Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Departement of Mathematics

Master Thesis | Spring 2024

Timon Reichert

Finding the Homology of Submanifolds from Samples

Submission Date: 13th May 2024

Advisor: Dr. Sara Kališnik Hintz

Abstract

Many Researchers today are confronted with large amounts of high dimensional data. This data often lies on or close to a manifold. The dimension of this submanifold is often much smaller than the dimension of the ambient Euclidean space. The goal of this thesis is to show that one can obtain geometric and topological information based on the given sample. In particular, we provide sufficient conditions on how dense the sample has to be, in order for the homology group of the submanifold to be determined from the sample. We also give a lower bound on the sample size in order for the density-conditions to be satisfied with high confidence. The density-conditions and the bounds on the sample size are obtained in terms of the condition number, which measures a shapes deviation from being flat. The thesis is based on the paper ≪Finding the Homology of Submanifolds with High Confidence from Random Samples≫ [13] and provides the necessary background knowledge to understand the results and gives detailed proofs of the statements.

Contents

1 Introduction

Due to technology with high computational capabilities and due to new experimental methods modern scientist are faced with a deluge of data. This data is usually collected in the form of very long vectors and hence lies in a very high dimensional space. However, the underlying space these points are drawn from often has a much lower dimension than the ambient space. Consider, for example, the data that consists of finitely many points that are drawn randomly from a circle, as illustrated in Figure 1. Such a collection of points is called a sample. We can see that the sample lies in $\mathbb{R}^2$, but the underlying circle locally looks like a line and therefore has dimension 1.

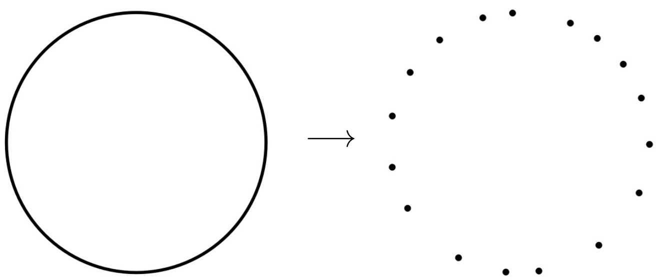

Figure 1: A sample of points drawn randomly from a circle.

Instead of a circle, one can consider an arbitrary manifold of dimension m, that is, a space that locally looks like a Euclidean space of lower dimension. For example, a circle is a manifold of dimension 1. Many researchers, such as Belkin and Niyogi [2], Donoho and Grimes [6], Roweis and Saul [14], Tenenbaum et al. [15] and Zomorodian and Carlsson [17] have worked on estimating geometric and topological properties of spaces from samples that come form submanifolds of Euclidean spaces. These kinds of questions belong to a class of problems called.

One property we are interested in is homology, which essentially gives information about the number of holes and high dimensional voids in a space. These numbers are called the *Betti numbers* β_k. For example, β_0 gives the number of 0-dimensional holes, which can be interpreted as the number of connected components of the space. The number β_1 then gives the number of holes, β_2 the number of voids and so on. For example, the unit circle S^1 has one connected component and one hole, hence $\beta_0 = \beta_1 = 1$. There are no voids of higher dimension in S^1, therefore the rest of the Betti numbers vanish. By identifying the homology of the submanifold we can therefore extract information about its geometric shape. Consider the sphere and the torus in Figure 2. Both have Betti numbers $\beta_0 = \beta_2 = 1$, since both are path connected and enclose a cavity. The torus has two 1-dimensional holes, indicated by the red and green path, and thus has $\beta_1 = 2$, while the sphere has $\beta_1 = 0$ since any 1-dimensional cycle is contractible to a point. Since both of them are 2-dimensional submanifolds of $\mathbb{R}^3$ (they locally look like $\mathbb{R}^2$) the numbers β_k vanish for $k \geq 3$.

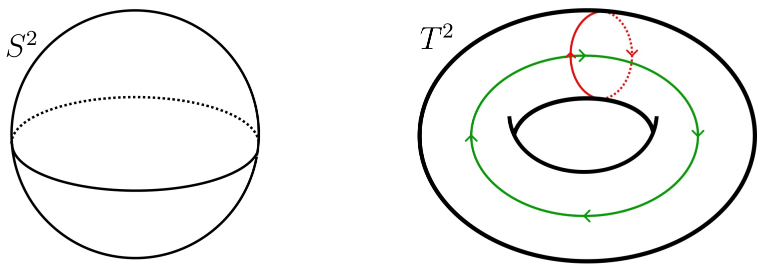

Figure 2: The sphere S^2 on the left and the torus T^2 are both 2-dimensional submanifolds of $\mathbb{R}^3$.

Consider now a sample $\overline{x}$ of points drawn from a compact submanifold M of $\mathbb{R}^N$. Regarding the points as a discrete topological space would enable us to compute its homology. However, this space will have Betti numbers $\beta_k = 0$ for all $k \geq 1$ and β_0 is just the number of points of the sample, which is not useful. We therefore have to approximate the point-cloud with a space that has non-trivial homology. For this we consider the set U given by the union of equally sized balls around each of the sample points, that is

$$U = \bigcup_{x \in \overline{x}} B_\varepsilon(x).$$

The challenge is to find the right ball-size ε, such that U has a similar shape as the point cloud, and yet is not just a collection of non-intersecting balls. In Figure 3 we can see the set U for the point-cloud from Figure 1 with three different radii. If the radius is too small, we are still left with non-intersecting balls and the homology only captures the number of points, as shown on the left. On the right, the radius is to large and the homology will be the one of a single contractible space. The figure in the middle shows a topological space that captures the shape from the underlying circle. Using standard constructions, such as simplicial complexes (see Bredon [4] or Hatcher [9]) the homology of the space U can easily be computed.

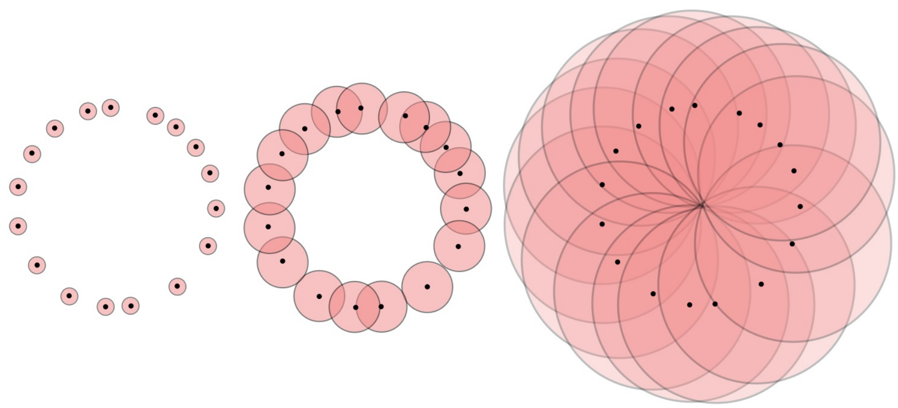

Figure 3: The set U with different radii for the point-cloud from Figure 1.

To prove that the homology of U is the same as that of M, we make use of the fact that two spaces that are homotopy equivalent have the same homology type. Intuitively, two spaces are homotopy equivalent, if one can be continuously deformed into the other, without cutting the space. For example, by smoothing out the edges, a square can be deformed into a circle without cutting.

If we can find the right ball-sizes ε, such that U and M are homotopy equivalent, then we can deduce the homology of M from the homology of U. In Figure 3 we have seen that the right ball-size has to be small enough, for U not to fill out any holes of M or create new ones. In fact, it is bounded from above by a number called the reach on M. Consider the space that is formed by thickening the circle, as shown in Figure 4. The reach R is

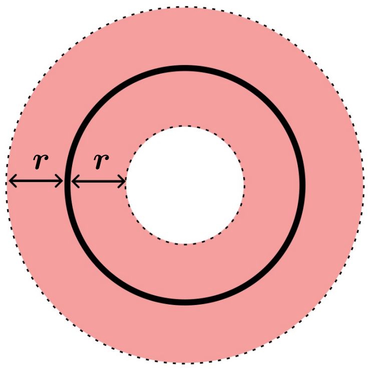

Figure 4: Thickening of the circle to get the tubular neighborhood of the circle.

the largest number r, such that this space can still be continuously deformed into M. On the other hand, the ball-size also has to be large enough, so that U does not just have the homology of a discrete set. This upper bound depends on the sample's density on the submanifold, since wide gaps between points require a larger radius, for the balls to overlap.

Our first result provides a range of radii and a density-condition on the sample, for which M is homotopy equivalent to U. In fact, we show that U deformation retracts to M, which implies homotopy equivalence. While the upper bound for the radii is given in terms of the condition number

$$\tau = \frac{1}{R}$$

of M, the lower bound on the density-condition is given in terms of the ball-size of U.

Theorem 4.1. *Let $\overline{x}$ be any finite collection of points $x_1, ..., x_n \in \mathbb{R}^N$ such that it is $\varepsilon/2$-dense in M. Then for any $\varepsilon < \sqrt{\frac{3}{5}}\tau$, we have that U strongly deformation retracts to M. Therefore the homology of U equals the homology of M.*

Clearly, the density of $\overline{x}$ in M depends on the number of points of the sample. However, the sample could be drawn in such a manner that the points are concentrated to one part of the manifold. Thus, a large sample is not a guarantee that the sample is dense enough in M. That is why, the next result considers a probabilistic setting, in which we provide a lower bound on the number of points of the sample, in order to guarantee the required density-condition with high probability. This lower bound depends on the ball-radius ε and on natural invariants of the submanifold M, such as the covering number (number of balls that are needed to cover M) and the packing number (number of balls that fit entirely into M without intersecting each other).

We also treat the situation where the data might be drawn from a probability distribution that is concentrated around the manifold. We introduce a model of noise, which considers samples that are drawn according to a probability measure, which is concentrated around the manifold M, rather on M. In particular, the sample points are allowed to lie within distance r of the manifold, that is they are contained in the tubular neighborhood of M with radius r. Considering again the set U, consisting of balls around each sample-point, we provide a range for the ball-size depending on r and the condition number as well as a density-condition in terms of r, for which U deformation retracts onto M.

Theorem 5.2. *Let* $0 < r < \left(\sqrt{9} - \sqrt{8}\right) \tau$ *and let* $\overline{x}$ *be a finite collection of points from within the distance* r *of* M*. If* $\overline{x}$ *is* r*-dense in* M*. Then for all*

$$\varepsilon \in \left(\frac{(r+\tau) - \sqrt{r^2 + \tau^2 - 6\tau r}}{2}, \frac{(r+\tau) + \sqrt{r^2 + \tau^2 - 6\tau r}}{2} \right),$$

the set U *consisting of the balls of radius* ε *centered at each sample-point strongly deformation retracts to* M*.*

As for samples without noise, we provide a lower bound on the sample-size in terms of the covering number, such that the density-condition is guaranteed with high probability.

In Chapter 2 we provide the differential geometric terminology starting with manifolds and tangent spaces. We will use the notions of connections and the second fundamental form to establish a relation between curvature and the condition number. The necessary probability preliminaries are provided in Chapter 3. We will give the basic notions of a probability space together with some examples. We define the support of a probability measure, which we use to define probability measures that are concentrated around a manifold. In Chapter 4 we first prove the theorem that provides sufficient conditions for the set U to deformation retract onto M. Then we prove the theorem that gives a lower bound on the number of points in the sample, such that the sample is dense enough on M with high confidence. The situation where the data is drawn from around the manifold is treated in Chapter 5. We define probability measures that are concentrated around a submanifold and We prove analogous theorems for the deterministic and probabilistic setting as in Chapter 4. At the end of Chapter 5 we briefly give thought to the situation where we weaken the density-condition on the sample with noise. We can observe, that the same or even stronger versions of our results can be derived this way.

2 Differential Geometry Preliminaries

In the introduction we have mentioned that a 1-dimensional submanifold of a Euclidean space is a subspace that locally looks like a line. The tangent space at a point of a 1-dimensional submanifold is the tangent line to the submanifold passing through that point. In this chapter we will extend these notions to the case of submanifolds of higher dimensions. In particular, at the end of this chapter we define a quadratic form on the tangent space, called the second fundamental form, which will help us establish a relation between the condition number and the curvature of the submanifold. For sections Section 2.1, Section 2.3 and Section 2.4, we mainly follow Lee [10], Merry [12] and Einsiedler and Wieser [8]. In Section 2.5 up to Section 2.9 the terminology is based on Lee [11] and Carmo [5]. Section 2.10 follows Niyogi et al. [13] and Lee [11].

2.1 Submanifolds and Smooth Maps

We introduce the main objects that we are working with in this thesis, namely smooth submanifolds of Euclidean spaces. In order compare and build relations between these objects, we define smooth maps between them. Smoothness is a notion that is well understood in the case where we consider subsets in the 1-dimensional case, that is of $\mathbb{R}$. We will extend the definitions to subsets of $\mathbb{R}^N$ for $N \geq 1$. All the notion mention above can of course even be extended to spaces that are not Euclidean, which is the theory of smooth manifolds and smooth maps. Since all the spaces we are working with are Euclidean, we will restrict our definitions to subspaces of Euclidean spaces. However, one can extend the notions to all smooth manifolds.

To be able to state the definition of smooth manifolds, first recall some definition from multivariable calculus.

Definition 2.1. *Let $U \subset \mathbb{R}^n$ an open subset and $f \colon U \to \mathbb{R}^m$ a continuous map. The **derivative** of f **along the vector** $v \in \mathbb{R}^n$ at $p \in U$ is defined by*

$$\partial_v f(p) := \frac{d}{dt}\Big|_{t=0} f(p + tv),$$

where $\frac{d}{dt}$ is the usual derivative with respect to t in all coordinates.

Denote by $\{e_i\}_{i=1}^n$ the standard basis of $\mathbb{R}^n$. This means that every vector $v \in \mathbb{R}^n$ can be written as

$$v = (v^1, ..., v^n) = v^1 e_1 + ... + v^n e_n.$$

For $j = 1, ..., N$, let x^j be the function that sends a vector v to its j-th coordinate. That is x^j is the real-valued function $x^j \colon \mathbb{R}^n \to \mathbb{R}$, defined by

$$x^j(v) = x^j(v^1, ..., v^n) = v^j \in \mathbb{R}^n.$$

Definition 2.2. *In the special case where $v = e_j$ for some $j \in \{1, ..., n\}$, we call the derivative along e_j of f at p*

$$\partial_{e_j} f(p) = \frac{d}{dt}\Big|_{t=0} f(p + te_j)$$

the ***partial derivative*** *of f at p in the j-th coordinate. We sometimes write $\frac{\partial f}{\partial x^j}(p)$, $\frac{\partial}{\partial x^i}\big|_p(f)$ or $\partial_j f(p)$.*

Example 2.3. *Consider the continuous map $f \colon \mathbb{R}^3 \to \mathbb{R}$ defined by*

$$f(x, y, z) = x(y^2 + \sin(z)).$$

The partial derivatives of f are

$$\partial_x f(x, y, z) = y^2 + \sin(z), \quad \partial_y f(x, y, z) = 2xy, \quad \partial_z f(x, y, z) = x \cos(z)$$

for $x, y, z \in \mathbb{R}$.

If $f = (f_1, ..., f_m) \colon U \subset \mathbb{R}^n \to \mathbb{R}^m$ is a continuous map with componen maps $f_1, ..., f_m$, we can compute the partial derivative of each component f_i at a point $p \in U$. The $(n \times m)$-matrix $J_f(p)$ whose (i, j)-th entry is the partial derivative $\partial_j f_i(p)$ is called the **Jacobi-matrix** of f.

Definition 2.4. *Let $f \colon U \subset \mathbb{R}^n \to \mathbb{R}^m$ be a function with components $(f_1, ..., f_m)$. The* ***derivative*** *of f at p is the linear map*

$$Df(p) \colon \mathbb{R}^m \to \mathbb{R}^n$$

given by the $(n \times m)$-Jacobi matrix, whose entries are the partial derivatives of f

$$Df(p) = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(p) & \frac{\partial f_1}{\partial x_2}(p) & \cdots & \frac{\partial f_1}{\partial x_m}(p) \\ \frac{\partial f_2}{\partial x_1}(p) & \frac{\partial f_2}{\partial x_2}(p) & \cdots & \frac{\partial f_2}{\partial x_m}(p) \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1}(p) & \frac{\partial f_n}{\partial x_2}(p) & \cdots & \frac{\partial f_n}{\partial x_m}(p) \end{pmatrix}.$$

We now have the definition of the derivative of maps between general Euclidean spaces. This enables us to define the notion of differentiable, k-differentiable and smoothness. But first we look at a concrete example of the derivative of a maps.

Example 2.5. *Let us consider the map $f \colon \mathbb{R}^2 \to \mathbb{R}^2$, given by $f(x, y) = (x + y^2, y + x^2)$. Then we have the components $f_1(x, y) = x + y^2$ and $f_2(x, y) = y + x^2$. Their partial derivatives are*

$$\partial_x f_1(x, y) = 1, \quad \partial_y f_1(x, y) = 2y$$

and

$$\partial_x f_2(x,y) = 2x, \quad \partial_y f_2(x,y) = 1.$$

The Jacobi matrix therefore is

$$Df(x,y) = \begin{pmatrix} 1 & 2y \\ 2x & 1 \end{pmatrix}.$$

Definition 2.6. *Let $f\colon U \subset \mathbb{R}^n \to \mathbb{R}^m$ be map on an open set U. We say that f is **differentiable** at p, if the partial derivatives $\partial_j f(p)$ exist at p for every $j \in \{1,...,n\}$. We say that f is **twice differentiable**, if it is differentiable and the second partial derivatives $\partial_j \partial_k f$ exist for all $j,k \in \{1,...,n\}$. In general, for $d \geq 2$, we say f is **d-differentiable**, if the partial derivatives $\partial_j f$ are $(d-1)$-differentiable. If f is d-differentiable for every $d \geq 1$, we say that f is **smooth**.*

We can now give the definition of a diffeomorphism between subsets of $\mathbb{R}^N$.

Definition 2.7. *Let $U, V \subset \mathbb{R}^m$. A map $f\colon U \to V$ is called a **diffeomorphism**, if it is smooth and bijective and if its inverse f^{-1} is also smooth.*

We are now able to define the notions of charts and transition maps, which we need in order to define submanifolds.

Definition 2.8. *Let M be a subspace of a Euclidean space. A **chart** (U,φ) on M consists of an open subset U of M and a diffeomorphism φ onto an open subset of some Euclidean space $\mathbb{R}^N$.*

If we have two charts (U,φ) and (V,ψ) on M into $\mathbb{R}^N$, we can observe that the map $\varphi \circ \psi^{-1}\colon \psi(U \cap V) \to \varphi(U \cap V)$ is in fact a map from an open set of $\mathbb{R}^N$ to $\mathbb{R}^N$. Thus, we can apply the notion of differentiability and smoothness from above to these maps.

Definition 2.9. *Let (U,φ) and (V,ψ) be charts on some subset M of a Euclidean space. The maps $\varphi \circ \psi^{-1}\colon \psi(U \cap V) \to \varphi(U \cap V)$ are called **transition maps.***

Definition 2.10. *Let $m \geq 1$ and $M \subset \mathbb{R}^N$ be a subspace. We say that M is a **submanifold of dimension** m **of** $\mathbb{R}^N$, if there is a collection of charts $\{(U_\alpha), \varphi_\alpha \mid \alpha \in A\}$ onto $\mathbb{R}^m$, such that $\{U_\alpha \mid \alpha \in A\}$ is an open cover of M and the transition maps*

$$\varphi_\alpha \circ \varphi_\beta^{-1}\colon \varphi_\beta(U_\alpha \cap U_\beta) \to \varphi_\alpha(U_\alpha \cap U_\beta)$$

are smooth for all $\alpha, \beta \in A$ with $U_\alpha \cap U_\beta \neq \emptyset$.

Let us have a look at some submanifolds of $\mathbb{R}^N$.

Example 2.11. *(a) Every open set $U \subset \mathbb{R}^N$ is an N-dimensional submanifold, with a single chart (U, id).*

(b) *Every discrete set $M \subset \mathbb{R}^N$ is a 0-dimensional submanifold. A discrete set is a subset M such that for every $p \in M$ there is $\varepsilon > 0$, such that $M \cap B_\varepsilon(p) = \{p\}$. For every point $p \in M$, we can consider $\varphi_p \colon B_\varepsilon(p) \to B_\varepsilon(0)$ to be the translation map $x \mapsto x - p$, for some small enough $\varepsilon > 0$. This yields a diffeomorphism $\varphi_p \colon \{p\} \to \{0\}$.*

(c) *Let S^2 be the 2-dimensional sphere*

$$S^2 = \{x \in \mathbb{R}^3 \mid \|x\| = 1\}.$$

This is a 2-dimensional submanifold of $\mathbb{R}^3$. Let $p_0 = (x_0, y_0, z_0) \in S^2$. In the following, we build charts around p_0, depending on the the values of its coordinates.

- *If $z_0 > 0$, let $U_+ \coloneqq B_1^2(0) \times (0, \infty)$, where $B_1^2(0)$ is the 2-dimensional open unit disc with radius 1 centered at the origin. The pair (φ_{p_0}, U_+), where φ_{p_0} is the map defined by*

$$\varphi_{p_0} \colon U_+ \to \varphi_{p_0}(U_+), (x, y, z) \mapsto (x, y, z - \sqrt{1 - x^2 - y^2}),$$

  gives a chart onto $\mathbb{R}^2$. Indeed, the map satisfies that $\varphi_{p_0}(x, y, z) \in \mathbb{R}^2 \times \{0\}$ if and only if $(x, y, z) \in U_+ \cap S^2$. We can see that φ_{p_0} is a diffeomorphism onto its image.

