PhD Dissertation

The Reach of Manifold Learning

Uniqueness Bounds for Latent Representations

Helene Hauschultz

Department of Mathematics
Aarhus University
2022

The Reach of Manifold Learning
Uniqueness Bounds for Latent Representations

PhD dissertation by
Helene Hauschultz

Department of Mathematics, Aarhus University
Ny Munkegade 118, 8000 Aarhus C, Denmark

Supervised by
Andrew du Plessis

Submitted to Graduate School of Natural Sciences, Aarhus, November 30, 2022

Dissertation was typesat in kpfonts with
pdfLATEX and the memoir class

AARHUS
UNIVERSITY
DEPARTMENT OF MATHEMATICS

Contents

Preface

This dissertation is the result of my PhD studies at the Department of Mathematics at Aarhus University. The dissertation consists of two papers together with an introduction to the research question and a chapter containing some background information on reach.

The papers are:

Paper A Pointwise normal reach. *Paper draft*.

Paper B Is an encoder within reach? Submitted to *The 26th Internation Conference on Artificial Intelligence and Statistics (AISTATS)*.

Paper B is the result of my research stay at DTU where I visited Professor Søren Hauberg. The targeted audience of the paper is the Machine Learning community, and the paper should be read as such. I have contributed significantly to both the theoretical analysis, to the experimental developments, as well as to the writing. The paper is identical except for layout to a version which has currently been submitted to the conference *AISTATS*, and appears on arXiv with id 2206.01552.

Paper A is a mathematical exploration of the pointwise normal reach explored in Paper B. The results in Paper A have been developed by me in collaboration with Andrew du Plessis. The paper is written solely by me, except for Appendix A.B which have been written by Andrew du Plessis, and edited by me in order to make it consistent with the rest of the paper.

In addition to the two papers a chapter about reach is included. The results in the chapter are previously known results and results done by Andrew du Plessis. It is included as it represents a large part of my time as a PhD student, however, in the end the project ended up in a different direction. As a consequence, most of the results in this chapter are not relevant for the results in the two papers.

As is the case with any good projects, this project had its share of unforeseen challenges to overcome. From the beginning the project was the first of its type at the department and there was no clear research question. In addition, it was hit with a global pandemic, so much of the time was carried out in isolation. For this reason, the project did not hit its stride until my stay at the DTU in the autumn of 2021. Looking back, I wish that the first years of the project had played out differently. My time as a PhD student has definitely been a learning experience. Ultimately I am very proud of what I have managed to achieve, and the project opens the door to many unanswered questions.

None of it had, of course, been possible without the support of a great many people! I would like to thank my advisor Andrew du Plessis for introducing me to reach and for his help in completing the mathematical results. This dissertation could not

have been completed without the guidance of Søren Hauberg who kindly welcomed me into his research group during my stay at the DTU. Our many discussions gave much needed direction to my otherwise aimless endeavors. Much thanks should be given to the entire Geometric Machine Learning group for making my stay at their group an incredible experience. I would like to especially thank my co-authors Pablo Moreno-Muñoz, Nicki Skafte Detlefsen and Rasmus Berg Palm, who in addition to great company, are much better at coding than me.

A big thank you goes out to my officemates Jacob Thøstesen, Ragnhild Laursen and Kenneth Borup, for many great talks and afternoon trips for coffee and cake. I also want to thank Rikke Eriksen for being just far enough ahead to answer all my questions about being a PhD student, and for great friendship for many years.

Finally I will say, I am evermore grateful to Thomas and Panda for being my home base, always cheering me on and supporting me at all times and in all matters.

Helene Hauschultz
Aarhus, November 2022

Summary

In manifold learning the aim is to represent data sets in lower dimension by fitting a manifold to the data. Here the word manifold refers to the more general idea of a subsets which can be described in fewer variables than the ambient space, as opposed to the stringent mathematical definition of a manifold. Lower dimensional representations can be found by projecting a data point to the (unique) nearest point on the manifold. However, for any non-linear manifold, such a unique projection does not exist everywhere.

One way to study where a unique projection does exist is through reach. Reach gives a global bound on how far one can move from the manifold while ensuring that a unique nearest point exists. However, this bound is too restrictive in a practical setting, as many points further away will still have a unique nearest point. Instead we introduce a new uniqueness bound called the *pointwise normal reach*, which gives a bound on how far one can move in a normal direction while ensuring unique projections. In addition to the pointwise normal reach, we introduce a related uniqueness bound on immersed manifolds. This bound is beneficial as it utilises standard analytical tools to compute the bound.

Finally, we use these bounds in practice, as we aim to understand the uniqueness of representation in the autoencoder algorithm. We create a test which asks if a datapoint is within reach of the assigned latent representation. A datapoint which fails this test is not ensured to have a unique best choice of latent representation. We employ Monte Carlo Sampling to estimate the pointwise normal reach of three autoencoders. We find that the algorithms fail this test for most data points. Thus, almost no data points are certain to have a unique best choice of representation. To improve this, we introduce a regularising term into the training algorithm. This significantly improves the amount of points which passes the reach test. However, the experimental setup is currently too expensive to employ for non-toy datasets.

Resumé

Formålet med manifold learning er at kunne repræsenterer et datasæt i lavere dimensioner ved at fitte en mangfoldighed til datasættet. Ordet mangfoldig henviser her til en mere generel idé om en delmængde der kan repræsenteres ved færre variable end det omgivende rum, frem for den stringente matematiske definition. Når man har fundet en mangfoldighed kan datasættet repræsenteres ved at projektere punkterne til det nærmeste punkt på mangfoldigheden. Vi ved dog at for enhver ikke-lineær mangfoldighed findes en sådan projektion ikke overalt.

I matematikken kan vi undersøge hvor en sådan projektion eksisterer ved hjælp af reach. Reach giver en begrænsning på hvor langt vi kan bevæge os fra en delmængde således at der stadig findes et entydigt nærmeste punkt. Det viser sig dog i praksis at denne grænse er for restriktiv, da mange punkter længere væk end reach stadig har et entydigt nærmeste punkt i mængden. Vi introducerer derfor pointwise normal reach, som er en lokal grænse for hvor langt man kan bevæge sig fra i punkt i en normal retning, samtidig med at man er sikker på at der findes et entydigt nærmest punkt. Derudover introducerer vi en relateret grænse for immersioner, hvor vi kan udnytte differentialet af immersionen i udregningen.

Vi udnytter disse grænser i praksis, da vi undersøger entydigheden af repræsentationer lært af en auto-encoder algoritme. Vi konstruerer en test som spørger hvilke datapunkter er indenfor pointwise normal reach af dens repræsentation. Hvis et datapunkt fejler denne test betyder det at vi ikke kan være sikker på at der er et entydigt bedste valg af repræsentation. Vi bruger Monte Carlo simulering til at estimerer pointwise normal reach i trænede auto-encoders. Vi finder at næsten alle datapunkter fejler testen. For at forbedre dette introducerer vi reach-regularisering under træningen af auto-encoderen. Med denne regularisering forbedres antallet af punkter der består testen betydeligt. Udfordringen er at reach-regulariseringen for nuværende er for dyr til at bruge på andet end meget små modeller.

Introduction

This chapter serves as an introduction to the topics and themes of the dissertation. The work presented in this dissertation falls in the intersection of two different fields, mathematics and machine learning. The motivation was to find ways where the theoretical knowledge of mathematics could give new insights into manifold learning, which is a type of machine learning which takes inspiration from geometry. The question we ended up tackling was: Can latent representations be expected to be unique?

A problem which can arise when the work is done in two different fields is that words have different meanings. This problem arises in this dissertation as well. Throughout this text, the word manifold is used not only to refer to the stringent mathematical definition of a manifold, but also to refer to the general idea of a subset which can be represented in fewer variables than the ambient space. This choice has been made since in the topic of manifold learning the word is used in this way. To make the confusion complete, at some point we need the assumption that a subset is a smooth submanifold. Thus any mention of smooth (sub)manifold will always refer to the mathematical definition (e.g. see [6]).

Machine learning

The goal of machine learning is to develop an algorithm which knows how to behave in a certain situation. For instance, given a picture of a handwritten digit, it should know which digit it is. Or a more complex situation; should a self-driving car turn, accelerate og slow down at any given time. To achieve this goal, one could try to encode our prior knowledge about the phenomenon by hand into the algorithm, but it turns out that this is not an effective method [3]. Instead a better approach is to let the program discover the expected behaviour by finding patterns in a set of example behaviours. This automatic pattern discovery in a data set is what is known as machine learning.

The fastest growing field of machine learning is deep learning. Here the desired behaviour from above is modelled as a mapping, F, sending an input state to the expected behaviour. This mapping is approximated by a neural network. A neural network is the composition of a sequence of functions. Traditionally, the sequence alternates between a linear transformation with an added bias term, and a simple non-linear function applied element wise. The latter is called an activation function. Thus, such an network is a composition of functions of the form

$$\tilde{F}(x) = l_1 \circ \cdots \circ l_k(x),$$

where $l_i(z) = \phi(W_i x + b_i)$ for some activation function ϕ and weights W_i and b_i. The dimensions of the weights and the activation function are determined ahead of the training. The weights are found during the training process by minimising some predetermined loss function on the data set. In the simplest situation, we minimise the difference between the output of F and the output of $\tilde{F}$, that is

$$W^*, b^* = \underset{W,b}{\operatorname{argmin}} \sum_{x \in \mathcal{D}} \|F(x) - \tilde{F}(x)\|,$$

where $\mathcal{D}$ is the set of training data.

Data representations

In machine learning it is important to have well suited representations of the data in order to have successful learning outcomes. Such representation can be found manually through a process called feature engineering, where human investigators use their prior knowledge of the data to say which features are significant. However, such an approach is time consuming and difficult to scale [3]. Instead, it is beneficial to *learn* good representations. This approach is known as representation learning. Representation learning covers many different approaches. However, for our endeavours we are only interested in one type; dimensionality reduction. Dimensionality reduction is a type of learning scheme where the goal is to reduce the dimension of the data. The benefits of this are three fold: Firstly, in practice most data sets are very high dimensional, thus the computational cost of training the models is very high. If we are able to reduce the dimension while keeping the same intrinsic information, we are able to reduce computation cost. Secondly, the expectation is that reducing the dimension also reduces the noise in the data by removing axes, which do not contribute significant information. Thirdly, by removing redundant information one can extract the most important information in the data. Thus making successful learning more likely.

In order for dimensionality reducing schemes to be beneficial, it needs to be feasible to representent the data in fewer dimensions while still maintaining most of the information. The reasoning behind this is known as the *manifold hypothesis*. The manifold hypothesis states that most real world occurring data sets lie near a lower dimensional submanifold of the ambient data space [3]. Because of this hypothesis, learning lower dimensional representations is often called *manifold learning*, as we aim to learn the submanifold the data lies near.

Unique latent representations

Under the manifold hypothesis it follows that data points can be represented in lower dimensions by projecting them onto the manifold. The most natural being that the new representation must correspond to the closest point on the manifold from the original data point. It is also clear that any submanifold with non-zero curvature somewhere, will result in points in the ambient space which does not have a unique projection. Thus we risk that a data point does not have a unique representation. In practice, dimensionality reducing algorithms produce singular latent representations, and do not take into account that multiple best choices of such representations exist.

The uses of such representations include visualisation, such as scatter plots which are used to form scientific hypotheses as well as latent space statistics such as finding clusters in the latent representations (see Paper B). Thus, there is a hidden assumption that the latent representations are unique, but no investigations into whether or not this assumption is valid are conducted.

To our knowledge, the work presented in this dissertation is the first to discuss the risk of non-unique representations in representation learning. However, in the general field of data science there is other work on such non-uniqueness. For instance, Laursen & Hobolt [5] studies non-unique representations using non-negative matrix factorization in cancer genomics.

Reach in manifold learning

In the previous section we discussed how manifold learning attempts to learn lower dimensional latent representations by fitting a lower dimensional manifold to the data. The study of manifolds is a big field in mathematics, thus there is a natural interest from mathematicians into manifold learning. The perhaps most famous example of this is the paper "Testing the Manifold Hypothesis" by Fefferman et.al. [2], in which the authors develop a statistical test for the manifold hypothesis. The method described in the paper tests if a given data set fits a manifold of dimension k and a certain reach. The reach of a subset is the maximal distance one can move from the set while ensuring that there exists a unique nearest point in the set. The reasons for including the reach as a parameter in their test is given as follows: "We consider a bound on the reach to be a natural constraint since if the data lie within a distance less than the reach of the manifold, it can be denoised by mapping data points to the nearest point on the manifold." [2, p.958]

The global reach has also been used in results on manifold estimation. Here the use of reach differs from the work of Fefferman et.al. (and the work in this dissertation) as the reach of the underlying manifold from the data is generated which is used, instead of the reach of the learned manifold which is fitted to the data. In [4] the reach of the underlying manifold is used as a regularity condition as a minimax bound on the estimation of the manifold is found. In [7] the global reach influences how densely the data points must be sampled on the underlying manifold to be able to accurately estimate the topology. This use of reach as a regularity condition lies closer to the original motivation for introducing reach, than how the use of reach to ensure unique projection.

Sets of positive reach were introduced by Federer in his 1959 paper "Curvature Measures" [1]. Prior to Federer's work, similar integral geometric theories had been developed for convex sets and differentiable manifolds, respectively. Concretely, Steiner's formula and Weyl's formula were two tube formulas with the same result for the two classes of sets, but with different assumptions, such that neither follows from the other [9]. Federer's aim was to unite the two into one general theory, something he managed with the introduction of sets with positive reach [1]. The intention of this dissertation is not to give a thorough background on the history of integral geometry. Instead we show how Federer's intent was not directly to understand which points in the ambient space have a unique projection onto the set. In fact, the global reach value which Federer introduced, finds the worst case situation. A point within reach distance of the set is guaranteed a unique nearest point, but many, in fact most, points further away still have one.

Pointwise normal reach and autoencoders

The global reach gives a very rough bound on which points have a unique nearest point on the manifold. To remedy this we introduce *pointwise normal reach*. The pointwise normal reach gives a local bound on how far one can move from a point on the manifold in a normal direction while ensuring that a unique nearest point still exists. A big benefit of the pointwise normal reach is that it is computable. Thus we are able to use it to analyse the uniqueness of latent representations.

