In general, the methods can be split up into **linear** and **non-linear** methods.

3.1 Linear manifold learning

Linear manifold learning works well if the given manifold is a hyperplane, that is, a plane with arbitrary dimension (a line, a two-dimensional plane, a cube, etc.). The crucial property of a manifold that is necessary for linear methods to yield "good", low-dimensional embeddings is that a local basis describing the manifold in a neighborhood of a point can be linearly extended to a global basis. These manifolds can often be identified with the Euclidean space they are embedded in: a (finite size, two-dimensional) plane is similar to two-dimensional Euclidean space, and a cube is similar to three-dimensional Euclidean space.

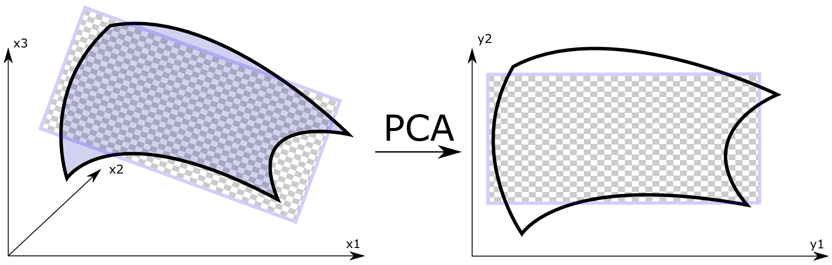

Figure 5: An embedding of a manifold in two-dimensional Euclidean space, using Principal Component Analysis. The manifold on the left is already embedded into three-dimensional Euclidean space, but PCA is able to find a two-dimensional embedding, because the manifold is almost planar.

3.2 Nonlinear manifold learning

Different from the linear methods, there is no canonical, non-linear method that works well for all manifolds. The reason for this is that non-linear manifolds cannot be described by a global coordinate basis in the way hyperplanes can. Sometimes, there is a parametrization of the manifold that allows an embedding in Euclidean space that has the same dimension than the manifold (see Fig.7, with a sphere example). However, the sphere (and the earth!) already provides an example where the embedding dimension (three) is higher than the manifold dimension (two, see Fig. 7).

One method for non-linear dimensionality reduction is called Diffusion Maps [1, 2]. The main idea is to construct an approximation of a linear operator on the given manifold, and then use eigenfunctions of that operator as coordinates for the points of the manifold. [Here, I still have to explain much more]

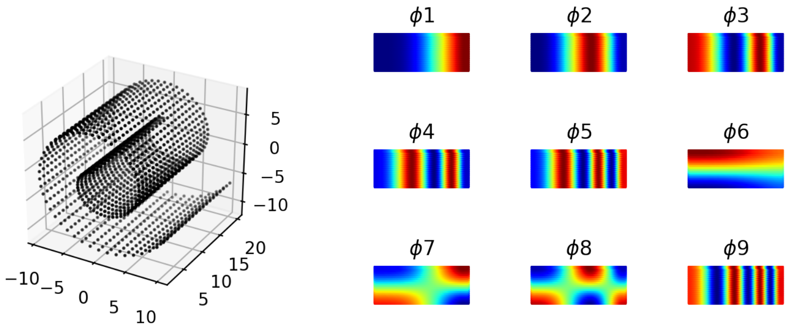

Figure 6: Left: the “swiss-roll” manifold, embedded in three-dimensional Euclidean space and sampled at the black dots. Right: Eigenfunctions of the Laplace-Beltrami operator, evaluated at the positions of the black dots on the left and plotted on the two “intrinsic” coordinates of the manifold. The color shows the function value. It is not possible to “unroll” the swiss-roll by fitting a plane in three-dimensional space (left), thus, normal PCA cannot embed this manifold properly in two dimensions.

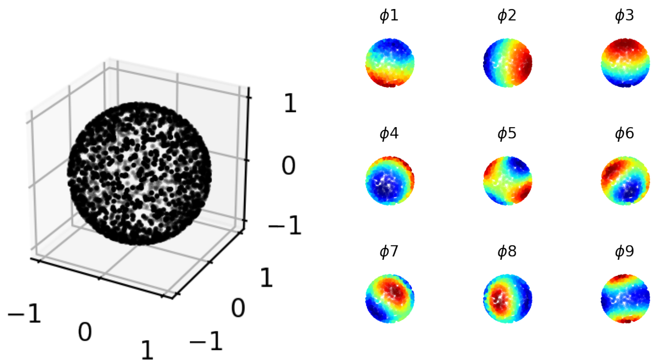

Figure 7: Left: a sphere, embedded in three-dimensional Euclidean space and sampled at the black dots. Right: Eigenfunctions of the Laplace-Beltrami operator, evaluated at the positions of the black dots on the left. The color shows the function value.

4 Mathematical literature

Here are some excellent introductions to the topic, for the mathematically inclined [5, 4, 3].

References

[1] Ronald R. Coifman, Stéphane Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for har-

monic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005.

[2] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 7 2006.

[3] John M. Lee. *Introduction to Smooth Manifolds*. Springer New York, 2012.

[4] Steven Rosenberg. *The Laplacian on a Riemannian Manifold*. Cambridge University Press, 1997.

[5] Loring W. Tu. *An Introduction to Manifolds*. Springer-Verlag GmbH, 2010.