- *If $z_0 < 0$, we consider the set $U_- \coloneqq B_1^2(0) \times (-\infty, 0)$ and the map*

$$\varphi_{p_0} \colon U_- \to \varphi_{p_0}(U_-), (x, y, z) \mapsto (x, y, z + \sqrt{1 - x^2 - y^2}),$$

  which is a diffeomorphism onto its image. Hence, (φ_{p_0}, U_-) is a chart on S^2 onto $\mathbb{R}^2$.

- *If $z_0 = 0$ and $y_0 > 0$, we build a chart $(\tilde{U}_+, \varphi_{p_0})$ by considering the set $\tilde{U}_+ \coloneqq \{(x, y, z) \in \mathbb{R}^3 \mid y > 0, x^2 + z^2 < 1\}$, (which is U_+ from above after exchanging y with z) and the map*

$$\varphi_{p_0} \colon \tilde{U}_+ \to \varphi_{p_0}(\tilde{U}_+), (x, y, z) \mapsto (x, z, y - \sqrt{1 - x^2 - z^2}).$$

  The same argumentation as above shows that this gives a diffeomorphism from $\tilde{U}_+ \cap S^2$ to $\mathbb{R}^2$.

- *We are left with the cases $z_0 = 0 \wedge y_0 < 0$ and $y_0 = z_0 = 0 \wedge x \in \{-1, 1\}$, for which the open set U_{p_0} and the diffeomorphism can be found analogously.*

In general, for $N \geq 3$, the $(N-1)$-dimensional sphere

$$S^{N-1} = \{x \in \mathbb{R}^N \mid \|x\| = 1\}$$

is a $(N-1)$-dimensional submanifold of $\mathbb{R}^N$.

(d) *The image of any curve $\gamma \colon (-1, 1) \to \mathbb{R}^n$ is a submanifold of dimension 1.*

In Example 2.11 *(b)*, we see that a point-cloud as a submanifold of $\mathbb{R}^N$ has dimension 0. That is why it does not make sense to consider the sample itself as a submanifold. Instead, we imagine the points as being collected from a submanifold of $\mathbb{R}^N$ and want to identify the geometrical shape of that submanifold.

Let us look at smooth maps between submanifolds of $\mathbb{R}^N$. The goal is to state the inverse function theorem, which says that a smooth map $F\colon M \to N$ between two submanifolds of Euclidean spaces is locally a diffeomorphism, if the derivative dF_p is invertible at every $p \in M$. For this, we first need to consider a setting, in which we can differentiate these maps.

Definition 2.12. *Let M and N be smooth manifolds of Euclidean spaces with dimension m and n, respectively. A continuous map $F\colon M \to N$ is called **smooth** if for each $p \in M$ and for some (and hence all) charts $\varphi\colon U \to \mathcal{O} \subset \mathbb{R}^m$ and $\psi\colon V \to \Omega \subset \mathbb{R}^n$ on M, respectively on N, the composite maps*

$$\psi \circ F \circ \varphi^{-1}\colon \varphi(U \cap F^{-1}(V)) \to \psi(F(U) \cap V)$$

are smooth.

Example 2.13. *Consider the two 2-dimensional submanifolds M and N of $\mathbb{R}^3$ as shown in Figure 5. The charts φ and ψ map an open neighborhood of p, respectively of $F(p)$ homeomorphically to an open set in $\mathbb{R}^2$. The map $\psi \circ F \circ \varphi^{-1}$ is a map from an open set of $\mathbb{R}^m$ to an open set of $\mathbb{R}^n$ and thus, we can apply the notion of smoothness to this map.*

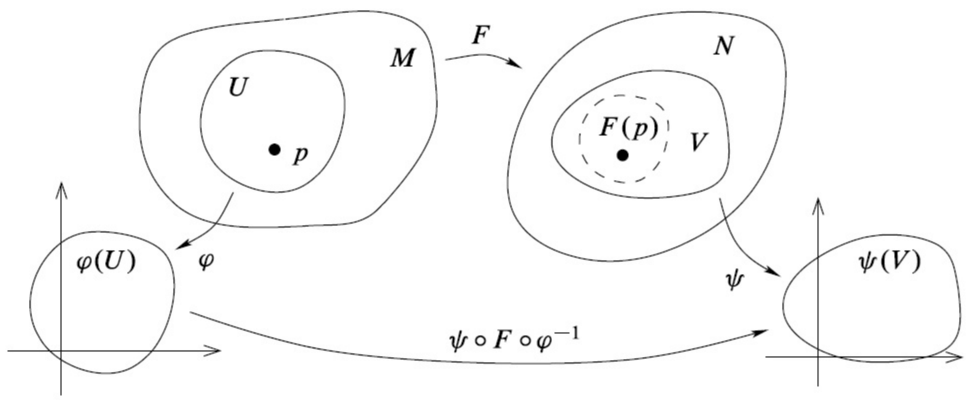

Figure 5: Smooth maps between manifolds with composite map, from Lee [10].

Definition 2.14. *Let M and N be submanifolds of $\mathbb{R}^N$ and $F\colon M \to N$ a smooth map. If F is a homeomorphism and its inverse $F^{-1}\colon N \to M$ is also smooth, then we say that F is a **diffeomorphism**.*

2.2 Homotopy and Strong Deformation Retract

To check if two spaces have the same geometrical shape, one can try to if somehow continuously deform one space into the other one. For example, by stretching, pulling or contracting. This brings us to the definition of a homotopy equivalence between two topological spaces.

Definition 2.15. *Let $f_0, f_1 \colon X \to Y$ be maps between two topological spaces. We say that f_0 and f_1 are* ***homotopic****, if there is a continuous map $F \colon X \times [0,1] \to Y$ with $F(x,0) = f_0(x)$ and $F(x,1) = f_1(x)$ for all $x \in X$. In that case write $f_0 \simeq f_1$.*

To say that ≪f_0 and f_1 are homotopic≫ is in fact an equivalence relation, that is, the following conditions are satisfied:

- $f \simeq f$, for every $f \colon X \to Y$,
- $f \simeq g \implies g \simeq f$, for all maps $f, g \colon X \to Y$, and
- $f \simeq g$ and $g \simeq h \implies f \simeq h$, for all maps $f, g, h \colon X \to Y$.

Definition 2.16. *A map $f \colon X \to Y$ between two topological spaces is called a* ***homotopy equivalence****, if there is a map $g \colon Y \to X$, such that $g \circ f \simeq \mathrm{id}_X$ and $f \circ g \simeq \mathrm{id}_Y$. In particular, if such a map f exists, we say that X and Y have the same* ***homotopy type****.*

Also the relation ≪X and Y have the same homotopy≫ type is an equivalence relation. Moreover, if $f \colon X \to Y$ is a homeomorphism, then f is clearly a homotopy equivalence, by taking $g = f^{-1}$.

Definition 2.17. *A space that has the same homotopy type as a single point $\{x\}$ is called* ***contractible****.*

There is even an equivalent definition of contractible spaces as shown in the following proposition.

Proposition 2.18. *A space X is contractible if and only if the identity map id_X is homotopic to a constant map $c \colon X \to X$.*

Example 2.19. *Let us look at a few examples of homotopy equivalences.*

(a) We show that for all $N \geq 1$, the Euclidean space $X = \mathbb{R}^N$ is contractible. Let $Y = \{y_0\}$ be the space with one element and $c \colon X \to Y$ be the constant map $c(x)$ for all $x \in X$. We claim that the inclusion map $i \colon \{0\} \to X$ is the homotopy inverse to c. In fact, we have $c \circ i = \mathrm{id}_Y$ and for the other homotopy equivalence, we consider the map $F \colon X \times [0,1] \to X$ defined by $F(x,t) = tx$. Then, the map F is continuous and it satisfies

$$F(x,0) = 0 = i \circ c(x) \text{ and } F(x,1) = x = \mathrm{id}_X(x).$$

Thus, we also have $i \circ c \simeq \mathrm{id}_X$, which shows that the constant map $c\colon X \to Y$ is a homotopy equivalence. But since Y is a single point, X is contractible.

(b) *Let X be an open annulus in $\mathbb{R}^2$, centered at the origin containing the unit circle S^1. Let $f\colon S^1 \to X$ be the inclusion map. Then it is easy to check that the map f*

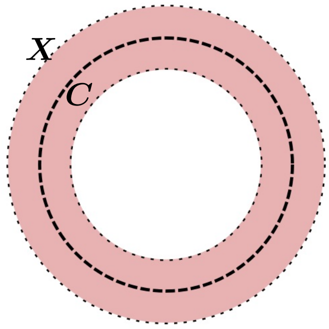

Figure 6: The annulus X is homotopy equivalent to a circle.

is a homotopy equivalence with homotopy inverse g, where $g\colon X \to S^1$ is the map that sends every point $x \in X$ to the unique point on S^1 that is closest to x.

In the second example we can see that the annulus was somehow „squished“ down onto the circle $S^1 \subset X$. We make this notion precise in the following definition.

Definition 2.20. *Let X be a topological space and $A \subset X$ a subspace. A continuous map $F\colon X \times [0,1] \to X$ is called a* ***strong deformation retract****, if for every $x \in X$ and every $a \in A$ the following hold*

- *$F(x,0) = x$;*
- *$F(x,1) \in A$;*
- *$F(a,t) = a$ for all $t \in [0,1]$.*

Example 2.21. (a) *In Example 2.19 (b) we considered the subspace $C \subset X$, which was an embedded circle in the annulus. This is in fact an example of a strong deformation retract where F is the map that moves every point x in the annulus along a straight line to the closest point on the embedded circle.*

(b) *Consider the subspace $S^2 \subset \mathbb{R}^3 \backslash \{0\}$, that is, the unit sphere as a subspace of the Euclidean space $\mathbb{R}^3$ pierced at the origin.*
One can check, that the map $F\colon \mathbb{R}^N \backslash \{0\} \times [0,1] \to \mathbb{R}^N \backslash \{0\}$ given by

$$F(x,t) := (1-t) \cdot x + t \cdot \frac{x}{\|x\|},$$

which sends a point x onto the unit sphere along the straight ray starting at the origin and passing through x (see Figure 7) is a strong deformation retract.

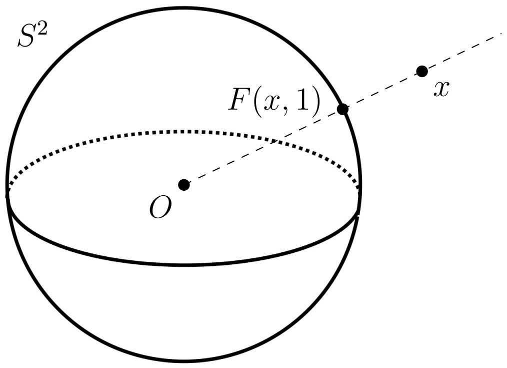

Figure 7: A deformation retract from the pierced Euclidean space to the sphere.

2.3 Tangent Vectors, Tangent Spaces and Normal Spaces

Consider a path $\gamma\colon (-1,1) \to M$ on some manifold M. We can compute its velocity vector at a given $t_0 \in (-1,1)$ by differentiating γ with respect to t at t_0, that is

$$\dot{\gamma}(t_0) = \frac{d}{dt}\gamma(t)|_{t=t_0}.$$

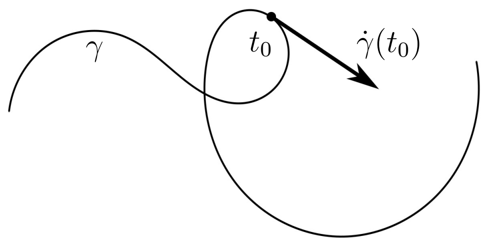

Figure 8: The velocity vector $\dot{\gamma}(t_0)$ tanget to γ at $\gamma(t_0)$.

In Figure 8 we see that this is a vector tangent to the 1-dimensional submanifold $\mathrm{im}(\gamma)$ at the point $\gamma(t_0)$. In fact, we say that $\dot{\gamma}(t_0)$ is an element of the tangent space of M at the point $\gamma(t_0)$. Instead of choosing M to be a 1-dimensional submanifold, we generalize the notion of the tangent space to all submanifolds of Euclidean spaces. The definition of the tangent space together with some of their properties will enable us to bring us one step closer to the condition number of a manifold. Moreover, in this section, we will show that we can identify the tangent space of $\mathbb{R}^N$ with $\mathbb{R}^N$ itself and we will provide a canonical basis of $T_p\mathbb{R}^N$. At the end of this section, we will equip the tangent space of a submanifold with an inner product to get Riemannian submanifolds.

Definition 2.22. *Let $M \subset \mathbb{R}^N$ be a submanifold and $p \in M$. The* ***tangent space*** *of M at p is the set*

$$T_pM = \{\dot{\gamma}(0) \mid \gamma\colon (-1,1) \to M \text{ is a smooth curve and } \gamma(0) = p\} \subset \mathbb{R}^N.$$

The tangent bundle of M is

$$TM = \{(p, v_p) \mid p \in M, v_p \in T_pM\}.$$

Recall the partial derivative operatores

Example 2.23. *Figure 9 shows the tangent vector space T_pS^2 at some point p on the unit sphere S^2.*

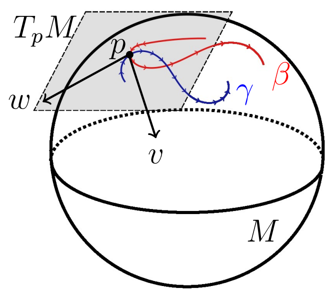

Figure 9: The two tangent vectors v and w are given by $v = \beta'(0)$ and $w = \gamma'(0)$.

We see that the tangent space of M at p is the plane tangent to the sphere passing through the point p. Since any plane in $\mathbb{R}^3$ is isomorphic to $\mathbb{R}^2$, we see that T_pM is isomorphic to $\mathbb{R}^2$.

This makes us wonder, if this is in general true. For the case, where $M = \mathbb{R}^N$ itself, we have the following statement.

Proposition 2.24. *For every $p \in \mathbb{R}^N$, the tangent space $T_p\mathbb{R}^N$ is canonically diffeomorphic to $\mathbb{R}^N$. In particular, the tangent bundle of $\mathbb{R}^N$ is canonically diffeomorphic to $\mathbb{R}^N \times \mathbb{R}^N$.*

Proof. Let $p \in \mathbb{R}^N$ and $v \in T_p\mathbb{R}^N$ with smooth curve $\gamma_v\colon (-1,1) \to \mathbb{R}^N$, such that $\gamma_v(0) = p$ and $\dot{\gamma}_v(0) = v$. Define the map $h_p\colon T_p\mathbb{R}^N \to \{p\} \times \mathbb{R}^N$ via

$$h_p(v) \coloneqq (\gamma_v(0), \dot{\gamma}_v(0)).$$

Since v is a tangent vector, $\gamma_v(0)$ and $\dot{\gamma_v}(0)$ do not depend on the choice of the curve γ_v. This map has an inverse given by

$$h_p^{-1}(p,x) := \frac{d}{dt}(p + t \cdot x)(0).$$

Indeed, $\gamma(t) = p + tx$ defines a smooth curve for $t \in (-1,1)$ with $\gamma(0) = p + tx|_{t=0} = 0$ and $\dot{\gamma}(0) = \frac{d}{dt}(p+tx)(0) = x \in \mathbb{R}^N$. In particular, h_p and h_p^{-1} are both smooth maps and therefore, h is a diffeomorphism. Since h_p does not depend on the choice of a basis of $T_p\mathbb{R}^N$, the map h_p is in fact a canonical diffeomorphism. Moreover, via the map

$$h\colon \mathbb{R}^N \times \mathbb{R}^N \to TM, (p,v) \mapsto (v, h_p(v)),$$

we get a canonical isomorphism between $\mathbb{R}^N \times \mathbb{R}^N$ and TM. □

Next we provide a basis of the tangent space $T_p\mathbb{R}^N$, using the partial derivative operators we defined earlier. Recall that for any point $p \in \mathbb{R}^N$ and any vector $v \in \mathbb{R}^N$, we have the partial derivative operator $\partial_v|_p$ defined by

$$\partial_v|_p\colon f \mapsto \partial_v|_p(f) = \frac{d}{dt}\Big|_{t=0} f(p + t \cdot v).$$

Recall that for $v = e_j$ we have the partial derivative operator $\partial_j|_p$ at p with respect to the j-th coordinate. Let $\mathrm{Par}(p)$ be the vector space generated by the partial derivative operator $\partial_j|_p$, for $j = 1, ..., N$.

Proposition 2.25. *The space* $\mathrm{Par}(p)$ *is canonically isomorphic to* $T_p\mathbb{R}^N$.

Proof. Let $p \in \mathbb{R}^N$. Define the map $\Phi_p\colon T_p\mathbb{R}^N \to \mathrm{Par}(p)$, by

$$\Phi_p(v) \mapsto \partial_v|_p.$$

First we show that Φ_p is linear. Let $a \in \mathbb{R}$ and $v, w \in T_p\mathbb{R}^N$. By the chain rule of the partial derivative, we have that Φ_p is linear in the v. Moreover, since $\partial_v|_p$ is uniquely determined by the directional vector v, the map Φ_p is injective.

Consider an arbitrary $\partial_p = a^1\partial_1|_p + ... + a^1\partial_N|_p \in \mathrm{Par}(p)$. For $v = a^1e_1 + ... + a^Ne_n$, we get $\Phi_p(v) = \partial_p$ due to the linearity in the variable v. But this means that Φ_p is also surjective. Hence, Φ_p is a linear isomorphism and since, the isomorphism does not depend on a chosen basis, the isomorphism is canonical. □

Identifying a tangent vector $v \in T_p\mathbb{R}^N$ with its partial derivative operator $\partial_v|_p$ enables us to view v as an operator on the smooth real-valued functions via

$$v(f) := \partial_v|_p(f).$$

In particular, we can write any vector $v \in T_p\mathbb{R}^N$ as a linear combination of the partial derivatives along the standard basis vectors, that is,

$$v = v^1\partial_1|_p + ... + v^N\partial_N|_p.$$

This is essentially a matter of labeling the tangent vectors that will become useful in the setting of vector fields, which are smooth maps from M to the tangent bundle TM, such that p is mapped into T_pM.

2.4 Derivatives of Smooth Maps and The Inverse Function Theorem

In this section we will state the inverse function theorem, which says that a smooth map with invertible derivative at some point is a diffeomorphism on a neighborhood of p. In order to state the theorem, we need some terminology about the derivative of a smooth map between two submanifold of $\mathbb{R}^N$.

Definition 2.26. *Let M and N be submanifolds of $\mathbb{R}^N$ of dimension m and n respectively. Let $F\colon M \to N$ be a smooth map and $p \in M$. If (U, φ) and (V, ψ) are charts around p and $F(p)$, the* ***derivative*** *of F at p is the linear map*

$$DF(p)\colon T_pM \subset \mathbb{R}^N \to T_{F(p)}N \subset \mathbb{R}^N$$

given by the matrix, whose entries are the partial derivatives of $\psi \circ F \circ \varphi^{-1}$ at p.

Using standard notions for a $(N \times N)$-matrix to be invertible form linear algebra, we can now state the inverse function Theorem:

Theorem 2.27 (**Inverse Function Theorem for Manifolds**)**.** *Suppose M and N are smooth manifolds and $F\colon M \to N$ is a smooth map. If $p \in M$ is a point such that $DF(p)$ is invertible, then there are connected neighborhoods U of p and V of $F(p)$ such that $F|_U\colon U \to V$ is a diffeomorphism.*

We will skip the proof of the Inverse Function Theorem and refer to Lee [10] for the proof.

2.5 Riemannian Manifolds

In Example 2.23 we have seen that the tangent space $T_pS^2 \subset \mathbb{R}^3$ at some point on the sphere is the plane that touches the sphere at p. In particular, the tangent space is a 2-dimensional submanifold of $\mathbb{R}^3$. The question arises, what happened to the third dimension. In fact, consider the affine space A that passes though p and is orthogonal to the tangent plane T_pS^2. Then we immediately get that $\mathbb{R}^3 = T_p\mathbb{R}^3$ is the direct sum of the tangent plane T_pS^2 and the affine space A. We will call A the normal space of the sphere S^2 at p. However, to make sense of orthogonality on the tangent space of a submanifold of $\mathbb{R}^N$ we need an inner product on the tangent space $\mathbb{R}^N$. This is a symmetric and positive-definite map that is linear in the first argument. If we equip a manifold with an inner product on each tangent space, we get a Riemannian manifold.

Definition 2.28. *Let V be a vector space. An **inner product** on V is a map*

$$\langle \cdot, \cdot \rangle : V \times V \to \mathbb{R}$$

that is

(a) symmetric: $\langle x, y \rangle = \langle y, x \rangle$ for all $x, y \in V$;

(b) linear in the first argument: $\langle ax + by, z \rangle = a\langle x, z \rangle + b\langle y, z \rangle$ for all $x, y, z \in V$ and $a, b \in \mathbb{R}$;

(c) positive-definite: $\langle x, x \rangle \geq 0$ and $\langle x, x \rangle > 0$ whenever $x \neq 0$ for all $x \in \mathbb{R}$.