We consider the type of dimensionality reducing algorithm called an autoencoder. In its basic form the algorithm is a composition of two functions f and g, where g maps from the ambient space of the data set into a lower dimensional space called the latent space, and f maps the latent space back into the data space. Then the optimal f and g are found by minimising the distance between the data points and the reconstructed points, that is

$$f^*, g^* = \min_{f,g\in\mathcal{F}\times\mathcal{G}} \sum_{x\in\mathcal{D}} \|x - f(g(x))\|.$$

Here $\mathcal{F}$ and $\mathcal{G}$ are the predetermined function classes of f and g. These are typically chosen to be some type of neural network. f is called the decoder, and g is called the encoder.

For our setup, we consider the image of the decoder f to be the manifold. The optimal encoder is then a function g such that the reconstructed point $f(g(x))$ is the closest point on the manifold to x. The use of the word manifold is not reasonable from a mathematical perspective. There is no guarantee that the image of f is a (smooth) manifold. A more reasonable assumption is that f is an immersion [8]. For this reason we consider how one can find a similar bound to the pointwise normal reach under this assumption.

In this dissertation we present two papers. In Paper A we develop the mathematical theory of pointwise normal reach, and how we can calculate bounds on the distance which ensures that a unique nearest point exists. In Paper B we investigate how these bounds can be used to analyse the uniqueness of the autoencoder algorithm. This paper covers both the theoretical motivation as well as an experimental section where we try to implement the ideas in practice. Paper B is written as a conference paper. It thus had a strict page limit. In addition, it is also written to target the machine learning community.

References

[1] Federer, H. (1959). Curvature Measures. eng. *Transactions of the American Mathematical Society* 93(3), 418–491.

[2] Fefferman, C., S. Mitter and H. Narayanan (2016). Testing the manifold hypothesis. eng. *Journal of the American Mathematical Society* 29(4), 983–1049.

[3] Goodfellow, I., Y. Bengio and A. Courville (2016). *Deep Learning*. http://www.deeplearningbook.org. MIT Press.

[4] Kim, A.K.H. and H.H. Zhou (2015). Tight minimax rates for manifold estimation under Hausdorff loss. *Electronic Journal of Statistics* 9(1), 1562–1582.

[5] Laursen, R. and A. Hobolth (2022). A Sampling Algorithm to Compute the Set of Feasible Solutions for NonNegative Matrix Factorization with an Arbitrary Rank. *SIAM Journal on Matrix Analysis and Applications* 43(1), 257–273.

[6] Lee, J.M (2013). *Introduction to smooth manifolds*. eng. 2. ed. Graduate texts in mathematics ; 218. New York: Springer.

[7] Niyogi, P., S. Smale and S. Weinberger (2008). Finding the Homology of Submanifolds with High Confidence from Random Samples. eng. *Discrete & Computational Geometry* 39(1-3), 419–441.

[8] Shao, H., A. Kumar and P.T. Fletcher (2018). "The Riemannian Geometry of Deep Generative Models". *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 428–4288.

[9] Thäle, C. (2008). 50 years sets with positive reach – a survey. eng. *Surveys in Mathematics and its Applications* 3, 123–165.

Global reach

This chapter contains background information on reach. Unless otherwise noted, the results in this chapter are based on lectures and lecture notes made by Andrew du Plessis. Some of the results also appeared in the dissertation of Helene Mathilde Svane [8].

1.1 Definitions and properties

Let A be an arbitrary subset of $\mathbb{R}^n$. Then we denote the set $\mathrm{Unp}(A)$ to be the set of points which have a unique nearest point in A. It follows that we can define a projection onto A by

$$P_A : \mathrm{Unp}(A) \to A$$

which associates a point $x \in \mathrm{Unp}(A)$ to its nearest point in A.

Definition 1.1. *A subset A of $\mathbb{R}^n$ has* reach r*, if r is the maximal distance such that all points x in $\mathbb{R}^n$ with $d(x,A) < r$ have a unique nearest point in A.*

Here $d(x,A) = \inf\{\|x-a\| : a \in A\}$. It is clear that if x has a unique nearest point, then there exists $a \in A$ such that $d(x,A) = \|x-a\|$. Consider a subset A which is not closed, then there exists a convergent sequence $a_1, a_2, \ldots$ in A with limit $a \notin A$. Then it is clear that a does not have a unique nearest point in A, as for any given point in A we can always find one nearer. Hence, this implies that if $\mathrm{reach}\, A > 0$ then A is closed. So from now on we only consider closed sets.

We can also define reach as a local feature, such that for a point $a \in A$ we denote

$$\mathrm{reach}(A,a) = \sup\{r > 0 : \{x : \|x-a\| < r\} \subset \mathrm{Unp}(A)\}$$

Then $\mathrm{reach}(A) = \inf_{a\in A} \mathrm{reach}(A,a)$.

Definition 1.2. *Let $A \subset \mathbb{R}^n$ and $a \in A$. We then say $u \in \mathbb{R}^n$ is a tangent vector to A at a if $u = 0$ or if for all $\varepsilon > 0$ there exists $b \in A$ such that*

$$0 < \|b-a\| < \varepsilon \quad \textit{and} \quad \left\|\frac{b-a}{\|b-a\|} - \frac{u}{\|u\|}\right\| < \varepsilon.$$

Let T_aA denote the space of tangent vectors of A at a, and call it the tangent cone. We define the normal cone of A at a to be the dual cone.

$$N_aA = \{u \in \mathbb{R}^n : \langle u, v\rangle \leq 0\, \forall v \in T_aA\}.$$

The normal cone is a convex cone, but this is not necessarily true for the tangent cone. We will not go into details about cones. The definition and more details are found in Federer's remark 4.5 [3].

The following theorem from Federer [2, Theorem 4.8] gives a number of important properties of spaces with positive reach. These results only require the tangent- and normal cones defined above, and not for them to be full vector spaces.

Theorem 1.3 ([3] Theorem 4.8). *Let $A \subset \mathbb{R}^n$ be a non-empty, closed subset of $\mathbb{R}^n$. Writing $d(\cdot, a) = d(\cdot)$, and $U = \operatorname{Unp}(A)$, the following statements holds:*

1. $|d(x) - d(y)| \leq |x - y|$ *for all* $x, y \in \mathbb{R}^n$,
2. *If $a \in A$ and $P = \{v \in \mathbb{R}^n : P_A(a+v) = v\}$, $Q = \{v \in \mathbb{R}^n : d(a+v) = \|v\|\}$. Then P and Q are convex and $P \subset Q \subset N_a A$.*
3. *Let $x \in \mathbb{R}^n \setminus A$ and suppose d is differentiable in x, then $x \in U$ and $\nabla d(x) = \frac{x - P_A(x)}{d(x)}$.*
4. *P_A is continuous.*
5. *d is continuously differentiable on the interior $\operatorname{Int}(U \setminus A)$ and d^2 is continuously differentiable on $\operatorname{Int} U$ with $\nabla d^2(x) = 2(x - P_A(x))$.*
6. *Let $a \in A$ and $v \in \mathbb{R}^n$ such that $0 < T = \sup\{t : P_A(a + tv) = a\} < \infty$, then $a + Tv \notin \operatorname{Int} U$.*
7. *Let $x \in U$ with $P_A(x) = a$ and $\operatorname{reach}(A, a) > 0$ and let $b \in A$, then $\langle x - a, a - b\rangle \geq -\frac{\|a-b\|^2 \|x-a\|}{2 \operatorname{reach}(A,a)}$.*
8. *Let $0 < r < q < \infty$. For $x, y \in U$ with and $d(x) < r$, $d(y) < r$ and $\operatorname{reach}(A, P_A(x)) > q$, $\operatorname{reach}(A, P_A(y)) > q$, then $\|P_A(x) - P_A(y)\| \leq \frac{q}{w-r}\|x - y\|$.*
9. *If $0 < s < r < \operatorname{reach}(A)$, then ∇d is Lipschitz on $\{x \in \mathbb{R}^n | s < d(x) < r\}$ and ∇d^2 is Lipschitz on $\{x \in \mathbb{R}^n | d(x) < r\}$.*
10. *If $a \in A$, $T_a A = \left\{u \in \mathbb{R}^n | \liminf_{t \to 0^+} t^{-1} d(a + tu) = 0\right\}$.*
11. *Let $a \in A$ and $\operatorname{reach}(A, a) > r > 0$, and let $u \in \mathbb{R}^n$ with $\langle u, v\rangle \leq 0$ whenever $P_A(a+v) = a$, and $\|v\| = r$, then $\lim_{t \to 0^+} t^{-1} d(a + tu) = 0$.*
12. *If $a \in A$ with $\operatorname{reach}(A, a) > r > 0$, then $N_a A = \{\lambda v | \lambda \geq 0, \|v\| = r, P_A(a + v) = a\}$. $T_a A$ is the convex dual cone to $N_a A$ and $\lim_{t \to 0^+} t^{-1} d(a + tu) = 0$ for $u \in T_a A$.*

Lemma 1.4. *Let $A \subset \mathbb{R}^n$ be a closed subset. Then the following are equivalent:*

1. $\operatorname{reach} A > r$,
2. *For all $x, y \in A$ with $\ell = \|x - y\| \leq 2r$, the angles between $N_x A$ and $N_y A$ and the line L between x and y are greater than $\cos^{-1}(\ell/2r)$.*

Let $V \subset \mathbb{R}^n$ be a vector space and $L \subset \mathbb{R}^n$ a line. We define the angle θ between V and L to be the acute angle between L and its projection onto V, in the plane spanned by the two lines. Let u be a unit vector in the direction of L, then it is clear that

$$\sin\theta = \frac{\|u - \pi_V u\|}{1} = \|u - \pi_V u\|.$$

This calculation also shows that θ is indeed the smallest possible angle, as the distance from u to any other point in V must be larger.

Proof. (1) $\Rightarrow$ (2): Let $x, y \in A$ with $\ell = \|x-y\| \leq 2r$. By translating the set we can assume without loss of generality that $x = 0$. Let $z \in N_x A$ with $\|z\| = r$. Now x is the unique nearest point to z in A, so $\|y-z\| > r$. The law of cosines tell us $\|z-y\|^2 = r^2 + \ell^2 - 2r\ell\cos\theta$ thus

$$\begin{aligned} r^2 < r^2 + \ell^2 - 2r\ell\cos\theta &\Leftrightarrow 2r\ell\cos\theta < \ell^2 \\ &\Leftrightarrow \cos\theta < \frac{\ell}{2r} \\ &\Leftrightarrow \theta > \cos^{-1}(\ell/2r). \end{aligned}$$ □

The last inequality holds as cosine is a decreasing function between 0 and $\pi/2$.

(2) $\Rightarrow$ (1): Let $z \in \mathbb{R}^n \setminus A$ and suppose $y_1, y_2 \in A$ satisfy $d(z, A) = \|z-y_1\| = \|z-y_2\| = s$. Next we need to prove $s > r$. Suppose by contradiction that $s \leq r$, then $\|y_1 - y_2\| \leq 2s \leq 2r$. Then we get that the angle θ between $z - y_i$ and L is greater than $\cos^{-1}(\ell/2r)$. Using the law of sines we get $\frac{1}{2}\ell = s\sin(\frac{1}{2}\alpha) = s\sin(\pi/2 - \theta) = s\cos(\theta) < \frac{s\ell}{2r}$, thus $r < s$.

Proposition 1.5. *Let $A \subset \mathbb{R}^n$ with* reach $A > r$. *Let $x, y \in A$ with $\|x - y\| \leq 2r$, and let L be the line segment between them. Then*

1. $L \subset \operatorname{Unp}(A)$ *and $P_A(L)$ is a simple curve in A from x to y.*
2. *Let B be an s-ball with $s \leq r$ and $x, y \in B$. Then $P_A(L) \subset B$.*

Proof. *(i):* Since $\|x - y\| \leq 2r$ all points on L must have distance less or equal to r to either x or y. So $L \subset \operatorname{Unp}(A)$.

$P_A\big|_L : L \to A$ gives rise to a parametrisation of a curve, but it is not necessarily an injective parametrisation. We want to show that we can find a parametrisation of $P_A(L)$ which is injective. First notice if $y_1, y_2 \in L$ with $P_A(y_1) = P_A(y_2) = a$. Then $y_1, y_2 \in a + N_a A$, and as $N_a A$ is a convex cone the line segment L' between the two points will also lie in $a + N_a A$. Hence $P_A(L') = \{a\}$. Furthermore Theorem 1.3 tells us P_A is continuous, hence $P_A^{-1}(a)$ will be closed, so especially will it be a closed interval.

Let $\rho : [0,1] \to A$ be a parametrisation of $P_A(L)$ given by $\rho(t) = P_A((1-t)x + ty)$. We can construct an equivalence relation as

$$t_1 \sim t_2 \quad \Leftrightarrow \quad \rho(t_1) = \rho(t_2).$$

This then induces a map

$$\bar{\rho} : {}^{[0,1]}\!\big/\!_{\sim} \to P_A(L)$$

which is bijective and continuous in the quotient topology.

Furthermore ${}^{[0,1]}\!/\!_{\sim}$ is compact and we show that it is Hausdorff. For that suppose $[a], [b] \in {}^{[0,1]}\!/\!_{\sim}$ with $[a] \neq [b]$. Then we have $a, b \in [0,1]$ with $\rho(a) \neq \rho(b)$. As $\rho(a), \rho(b) \in P_A(L) \subset \mathbb{R}^n$ we can find two open neighbourhoods $\rho(a) \in U_a \subset \mathbb{R}^n$ and $\rho(b) \in U_b \subset \mathbb{R}^n$ such that $U_a \cap U_b = \emptyset$. Now both U_a and U_b are relatively open in A so the preimages $\rho^{-1} U_a$, $\rho^{-1} U_b$ are open in $[0,1]$ then the quotients are open in the quotient topology. As we chose $U_a \cap U_b = \emptyset$, it follows that we have constructed two disjoint open neighbourhoods of $[a]$ and $[b]$.

We claim that ${}^{[0,1]}\!/\!_{\sim}$ is homeomorphic to an interval. Theorem 1.3 also tells us that P_A is Lipschitz, so the curve $P_A(L)$ must have finite length. Hence we can define

$s(t)$ to be the arc-length from x to $(1-t)x+y$. This gives a function $s:[0,1]\to[0,L]$ where L is the arc-length of $P_A(L)$. It is clear that s is surjective and continuous. The map induces a commutative diagram

$$\begin{array}{ccc} S:[0,1] & \longrightarrow & [0,L] \\ & \searrow \quad \nearrow & \\ & [0,1]/\sim & \end{array}$$

We claim that the map $\bar{s}: {}^{[0,1]}\!/_{L} \to [0,L]$ is injective. Notice that

$$s(a) > s(b) \quad \text{for } a > b, a \nsim b,$$

since if $a \nsim b$ then $P_A(a) \neq P_A(b)$ so there must be some distance between them. Thus $\bar{s}: {}^{[0,1]}\!/_{L} \to [0,L]$ is bijective, and as ${}^{[0,1]}\!/_{L}$ is compact and Hausdorff, we get that ${}^{[0,1]}\!/_{\sim}$ is homeomorphic to $[0,L]$. Thus, we get an injective parametrisation of $P_A(L)$.