Example 2.29. *The standard inner product on $\mathbb{R}^N$ of two vectors $v = (v^1, ..., v^n)$ and $w = (w^1, ..., w^n)$ is given by $\langle v, w \rangle = v^1 w^1 + v^2 w^2 + ... + v^n w^n$.*

Example 2.30. *An inner product on the vector space $C^1([0, 1])$ of continuous functions on $[0, 1]$ is defined as follows. Given two functions $f, g \in C^1([0, 1])$ we define*

$$\langle f, g \rangle := \int_0^1 f(x)g(x)dx.$$

The symmetry follows directly from the symmetry of the multiplication of real numbers. The linearity can be shown as follows.

$$\begin{aligned}
\langle af + bh, g \rangle &= \int_0^1 (af(x) + bh(x))g(x)dx \\
&= \int_0^1 (af(x)g(x) + bh(x)g(x))dx \\
&= a \int_0^1 f(x)g(x)dx + b \int_0^1 h(x)g(x)dx \\
&= a\langle f, g \rangle + b\langle h, g \rangle,
\end{aligned}$$

where we used linearity of the integral in the third equation. The positive-definiteness follows from the fact that $x \mapsto x^2$ is a non-negative function.

Definition 2.31. *A **Riemannian Manifold** (M, g) is a real smooth manifold equipped with an inner product $\langle \cdot, \cdot \rangle_p$ on the tangent space T_pM at each point $p \in M$.*

Example 2.32. *Consider the Euclidean space $\mathbb{R}^N$. We have seen that for every point $p \in \mathbb{R}^N$, the tangent space $T_p\mathbb{R}^N$ can be identified with $\mathbb{R}^N$ itself, by identifing the tangent* *vectors* $\frac{\partial}{\partial x^j}\big|_p$ *with* e_j *for every* $j = 1, ..., N$. *The inner product on* $T_p\mathbb{R}$ *then is defined on the basis*

$$\langle e_i, e_j \rangle = \delta_{ij} = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{otherwise.} \end{cases}$$

Another set of Riemannian submanifolds are immersed submanifolds of Riemannian submanifolds.

Definition 2.33. *Let* $M \subset \mathbb{R}^m$ *and* $N \subset \mathbb{R}^n$ *be submanifolds of Euclidean spaces. We say that a smooth map* $F: M \to N$ *is an* ***immersion*** *if the derivative*

$$DF(p): T_pM \subset T_p\mathbb{R}^m \to T_{F(p)}N \subset \mathbb{R}^m$$

is injective. If $F: M \to N$ *is an immersion, we say that* M *is an* ***immersed submanifold*** *of* N.

Example 2.34. *Let* $F: M \subset \mathbb{R}^m \to N \subset \mathbb{R}^n$ *be an immersion. If* N *is a Riemannian manifold,* F *induces a Riemannian structure on* M *by defining*

$$\langle u, v \rangle_p \coloneqq \langle DF(p)(u), DF(p)(v) \rangle_{F(p)},$$

for $u, v \in T_pM$. *The symmetry and linearity of* $\langle \cdot, \cdot \rangle_p$ *follow directly from the symmetry of* $\langle \cdot, \cdot \rangle_{F(p)}$ *and the linearity of both* $\langle \cdot, \cdot \rangle_{F(p)}$ *and the derivative* $DF(p)$. *Moreover, we have for all* $u \in T_pM$

$$\langle u, u \rangle_p = \langle DF(p)(u), DF(p)(u) \rangle_{F(p)} \geq 0.$$

since $\langle \cdot, \cdot \rangle_{F(p)}$ *is positive-definite. Moreover, since* $DF(p)(u) = 0$ *implies* $u = 0$, *by the injectivity of* $DF(p)$, *we therefore get the positive-definiteness of* $\langle \cdot, \cdot \rangle_p$.

Example 2.35. *By Example 2.34 we see that every submanifold* $M \subset \mathbb{R}^N$ *inherits a Riemannian structure by restricting the innerproduct of* $T_p\mathbb{R}^N$ *to* $T_pM \subset \mathbb{R}^N$.

2.6 Tubular Neighborhoods, the Reach and the Condition Number of a Submanifold

In this section we introduce the reach and the condition number of a submanifold of a Euclidean space. For this we first need the definition of normal spaces and tubular neighborhoods.

Definition 2.36. *Let* M *be a Riemannian submanifold of the Euclidean space* $\mathbb{R}^N$. *The* ***normal space*** *of* M *at* p *is the orthogonal complement of* T_pM *in* $T_p\mathbb{R}^N \subset \mathbb{R}^N \times \mathbb{R}^N$ *with respect to the inner product on* $\mathbb{R}^N \times \mathbb{R}^N$, *that is*

$$N_pM \coloneqq (T_pM)^\perp = \left\{ w \in \mathbb{R}^N \mid \langle v, w \rangle = 0, \forall v \text{ with } v_p \in T_pM \right\},$$

where $v_p \in T_pM$ corresponds to $v \in \mathbb{R}^N$. The ***normal bundle*** *is the set*

$$NM := \{(p, w) \in \mathbb{R}^N \times \mathbb{R}^N \mid p \in M, w_p \in N_pM\}.$$

We observe that for each point $p \in M$, the inner product on the tangent space splits $T_p\mathbb{R}^N$ into

$$T_p\mathbb{R}^N = T_pM \oplus N_pM.$$

Definition 2.37. *If $v \in T_p\mathbb{R}^N$ is a tangent vector, we can write $v = v^T + v^N$, where $v^T \in T_pM$ is the* ***tangential component*** *and $v^N \in N_pM$ is the* ***normal component*** *of v.*

Example 2.38. *Consider M to be the space that looks like two tori glued together, as drawn in Figure 10. The normal space N_pM is the 1-dimensional linear subspace passing through p that is orthogonal to the tangent space T_pM[1].*

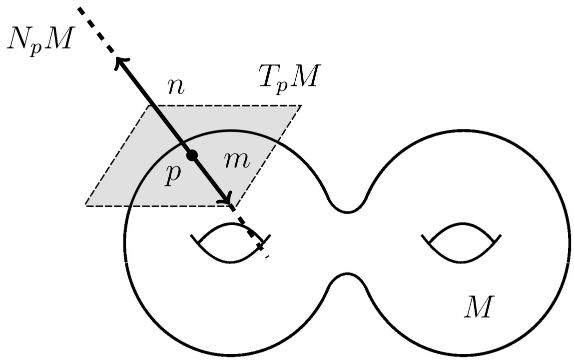

Figure 10: The normal space N_pM at p is the straight line passing through the point p that is orthogonal to the tangent space T_pM.

Consider the normal bundle NM of M. We can regard a normal vector v together with its base point p both as in $\mathbb{R}^N$ and thus they can be added together.

Definition 2.39. *Define $E\colon \mathbb{R}^N$ to be the map that adds the normal vector v to the point p in $\mathbb{R}^N$, that is $E(p, v) = p + v$.*

Definition 2.40. *A* ***tubular neighborhood*** $\mathrm{Tub}_r(M)$ *with radius $r > 0$ of M is a neighborhood U of M in $\mathbb{R}^N$ that is the diffeomorphic image under E of an open subset $V \subset NB$ of the form*

$$V = \{(p, v) \in NM \mid \|v\| < r\}.$$

[1]Actually, the plane drawn in Figure 10 is rather the set set $p + T_pM$, since T_pM is a linear vector space that always passes through the origin 0. In the same manner, we would need to write $p + N_pM$ for the line through p that is orthogonal to the grey plane. However, for simplicity reasons we will write $\mathbb{R}^N$ by identifying $v \in T_pM$ with $p + v$.

Remark 2.41. *Observe that the neighborhood $U_r = \{x \in \mathbb{R}^N \mid d(x, M) < r\}$ is a tubular neighborhood of M if and only if U is homotopic equivalent to M. Here*

$$d(x, M) = \min_{y \in M} \|x - y\|$$

is the distance of x to M. This can be illustrated by the example of a 1-dimensional submanifold that looks like a horseshoe drawn in Figure 11. If $r < d/2$ (see picture on the left), then E maps the open set $V = \{(p, v) \in NM \mid \|v\| < r\}$ diffeomorphically onto the the neighborhood U_r. In particular, both M and U_r are homotopy equivalent to a circle.

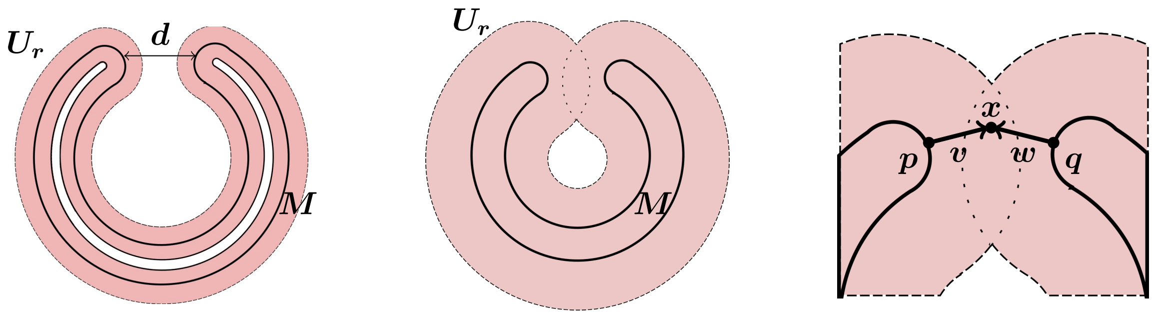

Figure 11: The point $x \in U_r$ is both $E(p, v)$ and $E(q, w)$.

If $r \geq d/2$ (see picture in the middle), then $E|_V : V \to U_r$ is not injective. If x is in the overlap of the tubular neighborhood (see picture on the right), there are two normal vectors that have its end point in x, each with base point on one side of the two arms of the manifold, that is $E(p, v) = x = E(q, w)$.

Example 2.42. *Consider the open set $M \subset \mathbb{R}^2$ as in Figure 12, that looks like the the symbol ∞.*

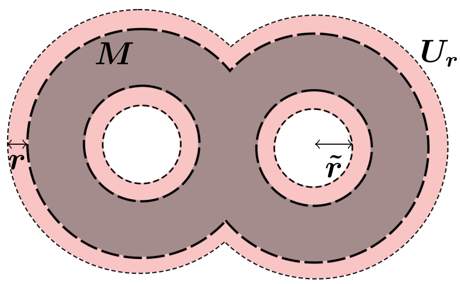

Figure 12: The tubular neighborhood U_r of the symbol ∞ with radius r.

If $U_r = \{x \in \mathbb{R}^2 \mid d(x, M) < r\}$ for $r < \tilde{r}$, then U_r has the same homotopy type as the manifold M since it deformation retracts onto M. Thus, U_r is a tubular neighborhood of M with radius r. If the radius r of the neighborhood U_r is greater then the radius $\tilde{r}$ of the holes, then U_r becomes contractible and is therefore homotopy equivalent to a point, while M has two loops and hence is not contractible. In that case U_r is not a tubular neighborhood.

We now associate the condition number τ to the submanifold M, in order to describe how large a manifold can thickened without changing its homotopy type. We will later use it to limit the curvature of the manifold.

Definition 2.43. *The **reach** of M, denoted by R, is the supremum of all $r > 0$, for which there is a tubular neighborhood of radius r.*

Definition 2.44. *Let R be the reach of M. The **condition number** τ of M is*

$$\tau := \begin{cases} 0, & \text{if } R = \infty, \\ \frac{1}{R}, & \text{otherwise.} \end{cases}$$

Remark 2.45. *By Remark 2.41, we can make the following two observations.*

1. *For any submanifold of $\mathbb{R}^N$ the condition number τ is finite.*
2. *We can also say that the reach of M is the supremum of all $r > 0$ for which the neighborhood $U_r = \{x \in \mathbb{R}^N \mid d(x, M) < r\}$ is homotopy equivalent to M.*

Definition 2.46. *Given M, we define the set*

$$G = \left\{x \in \mathbb{R}^N | \exists p \neq q \in M \text{ where } d(x, M) = \|x - p\| = \|x - q\|\right\}.$$

*We call the closure of G to be the **medial axis** of M.*

Example 2.47. *Let $M \subset \mathbb{R}^2$ be the horseshoe-like submanifold from Remark 2.41. Then the medial axis of M is the ray starting at the center C of M, that passes through the middle of the two arms of M. In fact, consider any point x on the vertical line as in Figure 13. The two points p and q on M, as in the definition of $G(M)$, are the base points of the two normal vectors v and w with $E(p, v) = x = E(q, w)$, as indicated. For any other point $y \in \mathbb{R}^2$ the minimal distance to M is unique.*

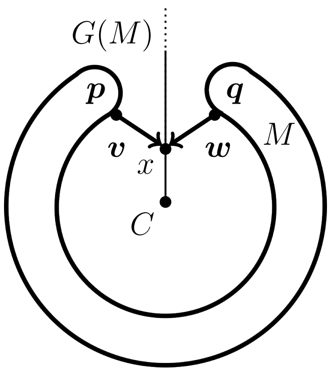

Figure 13: Middle axis of the manifold M. The point a shows the center of the two concentric open circles that bound M.

Definition 2.48. *For any point $p \in M$ we say that the* ***local feature*** *size $\sigma(p)$ is the distance of p to the medial axis.*

Then it follows directly from Remark 2.41 that $\tau = \inf_{p \in M} \sigma(p)$.

2.7 Vector Fields and Connections

Consider a curve γ on a submanifold of $\mathbb{R}^N$. We can think of the image of γ as a 1-dimensional submanifold of M. In particular, for every $t \in [0,1]$, we are provided with a point p on M. As described in Definition 2.22, we can determine the velocity vector of γ at any time t. Since the velocity vector $\dot{\gamma}(t)$ is a tangent vector at the point $p = \gamma(t)$, we can define a map that associates to every point $p = \gamma(t) \in M$ a tangent vector $v_p := \dot{\gamma}(t) \in T_pM$, as it can be seen in Figure 14.

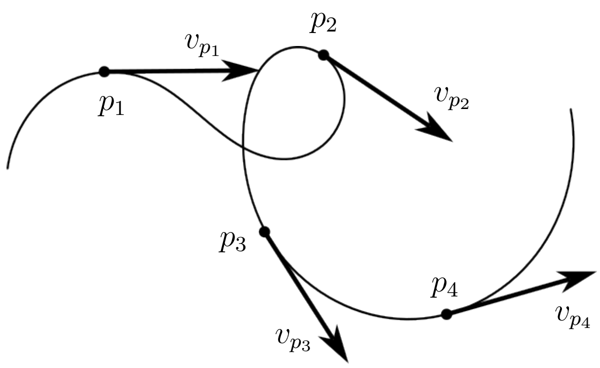

Figure 14: Velocity vector v_{p_i} at the point p_i, for $i = 1, 2, 3, 4$.

This way, we can think of a map $X \colon M \to TM$ from the manifold to its tangent bundle $TM = \{(p, v) \mid p \in M, v \in TM\}$ by sending p to the pair $(p, v_p) = (p, \dot{\gamma}(t))$, where $\gamma(t) = p$. Such maps are called vector fields on M. We now assume that M is a submanifold of $\mathbb{R}^N$ and we extend the notion of vector fields to all submanifolds of Euclidean spaces.

Definition 2.49. *A* ***vector field*** *X of M is a map $X \colon M \to TM$ that associates to every point $p \in M$ a tangent vector $X(p) = X_p \in T_pM$.*

Denote by $\mathcal{X}(M)$ the set of all vector fields on M. Let $X \in \mathcal{X}(M)$ be a vector field on M and $f \colon M \to \mathbb{R}$ be a smooth real-valued function. We can think of a new vector field fX using $(fX)_p := f(p)X_p \in T_pM$. This is not to be confused with the real-valued function Xf induced by the vector field as

$$(Xf)(p) := X_p(f) = \partial_{X_p} f(p) = \frac{d}{dt}\Big|_{t=0} f(p + tX_p),$$

we have seen in the proof of Proposition 2.25.

Consider the tangent vectors $\partial_j|_p = \frac{\partial}{\partial x^j}\big|_p$ that form a basis of $T_p\mathbb{R}^N$ and thus, also generate $T_pM \subset T_p\mathbb{R}^N$. We can therefore write the evaluation $X_p \in T_pM$ of a vector field as

$$X_p = (a_p^1\partial_1|_p + ... + a_p^N\partial_N|_p),$$

for some real numbers a_p^j that depend on the point p. Thus, by defining $X^j : M \to \mathbb{R}$ to be the function $X^j(p) := a_p^j$ and ∂_j to be the vector field on M via $p \mapsto \partial_j|_p$, we see that we can write the vector field $X \in \mathcal{X}(M)$ as

$$X = X^1\partial_1 + ... + X^N\partial_N,$$

which we abbreviate from now on by $X = X^1\partial_1 + ... + X^N\partial_N$.

Remark 2.50. *By the definition of the partial derivatives $\partial_j|_p$, for $j = 1, ..., N$, we see that ∂_j is just the operator that assigns to every point $p \in M$ the partial derivative operator of the j-th coordinate at the point p.*

Let $X, Y \in \mathcal{X}(M)$ be vector fields on M. We can write $Y = Y^1\partial_1 + ... + Y^N\partial_N$ for real-valued functions Y^j, for $j = 1, ..., N$ for the vector fields $\partial_j : p \mapsto \partial_j|_p$. We have also seen that for a real-valued function f, we can define another function by Xf, via $(Xf)(p) := \partial_{X_p} f(p)$. We can apply this to the functions Y^j to get functions (XY^j) for $j = 1, ..., N$. Using them as components with the generating vector fields ∂_j, we get a new vector field $(XY^1)\partial_1 + ... + (XY^N)\partial_N$.

Definition 2.51. *Let $X, Y \in \mathcal{X}(M)$ be vector fields on M. For $j = 1, .., N$, let Y^j be the component functions on Y, that is we have $Y = Y^1\partial_1 + ... + Y^N\partial_N$. Then we define XY to be the vector field defined componen-twise by $(XY)^j = XY^j$.*

We have seen that the map that associates to each point on a curve its velocity vector is a vector field on the curve. In Section 2.8, we introduce geodesics, which are a generalization of straight lines in $\mathbb{R}^N$. A curve on $\mathbb{R}^N$ is a straight line if and only if its acceleration is identically zero. However, on a arbitrary submanifold, we cannot just compute the acceleration of a path, since the velocity vectors do not even belong to the same tangent space. To make sense of the idea of differentiate the velocity vector field on an arbitrary submanifold, we introduce the notion of connections, which are essentially a coordinate-invariant set of rules for taking the directional derivatives of vector fields.

In the following, $C^\infty(M)$ is the class of smooth real-valued functions.

Definition 2.52. *A **linear connection** on M is a map $\nabla : \mathcal{X}(M) \times \mathcal{X}(M) \to \mathcal{X}(M)$, written $(X, Y) \mapsto \nabla_X Y$, satisfying the following properties:*

(a) $\nabla_X Y$ is linear over $C^\infty(M)$ in X:

$$\nabla_{fX_1+gX_2} Y = f\nabla_{X_1} Y + g\nabla_{X_2} Y \qquad \text{for } f, g \in C^\infty(M)$$

(b) $\nabla_X Y$ *is linear over* $\mathbb{R}$ *in* Y*:*

$$\nabla_X(aY_1 + bY_2) = a\nabla_X Y_1 + b\nabla_X Y_2 \qquad \textit{for } a, b \in \mathbb{R}$$

(c) $\nabla_X Y$ *satisfies the following product rule:*

$$\nabla_X(fY) = f\nabla_X Y + (Xf)Y \qquad \textit{for } f \in C^\infty(M)$$

Example 2.53. *An example of a connection is the* ***Euclidean connection*** $\overline{\nabla}$. *Let* X *and* $Y = Y^1\partial_1 + ... + Y^N\partial_N$ *be vector fields on* M. *Then we define* $\overline{\nabla}_X Y$ *to be the vector field*

$$\overline{\nabla}_X \left(Y^1\partial_1 + ... + Y^N\partial_N\right) := (XY^1)\partial_1 + ... + (XY^N)\partial_N.$$

In other words, $\overline{\nabla}_X Y$ *is the vector field whose components are the directional derivatives of the components of* Y *in direction of* X. *All the connections that we will use in this thesis are either* $\overline{\nabla}$ *itself or restriction of the Euclidean connections to submanifolds.*

Consider again the velocity vector field of a curve. Instead of mapping points on $\mathrm{im}(\gamma)$ to the velocity vectors, we can also map a given time to the velocity vector at that time. To generalize this idea we have the following definition.

Definition 2.54. *Let* $\gamma\colon I = [0,1] \to M$ *be a smooth curve. A* ***vector field along*** γ *is a smooth map* $V\colon I \to TM$ *such that* $V(t) \in T_{\gamma(t)}M$ *for all* $t \in I$. *Denote by* $\mathcal{T}(\gamma)$ *the space of all vector fields along* γ.

Example 2.55. *Let* $\gamma\colon I \to M$ *be a smooth curve. The most obvious example of a vector field along* γ *is its* ***velocity vector field*** $\dot\gamma \in \mathcal{T}(\gamma)$ *as shown in Figure 15.*

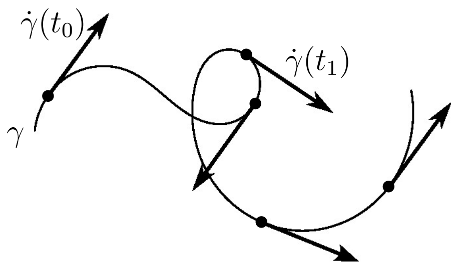

Figure 15: The velocity vector field along the curve γ. At each t the vector field gives the velocity vector $\dot\gamma(t)$ with base point $\gamma(t)$.

2.8 Geodesics and Parallel Transport of Vector Fields

In this section we define the notion of geodesics, which are essentially a generalization of straight lines. For example, if two points on a sphere have to be connected by a path on the sphere, the shortest path is not going to be a straight line, since the sphere does not contain any straight lines. In particular, if γ is any curve in $\mathbb{R}^N$ with velocity vector field $\dot\gamma$, we can check if γ is a straight line by checking if the acceleration $\ddot\gamma$ vanishes or not. However, in an arbitrary submanifold we cannot just compute $\ddot\gamma$, since $\dot\gamma(t)$ and $\dot\gamma(t')$, for $t \neq t' \in I$, live in different tangent spaces and the notion of subtracting has no meaning. This motivates a different way to define the acceleration of a curve.

Lemma 2.56. *Let ∇ be a linear connection on M. For each curve $\gamma\colon I \to M$, the connection ∇ determines a unique operator*

$$D_t\colon \mathcal{T}(\gamma) \to \mathcal{T}(\gamma)$$

satisfying the following properties:

(a) Linearity over $\mathbb{R}$:

$$D_t(aV + bW) = aD_tV + bD_tW, \qquad \text{for } a, b \in \mathbb{R}.$$

(b) Product rule:

$$D_t(fV) = \dot{f}V + fD_tV, \qquad \text{for } f \in C^\infty(I).$$

(c) If V is induced by a vector field $X \in \mathcal{X}(M)$, that is, for all $t \in I$ we have $V(t) = X(\gamma(t))$, then

$$D_tV(t) = \nabla_{d\gamma/dt}X.$$

For any $V \in \mathcal{T}(\gamma)$ we call D_tV the covariant derivative of V along γ.

We can use the notion of the covariant derivative to define the acceleration of a curve and what it means for a curve to be a geodesic.

Definition 2.57. *The **acceleration** of a smooth curve $\gamma\colon I \to M$ with respect to ∇ is $D_t\dot{\gamma}$, where $\dot{\gamma}$ is the velocity vector field of γ.*

If the connection ∇ in question is the connection of a submanifold of a Euclidean space relative to the standard metric, then, as expected, the acceleration of γ is

$$D_t\dot{\gamma}(t) := \frac{d}{dt}\dot{\gamma}(t) = \ddot{\gamma}(t),$$

that is the second derivative with respect to t in each component.

Definition 2.58. *A curve γ is called a **geodesic** if its acceleration vanishes, that is $D_t\dot{\gamma} \equiv 0$.*

The most obvious examples of geodesics are the ones in the Euclidean space $\mathbb{R}^N$. A geodesic between two points p and q in $\mathbb{R}^N$ is straight line connecting the two points. The following exampl e shows, that for different connections, one might has a path that is a geodesic relative to one connection, but not relative to the other connection.

Example 2.59. *Let $p \in S^1$ be a point on the unit circle. Using the standard parametrization of polar coordinates, indicated by square brackets, we have*

$$(1, \varphi) := (x(\varphi), y(\varphi)) = (\cos(\varphi), \sin(\varphi)), \ \text{for } \varphi \in [0, 2\pi),$$

we can write p in polar coordinates $p = (1, \varphi_0)$, for some $\varphi \in [0, 2\pi)$. Suppose that $0 < \varphi_0 \leq \pi$ and $\varphi \colon [0,1] \to S^1$ is the path in cartesian coordinates

$$\gamma(t) \coloneqq (\cos(t\varphi_0), \sin(t\varphi_0)).$$

Then, γ is the path in Figure 16 connecting the points $q = (1, 0)$ and p.

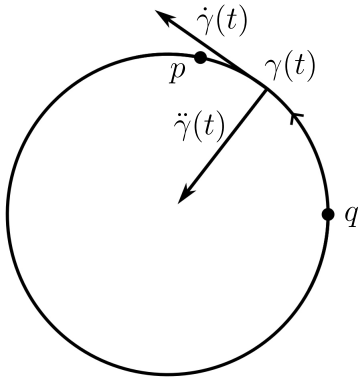

Figure 16: A geodesic γ between $q = (1, 0)$ and p on the sphere that is no geodesic in $\mathbb{R}^2$.

Its acceleration according to the Euclidean connection is just the second time derivative of the path with respect to the cartesian coordinates, that is,

$$\ddot\gamma(t) = \frac{d^2}{dt^2}\Big(\cos(t\varphi_0), \sin(t\varphi_0)\Big) = \Big(-\varphi_0^2\cos(t\varphi_0), \varphi_0^2\sin(t\varphi_0)\Big).$$

Since $\varphi_0 \neq 0$, we see that $\ddot\gamma \not\equiv 0$ and therefore, γ is not a geodesic in $\mathbb{R}^2$. However, regarding γ as a curve on S^1 with the polar coordinates, we see that

$$\gamma(t) = (1, t\varphi_0).$$

The covariant derivative on S^1 is the time derivative with respect to the polar coordinates. Therefore, we have

$$D_t\gamma(t) = \frac{d^2}{dt^2}(1, t\varphi_0) = \frac{d}{dt}(0, \varphi_0) = (0, 0).$$

Thus, with respect to the connection on S^1, we have $D_t\gamma(t) \equiv 0$ and hence, γ is a geodesic on S^1.

Example 2.60. *We consider the submanifold $S^2 \subset \mathbb{R}^3$ with the connection ∇ relative to the standard metric on the sphere. Let $p = (\varphi_1, \theta_1)$ and $q = (\varphi_2, \theta_2)$ be two points on the unit sphere, using spherical coordinates*

$$(\varphi, \theta) \mapsto \Big(\sin(\varphi)\cos(\theta),\ \sin(\varphi)\sin(\theta),\ \cos(\varphi)\Big),$$

for $\varphi \in \varphi \in [0,\pi], \theta \in [0, 2\pi)$*. Consider the curve* $\gamma\colon I \to S^2$*,*

$$\gamma(t) = (\varphi(t), \theta(t)) \coloneqq (1-t)\cdot(\varphi_1, \theta_1) + t \cdot (\varphi_2, \theta_2).$$

To see that γ *is a geodesic, we can calculate its acceleration* $D_t\dot{\gamma}$ *according to the spherical coordinates. In fact, we have*

$$D_t\dot{\gamma} = \frac{d}{dt}\Big(\dot{\gamma}(t)\Big) = \frac{d}{dt}\Big((\varphi_2, \theta_2) - (\varphi_1, \theta_1)\Big) = 0,$$

since p *and* q *do not depend on* t*. However, if we consider the acceleration with respect to the Euclidean connection* $\overline{\nabla}$*, we see that* $\overline{D}_t\dot{\gamma} \not\equiv 0$*, since the curve moves along the sphere and therefore has curvature, as it can be seen in Figure 17. From now on, We will refer to* $\overline{D}_t\dot{\gamma}$ *as* $\ddot{\gamma}$*.*

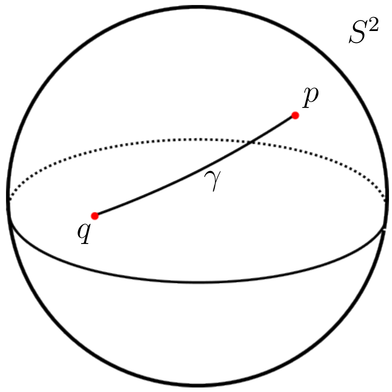

Figure 17: A geodesic curve γ on the unit sphere S^2.

Another notion that we need to simplify computation has to do with measuring the length of a path. Let $\gamma\colon I \to M$ be a smooth curve. Then

$$\text{length}(\gamma) = \int_0^1 \|\dot{\gamma}(t)\| dt.$$

We can parametrize γ by the length of the path.

Definition 2.61. *Let* $\gamma\colon I \to M$ *be a smooth curve with length* l*. We call the function* $s\colon I \to J = [0, l]$*,*

$$s(t) \coloneqq \int_0^t |\dot{\gamma}(r)| dr$$

the ***arc length function*** *of* γ*. We can solve for* t *to get the function* $t\colon J \to I$*. In particular, we can define a new curve* $\beta\colon J \to M$ *as* $\beta(s) \coloneqq \gamma \circ t(s) = \gamma(t(s))$*. We say that a curve of this form is said to be* ***parametrized by arc length****.*

An important feature of a curve that is parametrized by arc length is, that it has velocity vectors of length 1. To see this, let $\beta(s) \coloneqq (\gamma \circ t)(s)$ be a curve parametrized by arc length. Then

$$\|\beta'(s)\| = \left\|\frac{d}{ds}\gamma(t(s))\right\| = \left\|\dot{\gamma}(t(s)) \cdot \frac{dt}{ds}\right\| = \left\|\dot{\gamma}(t) \cdot \frac{1}{\dot{\gamma}(t)}\right\| = 1. \tag{1}$$

Definition 2.62. *Let M be a smooth manifold and $p, q \in M$ be two points. We define $d_M(p, q)$ to be the length of a geodesic that connects p and q.*

Observe that $d_M(p, q)$ is well-defined, since two geodesics connecting the same two points must have the same length.

A further notion that relates to the definition of geodesic is the notion of parallel vector fields along curves.

Definition 2.63. *Let ∇ be a linear connection on M. We say that a vector field V along a curve γ is **parallel** along γ with respect to ∇, if $D_t V(t) \equiv 0$.*

Remark 2.64. *Let γ be a smooth curve on M. and $V(t)$ the vector field along γ given by $V(t) = \dot{\gamma} = \frac{d\gamma}{dt}$. We observe that γ is a geodesic if and only if the velocity vector field $\dot{\gamma}$ is parallel along γ.*

For the notion of parallelism, we are provided with a theorem of uniqueness.

Theorem 2.65. *Given a curve $\gamma\colon I \to M$, $t_0 \in I$ and a vector $V_0 \in T_{\gamma(t_0)}M$, there exists a unique parallel vector field V along γ such that $V(t_0) = V_0$. This unique vector field V is called **parallel transport** of V_0 along γ.*

Example 2.66. *Let γ be any smooth curve on $\mathbb{R}^N$. The parallel vector fields $V(t)$ along γ with respect to the Euclidean connection are exactly the ones whose components are constant. Figure 18 shows that the parallel transport assigns vectors to every $t \in I$ that have exactly the same direction and the same length, namely the one of V_0.*

2.9 The Second Fundamental Form

We can now introduce the second fundamental form, which measures the curvature of a submanifold within of a Euclidean space. We will show that the second fundamental form is bounded from above by the condition number τ in all directions. For this, let M be a Riemannian submanifold of $\mathbb{R}^N$.

Consider the Euclidean connection $\overline{\nabla}$ on $\mathbb{R}^N$. If X and Y are vector fields on M, they extend to vector fields $\overline{X}$ and $\overline{Y}$ on $\mathbb{R}^N$, that is $\overline{X}$ and $\overline{Y}$ are vector fields on $\mathbb{R}^N$ with $\overline{X}|_M \equiv X$ and $\overline{Y}|_M \equiv Y$, respectively. Therefore we can define

$$\nabla_X Y := \left(\overline{\nabla}_{\overline{X}} \overline{Y}\right)^T,$$

that is, the tangential projection of the vector field $\overline{\nabla}_{\overline{X}} \overline{Y}$. It can be shown that this is in fact the Riemannian connection relative to the metric induced on M (see Carmo [5]).

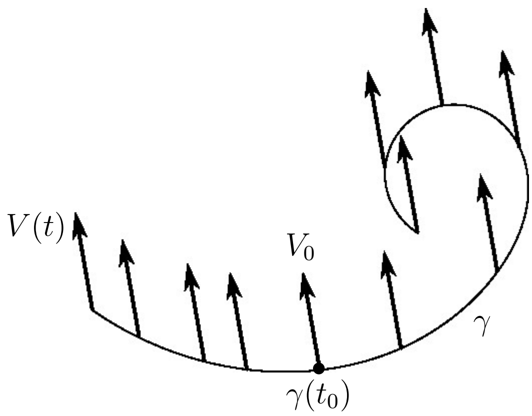

Figure 18: Parallel vector field $V(t)$ with respect to $\overline{\nabla}$ whose components are constant. In particular, $V(t)$ is the parallel transport of V_0 along γ.

Definition 2.67. *If $\mathcal{X}(M)$ is the space of vector fields on M, we can define a map $B\colon \mathcal{X}(M) \times \mathcal{X}(M) \to \mathcal{X}(M)^N$ via*

$$B(X,Y) \coloneqq \overline{\nabla}_{\overline{X}}\overline{Y} - \nabla_X Y = \left(\overline{\nabla}_{\overline{X}}\overline{Y}\right)^N,$$

where $\overline{X}$ and $\overline{Y}$ are any extensions of X and Y, respectively.

That this definition does not depend on the choice of the extensions, can be seen by the following: If $\overline{X}_1$ and $\overline{Y}_1$ are two different extensions of X and Y, then we have

$$B(X,Y) - B(X_1,Y) = \left(\overline{\nabla}_{\overline{X}}\overline{Y} - \nabla_X Y\right) - \left(\overline{\nabla}_{\overline{X}_1}\overline{Y} - \nabla_X Y\right) = \overline{\nabla}_{\overline{X}-\overline{X}_1}\overline{Y}$$

which vanishes on M, since $\overline{X} - \overline{X}_1 = 0$ on M. Similarly, we have

$$B(X,Y) - B(X,Y_1) = \left(\overline{\nabla}_{\overline{X}}\overline{Y} - \nabla_X Y\right) - \left(\overline{\nabla}_{\overline{X}}\overline{Y}_1 - \nabla_X Y\right) = \overline{\nabla}_{\overline{X}}\left(\overline{Y} - \overline{Y}_1\right).$$

This also vanishes on M, since $\overline{Y} - \overline{Y}_1 = 0$ on M. Hence, the map B is well-defined. Moreover, the map B is bilinear over $C^\infty(M)$ and symmetric. However, before we can prove this, we need to define the notion of Lie brackets, since we use them to prove the symmetry of B.

Definition 2.68. *For vector fields $X, Y \in \mathcal{X}(M)$, we define the* ***Lie bracket*** *of X and Y to be the vector field defined as $[X,Y] \coloneqq XY - YX$.*

The vector fields XY and YX are defined as in Definition 2.51. Thus, if we write $X = X^1\partial_1 + ... + X^N\partial_N$ and $Y = Y^1\partial_1 + ... + Y^N\partial_N$ for real-valued functions X^j and Y^j, for $j = 1, ..., N$, we see that

$$[X,Y] = XY - YX = \left(\left(XY^1\right)\partial_1 + ... + \left(XY^N\right)\partial_N\right) - \left(\left(YX^1\right)\partial_1 + ... + \left(YX^N\right)\partial_N\right).$$

Component-wise we have

$$[X,Y]^i = \left(X^j \partial_j X^i\right)\partial_i - \left(Y^j \partial_j X^i\right)\partial_i = \left(X^j \partial_j X^i - Y^j \partial_j X^i\right)\partial_i.$$

Since we are only interested on what happens on M and T_pM, we will omit the notation $\overline{X}$. If the vector field $X \in \mathcal{X}(M)$ appears in relation with $\overline{\nabla}$, this means that we apply $\overline{\nabla}$ to an extension of X.

Lemma 2.69. *The map $B\colon \mathcal{X}(M) \times \mathcal{X}(M) \to \mathcal{X}(M)^N$ is $C^\infty(M)$-bilinear and symmetric.*

Proof. The $C^\infty(M)$-linearity follows directly form the $C^\infty(M)$-linearity of $\overline{\nabla}$ and ∇ in both arguments.

To show that B is symmetric, we first prove that the Euclidean connection $\overline{\nabla}$ and thus, also the connection ∇, satisfies the following symmetry condition

$$\overline{\nabla}_X Y = \overline{\nabla}_Y X + [X,Y].$$

By Writing $X = X^1\partial_1 + ... + X^N\partial_N$ and $Y = Y^1\partial_1 + ... + Y^N\partial_N$, we get by definition of the Euclidean connection in each component

$$\begin{aligned}\left(\overline{\nabla}_X Y - \overline{\nabla}_Y X\right)^i &= \left(\overline{\nabla}_X Y\right)^i - \left(\overline{\nabla}_Y X\right)^i \\ &= \left(XY^i\right)\partial_i - \left(YX^i\right)\partial_i \\ &= \left(XY^i\right)\partial_i - \left(YX^i\right)\partial_i \\ &= \left(X^j \partial_j X^i - Y^j \partial_j X^i\right)\partial_i.\end{aligned}$$

Comparing the components, we see that these are exactly the ones of the Lie Bracket. Reordering the equation gives the desired symmetry. Thus, applying this symmetry condition to $B(X,Y)$, we get

$$\begin{aligned}B(X,Y) &= \overline{\nabla}_Y X + [X,Y] - \nabla_Y X - [X,Y] \\ &= \overline{\nabla}_{\overline{Y}}\overline{X} - \nabla_Y X \\ &= B(X,Y),\end{aligned}$$

where we used in the second equality that any extension of X restricted to M coincides with X. □

Before we continue with defining the second fundamental form, we will give the following two technical remarks.

Remark 2.70. *Let $\gamma\colon I \to M$ be a smooth curve, ∇ a linear connection on M. If $V, W \in \mathcal{T}(\gamma)$ are vector fields along γ that are induced by some vector fields $X, Y \in \mathcal{X}(M)$, we will write $\nabla_V W$ for $\nabla_X Y$.*

Remark 2.71. *If $\gamma\colon I \to \mathbb{R}^N$ is a curve and $V(t)$ is a vector field along γ, then we observe that the covariant derivative of V along γ coincides with $\frac{dV}{dt}(t)$, which is the vector field along γ, whose components are the time-derivative of the components of V. That is if $V(t) = V^1(t)\partial_N + ... + V^N(t)\partial_N$, for components $V^i\colon I \to \mathbb{R}$, then*

$$\overline{D}_t V(t) = \frac{dV^1}{dt}(t)\partial_1 + ... + \frac{dV^N}{dt}(t)\partial_N.$$

Let us get back to the map B defined above. Since B is bilinear, the value $B(X,Y)(p)$ only depends on the values $X(p)$ and $Y(p)$, hence B induces a map $B\colon T_pM \times T_pM \to N_pM$, which is bilinear and symmetric.

We are now in the position to define the second fundamental form. For $p \in M$ and $\eta \in N_pM$ with $\|\eta\| = 1$, the map $B_\eta\colon T_pM \times T_pM \to \mathbb{R}$, defined by

$$B_\eta(v,w) \coloneqq \langle B(v,w), \eta\rangle$$

is a symmetric bilinear form.

Definition 2.72. *Let $p \in M \subset \mathbb{R}^N$ and $\eta \in N_PM$ with $\|\eta\| = 1$. The* ***second fundamental form*** *at p along the normal vector η is the quadratic form $\mathbb{I}_\eta\colon T_pM \to \mathbb{R}$ defined by*

$$\mathbb{I}_\eta(u) \coloneqq B_\eta(u,u) = \langle \eta, B(u,u)\rangle.$$

Definition 2.73. *Let $\eta \in N_pM$ with $\|\eta\| = 1$. The* ***norm of the second fundamental form*** *in direction η is given by*

$$\lambda_\eta \coloneqq \sup_{u\in T_pM} \frac{\mathbb{I}_\eta(u)}{\langle u,u\rangle}.$$

Observe that B_η is associated to a linear self-adjoint[2] operator $L_\eta\colon T_pM \to T_pM$ by

$$\langle L_\eta u, v\rangle = B_\eta(u,v) = \langle \eta, B(u,v)\rangle.$$

Thus, we can also express λ_η using this operator by

$$\lambda_\eta = \sup_{u\in T_pM} \frac{\mathbb{I}_\eta(u)}{\langle u,u\rangle} = \sup_{u\in T_pM} \frac{\langle u, L_\eta u\rangle}{\langle u,u\rangle}.$$

In order to to compute the second fundamental form, it makes sense to find an expression for the operator L_η. We get the following proposition.

Proposition 2.74. *Let $p \in M$, $x \in T_pM$ and $\eta \in (T_pN)$ with $\|\eta\| = 1$. Let N be a local*

[2] A self-adjoint operator with respect to the inner product $\langle\cdot,\cdot\rangle$ is a linear operator L with $\langle L\cdot,\cdot\rangle = \langle\cdot, L\cdot\rangle$.

extension η normal to M and X a local extension of x on M. Then

$$L_\eta(x) = -\left(\overline{\nabla}_x N\right)^T = -\left(\overline{\nabla}_X N\big|_p\right)^T.$$

Proof. Let $y \in T_pM$ and let Y be a local extension of y on M. Then $\langle N, Y\rangle = 0$ and therefore

$$\langle L_\eta(x), y\rangle = \langle B(X,Y)(p), \eta\rangle.$$

By definition and the fact, that $\left\langle \left(\overline{\nabla}_X Y\right)^T, N\right\rangle$ vanishes, we get

$$\big\langle L_\eta(x), y\big\rangle = \langle \overline{\nabla}_X Y, N\rangle(p) = -\langle \overline{\nabla}_X N, Y\rangle(p) = \langle \overline{\nabla}_x N, y\rangle.$$

□

Example 2.75. *Let M be a plane in $\mathbb{R}^3$ of the form*

$$M\colon ax + by + cz = 0, \text{ for } a, b, c \in \mathbb{R}^3.$$

Let $v \in \mathcal{X}(M)$ be a tangent vector and

$$n = \frac{1}{\lambda}(a,b,c),$$

with $\lambda = \sqrt{a^2+b^2+c^2}$ be a unit normal vector to M. We want to compute $L_\eta(v)$. By Proposition 2.74, we have

$$L_\eta(v) = -\left(\overline{\nabla}_v \eta\right)^T = -\left(\overline{\nabla}_v\left(\frac{1}{\lambda}(a,b,c)\right)\right)^T.$$

By definition, this is the tangential component of the directional derivative of n along v, that is

$$L_\eta(v) = -\partial_v\left(\frac{1}{\lambda}(a,b,c)\right) = -\frac{1}{\lambda}\left(\partial_v a, \partial_v b, \partial_v c\right) = (0,0,0).$$

Since v was arbitrary in T_pM, we see that the norm of the second fundamental form of a plane is 0. This makes sense, since a plane in $\mathbb{R}^3$ has no curvature.