(2): Let $z \in L \setminus A$ with $P_A(z) = a$. We can then write $z = a + tn_a$, with $n_a \in N_aA$ and $t \in (0,r]$. Let B' be the ball with radius r and centre $a + rn_a$. Since reach $A > r$, $B' \cap A = \{a\}$.

Now, let P be the plane spanned by L and $z - a$, and let $C = B \cap P$, $D = B' \cap P$ and $M = D \cap L$. As the centre of B' is included in P, D is an r-disk. M is a line segment which is properly contained in L.

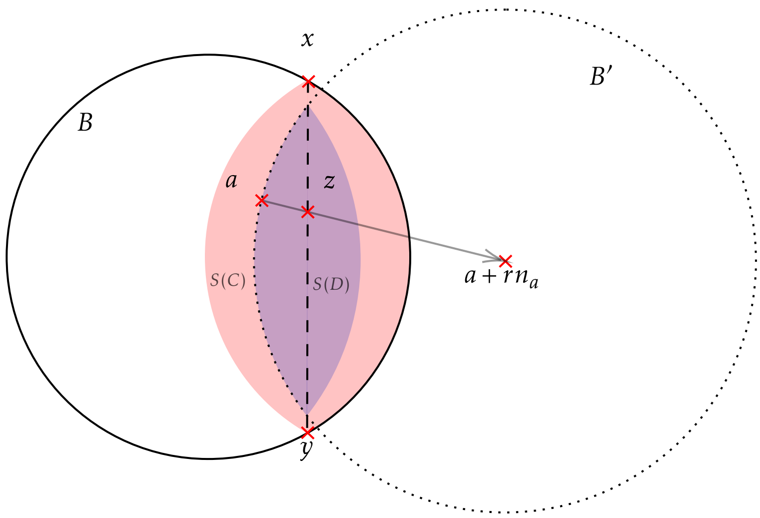

Figure 1: Illustration of the proof from Theorem 1.5 (II).

Now let ℓ denote the line containing L, and let

$$S(C) = C \cap \{C\text{'s reflection in } \ell\}$$
$$S(D) = D \cap \{D\text{'s reflection in } \ell\}.$$

Then $a \in S(D) \subset S(C) \subset C \subset B$. □

We can strengthen the proposition in the case that N_aA is a full vector space for all $a \in A$.

Corollary 1.6. *If in addition N_aA is a full vector space for all $a \in A$, then $P_A\big|_L$ is injective.*

Proof. If $P_A(z_1) = P_A(z_2) = a$ for two points on L with $z_1 \neq z_2$, it follows that $z_1, z_2 \in N_aA$. Then all of L must lie in N_aA, especially will the endpoints $x, y \in N_aA + a$, however this contradicts lemma 1.4. □

Theorem 1.7. *For $A \subset \mathbb{R}^n$ a closed subset,*

$$\operatorname{reach}(A) = \inf\left\{ \frac{\|x-y\|^2}{2d(y-x, T_xA)} : x, y \in A, y - x \notin T_xA \right\}.$$

Proof. Let $r > 0$. First assume that $x, y \in A$ with $\|x-y\| > 2r$. Note that $d(y-x, T_xA) \leq \|x-y\|$ as $x \in T_xA + x$. Hence we get

$$2r < \|x-y\| \leq \frac{\|x-y\|^2}{d(y-x, T_xA)}.$$

Now suppose $\|x-y\| \leq 2r$ and that the angle θ between N_xA and the line between x and y satisfy $\cos\theta < \frac{\|x-y\|}{2r}$. Using the fact that $\cos\theta = \frac{d(y-x,T_xA)}{\|x-y\|}$ we get that

$$\frac{d(y-x, T_xA)}{\|x-y\|} < \frac{\|x-y\|}{2r} \quad \Leftrightarrow r < \frac{\|x-y\|^2}{2d(y-x, T_xA)}.$$

This proves that the second statement of lemma 1.4 is equivalent to

$$r < \frac{\|x-y\|^2}{2d(y-x, T_xA)} \quad \forall x, y \in A\, y - x \notin T_xA.$$

Now applying lemma 1.4 gives us

$$r < \operatorname{reach} A \quad \Leftrightarrow \quad r < \frac{\|x-y\|^2}{2d(y-x, T_xA)} \quad \forall x, y \in A\, y - x \notin T_xA. \tag{1.1}$$

Thus we have that

$$\operatorname{reach}(A) = \inf\left\{ \frac{\|x-y\|^2}{2d(y-x, T_xA)} : x, y \in A, y - x \notin T_xA \right\}.$$ □

The next theorem we will state without proof. It shows that a topological manifold is ensured to be $C^{1,1}$.

Theorem 1.8. *[7] Let $A \subset \mathbb{R}^n$ be a subset with* $\operatorname{reach} A > r > 0$, *and suppose for $n > k \geq 1$ T_aA is a k–dimensional vector space for all $a \in A$. Then A is a $C^{1,1}$ manifold of dimension k.*

A function is $C^{1,1}$ if it is differentiable with Lipschitz continuous derivative. In fact, a $C^{1,1}$ function is 'almost' C^2, a fact which follows from Radermacher's Theorem below. Later we are going to utilise the fact that for a $C^{1,1}$ curve, the curvature exists almost everywhere.

Theorem 1.9 (Radermacher's Theorem). *[3, Theorem 3.1.6] If $f : \mathbb{R}^m \to \mathbb{R}^n$ is Lipschitz, then f is differentiable almost everywhere.*

1.2 Spindles and their properties

We now introduce the concept of spindles and how they relate to reach.

Definition 1.10. *Let L be a closed line-segment in $\mathbb{R}^n$ of length $|L| \leq 2r$. Then the* r-spindle *$S(L,r)$ generated by L is the intersection of all r-balls whose boundaries each contain both endpoints of L*

In the case that L has length $2r$, there is only one such ball, so $S(L,r)$ is this ball. Considering the case $n = 2$, there are two such balls, each the mirror image of the other in the line L. Hence, $S(L,r)$ is the intersection of these two balls. It is a well known fact that the angle from a chord subtended from two points on a circle is constant corresponding to the two arcs. So by computing this angle we get

$$S(L,r) = \left\{ x \in \mathbb{R}^2 : \text{ The angle } \angle axv \text{ is obtuse, with } \sin(\angle axb) \leq \frac{|L|}{2r} \right\} \tag{1.2}$$

Lemma 1.11. *Let $0 < s \leq r$, and let $M \subset N \subset \mathbb{R}^2$ be the closed line segments of length less than $2s$. Then $S(M,r) \subset S(N,s)$.*

Proof. Let a,b denote the endpoints of M and c,d denote the endpoints of N. Let $x \in S(M,r)$. Using (1.2) we get that the angle $\angle axb$ is obtuse, so the larger angle $\angle cxd$ is also obtuse, and

$$\frac{|M|}{2r} \leq \frac{|N|}{2s},$$

as $|M| \leq |N|$ and $r \geq s$. So $x \in S(N,s)$. □

Lemma 1.12. *Let L be a line in $\mathbb{R}^n$, $n \geq 3$ with $|L| \leq 2r$. Then we can describe $S(L,r)$ as the union of the two dimensional spindles $S_P(L,r)$ in P generated by $L \subset P$, $P \in \mathcal{P}$, where $\mathcal{P}$ consists of all the planes which contain L.*

Proof. Let $x \in S(L,r)$. If $x \in L$ it follows directly that $x \in S_P(L,r)$ for all $P \in \mathcal{P}$. Suppose that x does not lie on L, then there exists a unique plane $P_x \in \mathcal{P}$, such that it contains both L and x. So we have to show that $x \in S_{P_x}(L,r)$.

We noted earlier that $S_{P_x}(L,r)$ is the intersection of exactly two disks of radius r. An r-disk containing L must lie in an r-ball containing L, so it follows from the definition of the spindle that $x \in S_{P_x}(L,r)$.

Suppose conversely that $x \in S_{P_0}(L,r)$ for some $P_0 \in \mathcal{P}$. Let B be any ball of radius r which contains L. Now $B \cap P_0$ is not empty as both contain L, hence it must be a disk of radius $s \leq r$. Now using the definition of spindles and lemma 1.11 we get that

$$x \in S_{P_0}(L,r) \subseteq S_{P_0}(L,s) \subset B,$$

which proves the claim. □

We can now extend lemma 1.11 to arbitrary dimensions.

Corollary 1.13. *Let $0 < s \leq r$, and let $M \subset N \subset \mathbb{R}^n$ be the closed line segments of length less than $2s$. Then $S(M,r) \subset S(N,s)$.*

Proof. We consider the intersections with affine planes P containing M. We observe

$$S(M,r)\cap P = S_P(M,r),$$

and we know $S(M,r)=\bigcup_P S_P(M,r)$, for P affine planes containing M. Similarly, $S(N,s)$ is the union of $S_P(N,s)$. From the two dimensional case we know that $S_P(M,r)\subset S_P(N,s)$ for all P, so $S(M,r)\subset S(N,s)$. □

Lemma 1.14. *Let M be a closed line segment of length $|M|<2r$ in $\mathbb{R}^n$. Then $S(M,r)$ is the intersection of all closed balls of radius at most r which contains M.*

Proof. Let C be a closed s-ball which contains M and $s\le r$. Let ℓ be the affine line containing M. Then $C\cap\ell$ is a closed line segment N, which contains M. So from (1) we know that $S(M,r)\subset S(M,s)\subset C$, so $S(M,r)\subset S$. □

Lemma 1.15. *Let M be a closed line segment of length $|M|<2r$ in $\mathbb{R}^n$. Then $S(M,r)$ is the intersection of all closed balls of radius strictly less than r, which contains M.*

Proof. Let S denote the intersection of all balls with radius strictly less than r which contains M. Now, lemma 1.14 gives that $S\subset S(M,r)$.

Next let $x\in S$. We want to prove that $x\in S(M,r)$. First we observe that lemma 1.14 also gives that $x\in S(M,s)$ for all $s\in(\frac{1}{2}|M|,r)$. We can restrict the situation to the two-dimensional spindle in the plane containing M and x. Denote the end-points of M a,b we get

$$\sin(\angle axb)\le\frac{|M|}{2s}$$

for all $s\in(\frac{1}{2}|M|,r)$. Now taking the limit $s\to r$ gives

$$\sin(\angle axb)\le\frac{|M|}{2r}.$$

So x is contained in the two-dimensional spindle in the plane, and hence it is contained in $S(M,r)$.

□

Corollary 1.16. *Let $A\subset\mathbb{R}^n$ have $\operatorname{reach} A>r$. Let $x,y\in A$ with $\|x-y\|\le 2r$ and let L denote the line segment between them. Then $P_A(L)\subset S(L,r)$.*

Proof. This follows directly from Proposition 1.5 as $P_A(L)$ is contained in all balls with radius $s\le r$ which contain x and y. □

Proposition 1.17. *Let $A\subset\mathbb{R}^n$ be a closed subset, then the following are equivalent:*

1. *A has $\operatorname{reach} A>r$,*
2. *For all $x,y\in A$ with $0<\|x-y\|<2r$*

   $$S(L,r)\cap(A\setminus\{x,y\})\neq\emptyset,$$

   where L is the line segment from x to y.

Proof. The implication (1) $\Rightarrow$ (2) follows directly from Corollary 1.5. Conversely, we assume that A does not have reach $A > r$, then there exist $z \in \mathbb{R}^n \setminus A$ such that $d(z,A) = s \leq r$ and there is $x,y \in A$ with $\|x-z\| = \|y-z\| = s$. There must exist $w \in S(L,r) \cap (A \setminus \{x,y\})$ as it is non-empty. However we also have that $S(L,r) \setminus \{x,y\}$ is contained in the open ball $B(z,s)$, hence $d(z,A) < s$ giving the contradiction. □

Definition 1.18. *For a two dimensional spindle $S(L,r)$, the* spindle angle *is the acute angle between L and the boundary arcs of $S(L,r)$. An easy calculation shows that* $\sin\varphi = \frac{|L|}{2r}$. *The* spindle angle *of a higher dimensional spindle $S(L,r)$ is the spindle angle of $S_P(L,r)$ for any affine plane P containing L.*

Definition 1.19. *Let α be a curve. Suppose for all $x,y \in \alpha$ with $\|x-y\| \leq 2r$, the sub-arc of α from x to y is contained in the r-spindle with endpoints x and y. Then we say α has the* r-spindle property.

Proposition 1.20. *Let $\alpha \subset \mathbb{R}^n$ be a C^1 arc with the r-spindle property. Let ℓ be the line segment joining the endpoints of α and assume the length of ℓ is $\leq 2r$. Let φ be the spindle angle of the r-spindle S generated by ℓ. Then the angle between ℓ and any tangent vector to α is at most φ.*

Proof. The result follows directly for the tangent vectors at the endpoints x and y. Let z be another point in α, and let ℓ' be the line between x and z and let S' be the r-spindle generated by ℓ'. As $\alpha \subset S$, the sub-arc from x to z is contained in S, and hence $\ell' \subset S$. So in particular ℓ' is contained in any closed ball of radius at least r containing ℓ. So by the characterisation from lemma 1.14 (2), $S' \subset S$.