Example 2.76. *Let $M = S_r^2$ be the sphere with radius r and η the normal vector field that assigns to every $p \in M$ the normal vector at p pointing outward with unit norm. Fix a point $p \in S_r^2$ and let $v \in T_pS_r^2$. If $\gamma\colon [0,1] \to S_r^2$ is a curve with $\gamma(0) = p$ and $\dot\gamma(0) = v$, then we can compute*

$$L_\eta(v) = -\overline{\nabla}_v \eta,$$

which is the directional derivative of n along v. But since $v = \dot\gamma(0)$, this is the same as

$$\frac{d}{dt}\Big|_{t=0} n\circ\gamma(t).$$

But $n \circ \gamma(t)$ is the normal unit vector at $\gamma(t)$ pointing outward, which is given by $\frac{\gamma(t)}{\|\gamma(t)\|}$. Since $\gamma(t)$ is on S_r^2, we get

$$L_\eta(v) = \frac{d}{dt}\Big|_{t=0} \frac{\gamma(t)}{\|\gamma(t)\|} = \frac{d}{dt}\Big|_{t=0} \frac{\gamma(t)}{r} = \frac{1}{r}\dot{\gamma}(0) = \frac{1}{r}v.$$

By replacing v with $\frac{v}{\|v\|}$, we see that the norm of the second fundamental form of S_r^2 is $\lambda_\eta = \frac{1}{r}$. But since we know that for S_r^2 the reach is $R = r$, we see that $\lambda_\eta = \tau$, that is the norm of the second fundamental form of S_r^2 is just the condition number.

2.10 Curvature and the Condition Number

In Example 2.76 we have seen that the norm of the second fundamental form of a sphere is the same as its condition number τ. In this section we will show that, for an arbitrary submanifold of $\mathbb{R}^N$, the second fundamental form in all directions is bounded by the condition number, that is $\lambda_\eta \leq \tau$ for all normal unit vectors of the submanifold. Recall the norm of the second fundamental form in direction of η, which is defined by

$$\lambda_\eta = \sup_{u \in T_pM} \frac{\mathbb{I}_\eta(u)}{\langle u, u \rangle} = \sup_{u \in T_pM} \frac{\langle u, L_\eta u \rangle}{\langle u, u \rangle}.$$

Let us now state the following proposition.

Proposition 2.77. *If M is a submanifold of $\mathbb{R}^N$ with condition number τ, then the norm of the second fundamental form is bounded by τ in all directions. In other words, for all $p \in M$ and for all $\eta \in N_pM$ with $\|\eta\| = 1$, we have*

$$\lambda_\eta \leq \tau.$$

Proof. We proof by contradiction. Let R be the reach of the submanifold M. Let us assume that

$$\lambda_\eta > \tau = \frac{1}{R}.$$

This means that there is a point $p \in M$ and a tangent vector $u \in T_pM$, such that

$$\lambda_\eta \geq \frac{\langle \eta, B(u, u) \rangle}{\langle u, u \rangle} > \frac{1}{R}.$$

By replacing u with $u/\|u\|$, we can assume that u has unit norm, so

$$\lambda_\eta \geq \langle \eta, B(u, u) \rangle > \frac{1}{R}.$$

Consider a geodesic curve $\gamma\colon J \to M$ parametrized by arc length, such that $\gamma(0) = p$ and $\dot{\gamma}(0) = u$. Here $J = [0, l]$, where l is the arc length of γ. Consider the point $x := p + R \cdot \eta$, that is the point at distance R from p in the direction η. By definition of the reach of a submanifold, we see that p is the closest point on M to the center of the R-ball given by x, hence for all $t \in I$ we have

$$\|\gamma(t) - x\|^2 \geq \|R\eta\|^2 = R^2. \tag{2}$$

By using

$$\begin{aligned}\|\gamma(t) - x\|^2 &= \langle \gamma(t) - x, \gamma(t) - x \rangle \\ &= \langle \gamma(t), \gamma(t) \rangle - 2\langle \gamma(t), x \rangle + \langle x, x \rangle \\ &= \langle \gamma(t), \gamma(t) \rangle - 2R\langle \gamma(t), \eta \rangle + R^2,\end{aligned}$$

the inequality (2) becomes

$$\langle \gamma(t), \gamma(t) \rangle - 2R\langle \gamma(t), \eta \rangle \geq 0. \tag{3}$$

Now let us define the function $g \colon I \to \mathbb{R}$ via $g(t) = \langle \gamma(t), \gamma(t) \rangle - 2R\langle \gamma(t), \eta \rangle$. Then we have $g(0) = 0$, because $\gamma(0) = 0$. Moreover,

$$\begin{aligned}g'(t) &= \langle \dot\gamma(t), \gamma(t) \rangle + \langle \gamma(t), \dot\gamma(t) \rangle - 2R\langle \dot\gamma(t), \eta \rangle \\ &= 2\langle \dot\gamma(t), \gamma(t) \rangle - 2R\langle \dot\gamma(t), \eta \rangle.\end{aligned}$$

Since $\gamma(0) = 0$ and $\dot\gamma(t) \in T_pM$ and $\eta \in N_pM$, we have $g'(0) = 0$. Finally, we get

$$g''(t) = 2\langle \ddot\gamma(t), \gamma(t) \rangle + 2\langle \dot\gamma(t), \dot\gamma(t) \rangle - 2R\langle \ddot\gamma(t), \eta \rangle.$$

Again, since $\gamma(0) = 0$ and $\langle \dot\gamma(t), \dot\gamma(t) \rangle = 1$, due to the parametrization by arc length (see (1)), we get

$$g''(0) = 2 - 2R\langle \ddot\gamma(0), \eta \rangle.$$

By Remark 2.64, γ being a geodesic is equivalent to the velocity vector field $\dot\gamma = \frac{d\gamma}{dt}$ being parallel. We now consider

$$\begin{aligned}B\left(\frac{d\gamma}{dt}, \frac{d\gamma}{dt}\right) &= \overline{\nabla}_{\frac{d\gamma}{dt}} \frac{d\gamma}{dt} - \nabla_{\frac{d\gamma}{dt}} \frac{d\gamma}{dt} \\ &= \overline{D}_t \frac{d\gamma}{dt} - D_t \frac{d\gamma}{dt},\end{aligned}$$

where $\overline{D}_t$ is the covariant derivative with respect to the Euclidean connection $\overline{\nabla}$ and D_t is the covariant derivative with respect to the connection relative to the metric on M, that is

$$\nabla_{\frac{d\gamma}{dt}} \frac{d\gamma}{dt} = \left(\overline{\nabla}_{\frac{d\gamma}{dt}} \frac{d\gamma}{dt}\right)^T.$$

Since $\frac{d\gamma}{dt}$ is parallel, we have $D_t \frac{d\gamma}{dt} = 0$ and therefore

$$B\left(\frac{d\gamma}{dt}, \frac{d\gamma}{dt}\right) = \overline{D}_t \frac{d\gamma}{dt} = \ddot{\gamma}(t).$$

Thus, by our assumption we see, that

$$\langle \eta, B(u,u)\rangle = \left\langle \eta, B\left(\frac{d\gamma}{dt}, \frac{d\gamma}{dt}\right)\bigg|_{t=0} \right\rangle = \langle \eta, \ddot{\gamma}(0)\rangle > \frac{1}{R}.$$

Using this, we can compute

$$g''(0) = 2 - 2R\langle \dot{\gamma}(0), \eta\rangle > 2 - 2R\frac{1}{R} = 0.$$

Therefore, g has a maximum at $t = 0$ with $g(0) = 0$ and hence, there exists $t^* \in I$ with $g(t^*) < 0$, which is a contradiction to (3), that is, that $g(t) \geq 0$ for all $t \in I$. Thus, we have shown that $\lambda_\eta \leq \frac{1}{R} = \tau$. □

Remark 2.78. *Recall that, for a normal unit vector η, the map B from Definition 2.72 induces a linear operator L_η via*

$$\langle L_\eta u, v\rangle = \langle \eta, B(u,v)\rangle.$$

An important feature of this self-adjoint operator L_η is that $\|L_\eta\| \leq \tau$, where

$$\|L_\eta\| = \sup_{u,v \in T_pM\setminus\{0\}} \frac{\langle v, L_\eta u\rangle}{\|v\|\|u\|},$$

is the standard operator norm.

We have seen that for a compact submanifold the condition number is finite. Hence, Proposition 2.77 shows that the norm of the second fundamental form of our manifold M is bounded. Therefore, the manifold cannot curve too much locally. Thus, the angle between the tangent spaces of two different points that are nearby cannot be too large.

Let p and q be two points with associated tangent spaces T_pM and T_qM, where we consider them transported to the origin according to the Euclidean connection on $\mathbb{R}^N$. We can compare tangent vectors in each tangent space with each other. Thus, for any vectors $u \in T_pM$ and $v \in T_qM$, we define the angle θ between them, using the standard formula

$$\cos(\theta) = \frac{|\langle u', v'\rangle|}{\|u'\| \cdot \|v'\|},$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product in $\mathbb{R}^{\mathbb{N}}$ and u', v' are the vectors obtained by parallel transport of u and v, respectively, to the origin.

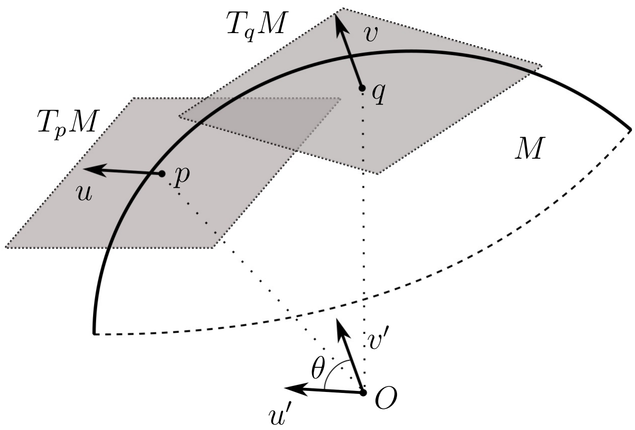

Figure 19: Angle between two tangent vectors u and v relative to different points p and q, after parallel transporting them to the origin O.

From now on we omit the prime notation, that is, when we write $\langle u, v\rangle$ for two tangent vectors in different tangent spaces, we mean the standard inner product between the parallel transport of u and v to the origin according to the Euclidean connection on $\mathbb{R}^N$. In order to define the angle between the two tangent spaces, we maximize over T_pM and minimize over T_qM, that is, the angle ϕ between T_pM and T_qM is defined by

$$\cos(\phi) = \min_{u\in T_pM} \max_{v\in T_qM} \frac{\langle u, v\rangle}{\|u\|\|v\|}.$$

Moreover, recall the $d_M(p, q)$, that is, the length of a geodesic connecting p and q.

Proposition 2.79. *Let M be a submanifold of $\mathbb{R}^N$ with condition number τ. Let p and q be two points on M. Let ϕ be the angle between the tangent spaces T_pM and T_qM. Then*

$$\cos(\phi) \geq 1 - \tau d_M(p, q).$$

Proof. Consider two points $p, q \in M$ connected by a geodesic curve $\gamma\colon J = [0, l] \to M$, that is parametrized by arc length, such that $\gamma(0) = p$ and $\gamma(L) = q$, where L is the arc length of the curve γ. Let $v_p \in T_pM$ with $\|v_p\| = 1$ and consider $V(t)$ to be the parallel transport of v_p along γ, with respect to the connection ∇ on M. Then we have $v_q := V(l) \in T_qM$. Moreover, since V is parallel, we have that $\langle V(t), V(t)\rangle$ is constant on all of J and since $\langle V(0), V(0)\rangle = \|v_p\|^2 = 1$, we have $\langle V(t), V(t)\rangle = 1$ for all $t \in J$.

Let us have a closer look at the vector field V along γ. We would like to express the vector $V(l) = v_q$ in terms of $V(0) = v_p$. To do this, we will use the fundamental theorem of calculus. According to Remark 2.71, the components of $\overline{D}_tV(t)$ are $\frac{dV^i}{dt}(t)$, for real-valued functions $V^i\colon J \to \mathbb{R}$. Since $\overline{D}_tV(t)$ is another vector field along the curve γ, we can integrate along γ. Using the fundamental theorem of calculus on the components we then get

$$\int_0^l \left(\overline{D}_t V\right)^i dt = \int_0^l \frac{dV^i}{dt}(t)dt = \left(V^i(l) - V^i(0)\right).$$

If we define w to be the vector with components $w^i = (V^i(l) - V^i(0))$, we can write $V(l) = V(0) + w$, where this addition is to be understood after parallel transportation to the origin. But then we have

$$\begin{aligned}\langle V(0), V(l)\rangle &= \langle V(0), V(l) + w\rangle \\ &= \langle V(0), V(0)\rangle + \langle V(0), w\rangle \\ &= 1 + \langle V(0), w\rangle.\end{aligned}$$

If θ is the angle between v_p and v_q, we then get

$$\begin{aligned}\cos(\theta) &= \|\langle V(0), V(l)\rangle\| \\ &= \|1 + \langle V(0), w\rangle\| \\ &\geq 1 - \|\langle V(0), w\rangle\|,\end{aligned}$$

since $\|v_p\| = \|v_q\| = 1$ and thus, we get

$$\cos(\theta) \geq 1 - \|w\|. \tag{4}$$

Observe that v_p was arbitrary. If we chose v_p, such that

$$\cos(\phi) = \min_{u \in T_pM} \max_{v \in T_qM} \frac{\langle u, v\rangle}{\|u\|\|v\|} = \max_{v \in T_qM} \frac{\langle v_p, v\rangle}{\|v_p\|\|v\|} = \max_{v \in T_qM} \langle v_p, v\rangle,$$

where we assume $\|v_p\| = \|v\| = 1$ in the last equality, this implies, that

$$\cos(\phi) \geq \cos(\theta) \geq 1 - \|w\|. \tag{5}$$

Observe that we have $\overline{D}_t V(t) = \overline{\nabla}_{\frac{d\gamma}{dt}} V(t)$ and since V is parallel along γ with respect to ∇, we see that

$$\left(\overline{\nabla}_{\frac{d\gamma}{dt}} V(t)\right)^T = \nabla_{\frac{d\gamma}{dt}} V(t) = 0.$$

Therefore, $\overline{\nabla}_{\frac{d\gamma}{dt}} V(t) = \left(\overline{\nabla}_{\frac{d\gamma}{dt}} V(t)\right)^N$. But note that we defined the bilinear map B to be the component of $\overline{\nabla}_{\frac{d\gamma}{dt}} V(t)$ in the normal direction, see (2.67). Hence, we get

$$\overline{D}_t V(t) = B\left(\frac{d\gamma}{dt}, V(t)\right).$$

Now consider $\eta := \frac{1}{\|\overline{D}_t V(t)\|} \cdot \overline{D}_t V(t)$, that is the normed vector in direction of $\overline{D}_t V(t)$. Then, we see that

$$\begin{aligned} \|\overline{D}_t V(t)\| &= \frac{1}{\|\overline{D}_t V(t)\|} \langle D_t V(t), D_t V(t) \rangle \\ &= \langle \eta, D_t V(t) \rangle \\ &= \langle \eta, B\left(\frac{d\gamma}{dt}, V(t)\right) \rangle \\ &= \langle \frac{d\gamma}{dt}, L_\eta V(t) \rangle \le \|\dot{\gamma}(t)\| \|L_\eta V(t)\|. \end{aligned}$$

But by Remark 2.78, the operator norm of L_η is bounded from above by τ and thus, we have

$$\overline{D}_t V(t) \le \|\dot{\gamma}(t)\| \|L_\eta\| \le \|\dot{\gamma}(t)\| \cdot \tau.$$

Using the definition of w, we get

$$\begin{aligned} \|w\| &= \| \int_0^l \overline{D}_t V(t) dt \| \\ &\le \int_0^l \|\overline{D}_t V(t)\| dt \\ &\le \tau \int_0^l \|\dot{\gamma}(t)\| dt \le \tau \cdot d_M(p, q), \end{aligned}$$

where the last inequality follows from the fact that γ is a geodesic from p to q. Combining (5) and what we just found, we get that

$$\cos(\phi) \ge 1 - \|w\| \ge 1 - \tau d_M(p, q),$$

which finishes the proof. □

The next proposition shows a relationship between the number $d_M(p, q)$ between two points and the Euclidean distance. This can then be used to give a bound on the angle between $T_p M$ and $T_q M$, that only depends on the condition number τ and the Euclidean distance $\|p - q\|_{\mathbb{R}^N}$, where we think of p and q as points in $\mathbb{R}^N$.

Proposition 2.80. *Let M be a submanifold of $\mathbb{R}^N$ with condition number τ. Let $d \le \frac{\tau}{2}$ and $p, q \in M$ be two points in M with Euclidean distance $\|p - q\|_{\mathbb{R}^N} = d$. Then $d_M(p, q)$ is bounded by*

$$d_M(p, q) \le \frac{1}{\tau} \left(1 - \sqrt{1 - 2d\tau}\right).$$

Proof. Let $d \le \frac{\tau}{2}$ and consider two points $p, q \in M$ with Euclidean distance $\|p - q\|_{\mathbb{R}^N} = d$ and let $\gamma\colon J = [0, l] \to M$ be a geodesic curve starting at p, parametrized by arc length, that connects the points p and q. If $s := d_M(p, q)$, then we have $\gamma(0) = p$ and $\gamma(s) = q$.

Since γ is a geodesic, by Remark 2.64, we see that the vector field $\frac{d\gamma}{dt}$ along γ is parallel with respect to ∇, and we have seen in the proof of Proposition 2.77 that

$$\ddot{\gamma} = B(\dot{\gamma}, \dot{\gamma}).$$

Using the bound from Proposition 2.77, we get that for any $\eta \in N_{\gamma(t)}M$ with $\|\eta\| = 1$, we have

$$\|\ddot{\gamma}\| = \|B(\dot{\gamma}, \dot{\gamma})\| \leq \lambda_\eta \leq \tau. \tag{6}$$

We now relate the Euclidean distance d to the geodesic distance s. Observe that by the fundamental theorem of calculus we have

$$\gamma(s) - \gamma(0) = \int_0^s \dot{\gamma}(t)dt,$$

where we consider p and q as points in $\mathbb{R}^N$. Moreover, using the fundamental theorem of calculus on $\dot{\gamma}$, we get

$$\dot{\gamma}(t) = \dot{\gamma}(0) + \int_0^t \ddot{\gamma}(r)dr.$$

By defining $u(t) := \int_0^t \ddot{\gamma}(r)dr$, where this integral is to be thought in every component of $\ddot{\gamma}(t)$, we get $\dot{\gamma}(t) = \dot{\gamma}(0) + u(t)$. In particular, we have for every $t \in J$, using (6), that

$$\|u(t)\| = \|\int_0^t \ddot{\gamma}(r)dr\| \leq \int_0^t \|\ddot{\gamma}(t)\|dr \leq \tau \cdot t.$$

But now we can bring everything together to get

$$\begin{aligned}
\|\gamma(s) - \gamma(0)\|_{\mathbb{R}^N} &= \|\int_0^s \dot{\gamma}(t)dt\| \\
&= \|\int_0^s (\dot{\gamma}(0) + u(t))\, dt\| \\
&\geq \|\int_0^s \dot{\gamma}(0)dt\| - \|\int_0^t \ddot{\gamma}(r)dr\| \\
&\geq s \cdot \|\dot{\gamma}(0)\| - \int_0^t \|\ddot{\gamma}(r)\|dr \\
&\geq s - \tau \cdot \int_0^s t\ dt,
\end{aligned}$$

where we used, that γ is parametrized by arc length in the last line. But this gives us

$$d = \|\gamma(s) - \gamma(0)\|_{\mathbb{R}^N} \geq s - \frac{1}{2}\tau s^2. \tag{7}$$

The inequality in (7) is satisfied only if

$$s \leq \frac{1}{\tau} - \frac{1}{\tau}\sqrt{1 - 2d\tau} \quad \text{or} \quad s \geq \frac{1}{\tau} + \frac{1}{\tau}\sqrt{1 - 2d\tau}.$$

However, if $d = 0$ that means that $p = q$ and thus $s = 0$ and we observe that only the first inequality applies to this case. But the first inequality gives us exactly

$$d_M(p, q) \leq \frac{1}{\tau}\left(1 - \sqrt{1 - 2d\tau}\right).$$

□

Example 2.81. *Let us again consider the two points $q = (1, 0)$ and $p = (1, \varphi_0)$ in polar coordinates on the unit circle from Example 2.59. We want to Proposition 2.79 to find a lower bound on the angle ϕ between T_qS^1 and T_pS^1. Proposition 2.79 gives us*

$$\cos(\phi) \geq 1 - \tau d_{S^1}(p, q).$$

For the unit circle we have $\tau = 1$. Moreover, we have seen in Example 2.59 that the curve $\gamma\colon [0, 1] \to S^1$ given by

$$\gamma(t) := (\cos(t\varphi_0), \sin(t\varphi_0)),$$

is a geodesic between p and q. Thus, we are left with computing its length.

$$\text{length}(\gamma) = \int_0^1 \|\dot{(\gamma)}(t)\| dt.$$

The norm of the velocity vector $\dot{\gamma}$ at t is given by

$$\|\dot{\gamma}(t)\| = \Big\|\big(- \varphi_0 \sin(t\varphi_0), \varphi_0 \cos(t\varphi_0)\big)\Big\| = \sqrt{\varphi_0} = \varphi_0,$$

since $\varphi_0 \geq 0$. Hence, ϕ is bounded by the relation

$$\cos(\phi) \geq 1 - \tau d_{S^1}(p, q) = 1 - \varphi_0.$$

Remark 2.82. *Combining Proposition 2.79 and Proposition 2.80 we get, that if p and q are two points in M with $d = \|p - q\| \leq \tau/2$, then*

$$\begin{aligned}\cos(\phi) &\geq 1 - \tau d_M(p, q) \\ &\geq 1 - \tau \cdot \frac{1}{\tau}\left(1 - \sqrt{1 - 2d\tau}\right) \\ &= \sqrt{1 - 2d\tau}.\end{aligned}$$

With this proposition we have finished our preparation in terms of background and in terms of differential geometry.