The angle between the tangent vector $\mathbf{t}$ at z to α is at most the spindle angle φ' of S'. So the angle between $\mathbf{t}$ and ℓ is at most φ' plus the angle θ between ℓ' and ℓ. However, $\theta + \varphi'$ is also the maximum angle between ℓ and the tangent vectors to S' at x since $S' \subset S$. So $\theta + \varphi' \leq \varphi$, which completes the proof. □

Proposition 1.21. *Let $\psi : I \to \mathbb{R}^n$ be a C^1 arc-length parametrisation of a curve C. Suppose C has the r-spindle property. Then the following holds*

1. $\|\psi'(t_1) - \psi'(t_2)\| \leq \frac{1}{r}\|t_1 - t_2\|$ *for all* $t_1, t_2 \in I$.
2. ψ *is twice differentiable almost everywhere, and the curvature is* $\leq \frac{1}{r}$ *almost everywhere.*

Proof. Suppose $t_1, t_2 \in I$ with $\|\psi(t_1) - \psi(t_2)\| \leq 2r$. Then $\psi([t_1,t_2]) \subset S([\psi(t_1),\psi(t_2)],r)$, so it follows from proposition 1.20 that the angles between $\psi'(t_1), \psi'(t_2)$ and the line $[\psi(t_1),\psi(t_2)]$ are less than the spindle angle ϕ for $S([\psi(t_1),\psi(t_2)],r)$. Hence the angle between $\psi'(t_1)$ and $\psi'(t_2)$ is $\leq 2r$. Then using the double-angle formula and the spindle angle calculation we get

$$\begin{aligned}\|\psi'(t_1) - \psi'(t_2)\|^2 =& 2 - 2\cos\theta \\ \leq& 2 - 2\cos 2\phi \\ =& 2 - 2(1 - 2\sin^2\phi) = (2\sin\phi)^2 \\ =& \left(\frac{1}{r}\|\psi(t_1) - \psi(t_2)\|\right)^2.\end{aligned}$$

Since ψ is arc-length parameterised, $\|t_1 - t_2\|$ is the arc-length from $\psi(t_1)$ to $\psi(t_2)$, hence

$$\|\psi'(t_1) - \psi'(t_2)\| \leq \frac{1}{r}\|\psi(t_1) - \psi(t_2)\| \leq \frac{1}{r}\|t_1 - t_2\|.$$

Now suppose that $\|\psi(t_1) - \psi(t_2)\| > 2r$. As $\|t_1 - t_2\| < \infty$, the arc-length of $\psi([t_1, t_2])$ is finite. Hence, we can create a sequence

$$t_1 = t^1 < t^2 < \cdots < t^n = t_2,$$

such that $\|\psi(t^i) - \psi(t^{i+1})\| \leq 2r$ for all $i = 1, \ldots, n-1$. Then

$$\begin{aligned}
\|\psi'(t_1) - \psi'(t_2)\| &\leq \sum_{i=1}^{n-1} \|\psi'(t^i) - \psi'(t^{i+1})\| \\
&\leq \sum_{i=1}^{n-1} \frac{1}{r}\|\psi(t^i) - \psi(t^{i+1})\| \\
&\leq \frac{1}{r}\sum_{i=1}^{n-1} \left\{\text{arc-length from } \psi(t^i) \text{ to } \psi(t^{i+1})\right\} \\
&= \frac{1}{r}\{\text{arc-length from } \psi(t_1) \text{ to } \psi(t_2)\} = \frac{1}{r}\|t_1 - t_2\|.
\end{aligned}$$

We have now shown that ψ' is Lipschitz, hence Radermachers theorem tells us that ψ' is differentiable almost everywhere. Suppose ψ' is differentiable in $t \in I$. Using the first part of the proof we get that

$$\|\psi'(t+h) - \psi'(t)\| \leq \frac{1}{r}\|h\|.$$

Dividing with $\|h\|$ and letting $h \to 0$ we get that $\|\phi''(t)\| \leq \frac{1}{r}$. □

Proposition 1.22. *Let $\alpha \subset \mathbb{R}^n$ be a closed arc with the r-spindle property. Then* reach $\alpha > r$.

Proof. Suppose that there exists $x \in \mathbb{R}^n$ and $z_1, z_2 \in \alpha$ such that $\|z_1 - x\| = \|z_2 - x\| = d = d(x, \alpha)$. Suppose for contradiction that $d \leq r$, then $\|z_1 - z_2\| \leq 2r$. We then have a closed ball B with radius d and centre x and $z_1, z_2 \in \partial B$. As $d < r$, the r-spindle S with endpoints z_1 and z_2 is contained in B by lemma 1.14 (2), and furthermore $S \cap \partial B = \{z_1, z_2\}$. As α has the r-spindle property, the sub-arc α' with endpoints z_1, z_2 is contained in S, but then there must be points in α' contained in the interior of B. So especially there must be $z_3 \in \alpha$ with $\|z_3 - x\| < d$. This contradicts the assumption that $d = d(x, \alpha)$, so we must have $d > r$. Then it follows that α must have reach r. □

1.3 Schur's Theorem for $C^{1,1}$ curves

Schur's Theorem is usually stated for C^2 curves, however, as a manifold with positive reach is only certain to be $C^{1,1}$ we here prove it for that situation.

Theorem 1.23 (Schur's Theorem). *Suppose that α is a closed, $C^{1,1}$ arc in $\mathbb{R}^n$ and that $\bar{\alpha}$ is a closed, simple C^2 plane closed, both arc-length parametrised. Suppose $\bar{\alpha}$, together with a line between the end-points, is the boundary of a convex set. Suppose that both are arc-length parametrised. Furthermore suppose*

- *$\bar{\alpha}$ has curvature $\bar{K}$ always positive,*
- *α has curvature K with $\bar{K}(s) \geq K(s)$ almost everywhere.*

Then
$$\|\alpha(L) - \alpha(0)\| \geq \|\bar{\alpha}(L) - \bar{\alpha}(0)\|.$$

Radermachers theorem tells us that a Lipschitz function is differentiable almost everywhere. Thus in our situation where the curve is $C^{1,1}$ the derivative is differentiable almost everywhere, which in turn implies that we can find the curvature almost everywhere. For the proof of Schur's theorem we need the following result on the length of a Lipschitz curve.

Lemma 1.24. *Let $f : [a,b] \to \mathbb{R}^n$ be a Lipschitz curve. Then there exists $s : [a,b] \to [0,L]$, where $s(t)$ is the arc-length from $f(a)$ to $f(t)$. Furthermore, for $a \leq \alpha \leq \beta \leq b$,*

$$\int_\alpha^\beta |f'(t)|\, dt = s(\beta) - s(\alpha).$$

Proof. As f is Lipschitz, it is rectifiable, that is, it has finite length. This is easy to verify, as the definition of the arc-length is the supremum of

$$\sum_{i=1}^{n} |f(t_i) - f(t_{i-1})|$$

over all partitions $a = t_0 < \ldots t_n = b$. Given such a partition,

$$\sum_{i=1}^{n} |f(t_i) - f(t_{i-1})| = \sum_{i=1}^{n} \left| \int_{t_{i-1}}^{t_i} f'(t)\, dt \right| \leq \sum_{i=1}^{n} \int_{t_{i-1}}^{t_i} \left| f'(t) \right| dt,$$

hence, it follows that

$$s(\beta) - s(\alpha) \leq \int_\alpha^\beta \left| f'(t) \right| dt. \tag{1.3}$$

We know that $|f(\beta) - f(\alpha)| \leq s(\beta) - s(\alpha)$, for all $\alpha \leq \beta$, so

$$\frac{|f(t+h) - f(t)|}{h} \leq \frac{s(t+h) - s(t)}{h} \leq \frac{1}{h} \int_t^{t+h} \left| f'(u) \right| du.$$

If we let $h \to 0$, the left hand side will go to $|f'(t)|$. Using the fundamental theorem of calculus for the Lebesgue integral, we get that the right hand side will go to $|f'(t)|$. Hence s' exists almost everywhere, and $s'(t) = |f'(t)|$ almost everywhere.

We know that s is a non-decreasing function. We can extend s to the interval $[\alpha - a, \beta + 1]$ by

$$s(t) = s(\alpha), \quad \text{for } t \in [\alpha - 1, \alpha],$$
$$s(t) = s(\beta), \quad \text{for } t \in [\beta, \beta + 1].$$

s is still non-decreasing, so s' is non-negative and Lebesgue integrable. Furthermore,

$$s'(t) = \lim_{n \to \infty} \frac{s(t + \frac{1}{n}) - s(t)}{\frac{1}{n}},$$

for almost all $t \in [\alpha, \beta]$.

Using Fatout's lemma we get that

$$\begin{aligned}\int_\alpha^\beta s'(t)\,dt &\leq \liminf_{n\to\infty} \int_\alpha^\beta \frac{s(t+\frac{1}{n})-s(t)}{\frac{1}{n}} \\ &= \liminf_{n\to\infty} n\left[\int_\alpha^\beta s\left(t+\frac{1}{n}\right)dt - \int_\alpha^\beta s(t)\,dt\right] \\ &= \liminf_{n\to\infty} n\left[\int_{\alpha+\frac{1}{n}}^{\beta+\frac{1}{n}} s(t)\,dt - \int_\alpha^\beta s(t)\,dt\right] \\ &= \liminf_{n\to\infty} n\left[\int_\beta^{\beta+\frac{1}{n}} s(t)\,dt - \int_\alpha^{\alpha+\frac{1}{n}} s(t)\,dt\right].\end{aligned}$$

However,

$$\int_\beta^{\beta+\frac{1}{n}} s(t)\,dt = \frac{1}{n}s(\beta) \quad \text{and} \quad \int_\alpha^{\alpha+\frac{1}{n}} s(t)\,dt \geq \frac{1}{n}s(\alpha),$$

so combining the above gives

$$\int_\alpha^\beta s'(t)\,dt \leq \liminf_{n\to\infty} n\left(\frac{1}{n}s(\beta) - \frac{1}{n}s(\alpha)\right) = s(\beta) - s(\alpha).$$

Combining this with equation (1.3) gives us

$$s(\beta) - s(\alpha) \leq \int_\alpha^\beta |f'(t)|\,dt = \int_\alpha^\beta s'(t)\,dt \leq s(\beta) - s(\alpha),$$

which completes the proof. □

Proof (Proof of Schur's Theorem). Radermachers theorem tells us that a Lipschitz function is differentiable almost everywhere, hence the curvature of α exists almost everywhere.

There must exist s_0 such that $\bar{\alpha}'(s_0)$ is parallel to the line ℓ between the end-points of $\bar{\alpha}$.

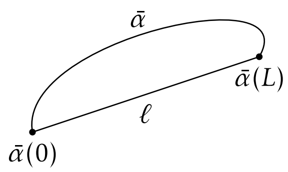

Considering $\bar{\alpha}$ as a curve in $\mathbb{R}^n$, we can, using a rigid motion, assume $\alpha(s_0) = \bar{\alpha}(s_0)$, and $\bar{\alpha}(s_0) = T_0$. Denote $\alpha'(s) = T(s)$ and $\bar{\alpha}'(s) = \bar{T}(s)$. Then

$$\begin{aligned}\|\alpha(L) - \alpha(0)\| &\geq \langle \alpha(L) - \alpha(0), T_0\rangle = \int_0^L \langle T(s), T_0\rangle\,ds, \\ \|\bar{\alpha}(L) - \bar{\alpha}(0)\| &= \langle \bar{\alpha}(L) - \bar{\alpha}(0), T_0\rangle = \int_0^L \left\langle \bar{T}(s), T_0\right\rangle ds,\end{aligned}$$

where the equality in the latter equation comes from the fact that $\bar{\alpha}(L)-\bar{\alpha}(0)$ is parallel to T_0. This tells us that it is enough to show that $\langle T(s), T_0\rangle \geq \left\langle \bar{T}(s), T_0\right\rangle$.

Both T and $\bar{T}$ are parametrised curves on S^{n-1}. As $\bar{\alpha}$ is a plane curve, $\bar{T}$ moves with positive speed $\bar{K}$ along a great circle. So the length of $\bar{T}$ from s_0 to s becomes

$$L(\bar{T}\big|_{[s_0,s]}) = \int_{s_0}^{s} \bar{K}(s)\,ds = \cos^{-1}\left\langle T_0, \bar{T}(s)\right\rangle.$$

As α is a $C^{1,1}$ curve, α' is Lipschitz and hence almost everywhere differentiable, per Radermachers theorem. So by letting K be the function which is α'' almost everywhere, we get from Lemma 1.24 that

$$L(T\big|_{[s_0,s]}) = \int_{s_0}^{s} K(s)\,ds \geq \cos^{-1}\langle T(s), T_0\rangle,$$

where the inequality comes from the fact that $\cos^{-1}\langle T(s), T_0\rangle$ is the length of the great circle arc which is the shortest path between the two points.

Now

$$\cos^{-1}\left\langle \bar{T}(s), T_0\right\rangle = \int_{s_0}^{s} \bar{K}(s)\,ds \geq \int_{s_0}^{s} K(s)\,ds \geq \cos^{-1}\langle T(s), T_0\rangle.$$

$\cos^{-1} : [0,1] \to [0,\pi]$ is decreasing, so $\left\langle \bar{T}(s), T_0\right\rangle \leq \langle T(s), T_0\rangle$, which completes the proof. □

Lemma 1.25. *If* $\|\alpha(L)-\alpha(0)\| = \|\bar{\alpha}(L)-\bar{\alpha}(0)\|$ *then we obtain* α *from* $\bar{\alpha}$ *via a rigid motion.*

Proof. The equality combined with the first inequalities from the previous proof gives us

$$\int_0^L \left\langle \bar{T}(s), T_0\right\rangle ds \geq \int_0^L \langle T(s), T_0\rangle\, ds,$$

but since $\left\langle \bar{T}(s), T_0\right\rangle \leq \langle T(s), T_0\rangle$ we get equality in both places. Thus

$$\|\alpha(L)-\alpha(0)\| = \int_0^L \langle T(s), T_0\rangle\, ds.$$

and $\cos^{-1}\left\langle T_0, \bar{T}(s)\right\rangle = \cos^{-1}\langle T_0, T(s)\rangle$. Once again considering the lengths

$$L(\bar{T}\big|_{[s_0,s]}) = \int_{s_0}^{s} \bar{K}(s)\,ds = \cos^{-1}\left\langle T_0, \bar{T}(s)\right\rangle$$
$$L(T\big|_{[s_0,s]}) = \int_{s_0}^{s} K(s)\,ds \geq \cos^{-1}\langle T_0, T(s)\rangle,$$

we get that

$$\int_{s_0}^{s} K(s)\,ds \geq \int_{s_0}^{s} \bar{K}(s)\,ds$$

which implies $K \equiv \bar{K}$, and

$$\int_{s_0}^{s} K(s)\,ds = \cos^{-1}\langle T_0, T(s)\rangle.$$ □

Thus T lies on a great circle, so α is a plane curve. Now the result follows from the fundamental theorem of plane curves.

Lemma 1.26. *α is a simple curve.*

Proof. It is clear that $\bar{\alpha}$ is a simple curve. Let $x, y \in Im\,\alpha$ with two different parameter values. Let $\bar{x}, \bar{y} \in Im\,\bar{\alpha}$ be the points corresponding to the same parameter values. Let $\beta \subset \alpha$ be the arc with x, y as endpoints, and $\bar{\beta} \subset \bar{\alpha}$ be the arc with $\bar{x}, \bar{y}$ as endpoints. We know $\bar{x} \neq \bar{y}$ as $\bar{\alpha}$ is simple, and then $x \neq y$ since

$$\|x - y\| \geq \|\bar{x} - \bar{y}\|.$$ □

Corollary 1.27 (Schwartz' Theorem). *Let α be a closed $C^{1,1}$ arc in $\mathbb{R}^n$ with curvature $\leq \frac{1}{r}$ almost everywhere, such that the distance between the endpoints x, y is $\leq 2r$. We write L_α as the length of α. Let C be a circle of radius r which contains x and y. Let A_0, A_1 be arcs on the circle with x and y as endpoints, with length L_0, L_1. Assume $L_1 \geq L_0$. Then $L_\alpha \geq L_1$ or $L_\alpha \leq L_0$.*

Proof. Let θ be the angle in the centre of C between $\alpha(0)$ and $\alpha(L_\alpha)$.

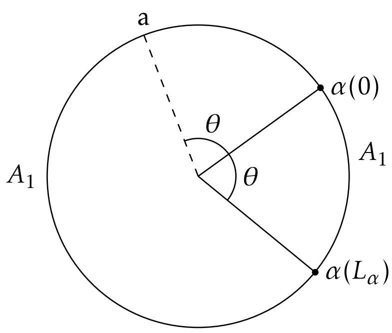

Let $\bar{\alpha}$ be an arc on C with $\bar{\alpha}(0) = \alpha(0)$ and length L_α. We observe that we are in the situation from Schur's theorem, hence we find that the distance between the endpoints of $\bar{\alpha}$ is less than the distance between $\alpha(0)$ and $\alpha(L_\alpha)$. Thus we get that $\bar{\alpha}(L_\alpha)$ lies in the arc between $\alpha(L_\alpha)$ and a. Hence it follows that if $L_\alpha < L_1$ then $L_\alpha \leq L_0$ and otherwise $L_\alpha \geq L_1$.

1.4 Reach and curvature

Intuitively it is clear that there is a connection between reach and curvature. In this section we delve further into this relationship.

Lemma 1.28. *Let $\mathcal{M} \subset \mathbb{R}^n$ be a smooth manifold of dimension $k \geq 2$. $\mathcal{M}$ has a Riemannian structure from the embedding in $\mathbb{R}^n$ with the usual Euclidean structure. Let $p \in \mathcal{M}$ and assume that the absolute value of all the normal curvatures of $\mathcal{M}$ at p is bounded by κ. Then all the sectional curvatures of $\mathcal{M}$ at p lie in the interval $[-\min(n-k)\kappa^2, \kappa^2]$.*

Proof. Let $v \in T_p\mathcal{M}$ be a unit vector, and let $\gamma : J \to \mathcal{M}$ be a smooth curve defined in an open neighbourhood of $0 \in \mathbb{R}$, such that $\gamma(0) = p$ and $\gamma'(0) = v$. For a unit normal vector $\mathbf{n}$, the $\mathbf{n}$-*normal curvature* is given by $\gamma''(0) \cdot n$.

The normal curvature is related to the sectional curvature through the second fundamental form of $\mathcal{M}$at p,

$$\mathrm{II} : T_p\mathcal{M} \times T_p\mathcal{M} \to N_p\mathcal{M}. \tag{1.4}$$

We need the following properties of the second fundamental form [5],

- II is symmetric and bilinear,
- the $\mathbf{n}$-normal curvature in p in the direction v is $\mathrm{II}(v,v)\cdot\mathbf{n}$,
- the sectional curvature K_P of $\mathcal{M}$at p with respect to the plane $P \subset T_p\mathcal{M}$ spanned by an orthonormal basis $\{v,w\}$ is equal to

$$\mathrm{II}(v,v)\cdot\mathrm{II}(w,w) - \|\mathrm{II}(v,w)\|^2.$$

Let $\mathbf{n}_0 \in N_p\mathcal{M}$ be a unit vector. Let $\kappa_v = \|\mathrm{II}(v,v)\|$, $\kappa_w = \|\mathrm{II}(w,w)\|$ and

$$\mathbf{n}_v = \begin{cases} \frac{1}{\kappa_v}\mathrm{II}(v,v), & \kappa_v \neq 0 \\ \mathbf{n}_0, & \text{otherwise} \end{cases},$$

$$\mathbf{n}_w = \begin{cases} \frac{1}{\kappa_v}\mathrm{II}(w,w), & \kappa_w \neq 0, \\ 0, & \text{otherwise} \end{cases}.$$

This implies that $\kappa_v = \mathrm{II}(v,v)\cdot\mathbf{n}_v$ and $\kappa_w = \mathrm{II}(w,w)\cdot\mathbf{n}_w$ are normal curvatures, so $\kappa_v, \kappa_w \leq \kappa$. Hence

$$\begin{aligned} K_P &= \mathrm{II}(v,v)\cdot\mathrm{II}(w,w) - \|\mathrm{II}(v,w)\|^2 \\ &\leq \mathrm{II}(v,v)\cdot\mathrm{II}(w,w) \\ &= \kappa_v\kappa_w\,\mathbf{n}_v\cdot\mathbf{n}_w \\ &\leq \kappa^2. \end{aligned}$$

If $\mathrm{II}(v,w) = 0$, then

$$K_P = \kappa_v\kappa_w\,\mathbf{n}_v\cdot\mathbf{n}_w \geq -\kappa^2.$$

Alternatively let $\lambda = \|\mathrm{II}(v,w)\|$, and $\mathbf{m} = \frac{1}{\lambda}\mathrm{II}(v,w)$. Let $\mathbf{p} = \frac{1}{\sqrt{2}}(v+w)$ and $\mathbf{q} = \frac{1}{\sqrt{2}}(v-w)$. Now $\mathbf{p}$ and $\mathbf{q}$ are unit vectors, and since II is symmetric and bilinear,

$$\mathrm{II}(v,w) = \frac{1}{2}\left(\mathrm{II}(p,p) - \mathrm{II}(q,q)\right).$$

So $\lambda = \mathrm{II}(v,w)\cdot m = \frac{1}{2}\left(\mathrm{II}(p,p)\cdot m - \mathrm{II}(q,q)\cdot m\right)$. However, $\mathrm{II}(p,p)\cdot m$ and $\mathrm{II}(q,q)\cdot m$ are normal curvatures, so $|\lambda| \leq \frac{1}{2}(\kappa+\kappa) = \kappa$. Hence

$$\begin{aligned} K_P &= \mathrm{II}(v,v)\cdot\mathrm{II}(w,w) - \|\mathrm{II}(v,w)\|^2 \\ &= \kappa_v\kappa_w\,\mathbf{n}_v\cdot\mathbf{n}_w - |\lambda|^2 \\ &\geq -\kappa^2 - \kappa^2 = -2\kappa^2. \end{aligned}$$

What is left is to consider the special situation of $n-k=1$. In that case $N_p\mathcal{M}$ has dimension 1, and the normal space is spanned by $\mathbf{n}_0$. Then the inner product between II and $\mathbf{n}_0$ is an quadratic form on $T_p\mathcal{M}$, so it can be orthogonally diagonalised. Hence it can be orthogonal diagonalised, which means we can find an orthonormal basis $\{a,b\}$ for $T_p\mathcal{M}$ such that $\mathrm{II}(a,b)\cdot\mathbf{n}_0$, which implies $\mathrm{II}(a,b) = 0$. □

Proposition 1.29. *Let $\mathcal{M}$ be a C^2 manifold with* reach $\mathcal{M} > r$. *Let $x, y \in \mathcal{M}$, and $\gamma : [0, L] \to \mathcal{M}$ a geodesic from x to y. Then the curvature of γ is less than $\frac{1}{r}$ everywhere.*

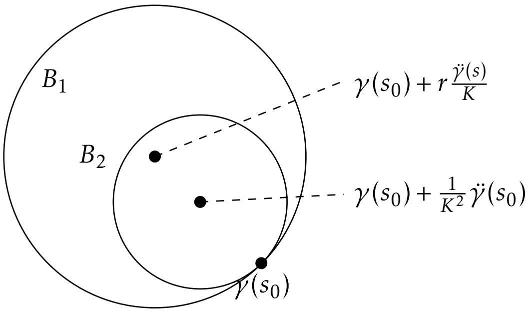

Proof. Suppose γ is arc-length parametrised. Let $s_0 \in [0, L]$. We want to show that

$$K = \|\ddot{\gamma}(s_0)\| \leq \frac{1}{r}.$$

We know that $\ddot{\gamma}(s_0) \perp T_{\gamma(s_0))}\mathcal{M}$, so $\ddot{\gamma}(s_0) \in N_{\gamma(s_0)}\mathcal{M}$, hence

$$P_{\mathcal{M}}\left(\gamma(s_0) + t\frac{\ddot{\gamma}(s_0)}{K}\right) = \gamma(s_0)$$

for all $t \leq r$.

Now, suppose that $K > \frac{1}{r}$, equivalently $r > \frac{1}{K}$. Let B_1 be the ball with radius r and centre $\gamma(s_0) + r\frac{\ddot{\gamma}(s)}{K}$ and B_2 the ball with radius $\frac{1}{K}$ and centre $\gamma(s_0) + \frac{1}{K}\ddot{\gamma}(s_0)$. By the definition of the curvature K there exists $\varepsilon > 0$ such that

$$\gamma(s_0 - \varepsilon, s_0 + \varepsilon) \setminus \gamma(s_0) \subset Int(B_1),$$

but since we assumed that M has reach greater than r we have a contradiction. □

In fact the relationship between reach and curvature goes both ways, but we need an additional requirement apart from the curvature.

Theorem 1.30. *Let $\mathcal{M} \subset \mathbb{R}^n$ be a closed smooth submanifold. Then, for any $r > 0$, the following are equivalent*

1. *$\mathcal{M}$ has reach $> r$,*
2. *All normal curvatures on $\mathcal{M}$ have absolute value $< \frac{1}{r}$, and for any $p \in M$, $B(p, 2r) \cap \mathcal{M}$ is connected.*

We need to ensure that $B(p, 2r) \cap \mathcal{M}$ is connected, as we risk that while the curvature of the manifold stays low, it slowly bends back and ends up near itself again.

Proof. *(i) ⇒ (ii)*: Let $r' \in (r, \tau_{\mathcal{M}})$, where $\tau_{\mathcal{M}} =$ reach $\mathcal{M}$. Let $p \in \mathcal{M}$ and let $\mathbf{t}$ be an arbitrary unit tangent vector to $\mathcal{M}$ at p, and let $\mathbf{n}$ be an arbitrary unit normal vector to $\mathcal{M}$ at p.

Now, let $\gamma : (-\varepsilon, \varepsilon) \to \mathcal{M}$ be a smooth curve with $\gamma(0) = p$ and $\gamma'(0) = \mathbf{t}$. Let $F : (-\varepsilon, \varepsilon) \to [0, \infty)$ be given by

$$F(t) = \|\gamma(t) - (p + r'\mathbf{n})\|^2, \quad t \in (-\varepsilon, \varepsilon).$$

Since $\mathcal{M}$has reach greater than r' we know that $P_{\mathcal{M}}(p + r'\mathbf{n}) = p$, so the sphere with radius r' and centre $p + r'\mathbf{n}$ will intersect $\mathcal{M}$ only in $\{p\}$. This means that F must have a minimum in 0. So

$$F'(0) = (\gamma(0) - (p + r'\mathbf{n})) \cdot \gamma'(0) = r'\mathbf{n} \cdot \mathbf{t} = 0$$
$$F''(0) = \|\gamma'(0)\|^2 - (\gamma(0) - (p + r'\mathbf{n})) \cdot \gamma''(0) = 1 - r'\mathbf{n} \cdot \gamma''(0) \geq 0.$$

Hence

$$\mathbf{n} \cdot \gamma''(0) \leq \frac{1}{r'} < \frac{1}{r}.$$

By changing $\mathbf{n}$ to $-\mathbf{n}$ in the calculations we get that $-\mathbf{n} \cdot \gamma''(0) < \frac{1}{r}$, so $\mathbf{n} \cdot \gamma''(0) > -\frac{1}{r}$.

Since p, $\mathbf{t}$ and $\mathbf{n}$ were chosen arbitrarily, it follows that the absolute values of all the normal curvatures on M must be less than $\frac{1}{r}$.

Let $p \in \mathcal{M}$ and let $p' \in \mathcal{M} \cap B(p, 2r)$. Then $\|p - p'\| = 2s \leq 2r$. Let L be the line segment between two points. Then it follows from Proposition 1.5 that $P_{\mathcal{M}}(L) \subset B(\frac{p+p'}{2}, 2)$, and thus

$$P_{\mathcal{M}}(L) \subset B(\frac{p + p'}{2}, s) \subset B(p, 2r).$$

Hence $\mathcal{M} \cap B(p, 2r)$ is connected, and as p was chosen arbitrarily it holds for all $p \in \mathcal{M}$.

(ii)$\Rightarrow$ *(i)*: First we recall for a Riemannian manifold with sectional curvature less than $K > 0$ everywhere, we have the following estimate of the injectivity radius by Klingenberg [4]:

$$inj(\mathcal{M}) \leq \min\left\{\frac{\pi}{\sqrt{K}}, \frac{1}{2}\{\text{the length of the shortest geodesic loop though } p\}\right\}.$$

Returning to our situation, $\mathcal{M} \subset \mathbb{R}^n$ inherits a Riemannian structure from $\mathbb{R}^n$. As all normal curvatures have absolute value less than $\frac{1}{r}$, it follows that all sectional curvatures are less than $\left(\frac{1}{r}\right)^2$. The condition on the normal curvatures also gives that the curvature of an arbitrary geodesic arc is less than $\frac{1}{r}$ everywhere.

Then it follows from Schurs' Theorem (1.23) that if a geodesic arc has length less than $2\pi r$, then the distance between the endpoints d is greater than the distance between the endpoints of a circular arc of the same length. Hence $d > 0$, and thus it follows that a geodesic loop on $\mathcal{M}$ must have length strictly greater than $2\pi r$. Now Klingenberg's estimate gives that the injectivity radius of $\mathcal{M}$ at each point p must be $> \pi r$.