3 Probability Preliminaries

In this chapter we give the necessary terminology about probability spaces to prove these bounds and in order to prove the following lemma.

Theorem 3.3. *Let $\alpha > 0$ and $\delta > 0$. Let $\{A_i\}_{i=1}^{l}$ be a finite collection of measurable sets and let μ be a probability measure on $\bigcup_{i=1}^{l} A_i$ such that for all $1 \leq i \leq l$, we have $\mu(A_i) \geq \alpha$. Let $\overline{x} = \{x_1, ..., x_n\}$ be a set of i.i.d. draws according to μ. Then if*

$$n \geq \frac{1}{\alpha}\left(\log l + \log \frac{1}{\delta}\right),$$

we have

$$\mu\left(\overline{x} \cap A_i \neq \emptyset, \textit{for all } i\right) \geq 1 - \delta.$$

The bounds on the sample size for capturing the homology with high confidence will be shown in terms of natural invariants of the underlying submanifold, such as the covering number and the packing number. In the second part of this chapter we will define these invariants and show how they relate to each other. In Section 3.1 we follow the notions of Durrett [7] and in Section 3.2 the notion follow the one in Wainright [16] and Bölcskei [3].

3.1 Probability Measure and its Support

Let us start with the definition of probability measures.

Definition 3.1. *A* ***probability space*** *is a triple $(\Omega, \mathcal{F}, \mu)$, consisting of the set of outcomes Ω, the set of events $\mathcal{F}$ and the probability measure μ that satisfy the following conditions.*

- *$\Omega \subset \mathbb{R}^N$ is a non-empty subset.*
- *$\mathcal{F}$ is a σ-**algebra**, that is a non-empty collection of subsets of Ω satisfying:*
  - *(a) $\emptyset, \Omega \in \mathcal{F}$,*
  - *(b) $A \in \mathcal{F} \implies \Omega \setminus A \in \mathcal{F}$ and*
  - *(c) if $\{A_i\}_{i\geq 0}$ are countably many disjoints sets in $\mathcal{F}$, then $\bigcup_{i\geq 0} A_i \in \mathcal{F}$.*
- *$P \colon \mathcal{F} \to [0, 1]$ is a non-negative function, such that*
  - *(a) $\mu(A) \geq \mu(\emptyset) = 0$, for all $A \in \mathcal{F}$ and*
  - *(b) if $\{A_i\}_{i\geq 0}$ are countably many disjoint sets in $\mathcal{F}$, then*

$$\mu\left(\bigcup_{i\geq 0} A_i\right) = \sum_i \mu(A_i).$$

Let us look at some basic examples of probability spaces.

Example 3.2 (Discrete Probability Measure). *Let Ω be a discrete set, that is, Ω is finite or countably infinite. Let $\mathcal{F}$ be the set of all subsets of Ω, that is $\mathcal{F} = \mathcal{P}(\Omega)$. Let $p\colon \Omega \to [0,1]$ be a function with*

$$p(\omega) \geq 0, \forall \omega \in \Omega \text{ and } \sum_{\omega \in \Omega} p(\omega) = 1.$$

Define $\mu\colon \mathcal{F} \to [0,1]$ as $\mu(A) = \sum_{\omega \in A} p(\omega)$. Then $(\Omega, \mathcal{F}, \mu)$ is a probability measure. In many case, where Ω is finite, one considers the function $p(\omega) = 1/|\Omega|$.

Definition 3.3. *The **Borel σ-algebra** $\mathcal{B}(M)$ of a subset of $\mathbb{R}^N$ is the smallest σ-algebra containing the open sets of M. We call $B \in \mathcal{B}(M)$ a **Borel set**.*

Example 3.4 (Measures on the real line). *Let $\Omega := [a,b]$ be a compact interval and $\mathcal{F}$ the Borel sets on $\mathbb{R}$. Since every open set can be approximated by a union of countably many intervals of the form $(c,d]$, for $a \leq c < d \leq b$, it suffices to define the probability measure on these intervals. We define $\mu\colon \mathcal{B}([a,b]) \to [0,1]$ to be*

$$\mu\Big((c,d]\Big) := \frac{d-c}{b-a}.$$

This function defines a probability measure on the interval Ω. In the case, where $N \geq 2$ and $\Omega = M \subset \mathbb{R}^N$ is a compact subset, we can again consider the Borel sets $\mathcal{B}(M)$ and the non-negative function $\mu\colon \mathcal{B}(M) \to [0,1]$ defined as

$$\mu_N(B) := \frac{\operatorname{vol}(B)}{\operatorname{vol}(M)},$$

*where $\operatorname{vol}(\cdot)$ is the N-dimensional volume on $\mathbb{R}^N$. For $N = 1$, the measure μ_N is the same as the measure μ defined on the compact interval. These measures are called **uniform probability measures**.*

Lemma 3.5. *Let $A_i \subset M$ be measurable sets and $A \subset \bigcup_{i=1}^m A_i$. Then*

$$\mu(A) \leq \sum_{i=1}^m \mu(A_i).$$

A further notion in order to make sense of „drawing random“ from a submanifold is independence.

Definition 3.6. *Let $(\Omega, \mathcal{F}, \mu)$ be a probability space. We say that events $A_i \in \mathcal{F}$ for $i = 1, ...l$ are **independent** if and only if*

$$\mu\left(\bigcap_{i=1}^l A_i\right) = \prod_{i=1}^l \mu(A_i).$$

We say that events are **independent and identically distributed** (short **i.i.d.**, if they are independent and have the same probability distribution, that is, they belong to the same probability space.

We now have enough to state the following lemma, which, once expressed by information about our submanifold M, will provide an estimate on the number of data points needed to capture the homology of M with high confidence.

Lemma 3.7. *Let $\alpha > 0$ and $\delta > 0$. Let $\{A_i\}_{i=1}^l$ be a finite collection of sets and let μ be a probability measure on $\bigcup_{i=1}^l A_i$, such that we have $\mu(A_i) \geq \alpha$, for all $1 \leq i \leq l$. Let $\overline{x} = \{x_1, ..., x_n\}$ be a set of i.i.d. draws according to μ. Then, if*

$$n \geq \frac{1}{\alpha}\left(\log l + \log \frac{1}{\delta}\right),$$

we have

$$\mu\left(\overline{x} \cap A_i \neq \emptyset, \textit{for all } i\right) \geq 1 - \delta.$$

Proof. For $i = 1, ..., l$, let E_i be the event „$\overline{x} \cap A_i$ is empty“. Then, by the independence of the events A_i, we have

$$\begin{aligned} \mu(E_i) &= \mu\left(\bigcap_{j=1}^n \{x_j \notin A_i\}\right) \\ &= \prod_{j=1}^n \mu(x_j \notin A_i) \\ &= (1 - \mu(A_i))^n \geq (1-\alpha)^n. \end{aligned}$$

Therefore, we have

$$\begin{aligned} \mu\left(\bigcup_{i=1}^l E_i\right) &\leq \sum_{i=1}^l \mu(E_i) \\ &\leq \sum_{i=1}^l (1-\alpha)^n = l(1-\alpha)^n. \end{aligned}$$

Let now $n \geq \frac{1}{\alpha}\left(\log l + \log \frac{1}{\delta}\right)$. We want to show that $l(1-\alpha)^n \leq \delta$. For this, we consider the function

$$f(x) = x \cdot e^x - e^x + 1, \text{ for } x \geq 0.$$

We observe that $f(0) = 0$ and $f'(x) = x \cdot e^x \geq 0$, when $x \geq 0$. Hence, f is monotonically increasing, therefore, for all $x \geq 0$, we have $f(x) \geq 0$. Applied to $x = \alpha$, we see that $f(\alpha) \geq 0$ implies

$$\alpha e^\alpha - e^\alpha + 1 \geq 0,$$

which simplifies to

$$e^{-\alpha} \geq (1-\alpha).$$

We can use this to show that

$$\begin{aligned} l(1-\alpha)^n &\leq le^{-\alpha n} \\ &\leq le^{-\alpha \cdot \frac{1}{\alpha} \cdot \left(\log l + \log \frac{1}{\delta}\right)} = \delta. \end{aligned}$$

Putting all together, we get that

$$\begin{aligned} \mu\left(\overline{x} \cap A_i \neq \emptyset, \forall i\right) &= \mu\left(\bigcap_{i=1}^{l} \{\overline{x} \cap A_i \neq \emptyset\}\right) \\ &= 1 - \mu\left(\bigcup_{i=1}^{l} E_i\right) \\ &\geq 1 - l(1-\alpha)^n \geq 1 - \delta. \end{aligned}$$

□

In Section 5.2 we will consider a probability measure that is concentrated around the manifold. What we mean is, that it has support on a neighborhood of M.

Definition 3.8 (Ambrosio et al. [1])**.** *The* ***support*** *of the measure μ is the set defined by*

$$\operatorname{supp}(\mu) := \overline{\left\{x \in \mathbb{R}^N \colon \mu(U) > 0 \text{ for each open neighborhood } U \text{ of } x\right\}}.$$

3.2 Covering Number and Packing Number

In this section we will give the definitions of the covering number of a manifold, which is the minimal number of balls that covers the manifold, as well as the packing number, which is the maximal number of balls that fit into the manifold.

Definition 3.9. *Let $\varepsilon > 0$. An ε-covering of $M \subset \mathbb{R}^N$ is a set $\{x_1, ..., x_l\} \subset M$, such that for every $x \in M$, there exists $1 \leq i \leq l$, such that $x \in B_\varepsilon(x_i)$. The ε-covering number $C(\varepsilon)$ is the cardinality of the smallest ε-covering.*

Definition 3.10. *Let $\varepsilon > 0$. An ε-packing of $M \subset \mathbb{R}^N$ is a set $\{x_1, ..., x_l\} \subset M$, such that $\|x_i - x_j\| > 2\varepsilon$, for any $1 \leq i < j \leq l$. The ε-packing number $P(\varepsilon)$ is the cardinality of the largest ε-packing.*

Example 3.11. *An ε-covering of M, see for example Figure 20 on the left, is a collection of balls centered at elements in M with radius ε that covers all of M, that is*

$$M \subset \bigcup_{i=1}^{l} B_\varepsilon(x_i).$$

An ε-packing, see for example Figure 20 on the right, is a collection of non-intersecting balls with radius ε and centered in M.

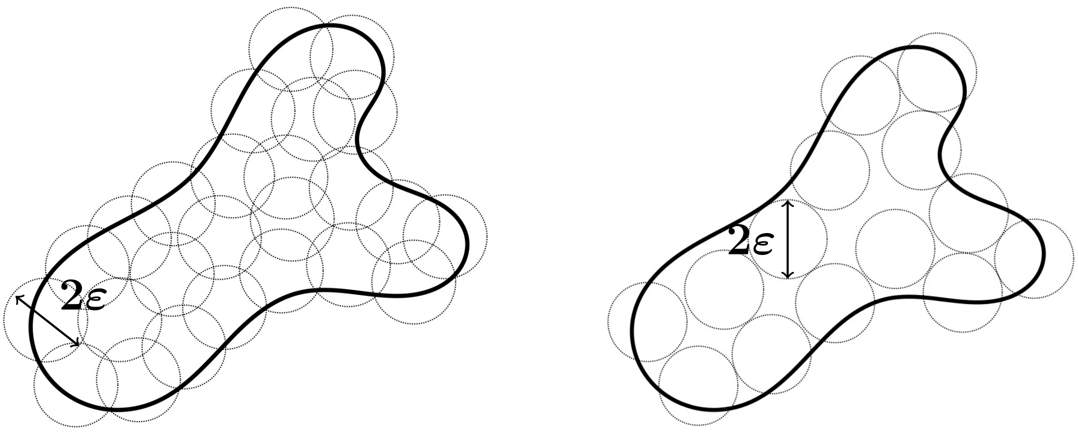

Figure 20: **Left**: An ε-covering of the compact submanifold M. All the balls have radius ε and center in M. Thus, in this case $C(\varepsilon) \leq 25$. **Right**: An ε-packing of the compact submanifold M. All the balls have radius ε and center in M. Thus, in this case $P(\varepsilon) \geq 14$.

The following lemma puts the covering number in relation to the packing number and shows that they essentially provide the same measure of size of a set.

Lemma 3.12. *Let M be as above and $\varepsilon > 0$. Then $P(2\varepsilon) \leq C(2\varepsilon) \leq P(\varepsilon)$.*

Proof. For the first inequality consider a minimal 2ε-covering and a maximal 2ε-packing of M. By definition of the 2ε-packing no two centers can lie in the same ball of the 2ε-covering. Thus, there are at least as many balls in the 2ε-covering as centers in the 2ε-packing, which shows that $P(2\varepsilon) \leq C(2\varepsilon)$.

For the second inequality consider a maximal ε-packing $\{x_1, ..., x_l\}$, that is $l = P(\varepsilon)$. Suppose that $x_1, ..., x_l$ does not form a minimal 2ε-covering, that is there is $x \in M$ such that $\|x - x_i\| > \varepsilon$, for all $1 \leq i \leq l$. But then we have $B_\varepsilon(x) \cap B_\varepsilon(x_i) = \emptyset$ for all $1 \leq i \leq l$, hence $\{x_1, ..., x_l, x\}$ forms an ε-packing of size $l + 1$. But this contradicts the fact that $l = P(\varepsilon)$. Therefore, $\{x_1, ..., x_l\}$ is a 2ε-covering, which proves $C(2\varepsilon) \leq P(\varepsilon)$. □

4 Capturing the Homology of a Submanifold from Samples

In the first part of this chapter we will discuss the conditions on the data sample, in order to capture the homology group of the underlying submanifold from the union of balls centered at the data points. In the second part we will use the results from Section 2.10 to provide an upper bound on the number of points in the sample, such that the homology group of the unerlying submanifold can be found with high confidence.

4.1 Deterministic Setting

Let M be a compact Riemannian submanifold of the Euclidean space $\mathbb{R}^N$. Consider a finite collection of points $\overline{x} = \{x_1, ..., x_n\} \subset M$, which illustrates the sample, and Euclidean balls $B_\varepsilon(x_i)$ of radius ε centered at x_i. We define the open set $U \subset by\mathbb{R}^N$

$$U := \bigcup_{x \in \overline{x}} B_\varepsilon(x).$$

We want to prove the following theorem.

Theorem 4.1. *Let $\overline{x}$ be any finite collection of points $x_1, ..., x_n \in \mathbb{R}^N$, such that it is $\varepsilon/2$-dense in M. Then, for any $\varepsilon < \sqrt{3/5}\tau$, we have that U strongly deformation retracts onto M. Therefore the homology of U equals the homology of M.*

We start with the definition of an ε-dense set in M.

Definition 4.2. *Let $\varepsilon > 0$. A subset $S \subset M$ is said to be ε-**dense**, if for every $p \in M$ there is $x \in S$, such that $\|x - p\| < \varepsilon$.*

Consider the canonical map $\pi\colon U \to M$, which is given by the restriction of the canonical projection $\pi_0\colon \mathrm{Tub}_\tau(M) \to M$, that is

$$\pi(x) = \arg\min_{p \in M} \|x - p\|.$$

We observe that the fibres $\pi^{-1}(p)$ are given by $N_pM \cap U \cap B_\tau(p)$. We need the intersection with $B_\tau(p)$, in order to eliminate distant regions that may intersect with N_pM.

Therefore,

$$\pi^{-1}(p) = \bigcup_{x \in \overline{x}} B_\varepsilon(x) \cap N_pM \cap B_\tau(p).$$

We now define a set $\mathrm{st}(p)$ that turns out to be star shaped relative to p and equal to the fibre $\pi^{-1}(p)$ proving that $\pi^{-1}(p)$ contracts to p. Let us define $\mathrm{st}(p)$ as

$$\mathrm{st}(p) = \bigcup_{\{x \in \overline{x}; x \in B_\varepsilon(p)\}} B_\varepsilon(x) \cap N_pM \cap B_\tau(p).$$

Clearly, $\mathrm{st}(p) \subseteq \pi^{-1}(p)$. Figure 21 shows, how the intersections $B_\varepsilon(x) \cap N_pM \cap B_\tau(p)$ are formed, in the case where the submanifold M is the unit circle S^1. Moreover, it shows why we need the intersection with $B_\tau(p)$.

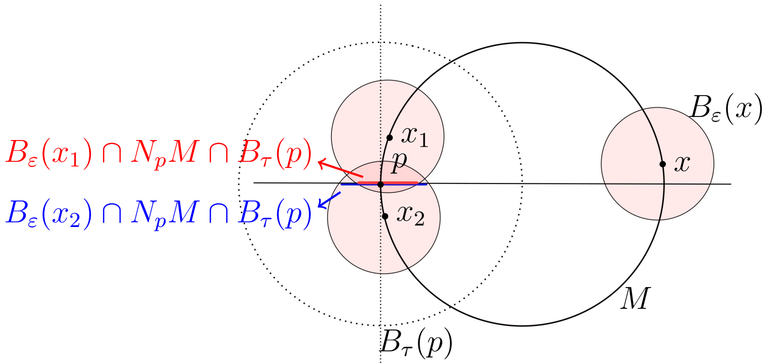

Figure 21: The two colored lines represent two intersections of $\mathrm{st}(p)$. For the fibre $\pi^{-1}(p)$ the corresponding ball $B_\varepsilon(x_i)$ in the intersection $B_\varepsilon(x_i) \cap N_pM \cap B_\tau(p)$ does not necessarily have to contain p itself.

We can see in Figure 21, why the intersection with $B_\tau(p)$ is necessary. Without it, the line segment of N_pM that intersects $B_\varepsilon(x)$ on the right of the picture would be in $\pi^{-1}(p)$ as well, however π does not map this line segment to p, but to its antipodal point instead. In particular, the following lemma is true.

Lemma 4.3. *Let $p \in M$. Then the set $\mathrm{st}(p)$ is star shaped relative to p and therefore contracts to p.*

Proof. Let $q \in \mathrm{st}(p)$ be arbitrary. We need to show that the line segment $\overline{pq}$ is entirely contained in $\mathrm{st}(p)$. Since p and q are both in N_pM, which is an affine subspace of $\mathbb{R}^N$, the line segment $\overline{pq}$ is in N_pM. Moreover, by definition of $\mathrm{st}(p)$, there is an $x \in \overline{x}$ with $x \in B_\varepsilon(p)$, such that $q \in B_\varepsilon(x)$. But since $x \in B_\varepsilon(p)$, we also have $p \in B_\varepsilon(x)$ and by convexity of $B_\varepsilon(x)$ the line segment $\overline{pq}$ is also contained in $B_\varepsilon(x)$. Last, the line segment is contained in $B_\tau(p)$, due to its convexity and thus

$$\overline{pq} \subseteq B_\varepsilon(x) \cap N_pM \cap B_\tau(p),$$

for some $x \in \overline{x}$ with $x \in B_\varepsilon(p)$, which shows that $\overline{pq} \subseteq \mathrm{st}(p)$. □

The next lemma shows that the set $\mathrm{st}(p)$ and the fibre $\pi^{-1}(p)$ are in fact the same.

Lemma 4.4. *Let $p \in M$, then $\mathrm{st}(p) = \pi^{-1}(p)$.*

Proof. Since we already know that $\operatorname{st}(p) \subset \pi^{-1}(p)$, we need to show the other inclusion. Thus, we consider an arbitrary $v \in B_\varepsilon(q) \cap N_pM \cap B_\tau(p)$, where $q \in \overline{x}$ and $q \notin B_\varepsilon(p)$. For such a v, we can consider the worst case scenario, shown in Figure 22a on the left, where p is such that $\sigma(p) = \tau$.

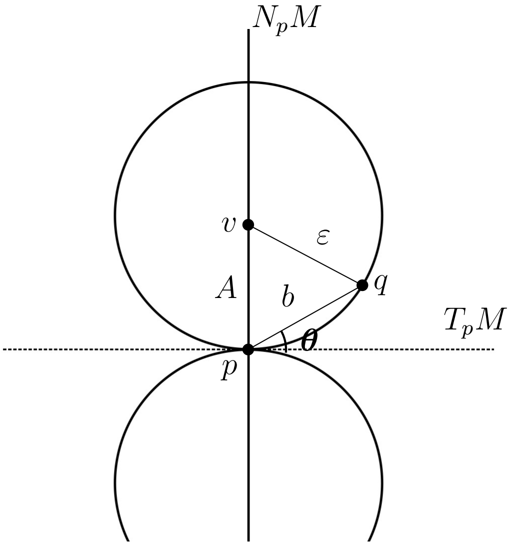

(a) The point q could potentially lie anywhere outside the two circles, but not inside, since otherwise we would have $\sigma(q) < \sigma(p)$. However, the distance $\|v-q\|$ is greatest, when q is on one of the two circles with radius τ. Therefore, this is the worst case. Without loss of generalities, one may consider q lying on the top circle.