Using Schurs' Theorem again, we see that if a geodesic in $\mathcal{M}$has length πr, then the distance between the endpoints must be greater than $2r$. Now let $p \in \mathcal{M}$ and let $S \subset \mathcal{M}$ be the geodesic sphere of radius πr centred at p. Then all points in S must have distance greater than $2r$ from p. So $S \cap B(p, 2r)$ is empty, and especially is $S \cap (B(p, 2r) \cap \mathcal{M})$ empty.

We now know that $\mathcal{M} \cap S$ is not connected, and that one of the connected components must be the open geodesic ball $\mathcal{O}$ centred at p. Since $B(p, 2r) \cap \mathcal{M}$ is connected and intersects $\mathcal{O}$ at p,

$$B(p, 2r) \cap \mathcal{M} \subset \mathcal{O}.$$

We are now ready to prove that $\mathcal{M}$has reach greater than r. Suppose otherwise. Then there must be a point $z \in \mathbb{R}^n$ with $d(z, \mathcal{M}) = s \leq r$, and two points $p_1, p_2 \in \mathcal{M}$ such that

$$\|z - p_1\| = \|z - p_2\| = s.$$

Then $\|p_1 - p_2\| \leq 2s$, so p_2 must lie in the open geodesic ball with radius πr and centre p_1. Thus there is a geodesic arc A from p_1 to p_2 of length $< \pi r$. As the normal curvature is strictly less than $\frac{1}{r}$ everywhere, we can find $r' > r$, for instance by letting $\frac{1}{r'}$ be the maximum curvature on A, such that A is contained in the r'-spindle with endpoints p_1 and p_2. This follows from Proposition 1.16. But since $s \leq r < r'$, the r'-spindle is contained in the interior of the s-ball with centre z, except for the endpoints. And thus the same holds for A. Since $A \subset \mathcal{M}$ this contradicts $s = d(z, \mathcal{M})$, so reach $\mathcal{M} > r$. □

1.5 Reach and shortest paths

Theorem 1.31. *[6] Let $A \subset \mathbb{R}^n$ be closed. Suppose there is a continuous rectifiable curve from x to y in A. Then there exists a continuous rectifiable curve of shortest length joining x and y.*

Theorem 1.32. *[1] Let $A \subset \mathbb{R}^n$ be closed. Then A has reach r if and only if for each pair $x, y \in A$ with $\|x - y\| < 2r$ there exists a continuous rectifiable curve joining x and y of length less or equal to $2r \arcsin\left(\frac{\|x-y\|}{2r}\right)$.*

Observe that $2r \arcsin\left(\frac{\|x-y\|}{2r}\right)$ is the length of the minor arc of a circle of radius r and endpoints x and y. The angle subtended at the centre of the circle is thus $2 \arcsin\left(\frac{\|x-y\|}{2r}\right)$. We call this angle the *r-associated angle* of the arc $[x, y]$.

Proof. Suppose that we have rectifiable curves with the stated bounds on the length. We want to show that A has reach r. Suppose on the contrary that there exists $z \in \mathbb{R}^n \setminus A$ and $x, y \in A$ with

$$d(z, A) = \|x - z\| = \|y - z\| = s < r.$$

Then $\|x - y\| \leq 2s < 2r$. Using the assumption and Theorem 1.31 together, we know that there exists a curve of shortest length $\gamma : [0,1] \to A$ from x to y, and the the length of γ must be less or equal to $2r \arcsin\left(\frac{\|x-y\|}{2r}\right)$.

We now claim that $Im\gamma \subset S(L, r)$, where L is the line segment from x to y. To show this, it is enough to prove that $Im\gamma \subset B$ for an arbitrary ball B of radius r which contains x and y.

Suppose that is not the case. Then there exists $t_0 < t_1$ in $[0,1]$, such that $\gamma(t_0), \gamma(t_1) \in B$ and

$$\gamma\left((t_0, t_1)\right) \subset \mathbb{R}^n \setminus B.$$

The radial projection onto the boundary of B is a smooth map $\rho : \mathbb{R}^n \setminus B \to \delta B$ which decreases length, so

$$\begin{aligned}
2r \arcsin\left(\frac{\|\gamma(t_0) - \gamma(t_1)\|}{2r}\right) \geq & \operatorname{length}\left(\gamma\big|_{[t_0,t_1]}\right) \\
> & \operatorname{length}\left(\rho \circ \gamma\big|_{[t_0,t_1]}\right) \\
\geq & \operatorname{length}(\text{the short geodesic on } \partial B \text{ from } \gamma(t_0) \text{ to } \gamma(t_1)) \\
= & 2r \arcsin\left(\frac{\|\gamma(t_0) - \gamma(t_1)\|}{2r}\right),
\end{aligned}$$

which is a contradiction, hence $Im\gamma \subset S(L, r)$.

However $S(L,r)\setminus\{x,y\}\subset B^{o}(z,s)$, so there must exists an element $\hat{x}\in Im\gamma\subset A$ with $\|\hat{x}-z\|<s$, which is also a contradiction. Hence A must have reach γ.

Now assume that A has reach r. Let $x,y\in A$ with $\|x-y\|<2r$. Let θ denote the r-associated angle for $[x,y]$. Let H be the hyperplane through $\frac{x+y}{2}$ which is orthogonal to $[x,y]$. As $P_A([x,y])$ is connected and contained in $S([x,y],r)$, there must be an element

$$m\in A\cap S([x,y],r)\cap H.$$

Let c be the points where the line starting in $\frac{x+y}{2}$ and passing through m intersects $\partial S([x,y],r)$. Then

$$\|x-m\|=y-m\leq\|x-c\|=\|y-c\|.$$

The line segments $[x,c]$ and $[y,c]$ both have r-associated angle $\frac{1}{2}\theta$, to the segments $[x,m]$ and $[y,m]$ has r-associated angle $\leq\frac{1}{2}\theta$.

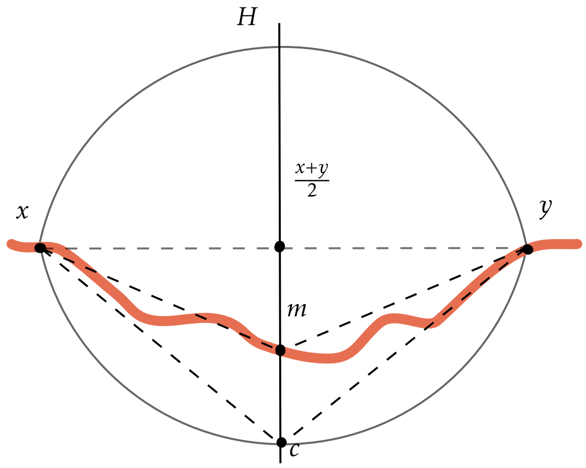

We can apply the same procedure to the segments of $P_A([x,y])$ from x to m and from m to y. Then we find new points and line segments which have r-associated angle less than $\frac{1}{r}\theta$. By continuing this process k times for an arbitrary $k\in\mathbb{N}$, we get a sequence S consisting of $2k+1$ points, $p_1,\dots,p_{2k+1}$ in $A\cap S([x,y],r)$, which starts in x and ends in y, such that each segment $[p_i,p_{i+1}]$ has r-associated angle $\leq\frac{1}{2^k}\theta$.

Consider the minor arc of a circle of radius r joining x and y. Since $[x,y]$ has r-associated angle θ, then the length of this arc is $r\theta$. Divide the arc into 2^k arcs of the same length $\frac{r}{2^k}\theta$. Let T be the collection of line segments between the endpoints of the arcs. Then each segment in T has r-associated angle $\frac{1}{2^k}\theta$.

Then

$$\begin{aligned}r\theta\geq&\text{sum of the lengths of the segments in T}\\ \geq&\text{sum of the lengths of the segments in S.}\end{aligned}$$

Next we want to apply P_A to the line segments from S. First consider, for an arbitrary line segment L of length $\ell<2r$ and en points in A, then

$$\text{length}\,(P_A(L))\leq\frac{r}{r-\frac{1}{2}\ell}\ell.$$

This follows from Theorem 1.3 (8). If L is a segment from S, then $\ell = 2r\sin(\frac{1}{2}\theta_L) \leq 2r\sin(\frac{1}{2^k}\theta)$, where θ_L is the r-associated angle for L. So it follows that the length of $P_A(\cup_{L\in S} P_A(L))$ is $\leq \frac{1}{1-2\sin(\frac{1}{2^k}} r\theta$. Letting $k \to \infty$, we see that

$$\inf\{\text{length}(\gamma) : \gamma \text{ is a curve joining} x \text{ and } y\} \leq r\theta. \qquad \square$$

From Theorem 1.31 it follows that this infimum is attained. Hence, there is a curve of length $\leq r\theta$ joining x and y. This completes the proof.

Corollary 1.33. *Suppose $A \subset \mathbb{R}^n$ with* $\operatorname{reach} A > r$. *Suppose C is a curve in A of shortest length between two points in A. Then C has the r-spindle property.*

Proof. Let $x, y \in C$ with $\|x-y\| \leq 2r$. Let B be a ball with radius r and x and y in the boundary of B. Let $\gamma : [0,1] \to C$ be a parametrisation of the curve between x and y. Suppose that $\gamma([0,1])$ is not included in B. Then there exists $t_1, t_2 \in [0,1]$ such that $\gamma(t_1), \gamma(t_2) \in \partial B$ and $\gamma(t_1, t_2) \subset B^c$.

As C is a curve of shortest length, $\gamma([t_1, t_2])$ must be the shortest path from $\gamma(t_1)$ to $\gamma(t_2)$. Hence by Theorem 1.32 $length(\gamma([t_1, t_2])) \leq 2r\arcsin\left(\frac{\|\gamma(t_1)-\gamma(t_2)\|}{2r}\right)$. However, this is also the length of the short curve of the boundary of B from $\gamma(t_1)$ to $\gamma(t_2)$, but this curve is also the radial projection of $\gamma([t_1, t_2])$ down to B. However, this is a contradiction, as the radial projection would strictly shorten the length of the curve. $\square$

Corollary 1.34. *Let $A \subset \mathbb{R}^n$ with* $\operatorname{reach} A = r$. *Suppose C is a path from x to y in A of shortest length. Then C has reach r.*

Proof. Combining Corollary 1.33 and Proposition 1.22 gives the result. $\square$

References

[1] Boissonnat, J.-D., A. Lieutier and M. Wintraecken (2019). The reach, metric distortion, geodesic convexity and the variation of tangent spaces. eng. *Journal of Applied and Computational Topology* 3(1), 29–58.

[2] Federer, H. (1959). Curvature Measures. eng. *Transactions of the American Mathematical Society* 93(3), 418–491.

[3] Federer, H. (1969). *Geometric measure theory.* eng. Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete ; 153. Berlin: Springer.

[4] Klingenberg, W. (1959). Contributions to Riemannian Geometry in the Large. eng. *Annals of mathematics* 69(3), 654–666.

[5] Lee, J.M (1997). *Riemannian manifolds : an introduction to curvature.* eng. Graduate texts in mathematics ; 176. New York: Springer.

[6] Myers, S.B. (1945). Arcs and Geodesics in Metric Spaces. eng. *Transactions of the American Mathematical Society* 57(2), 217–227.

[7] Rataj, J. (2007). *Three lectures on geometric measure theory*. URL: http://thiele.au.dk/fileadmin/www.thiele.au.dk/Events/2007/WS/Rataj.pdf, Last visited: 29/03/2022.

[8] Svane, H.M (2019). "Reconstructing R-regular Objects from Trinary Digital Images". PhD dissertation. Aarhus Universitet, Institut for Matematik.

PAPER A

Pointwise normal reach

Helene Hauschultz and Andrew du Plessis

A.1 Introduction

In modern machine learning many data sets are very high dimensional. The goal of manifold learning is to find lower dimensional representations of such high dimensional data. This is done by fitting a submanifold of lower dimension than the ambient space to the data. The data can then be represented by the projection onto the learned manifold. As these representations are used to further analyse the data, it is, from a theoretical perspective, important that the representations are unique. However, in practice manifold learning algorithms never consider that such a projection is actually unique.

In the influential paper *Testing the manifold hypothesis* by Fefferman et.al. [3] the authors tackle the problem of uniqueness by asking if the given data set fit a manifold of a certain reach. The reach of a subset is the maximal distance you can move from the subset while ensuring a unique nearest point. Aamari et.al [1] in their paper tackle the estimation of the reach on smooth closed manifolds. Reach is a concept widely used in geometric measure theory, and sets with positive reach behaves nicely, so the use of reach is also beneficial in the theory of manifold estimation [6, 8]. This paper originates from the idea of using reach in concrete manifold learning algorithms, to ensure that the data points have a unique representation. However, from this perspective, the global reach is too restrictive, and we found that a finer-grained analysis was needed.

For a given subset $\mathcal{M}$ of $\mathbb{R}^n$ we denote $\mathrm{Unp}(\mathcal{M})$ the set of points in $\mathbb{R}^n$ which has a unique nearest point in $\mathcal{M}$. To learn trustworthy representations we should require the data to lie in the set with unique nearest points to the learned manifold. However, for arbitrary subsets, or even smooth manifolds, the set $\mathrm{Unp}(\mathcal{M})$ is not easily computed. Introduced by Federer [2], $\mathrm{reach}(\mathcal{M})$ is the largest number such that any point at distance less than $\mathrm{reach}(\mathcal{M})$ to $\mathcal{M}$ lies inside $\mathrm{Unp}(\mathcal{M})$. It is clear that enforcing that the data stays within the reach distance of the manifold ensures unique representation. However, many points with distance larger than the reach still have a unique nearest point. Thus we develop a new local reach value, called the *pointwise normal reach*. The pointwise normal reach gives a bound on how far we can move in a normal direction from a given point while ensuring unique projection onto the subset.

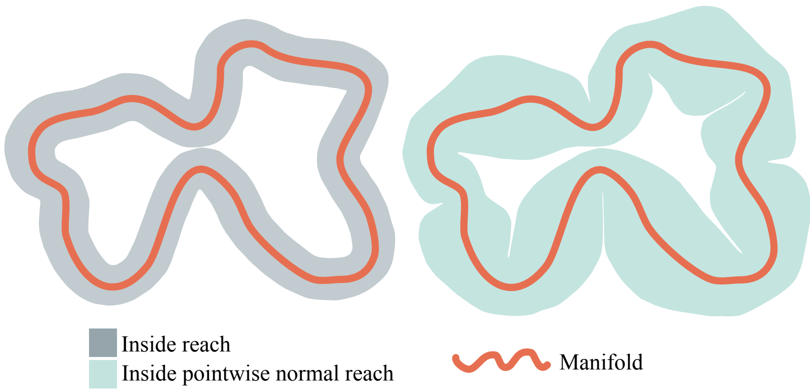

As our efforts come from applying the concept to actual algorithms, the pointwise normal reach can be directly calculated as

$$r_N(x) = \inf_{y \in \mathcal{M}} \frac{\|y-x\|^2}{2d(y-x, T_x\mathcal{M})},$$

where $T_x\mathcal{M}$ denotes the general tangent cone as defined by Federer [2]. If $\mathcal{M}$ is a manifold, $T_x\mathcal{M}$ is the standard tangent space. We will prove that the pointwise normal reach is continuous on smooth submanifolds which are closed subsets of $\mathbb{R}^n$. The need for continuity also comes from concrete applications. As learning algorithms often embed noise, we might end up calculating the pointwise normal reach in a point close to, but not exactly the nearest point to a given data point. By ensuring continuity, we do ensure that we are at least close to the correct value.

In practice, it is not feasible to obtain the general tangent cones of the learned manifolds. Instead we would like to utilise standard analytical tools, such as the Jacobians. As done by Shao, Kumar & Fletcher [9], we assume that our learned manifold is the image of an immersion $f : \mathbb{R}^m \to \mathbb{R}^n$, $m < n$. That is, f is a smooth map with injective differential at each point. We then develop a new bound for ensuring uniqueness, which uses the Jacobian matrix to estimate the distance to the tangent space.

A.2 Background

In this paper we consider $\mathbb{R}^n$ with the standard Euclidean inner product $\langle \cdot, \cdot \rangle$ and norm $\|\cdot\|$. We denote by $B(z, r)$ the open ball with centre $z \in \mathbb{R}^n$ and radius r. For any given subset $\mathcal{A} \subset \mathbb{R}^n$ we define the distance from a point $z \in \mathbb{R}^n$ to $\mathcal{A}$ as $d(z, A) = \inf_{x \in \mathcal{A}} \|x - z\|$. For a vector or vector space V we denote by π_V orthogonal projection onto V.

As done by Federer [2], for an arbitrary subset $\mathcal{A} \subset \mathbb{R}^n$ we consider the set $\mathrm{Unp}(\mathcal{A}) = \{z \in \mathbb{R}^n : z \text{ has a unique nearest point in } \mathcal{A}\}$. We define the function $P_\mathcal{A}$: $\mathrm{Unp}(\mathcal{A}) \to \mathcal{A}$ as the function sending z in $\mathrm{Unp}(\mathcal{A})$ to its unique nearest point in $\mathcal{A}$. For a point x in $\mathcal{A}$ the *local reach* of $\mathcal{A}$ at x is $\mathrm{reach}(\mathcal{A}, x) = \sup\{r > 0 : B(x,r) \subset \mathrm{Unp}(\mathcal{A})\}$, and the *global reach* of $\mathcal{A}$ is $\mathrm{reach}(\mathcal{A}) = \inf_{x\in\mathcal{A}} \mathrm{reach}(\mathcal{A}, x)$. It is easy to see that if $\mathcal{A}$ is not closed then $\mathrm{reach}(\mathcal{A}) = 0$.

For a subset $\mathcal{A} \subset \mathbb{R}^n$ and a point $x \in \mathcal{A}$, we say that $\boldsymbol{t} \subset \mathbb{R}^n$ is a tangent vector if for all $\varepsilon > 0$ there is a point $y \in \mathcal{A}$ such that

$$0 < \|y - x\| < \varepsilon \quad \text{and} \quad \|\frac{y-x}{\|y-x\|} - \frac{\boldsymbol{t}}{\|\boldsymbol{t}\|}\| < \varepsilon.$$

We denote $T_x\mathcal{A}$ the set of all tangent vectors of $\mathcal{A}$ at x. We denote the set of all normal vectors as the dual of the tangent cone $T_x\mathcal{A}$, $N_x\mathcal{A} = \{\boldsymbol{n} \in \mathbb{R}^n : \langle \boldsymbol{n}, \boldsymbol{t}\rangle < 0 \text{ for all } \boldsymbol{t} \in T_x\mathcal{A}\}$. Note that if $\mathcal{A}$ is a manifold, the tangent and normal cones correspond to the standard tangent and vector spaces.

Next we recall some useful results from Federer.

Lemma A.1 ([2] Theorem 4.8). *Let $\mathcal{A}$ be a closed subset of $\mathbb{R}^n$ and let $x \in \mathcal{A}$.*

1. *Let $P_\mathcal{A} : \mathrm{Unp}(\mathcal{A}) \to \mathcal{A}$ denote the projection onto the unique nearest point and let*

$$P = \{v \in \mathbb{R}^n : P_\mathcal{A}(x+v) = x\}, \quad Q = \{v : d(x+v, \mathcal{A}) = \|v\|\}.$$

   Then P and Q are convex and $P \subset Q \subset N_x\mathcal{A}$.

2. *If $z \in \mathrm{Unp}(\mathcal{A})$, $x = P_\mathcal{A}(z)$, $\mathrm{reach}(\mathcal{A}, x) > 0$ and $y \in \mathcal{A}$, then*

$$\langle z - x, y - x\rangle \le \frac{\|y-x\|^2\|x-z\|}{2\,\mathrm{reach}(\mathcal{A}, x)}.$$

From Federer we also get an explicit computation of the global reach.

Theorem A.2 ([2] Theorem 4.18). *For a closed subset $\mathcal{A}$ in $\mathbb{R}^n$,*

$$\mathrm{reach}(\mathcal{A})^{-1} = \sup_{y\in\mathcal{A}, y\neq x} \frac{2d(y-x, T_x\mathcal{A})}{\|y-x\|^2}.$$

Here we say if $\mathrm{reach}(\mathcal{A})^{-1} = 0$ *then* $\mathrm{reach}(\mathcal{A}) = \infty$ *and vice versa.*

A.3 Introducing pointwise normal reach

The motivation behind this work is to ensure unique representations in dimensionality reducing algorithms. Unique representations can be ensured by requiring data points to be within reach, by which we mean the distance from the point to the fitted manifold is less than the reach of the manifold. Thus the global reach provides a bound which ensures uniqueness. However, this bound is too rough, and many points which do have a unique nearest point on the manifold are discarded. The local reach defined by Federer does give a finer-grained analysis of $\mathrm{Unp}(\mathcal{M})$, but $\mathrm{reach}(\mathcal{M}, x)$ is generally not computable except in very simple cases. Thus the goal is to develop tools to give a computable finer-grained analysis of the set $\mathrm{Unp}(\mathcal{M})$.

Definition A.3. *Let $\mathcal{M} \subset \mathbb{R}^n$ be a closed subset. Let x be a point in $\mathcal{M}$. Then, for $y \in \mathcal{M}$ with $y - x \notin T_x\mathcal{M}$ we define*

$$R(x,y) = \frac{\|y-x\|^2}{2d(y-x, T_x\mathcal{M})}.$$

If $y - x \in T_x\mathcal{M}$ we say $R(x,y) = \infty$. We define the pointwise normal reach *to be*

$$r_N(x) = \inf_{y \in \mathcal{M}} R(x,y). \tag{A.1}$$

Thus pointwise normal reach is the pointwise minimisation of the computation from Theorem A.2. We can interpret $R(x,y)$ as the radius of the sphere tangent to $\mathcal{M}$ at x and passing through y.

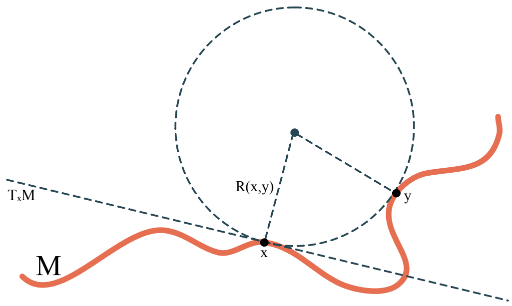

Figure 2: For two points $x, y \in \mathcal{M}$ the fraction $R(x,y)$ is the radius of ball tangent to $\mathcal{M}$ at x and passing through y.

In the following theorem we shall see the motivation behind the name the *pointwise normal reach*. The theorem shows that $r_N(x)$ gives a bound on how far we can move away from x in a normal direction, while staying inside $\text{Unp}(\mathcal{M})$. In other words, if for some $z \in \mathbb{R}^n$ the distance from z to $\mathcal{M}$ satisfies $d(z,\mathcal{M}) = \|z - x\|$ for some $x \in \mathcal{M}$, with $\|z - x\| < r_N(x)$, then $z \in \text{Unp}(\mathcal{M})$.

Theorem A.4. *Let $\mathcal{M} \subset \mathbb{R}^n$ be closed. For all $x \in \mathcal{M}$*

$$B(x, r_N(x)) \cap (N_x\mathcal{M} + x) \subset \text{Unp}(\mathcal{M}). \tag{A.2}$$

In order to prove the theorem we first prove the following technical lemma.

Lemma A.5. *Let $\mathcal{M} \subset \mathbb{R}^n$ be closed and let $x \in \mathcal{M}$. Then the following holds true:*

1. *For $v \in N_x\mathcal{M}$ and $u \in \mathbb{R}^n$ with $\langle v, u \rangle \geq 0$, then $\|\pi_v(u)\| \leq d(u, T_x\mathcal{M})$, where π denotes the orthogonal projection.*

2. *Suppose $z \in \mathbb{R}^n$ with $z - x \in N_x\mathcal{M}$, and suppose there exists $y \in \mathcal{M}$, $y \neq x$, such that $\|z - y\| \leq \|z - x\|$, then $R(x,y) \leq \|z - x\|$.*

Proof. To prove 1. we first prove that $d(u, \mathrm{dual}(v)) = \|\pi_v(u)\|$. Note that we can write $u = u_v + u_{v^\perp}$, for $u_v = \pi_v(u)$ and $u_{v^\perp} \in v^\perp$.

$$\begin{aligned} d(u, \mathrm{dual}(v)) = &\inf_{w \in dual(v)} \|u - v\| = \inf_{w \in \mathrm{dual}(v)} \|u_v + u_{v^\perp} - w\| \\ =&\|u_v\| + \|u_{v^\perp} - w\| - 2\langle u_v, w\rangle. \end{aligned}$$

As $\langle u_v, w\rangle \leq 0$, it follows that the infimum is achieved when $w = u_{v^\perp}$. By the definition of the dual, if follows that

$$\mathrm{dual}(v) \supset \mathrm{dual}(N_x\mathcal{M}) \supset T_x\mathcal{M}.$$

Hence, $d(u, T_x\mathcal{M}) = \inf_{w \in T_x\mathcal{M}} \|u - w\| \geq \inf_{w \in \mathrm{dual}(v)} \|u - w\| = \|\pi_v(u)\|$.

To prove 2. note that, as $\|z - y\| \leq \|z - x\|$, it follows that $\langle z - x, y - x\rangle \geq 0$. Thus we can consider the triangle with vertices x, y, z, where the angle θ at the vertex x is acute. Combining the inequality with the law of cosine gives $\|z - x\|^2 \geq \|z - y\|^2 = \|z - x\|^2 + \|y - x\|^2 - 2\|z - x\|\|y - x\|\cos\theta$, which is equivalent to

$$\|z - x\| \geq \frac{\|y - x\|^2}{2\|y - x\|\cos\theta} = \frac{\|y - x\|^2}{2\|\pi_{z-x}(y - x)\|}.$$

Following 1. we have that $\|\pi_{z-x}(y - x)\| \leq d(y - x, T_x\mathcal{M})$, hence

$$\frac{\|y - x\|^2}{2d(y - x, T_x\mathcal{M})} \leq \frac{\|y - x\|^2}{2\|\pi_{z-x}(y - x)\|}.$$

Combining the two inequalities, we get that $R(x, y) \leq \|z - x\|$. □

Proof (Proof of theorem A.4). Suppose for the sake of contradiction that there exists $z \in (B(x, r_N(x)) \cap (N_x\mathcal{M} + x))$ with $z \notin \mathrm{Unp}(M)$. That is, there exists $y_1, y_2 \in \mathcal{M}$ such that $d(z, \mathcal{M}) = \|y_1 - z\| = \|y_2 - z\|$. In particular, there must be at least one $y \in \mathcal{M}$ with $y \neq x$ such that $\|y - z\| \leq \|x - z\| < r_N(x)$. Applying Lemma A.4 then gives $R(x, y) \leq \|z - x\| < r_N(x)$, which gives a contradiction. □

Relation to local reach

In the following section we discuss how the pointwise normal reach relates to the local reach defined by Federer.

Proposition A.6. *For any* $x \in \mathcal{M}$, $\mathrm{reach}(\mathcal{M}, x) \leq r_N(x)$.

Proof. If $\mathrm{reach}(\mathcal{M}, x) = 0$, the result follows directly. Hence, suppose $\mathrm{reach}(\mathcal{M}, x) > 0$. Let $y \in \mathcal{M}$. If $y - x \in T_x\mathcal{M}$, then $R(x, y) = \infty$, otherwise we can find a unit normal vector $w \in N_x\mathcal{M}$ such that $\langle w, y - x\rangle \geq 0$. In that case

$$\langle w, y - x\rangle = \|w\|\|y - x\|\cos\theta = \|\pi_w(y - x)\|,$$

where θ is the acute angle between w and $y - x$ so $\cos\theta = \|w\|\frac{\|\pi_w(y-x)\|}{\|y-x\|}$.

Applying Lemma A.1 (2) then gives us

$$\mathrm{reach}(\mathcal{M}, x) \leq \frac{\|y - x\|^2\|w\|}{2\langle w, y - x\rangle} \leq \frac{\|y - x\|^2}{2\|\pi_w(y - x)\|} \leq R(x, y). \qquad \square$$

Corollary A.7. *For any* $x \in \mathcal{M}$, $\operatorname{reach}(\mathcal{M},x) \leq \inf_{y\in\mathcal{M}} r_N(y) + \|x-y\|$.