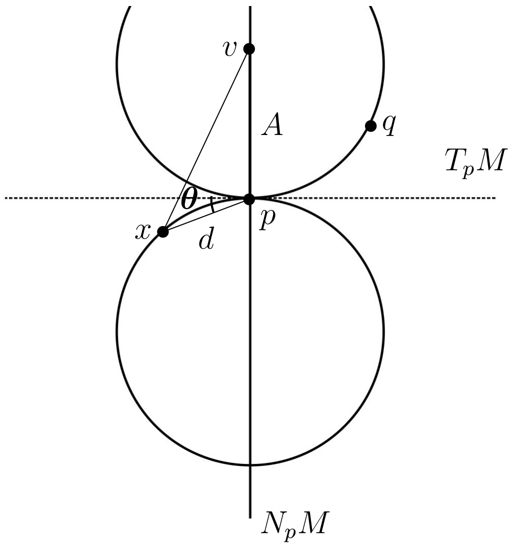

(b) The point x lies within $\varepsilon/2$-distance of p, but could potentially lie anywhere outside of the two circles of radius τ. With the assumption that q is on the top circle, however, the worst case for x is, when it lies on the bottom circle, as shown in the picture.

In order to show that $v \in \operatorname{st}(p)$, we need to find $x \in \overline{x}$ such that $v \in B_\varepsilon(x) \cap N_pM$. Since $\overline{x}$ is $\varepsilon/2$-dense in M, we can find $x \in \overline{x}$ within $\varepsilon/2$-distance of p. We consider again the worst case scenario for the position of x, which is shown in Figure 22b. To show that such x exists, we first show that the distance $A = \|v - p\|$ does not exceed ε^2/τ. Using this bound, we show that for the most unfavourable position of x, as in Figure 22b, we have $\|v - x\| < \varepsilon$ if and only if $\varepsilon < \sqrt{\frac{3}{5}}\tau$.

Thus, let us first prove the following lemma.

Lemma 4.5. *The distance $A = \|v - p\|$ is at most ε^2/τ.*

Proof. We consider the configuration of v, q and p that makes the distance $\|v - p\|$ as large as possible. It suffices to reason about this in the plane passing through these points. Following Figure 22a, we have

$$A = b\sin(\theta) + \sqrt{\|v-q\|^2 - b^2\cos^2(\theta)} \leq b\sin(\theta) + \sqrt{\varepsilon^2 - b^2\cos^2(\theta)},$$

where we used that $v \in B_\varepsilon(q)$ and where $b = 2\tau\sin(\theta)$. This gives us

$$A \leq 2\tau\sin^2(\theta) + \sqrt{\varepsilon^2 - 4\tau^2\sin^2(\theta)\cos^2(\theta)}.$$

To find the maximal A, we calculate $\frac{dA}{d\theta}$, which is given by

$$\frac{dA}{d\theta} = 2\tau\sin(\theta) - \frac{4\tau^2\sin(2\tau)\cos(2\tau)}{2\sqrt{\varepsilon^2 - \tau^2\sin^2(2\tau)}} = 2\tau\sin(2\theta)\left(1 - \frac{\tau\cos(2\theta)}{\sqrt{\varepsilon^2 - \tau^2\sin^2(2\tau)}}\right).$$

If $\varepsilon < \tau$, we have

$$\begin{aligned}\frac{dA}{d\theta} &< 2\tau\sin(2\theta)\left(1 - \frac{\tau\cos(2\theta)}{\sqrt{\tau^2 - \tau^2\sin^2(2\tau)}}\right)\\ &= 2\tau\sin(2\theta)\left(1 - \frac{\tau\cos(2\theta)}{\sqrt{\tau^2(1-\sin^2(2\tau))}}\right) = 0.\end{aligned}$$

Thus, A is monotonically decreasing with θ and in order to get the maximal A, the angle θ must be as small as possible. In particular, θ is minimal when $b = 2\tau\sin(\theta) = \varepsilon$. With this value of θ we get

$$\begin{aligned}A &\leq 2\tau\sin^2(\theta) + \sqrt{\varepsilon^2 - 4\tau^2\sin^2(\theta)\cos^2(\theta)}\\ &= 2\tau\sin^2(\theta) + \sqrt{4\tau^2\sin^2(\tau)(1-\cos^2(\tau))}.\end{aligned}$$

But this simplifies to $A \leq 4\tau\sin^2(\theta) = \varepsilon^2/\tau$. Since this is the largest value that A can take, we have shown, that A is at most ε^2/τ. □

What is left to show is, that there exists $x \in \overline{x}$ with $v \in B_\varepsilon(x) \cap N_pM$. By the $\varepsilon/2$-density, there is an $x \in \overline{x}$ such that $x \in B_{\varepsilon/2}(p)$. Consider Figure 22b for the most unfavorable position of x. By the same argument as in the proof of Lemma 4.5, we see that

$$A = \sqrt{\|v-x\|^2 - b^2\cos^2(\theta)} - b\sin(\theta),$$

where $d = 2\tau\sin(\tau) = \varepsilon/2$. Using this value, we get

$$A = \sqrt{\|v-x\|^2 - \frac{\varepsilon^2}{4}\left(1 - \frac{\varepsilon^2}{16\tau^2}\right)} - 2\tau\frac{\varepsilon^2}{16\tau^2}.$$

From the picture we see, that the worst case happens, when A takes the largest value possible, which is at most ε^2/τ due to Lemma 4.5. Plugging in this value for A gives us

$$\sqrt{\|v-x\|^2 - \frac{\varepsilon^2}{4}\left(1-\frac{\varepsilon^2}{16\tau^2}\right)} - 2\tau\frac{\varepsilon^2}{16\tau^2} = \frac{\varepsilon^2}{\tau},$$

which leads to

$$\sqrt{\|v-x\|^2 - \frac{\varepsilon^2}{4}\left(1-\frac{\varepsilon^2}{16\tau^2}\right)} = \frac{9}{8}\frac{\varepsilon^2}{\tau}.$$

Squaring both sides and solving for $\|v-x\|^2$ gives

$$\|v-x\|^2 = \frac{81}{64}\frac{\varepsilon^4}{\tau^2} - \frac{\varepsilon^2}{4}.$$

But since this expression represents the largest value that $\|v-x\|^2$ can take, we see that $\|v-x\| < \varepsilon$ if and only if

$$\frac{81}{64}\frac{\varepsilon^4}{\tau^2} - \frac{\varepsilon^2}{4} < \varepsilon^2,$$

which simplifies to

$$\frac{\varepsilon^2}{\tau^2} < \frac{3}{5}.$$

Hence, we have shown that $\|v-x\| < \varepsilon$ if and only if $\varepsilon < \sqrt{\frac{3}{5}}\tau$. This finishes the proof of Lemma 4.4. □

Proof of Theorem 4.1. We can now construct a strong deformation retract form U to M by

$$F\colon U\times[0,1]\to U,\ F(x,t) := t\cdot\pi(x) + (1-t)\cdot x.$$

The map F is well-defined and continuous, because $\mathrm{st}(p) = \pi^{-1}(p) \subset U$ is star shaped for all $p \in M$. Moreover we have

- $F(x,0) = x$ for every $x \in U$,
- $F(x,1) = \pi(x) \in M$ for every $x \in U$,
- $F(a,t) = t\cdot\pi(a) + (1-t)a = t\cdot a + (1-t)\cdot a = a$, for every $a \in M$,

where the last condition follows from the fact, that $\pi(a) = a$ for every $a \in M$. Thus, we see that F is a strong deformation retract form U to M and thus, the homology of U equals the homology of M, which finishes the proof. □

4.2 Probabilistic Setting

Let us again consider a compact submanifold M of the Euclidean space $\mathbb{R}^N$ with condition number τ. In this section we will show the following theorem.

Theorem 4.6. *Let M be a compact submanifold of $\mathbb{R}^N$ and let $\varepsilon < \sqrt{3/5}\tau$. Consider a finite collection of points $\bar{x} = \{x_1, ..., x_n\}$ of M, drawn in an i.i.d. fashion according to the uniform probability measure on M and let $\delta > 0$. If*

$$n \geq \frac{1}{\alpha}\left(\log l + \log \frac{1}{\delta}\right),$$

where

$$l = C(\varepsilon/4) \text{ and } \alpha = \frac{1}{P(\varepsilon/4)},$$

then we have with probability greater than $1 - \delta$ that M is $\varepsilon/2$ dense.

We start with a series of lemmas as preparation for the proof of the theorem. For this, consider the tangent space T_pM of M at the point p and let $f\colon \mathbb{R}^N \to T_pM$ be the canonical projection. We begin with the following lemma.

Lemma 4.7. *Let $\varepsilon > 0$ and $A = M \cap B_\varepsilon(p)$. The derivative of df is non-singular at all points $q \in A$.*

Proof. We prove by contradiction and assume, that there is $q \in A$, so that df is singular. Then T_qM is oriented so that the vector with origin q and end $f(q)$ lies in T_qM.

Since $q \in B_\varepsilon(p)$, we have $d = \|q - p\|_{\mathbb{R}^N} < \tau/2$. By combining Proposition 2.79 and Proposition 2.80, we get, using Remark 2.82, that

$$\cos(\phi) \geq \sqrt{1 - \frac{2d}{\tau}} > 0,$$

where ϕ is the angle between the tangent spaces T_qM and T_pM. But $\cos(\phi) > 0$ implies that $|\phi| < \frac{\pi}{2}$, which contradicts the fact that the vector from q to $f(q)$ is orthogonal to T_pM and lies in T_qM.

□

Lemma 4.8. *Let $p \in M$ be a point and k the dimension of M. Let $\varepsilon > 0$ and consider $A = M \cap B_\varepsilon(p)$. If $\theta = \arcsin\left(\frac{\varepsilon}{2\tau}\right)$, then, in terms of k-dimensional volumes, we have*

$$\operatorname{vol}(A) \geq \left(\cos(\theta)\right)^k \operatorname{vol}\left(B^k_\varepsilon(p)\right),$$

where $B^k_\varepsilon(p)$ is the k-dimensional ball in T_pM centered at p.

Proof. Let $B^k_r(p)$ be the k-dimensional ball of radius $r = \varepsilon \cdot \cos(\theta)$, where $\theta = \arcsin(\varepsilon/2\tau)$ on the tangent space T_pM, which is illustrated in Figure 24. We can see, that the image of

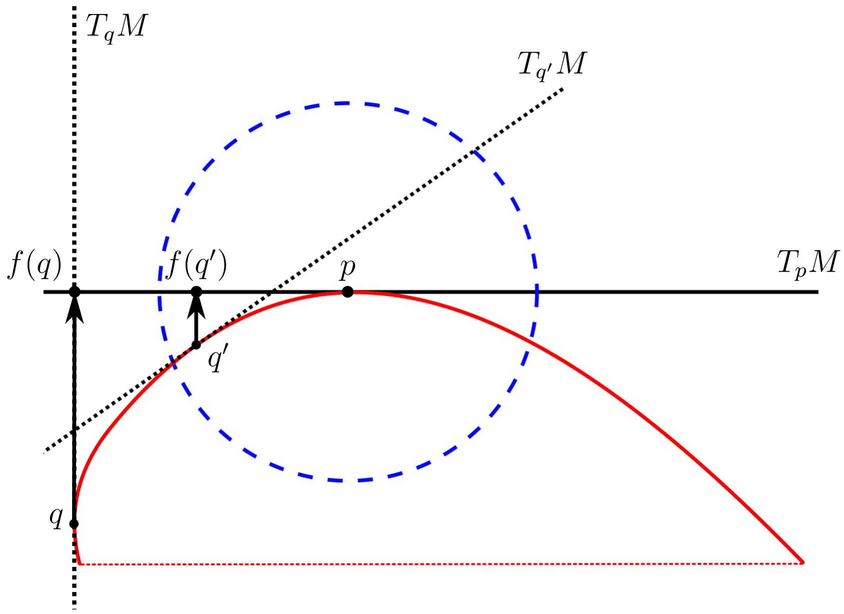

Figure 23: Singularity of df at the point q and non-singularity within A, illustrated by the red line segment within the blue circle.

A, illustrated by the red line within the blue dashed circle, in the tangent space contains $B_r^k(p)$, since the curvature of M cannot be greater than the one of the big circle with radius τ.

Consider the image of A under f. If we can show, that $B_r^k(p) \subset f(A)$, we are done, because then

$$\operatorname{vol}(A) \geq \operatorname{vol}(f(A)) \geq \operatorname{vol}\left(B_r^k(p)\right) = \left(\cos(\theta)\right)^k \operatorname{vol}\left(B_\varepsilon^k(p)\right),$$

where we used

$$\operatorname{vol}\left(B_{d\cdot\varepsilon}^k(p)\right) = d^k \cdot \operatorname{vol}\left(B_\varepsilon^k(p)\right),$$

for $d \geq 0$. To see that $B_r^k(p) \subset f(A)$, we notice that f is an open map since it is a projection map onto an open subset and moreover, by Lemma 4.7, its derivative is non-singular at every point $q \in A$, thus $df(q)$ is invertible at every point in A, in particular at $p \in A$. Therefore, we can apply the inverse function theorem (Theorem 2.27) to get that f is locally invertible at p and hence, there is a ball $B_r^k(p)$, such that

$$f^{-1}\left(B_s^k(p)\right) \subset A.$$

We can increase s and denote by s' the smallest s, for which

$$f^{-1}\left(\overline{B_{s'}^k(p)}\right) \not\subset A \text{ and } f^{-1}\left(B_{s'}^k(p)\right) \subset A.$$

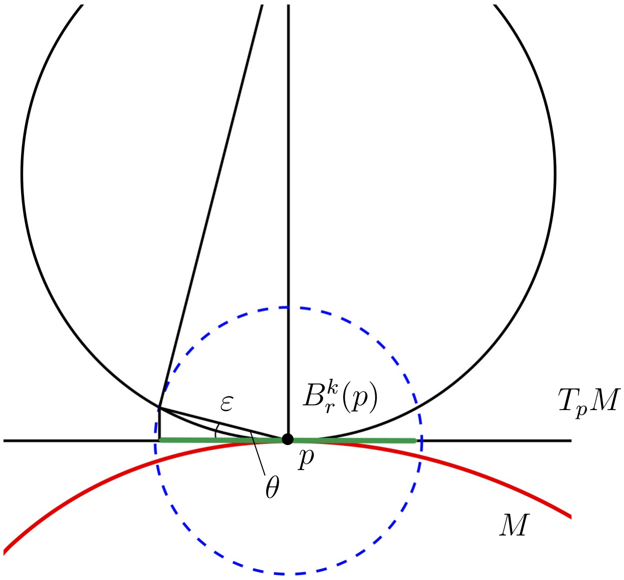

Figure 24: The green ball $B_r^1(p)$ is contained in the projection of the set $A = M \cap B_\varepsilon(p)$ to T_pM.

Here, $\overline{B_{s'}^k(p)}$ denotes the closure of the k-dimensional ball of radius s with center p. This implies, that there is $x \in \partial B_{s'}^k(p)$ with $q = f(x) \notin A$ and $q \in \overline{A}$, since f is an open map. Therefore, q is in the boundary of A and thus, $\|q - p\|_{\mathbb{R}^N} = \varepsilon$. We see that $s' = \varepsilon \cdot \cos(\varphi)$, where φ is the angle between the lines $\overline{qp}$ and $\overline{f(q)p}$. By the curvature bound implied by τ, we see that $|\varphi| \leq |\theta|$. This gives $s' = \varepsilon \cos(\varphi) \geq \varepsilon \cos(\theta) = r$.

The way we defined s', we now see, that $f^{-1}\left(B_r^k(p)\right) \subset A$ and thus, $B_r^k(p) \subset f(A)$, which finishes the proof. $\square$

Proof of Theorem 4.6. We want to apply Lemma 3.7 to our setting form Section 4.1. For this we have to define the numbers l, the measurable sets A_i for $1 \leq i \leq l$ and the bound α. For our setting, we consider a minimal cover of the manifold M by balls of radius $\varepsilon/4$. Let $\{y_i \colon 1 \leq i \leq l\}$ be the centers of balls that constitutes such minimal cover. Hence $l = C(\varepsilon/4)$ is the $\varepsilon/4$-covering number. For $1 \leq i \leq l$, we choose the sets A_i to be $A_i := B_{\varepsilon/4}(y_i) \cap M$. Since μ is the uniform probability measure on M, we have

$$\mu(A_i) \geq \frac{\operatorname{vol}(A_i)}{\operatorname{vol}(M)},$$

for all $i \leq i \leq l$. With $\theta = \arcsin(\varepsilon/8\tau)$, let us define the value α to be

$$\alpha := \frac{\left(\cos^k(\theta)\right)\operatorname{vol}\left(B_{\varepsilon/4}^k(p)\right)}{\operatorname{vol}(M)}.$$

Then, for all $1 \leq i \leq l$, Lemma 4.8 provides that $\mu(A_i) \geq \alpha$. Our setting now satisfies the requirements from Lemma 3.7, so for any $\delta > 0$, if

$$n \geq \frac{1}{\alpha}\left(\log l + \log \frac{1}{\delta}\right),$$

we have with probability greater than $1-\delta$ that for each $1 \leq i \leq l$, that $\overline{x} \cap A_i$ is not empty. To show that this implies $\varepsilon/2$-density, we pick $p \in M$ arbitrary. Since $\{y_i : 1 \leq i \leq l\}$ constitute a covering, there is such an y_i with $p \in B_{\varepsilon/4}(y_i)$. Because we have with high probability that $A_i \cap \overline{x} \neq \emptyset$, there is $x \in \overline{x}$, such that $x \in B_{\varepsilon/4}(y_i)$. Combining them yields, that for this particular x we have, using the triangular inequality, that

$$\|p - x\| \leq \|p - y_i\| + \|y_i - x\| < \frac{\varepsilon}{4} + \frac{\varepsilon}{4} = \frac{\varepsilon}{2}.$$

But this is exactly saying, that with probability greater than $1 - \delta$ the data sample $\overline{x}$ lies $\varepsilon/2$-dense in M. □

In particular, using Lemma 4.8, we immediately get a simple bound on the packing number

$$P(\varepsilon) \leq \frac{\operatorname{vol}(M)}{\operatorname{vol}(M \cap B_\varepsilon(p))} \leq \frac{\operatorname{vol}(M)}{(\cos^k(\theta)) \operatorname{vol}(B^k_\varepsilon(p))}, \tag{8}$$

where $p \in M$ is a point and $\theta = \arcsin(\varepsilon/2\tau)$. Hence, using $\varepsilon/4$ instead, we get

$$P\left(\frac{\varepsilon}{4}\right) \leq \frac{\operatorname{vol}(M)}{(\cos^k(\theta)) \operatorname{vol}\left(B^k_{\varepsilon/4}(p)\right)} = \frac{1}{\alpha}.$$

Putting everything together, we get the theorem.

Theorem 4.1 and Theorem 4.6 together show that if we were to draw enough sample points from the submanifold M, we get the necessary density condition with high probability and thus, we can compute the homology with high confidence by computing the homology of U. In particular, using Lemma 3.12 and the bound on the packing number found in (8), we obtain the lower bound on the sample size entirely in terms of the condition number.

Unfortunately, in many problems, the sample might not exactly lie on the manifold itself, for the draws happen with a certain level of noise. The next section discusses exactly this problem and makes sure, that also with noise, we can still compute the homology of the manifold M with high confidence.

5 Capturing the Homology of a Submanifold from Samples with Noise

In this chapter we introduce noisy data, that is, a sample is drawn according to a probability distribution from around the manifold, rather from the manifold. Similar as in Section 4.1, we provide necessary conditions, such that the set U deformation retracts onto the submanifold. In the second part of this chapter we discuss again the probabilistic setting and apply Lemma 3.7 to provide an upper bound on the number of points in the sample, such that the set U captures the homology group of M with high confidence. In the last section, we discuss the setting of samples with noise with a weaker density condition. Applied for the special setting in Theorem 4.1, we will be able to improve the bound on the ball size of U.

5.1 Deterministic Setting

We will proceed in a similar fashion as in Section 4.1. The main difference is, that our sample set does not need to be contained in M. In the following, we will prove this theorem.

Theorem 5.1. *Let $0 < r < \left(\sqrt{9} - \sqrt{8}\right)\tau$ and let $\overline{x}$ be an r-noisy set of finite points that is r-dense in M. Then for all*

$$\varepsilon \in \left(\frac{(r+\tau) - \sqrt{r^2 + \tau^2 - 6\tau r}}{2}, \frac{(r+\tau) + \sqrt{r^2 + \tau^2 - 6\tau r}}{2} \right),$$

the set U defined by

$$U := \bigcup_{x \in \overline{x}} B_\varepsilon(x)$$

strongly deformation retracts to M.

Before we begin with the prove, we define the notion of r-noisy sets.

Definition 5.2. *Let $0 < r < \tau$. A finite set $\overline{x}$ is called r-**noisy**, if it is contained in the tubular neighborhood $\mathrm{Tub}_r(M)$.*

Proof of Theorem 5.1. First, we show, that for the choice of ε, we have $M \subset U$. For this, we observe, since $r < \tau$, that we have

$$\frac{(r+\tau) - \sqrt{r^2 + \tau^2 - 6\tau r}}{2} \geq \frac{(r+\tau) - \sqrt{r^2 + \tau^2 - 2\tau r}}{2},$$

which simplifies to

$$\frac{(r+\tau) - \sqrt{(r-\tau)^2}}{2} = \frac{(r+\tau) - (\tau - r)}{2} = r.$$

Therefore, for the choice of ε from the statement, we certainly have $\varepsilon > r$. Let now $p \in M$, then there is $x \in \overline{x}$, such that $p \in B_r(x)$, since $\overline{x}$ is r-dense in M. By the fact that $\varepsilon > r$, we see that

$$p \in B_r(x) \subset B_\varepsilon(x) \subset U$$

and thus, $M \subset U$.

Consider the projection $\pi_0\colon \mathrm{Tub}_r(M) \to M$ from the tubular neighborhood to M and let $\pi\colon U \to M$ be its restriction to U. We claim that for each $p \in M$, the fibre $\pi^{-1}(p)$ contracts to the point p. We prove this by showing that $\pi^{-1}(p)$ is star-shaped with respect to p. Let $v \in \pi^{-1}(p)$ and consider the line segment $\overline{vp}$. Let $q \in \overline{x}$, such that $v \in B_\varepsilon(q)$, which exists since $v \in \pi^{-1}(p) \subset U$. If $q \in B_\varepsilon(p)$, we have $\overline{vp} \subset B_\varepsilon(p) \cap B_\varepsilon(q)$ due to the convexity of balls in $\mathbb{R}^N$. But this means that $\overline{vp} \subset \pi^{-1}(p)$ and therefore $\pi^{-1}(p)$ is star-shaped and hence, it contracts to p.

Let us assume that $q \notin B_\varepsilon(p)$. Since $\overline{x}$ is r-dense in M, there is $x \in B_r(p) \subset B_\varepsilon(x)$. We want to show that for this x, we have $v \in B_\varepsilon(x)$, so we can use the same argumentation as before to see, that the fibre of p contracts to p.

As in the proof of Theorem 4.1, it is enough to reason in the plane passing through the point v, p and q, as shown in Figure 25, where T_pM and N_pM intersect with this plane in the point p.

We want to find a sufficient condition, so that $v \in B_\varepsilon(x)$. In Figure 25 we see that $v \in B_\varepsilon(x)$ as long as $\|v - p\| < \varepsilon - r$. Rearranging and squaring gives

$$\Big(\tau - (\varepsilon - r)\Big)^2 < A^2 \leq (\tau - r)^2 - \varepsilon^2, \tag{9}$$

where the latter inequality follows by using the Pythagorean theorem on the triangle drawn in the upper circle. Expanding the squares, it becomes

$$\varepsilon^2 - \varepsilon(\tau + r) + 2\tau r < 0.$$

This quadratic inequality is satisfied, whenever

$$\begin{aligned}\varepsilon &\in \left(\frac{(\tau + r) - \sqrt{(r+\tau)^2 - 8\tau r}}{2}, \frac{(r+\tau) + \sqrt{(r+\tau)^2 - 8\tau r}}{2}\right)\\ &= \left(\frac{(r + \tau) - \sqrt{r^2 + \tau^2 - 6\tau r}}{2}, \frac{(r+\tau) + \sqrt{r^2 + \tau^2 - 6\tau r}}{2}\right),\end{aligned}$$

provided that the term under the square root is strictly greater than 0, that is

$$r^2 + \tau^2 - 6\tau r > 0.$$

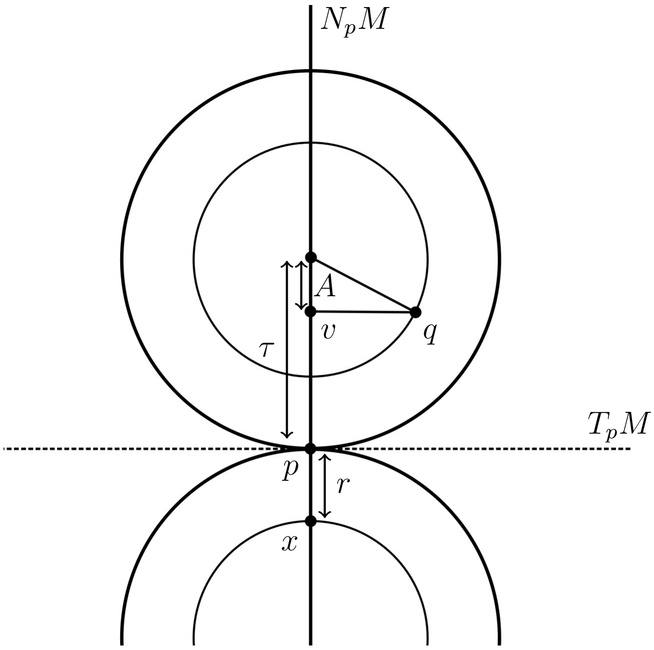

Figure 25: The picture shows the configuration of the points p, v, q and x that make the distance $\|v - x\|$ as large as possible. The two pair of concentric circles of radius τ and $\tau - r$ indicate, that q cannot lie inside the smaller of the two circles, due to the curvature of M. However, the point q could potentially lie anywhere outside the circle with radius $\tau - r$, but similar as in Figure 22a, the distance A is maximal when q is on the circle.

However, this quadratic inequality in r is satisfied as long as

$$r < \frac{6\tau - \sqrt{36\tau^2 - 4\tau^2}}{2} = (3 - 2\sqrt{2})\tau = \left(\sqrt{9} - \sqrt{8}\right)\tau,$$

or

$$r < \frac{6\tau + \sqrt{36\tau^2 - 4\tau^2}}{2} = (3 + 2\sqrt{2})\tau = \left(\sqrt{9} + \sqrt{8}\right)\tau.$$

Since $r > \tau$ implies that $\mathrm{Tub}_r(M)$ has different homology than M, the second condition cannot be taken into account.

Altogether, we have shown that, as long as $r < \left(\sqrt{9} - \sqrt{8}\right)\tau$ and

$$\varepsilon \in \left(\frac{(r+\tau) - \sqrt{r^2 + \tau^2 - 6\tau r}}{2}, \frac{(r+\tau) + \sqrt{r^2 + \tau^2 - 6\tau r}}{2}\right),$$

we have that $\overline{vp} \in B_\varepsilon(x)$ and thus $\pi^{-1}(p)$ contracts to p.

The map $F\colon U \times [0,1] \to U$ defined by

$$F(x,t) = (1-t)x + t\pi(x)$$

gives a strong deformation retract from U to M. □

5.2 Probabilistic Setting

Let M be a submanifold of the Euclidean space $\mathbb{R}^N$ and let μ be a probability measure on $\mathbb{R}^N$. Let $0 < r < \tau$ be a positive value. Different then in the previous chapters, our probability measure is not concentrated on the manifold M itself, but rather around it. To be precise, we assume that the probability measure μ has support within $\mathrm{Tub}_r(M)$. We prove the following theorem.

Theorem 5.3. *Let M be a compact submanifold of $\mathbb{R}^N$ and let $0 < r < \left(\sqrt{9} - \sqrt{8}\tau\right)$. Suppose that $l \coloneqq C(r/2)$ is the $r/2$-covering number of M. For any r-noisy data set $\overline{x} = \{x_1, ..., x_n\}$, drawn in an i.i.d. fashion according to a probability measure μ that is r-concentrated around M and for any $\delta > 0$, if*

$$n > \frac{1}{k_{r/2}} \left(\log(l) + \log\left(\frac{1}{\delta}\right)\right),$$

then we have with probability greater then $1 - \delta$ that $\overline{x}$ is r-dense in M.

First, we define the setting of a measure that also respects draws within a tubular neighborhood of M.

Definition 5.4. *We say that the probability measure μ is r-**concentrated** around M, if it satisfies the following two conditions:*

(a) We have that $\mathrm{supp}(\mu) \subset \mathrm{Tub}_r(M)$.

(b) For every $0 < s < r$, there is $k_s > 0$, independent of p, such that

$$\inf_{p \in M} \mu\left(B_s(p)\right) > k_s.$$

Before we state the result on the number of data points we need, in order to meet the requirements from Theorem 5.1, that is, that $\overline{x}$ is r-dense in M, we discuss an example of a measure that is r-concentrated around M.

Example 5.5. *Let $0 < r < \tau$. Consider the uniform probability measure μ_r on the tubular neighborhood* $\mathrm{Tub}_r(M)$, *that is for every measurable set $A \subset \mathrm{Tub}_r(M)$, we have*

$$\mu_r(A) = \int_A \frac{x}{\mathrm{vol}\left(\mathrm{Tub}_r(M)\right)} d\mu_r.$$

Observe that we need the factor $1/\operatorname{vol}(\operatorname{Tub}_r(M))$*, in order to achieve* $\mu_A(\operatorname{Tub}_r(M)) = 1$*. Define the measure* μ *on the Borel sets of* $\mathbb{R}^N$ *by*

$$\mu(B) := \mu_r(B \cap \operatorname{Tub}_r(M)), \text{ for all Borel sets } B \subset \mathbb{R}.$$

Clearly, $\operatorname{supp}(\mu) \subset \operatorname{Tub}_r(M)$ *and since for every* $p \in M$ *and for every* $0 < s < r$*, the ball* $B_s(p)$ *is entirely contained in* $\operatorname{Tub}_r(M)$*, we see that*

$$\mu(B_s(p)) = \operatorname{vol}(B_s(p)) > 0,$$

which does not depend on p*, but only on the number* s*. This shows that* μ *is indeed a* r*-concentrated measure around* M*.*

We can now prove Theorem 5.3.

Proof of Theorem 5.3. For the proof we choose $p \in M$ to be be arbitrary, $\delta > 0$ and we set $l := C(r/2)$. Suppose that

$$n > \frac{1}{k_{r/2}}\left(\log(l) + \log\left(\frac{1}{\delta}\right)\right).$$

Let $\{y_1, ..., y_l\}$ be a minimal $r/2$-covering. We define

$$A_i := B_{r/2}(y_i), \text{ for } i = 1, ..., l.$$

By definition of an $r/2$-covering, there is $j \in \{1, ..., l\}$, such that $p \in A_j$. In particular, by assumption on μ, we have for every $1 \leq i \leq l$ that

$$\mu(A_i) \geq \inf_{q \in M} \mu\left(B_{r/2}(p)\right) > k_{r/2}.$$

We can apply Lemma 3.7, to see that with probability greater than $1 - \delta$ we have

$$\overline{x} \cap A_i \neq \emptyset, \text{ for every } 1 \leq i \leq l.$$

Hence, there is $x \in \overline{x}$, such that $x \in A_j$ for the j from above. But for this x we have, using the triangle inequality, that

$$\|x - p\| \leq \|x - y_j\| + \|p - y_j\| < \frac{r}{2} + \frac{r}{2} = r.$$

Since we chose p to be arbitrary, this shows that $\overline{x}$ is in fact r-dense with probability greater than $1 - \delta$. □

Similar as in Chapter 4, if we put Theorem 5.1 and Theorem 5.3 together, we can see that, if we were to draw enough sample points from around the submanifold M, we can compute the homology with high confidence by computing the homology of U.

5.3 Improving Bound for Deterministic Setting without Noise

In Section 5.1 we assumed the sample $\overline{x}$ to be r-dense in M, where r is the level of noise allowed on the sample. The following result gives conditions on the ball-size, in order that U captures the homology group of M, given that $\overline{x}$ is s-dense in M for an arbitary $0 < s < r$.

Theorem 5.6. *Let $0 < r < \left(\sqrt{9} - \sqrt{8}\right)\tau$ and let $\overline{x}$ be an r-noisy set of finite points that is r-dense in M. For $0 < s < r$, suppose that $\overline{x}$ is s-dense in M. Provided that*

$$(\tau - w)^2 < (\tau - r)^2 - \varepsilon^2 \text{ and } w = \sqrt{\beta^2 - (s^2 - \varepsilon^2)} - \beta,$$

where $\beta = \frac{s^2}{2\tau} + r - \frac{r^2}{2\tau}$ the set U defined by

$$U := \bigcup_{x \in \overline{x}} B_\varepsilon(x)$$

strongly deformation retracts to M.

Proof. Once again we consider the worst possible case on the position of some arbitrary $v \in B_\varepsilon(q) \cap T_P M \cap B_\tau(p)$, where $q \notin B_\varepsilon(p)$. Following the same argumentation in the plane as in the previous chapter, we can see the worst case shown in Figure 26. The worst case is achieved, when the distance $\|v - q\|$ and $\|x - p\|$ are maximal within the provided bound. Since $v \in B_\varepsilon(q)$, we draw the worst case $\|v - q\| = \varepsilon$ and since $\overline{x}$ is s-dense, we draw $\|x - p\| = s$.

Following Figure 26, we find two equations describing the relations of w and β to the rest of the parameters r, s, ε and τ. By describing b^2 in two different ways, we get first equation

$$(\tau - r)^2 - (\tau - \beta)^2 = s^2 - \beta^2. \tag{10}$$

The other equation follows from the Pythagorean theorem on the rectangular triangle with side lengths ε, b and $w + \beta$, that is

$$s^2 - \beta^2 + (\beta + w)^2 = \varepsilon^2, \tag{11}$$

where we used that $b = s^2 - \beta^2$. Solving (10) for β gives us $\beta = \frac{s^2}{2\tau} + r - \frac{r^2}{2\tau}$, whereas simplifying (11) gives us

$$w^2 + 2\beta w + (s^2 - \varepsilon^2) = 0.$$

This equation possibly has two solutions, namely

$$w = \sqrt{\beta^2 - (s^2 - \varepsilon^2)} - \beta \text{ or } w = -\sqrt{\beta^2 - (s^2 - \varepsilon^2)} - \beta,$$

whereas the latter is impossible since w is a distance and thus $w \geq 0$. Hence we get

$$\beta = \frac{s^2}{2\tau} + r - \frac{r^2}{2\tau} \text{ and } w = \sqrt{\beta^2 - (s^2 - \varepsilon^2)} - \beta.$$

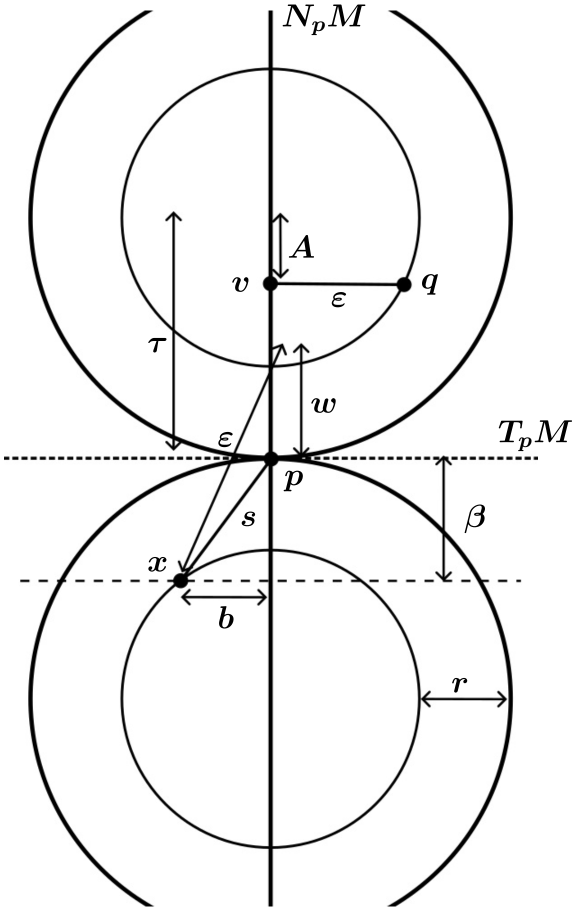

Figure 26: A picture showing the worst case of the position of x. In order for v to be in $B_\varepsilon(x)$, the distance A has to be greater than $\tau - w$.

In order to get the homotopy type between U and M, we want to show that the fibres $\pi^{-1}(p)$ as described in Section 4.1 retract to the point p. This is true, provided that v lies within ε-distance of the data point x. But following Figure 26, this is the same as requiring that the distance A is greater than $\tau - w$. Hence, we have the same homotopy type between U and M, provided that

$$(\tau - w)^2 < (\tau - r)^2 - \varepsilon^2, \tag{12}$$

with the above relation on β and w. □

The following examples treat two special cases, namely when $s = r$ and when $r = 0$. Following the equations (10)-(12), we can observe that the first case reduces to the results found in 5.1.

Example 5.7. *Let us set $s = r$. From (10), we get that*

$$(\tau - r)^2 - (\tau - \beta)^2 = r^2 - \beta^2,$$

which simplifies to $2\tau(r-\beta)=0$. *From (10), we get that also* $\beta = r$. *Putting* $s=\beta=r$ *in (11), provides*

$$r^2 - r^2 + (r+w)^2 = \varepsilon^2,$$

from which it follows that $w=\varepsilon-r$. *At last ,we insert* $w=\varepsilon-r$ *in (12) and get the inequality*

$$(\tau-(\varepsilon-r))^2 < (\tau-r)^2-\varepsilon^2,$$

which is exactly what we found in the inequality (9) in Section 5.1.

Example 5.8. *Consider the second case to be* $r=0$. *Using (12), we get*

$$(\tau-w)^2 < \tau^2-\varepsilon^2,$$

which is the same as requiring that $w^2-2\tau w+\varepsilon^2<0$. *This shows that we require*

$$w > \tau - \sqrt{\tau^2-\varepsilon^2}. \tag{13}$$

Solving (10) for β *gives us* $\beta = s^2/2\tau$, *which we can substitute in (11) to get the quadratic equation in* w

$$s^2 - \frac{s^4}{4\tau^2} + \left(\frac{s^2}{2\tau}+w\right)^2 = \varepsilon^2.$$

This equation has the positive solution

$$w = -\frac{s^2}{2\tau} + \sqrt{\varepsilon^2 - s^2 + \frac{s^4}{4\tau^2}},$$

which together with (13) gives rise to the following condition:

$$-\frac{s^2}{2\tau} + \sqrt{\varepsilon^2 - s^2 + \frac{s^4}{4\tau^2}} > \tau - \sqrt{\tau^2-\varepsilon^2}. \tag{14}$$

The inequality (14) provides a range for s *and* ε, *for which we achieve a homotopy equivalence between* U *and* M. *For the special case where* $s=\varepsilon/2$ *(compare with the density-choice we made in Theorem 4.1), the inequality (14) simplifies to*

$$\varepsilon^4 + 51\varepsilon^2\tau^2 - 48\tau^4 < 0. \tag{15}$$

which is satisfied whenever $0<\varepsilon^2<0.9244\tau^2$ *or* $0<\varepsilon<0.96\tau$.

Observe that in Theorem 4.1 we required ε to be in $(0,\sqrt{3/5}\tau)\approx(0,0.77\tau)$, whereas in Example 5.8 we found that ε can also be in the interval $(0,0.96\tau)$. Thus, Example 5.8 gives a slightly stronger version of Theorem 4.1 in terms of the balls-size ε. This is due to the fact, that we require ε to be in such a way that $\mathrm{st}(p)$ is equal to $\pi^{-1}(p)$, which is a stronger condition than simply assuming $B_\varepsilon(q)\cap B_\varepsilon(x)\cap\overline{vp}\neq\emptyset$.

Furthermore, if we assume the condition number τ and the noise factor r to be out of our control, the sample complexity entirely depends upon the number s, which gives information about the density of the data sample. In order to draw the fewest number of sample points, it is necessary to maximize s, which can be achieved by taking the largest s possible that satisfies Equation (14).

Bibliography

[1] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows*. 1st ed. Lectures in Mathematics. ETH Zürich. Basel: Birkhäuser, 2005.

[2] Mikhail Belkin and Partha Niyogi. "Semi-Supervised Learning on Riemannian Manifolds". In: *Machine Learning* 56.1 (July 2004), pp. 209–239. ISSN: 0885-6125.

[3] Helmut Bölcskei. "Lecture Notes Neural Network Theory". In: (2021). (Visited on 03/20/2024).

[4] Glen E. Bredon. *Topology and Geometry*. 1st ed. Graduate Texts in Mathematics. New York, NY: Springer, 1993.

[5] Manfredo P. do Carmo. *Riemannian Geometry*. 1st ed. Mathematics: Theory & Applications. Boston, MA: Birkhäuser, 1992.

[6] David L. Donoho and Carrie Grimes. "Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data". In: *Proceedings of the National Academy of Sciences* 100.10 (May 13, 2003), pp. 5591–5596.

[7] Rick Durrett. *Probability: Theory and Examples*. 1st ed. North Carolina: Duxbury Press, 1996.

[8] Manfred Einsiedler and Andreas Wieser. "Vorlesungsskript 2021/2022". In: (2021). URL: https://people.math.ethz.ch/~einsiedl/Analysis-Skript.pdf (visited on 03/20/2024).

[9] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.

[10] John M. Lee. *Introduction to Smooth Manifolds*. 2nd ed. Graduate Texts in Mathematics. New York, NY: Springer, 2012.

[11] John M. Lee. *Riemannian Manifolds*. 1st ed. Graduate Texts in Mathematics. New York, NY: Springer, 1997.

[12] Will Merry. "Lecture Notes Differential Geometry". In: (2020). URL: https://www2.math.ethz.ch/will-merry/files/Merry%20-%20Differential%20Geometry%20(2021).pdf (visited on 03/20/2024).

[13] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. "Finding the Homology of Submanifolds with High Confidence from Random Samples". In: *Discrete & Computational Geometry* 39.1 (Mar. 1, 2008), pp. 419–441.

[14] Sam T. Roweis and Lawrence K. Saul. "Nonlinear Dimensionality Reduction by Locally Linear Embedding". In: *Science* 290.5500 (Dec. 22, 2000), pp. 2323–2326.

[15] Joshua B. Tenenbaum, Vin De Silva, and John C. Langford. "A Global Geometric Framework for Nonlinear Dimensionality Reduction". In: *Science* 290.5500 (Dec. 22, 2000), pp. 2319–2323.

[16] Martin J. Wainright. *High Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.

[17] Afra Zomorodian and Gunnar Carlsson. "Computing Persistent Homology". In: *Discrete and Computational Geometry* 33.2 (Feb. 2005), pp. 249–274.