Proof. Let $x \in \mathcal{M}$ and consider an arbitrary $y \in \mathcal{M}$. There is a sequence $\{z_k\} \subset \operatorname{Unp}(\mathcal{M})^c$ such that the distance to y converges down to $\operatorname{reach}(\mathcal{M},y)$ as k goes to infinity. For any $k \in \mathbb{N}$ we have that $\|x-z_k\| \leq \|y-z_k\| + \|x-y\|$. Thus $\|x-z_k\| \to \operatorname{reach}(\mathcal{M},y) + \|x-y\|$ as $k \to \infty$. Thus it follows that

$$\operatorname{reach}(\mathcal{M},x) \leq \inf_{y\in\mathcal{M}} \operatorname{reach}(\mathcal{M},y) + \|x-y\|.$$

Combining this result with Proposition A.6 gives the result. □

Proposition A.8. *For any* $x \in \mathcal{M}$,

$$\operatorname{reach}(\mathcal{M},x) \geq \inf_{y\in B(x,2r_N(x))\cap\mathcal{M}} r_N(y).$$

Proof. By the definition $\operatorname{reach}(\mathcal{M},x) = \inf_{z\in\operatorname{Unp}(\mathcal{M})^c}\|z-x\|$, and from the proceding proposition we know that $\operatorname{reach}(\mathcal{M},x) \leq r_N(x)$. So suppose there exists $z \in B(x, r_N(x))$ with $z \notin \operatorname{Unp}(\mathcal{M})$. Then there exists $y_1, y_2 \in \mathcal{M}$ such that $d(z,\mathcal{M}) = \|y_1 - z\| = \|y_2 - z\| \leq \|x-z\|$. It follows that $\|y_i - x\| \leq 2r_N(x)$ for $i = 1,2$, and that $\|y_i - z\| \geq r_N(y_i)$ for $i = 1,2$. Hence it follows that $\|x-z\| \geq r_N(y_i) \geq \inf_{y\in B(x,2r_N(x))\cap\mathcal{M}} r_N(y)$, which proves the claim. □

The lower bound on $\operatorname{reach}(\mathcal{M},x)$ is, in fact, quite weak. Points near the edge of the $2r_N(x)$–radius with a small normal reach r_N will have a large impact on the bound. However, we can strengthen the bound in the following way:

Corollary A.9.

$$\operatorname{reach}(\mathcal{M},x) \geq \inf_{y\in B(x,2r_N(x))\cap\mathcal{M}} \max\left\{r_N(y), \frac{1}{2}\|y-x\|\right\}$$

Proof. Let $z \in \mathbb{R}^n$ and suppose $\|y-z\| \leq \|x-z\|$. Using the triangle inequality $\|y-x\| \leq \|y-z\| + \|z-x\| \leq 2\|x-z\|$, which implies $\|x-z\| \geq \frac{1}{2}\|y-x\|$. □

Restricting to convex tangent cones

We now restrict the study of the pointwise normal reach to the case where the tangent cones are convex. Indeed, we see that the normal reach is not necessarily informative otherwise.

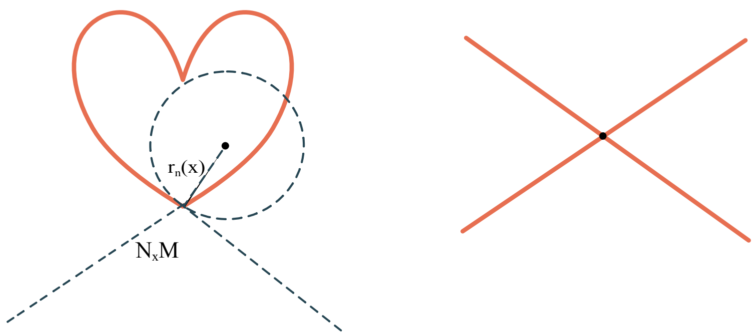

In the case of the cross, the pointwise normal reach in the centre will be ∞. However, this is not very informative as none of the points in $\mathbb{R}^n$ will project onto the centre, and the normal cone is empty. On the other hand, in the case of the heart, the pointwise normal reach in the lower vertex, will be bounded. However, we can move infinitely far away in any normal direction while staying inside $\mathrm{Unp}(\mathcal{M})$.

In the following lemma we prove that if the tangent cones are convex, the pointwise normal reach does indeed tell us the maximal distance we can move in any normal direction from $x \in \mathcal{M}$ while still ensuring the unique nearest point is x.

Lemma A.10. *Let $\mathcal{M}$ denote a closed subset of $\mathbb{R}^n$ and let $x \in \mathcal{M}$ with $T_x\mathcal{M}$ convex. Let $\tau_v = \sup\{t > 0 : P_{\mathcal{M}}(x + tv) = x\}$ for each unit vector $v \in N_x\mathcal{M}$, then*

$$r_N(x) = \inf_{v \in N_x\mathcal{M}, \|v\|=1} \tau_v.$$

Proof. Let $\tau = \inf_{v \in N_x\mathcal{M}} \tau_v$. Suppose there exists a unit vector $v \in N_x\mathcal{M}$ and $t < r_N(x)$ such that $\xi(x + tv) \neq x$. Then there exists $y \in \mathcal{M}$ with $\|x + tv - y\| \leq \|x + tv - x\| = t$. However, then Lemma A.5 gives $R(x, y) \leq t < r_N(x)$, which is a contradiction. Hence, $r_N(x) \leq \tau$.

Next let $t < \tau$. Let v be a unit vector in $N_x\mathcal{M}$, such that $\xi(x + tv) = x$. Hence for all $y \in \mathcal{M}$ $\|x + tv - y\| > t$. Using the same idea as in Lemma A.5, but with the opposite inequality, we get that

$$t^2 < \|x + tv - y\|^2 = \|x - y\|^2 + t^2 - 2t\|\pi_v(y - x)\|,$$

which is equivalent to

$$t < \frac{\|x - y\|}{2\|\pi_v(y - x)\|}$$

for all $y \in \mathcal{M}$ with $\langle v, y - x\rangle \geq 0$. Especially, for an arbitrary $y \in \mathcal{M} \setminus T_x\mathcal{M}$, $\frac{P_{N_x\mathcal{M}(y-x)}}{\|P_{N_x\mathcal{M}}(y-x)\|}$ will be a unit normal vector. Hence

$$t < \frac{\|y - x\|^2}{2\|P_{N_x\mathcal{M}}(y - x)\|} = R(x, y).$$

Hence $\tau \leq r_N(x)$, which proves the lemma. □

Normal reach in smooth manifolds

As we work towards proving that the pointwise normal reach is continuous, we tighten the assumptions and assume that $\mathcal{M}$ is a smooth manifold, in addition to a closed subset. In the following example we show that the C^1 is not sufficient to ensure continuity.

Example A.11. We consider a function $h : \mathbb{R}^2 \to \mathbb{R}^3$, and let $\mathcal{M} = h(\mathbb{R}^2)$, where h is given by

$$h(x, y) = \begin{cases} 2\sqrt{r^2 - \left(\frac{y}{2}\right)^2} - \sqrt{r^2 - x^2} - r & \text{if } 0 \leq |x| \leq \frac{y}{2} \\ \sqrt{r^2 - (x - y)^2} - r & \text{if } \frac{y}{2} \leq x \leq y \\ \sqrt{r^2 - (x + y)^2} - r & \text{if } -y \leq x \leq -\frac{y}{2} \\ 0 & \text{otherwise} \end{cases}$$

In Appendix A.A we show that the pointwise normal reach is in fact not continuous

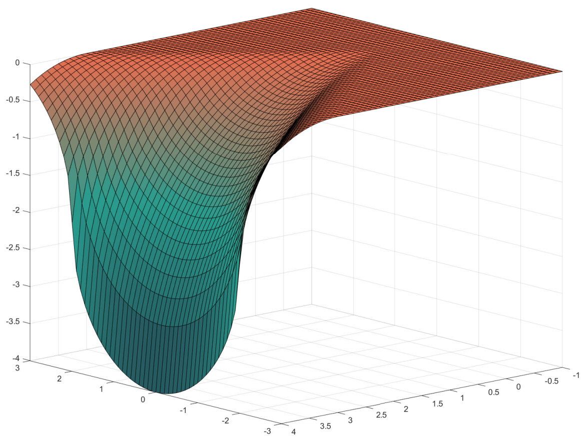

Figure 3: The graph of the function h from Example A.11.

in the point $(0,0,0)$. As $\mathcal{M}$ is $C^{1,1}$ we see that it is not sufficient to ensure the continuity. This is perhaps natural as the pointwise normal reach is closely related to the curvature of a manifold, which is a C^2 property.

Proposition A.12. *Suppose $\mathcal{M} \subset \mathbb{R}^n$ is closed and is a smooth submanifold. Then $r_N(x) > 0$ for all $x \in \mathcal{M}$.*

Proof. This follows from the fact that around any point in $\mathcal{M}$ we can find a neighbourhood in the normal bundle, which can be embedded into $\mathbb{R}^n$ ([4], Theorem 5.1). □

Suppose $\mathcal{M} \subset \mathbb{R}^n$ is a smooth submanifold and let $x \in \mathcal{M}$. Let $\boldsymbol{t}$ be a unit tangent vector for $\mathcal{M}$ at x, and let $\boldsymbol{n}$ be a unit normal vector for $\mathcal{M}$ at x. Suppose $\gamma : I \to \mathcal{M}$ is a smooth curve from an open neighbourhood I of 0 in $\mathbb{R}^n$, such that $\dot{\gamma}(0) = \boldsymbol{t}$ and $\gamma(0) = x$. Then the $\boldsymbol{n}$-normal curvature at x in the direction $\boldsymbol{t}$ is given by $\ddot{\gamma}(0) \cdot \boldsymbol{n}$. This is equivalent to $\mathrm{I\!I}(\boldsymbol{t},\boldsymbol{t}) \cdot \boldsymbol{n}$, where $\mathrm{I\!I}(\cdot,\cdot)$ denotes the second fundamental form. This implies that the normal curvature is continuous.

Proposition A.13. *Let $\mathcal{M} \subset \mathbb{R}^n$ be a smooth submanifold and let $x \in \mathcal{M}$. Suppose for some $r > 0$ that $r_N(x) > r$ then all normal curvatures at x in $\mathcal{M}$ have absolute value less than $\frac{1}{r}$.*

Proof. Let $r' \in (r, r_N(x))$. Let $\boldsymbol{t}$ be an arbitrary unit tangent vector to $\mathcal{M}$ at x, and let $\boldsymbol{n}$ be an arbitrary unit normal vector to $\mathcal{M}$ at x. Now, let $\gamma : (-\varepsilon, \varepsilon) \to \mathcal{M}$ be a smooth curve with $\gamma(0) = x$ and $\dot{\gamma}(0) = \boldsymbol{t}$. We then define a function $F : (-\varepsilon, \varepsilon) \to [0, \infty)$ to be given by

$$F(t) = \|\gamma(t) - (x + r'\boldsymbol{n})\|^2, \quad t \in (-\varepsilon, \varepsilon).$$

Since $r_N(x) > r'$, we know that $P_{\mathcal{M}}(x + r'\boldsymbol{n}) = x$, so the sphere with centre $x + r'\boldsymbol{n}$ and radius r' will intersect $\mathcal{M}$ only in x. This implies that F has a minimum in 0. Thus we have that

$$F'(0) = (\gamma(0) - (x + r'\boldsymbol{n})) \cdot \dot{\gamma}(0) = r'\boldsymbol{n} \cdot \boldsymbol{t} = 0,$$
$$F''(0) = \|\dot{\gamma}(0)\|^2 - (\gamma(0) - (x + r'\boldsymbol{n})) \cdot \ddot{\gamma}(0) = 1 - r'\boldsymbol{n} \cdot \ddot{\gamma}(0) \geq 0.$$

Thus $\boldsymbol{n} \cdot \ddot{\gamma}(0) \leq \frac{1}{r'} < \frac{1}{r}$. By changing $\boldsymbol{n}$ to $-\boldsymbol{n}$ in the calculation we get that $-\boldsymbol{n} \cdot \ddot{\gamma}(0) < \frac{1}{r}$, so $\boldsymbol{n} \cdot \ddot{\gamma}(0) > -\frac{1}{r}$. □

Next we need to recall the definition of spindles:

Definition A.14. *Let L be a closed line-segment in $\mathbb{R}^n$ of length $|L| \leq 2r$. Then the r-spindle $S(L,r)$ generated by L is the intersection of all r-balls whose boundaries each contain both endpoints of L.*

In fact we have that the r-spindle $S(L,r)$ is the intersection of all closed balls of radius at most r which contains L. Suppose x and y are points in $\mathcal{A}$, then the spindle with endpoints x and y, is the spindle generated by the line between these two points.

Before moving on we recall the following result which follows from Corollary A.57 in [10]:

Lemma A.15. *Suppose $x, y \in \mathcal{M}$ are points on a smooth manifold with $\|x - y\| \leq 2r$ and suppose there is a curve α between them with curvature $\leq \frac{1}{r}$ everywhere and length $\leq r\pi$. Then α is contained in the r-spindle with endpoints x and y.*

Lemma A.16. *Suppose $\mathcal{M}$ is a smooth submanifold and closed subset of $\mathbb{R}^n$. Then*

$$B(x, r_N(x)) \cap N_x\mathcal{M} \subset \mathring{\mathrm{Unp}}(\mathcal{M}).$$

Proof. Let $x \in \mathcal{M}$ with $r_N(x) > 0$. Suppose there exists $z \in B(x, r_N(x)) \cap (N_x\mathcal{M} + x)$ with $z \notin \mathring{\mathrm{Unp}}(\mathcal{M})$. Then there must exists a sequence $z_k \to z$ with $z_k \notin \mathrm{Unp}(\mathcal{M})$ for all $k \in \mathbb{N}$. For each k, z_k has two distinct associated points $p_k, q_k \in \mathcal{M}$, $p_k \neq q_k$, and

$$d(z_k, \mathcal{M}) = \|z_k - p_k\| = \|z_k - q_k\|.$$

As these sequences are bounded we can choose appropriate subsequences such that $p_k \to p$ and $q_k \to q$ for $k \to \infty$. We now claim that $p = q = x$. As $\mathcal{M}$ is closed, $p, q \in \mathcal{M}$. Suppose $p \neq x$. First notice that $\|z_k - p_k\| \leq \|z_k - x\| \leq \|z_k - z\| + \|z - x\|$.

$$\begin{aligned} \|z - p\| &\leq \|z - z_k\| + \|z_k - p\| \\ &\leq \|z - z_k\| + \|z_k - p_k\| + \|p_k - p\| \\ &\leq \|z - x\| + 2\|z_k - z\| + \|p_k - p\|. \end{aligned}$$

As $\|z_k - z\| \to 0$ and $\|p_k - p\| \to 0$ for $k \to \infty$, we have that $\|z - p\| \leq \|z - x\|$, but then we have a contradiction.

Now let $r = \|z - x\|$ and $r_N(x) = r + a$. Then $\frac{1}{\tau} \geq r + a$, for all normal curvatures τ for $\mathcal{M}$ at x. As the normal curvature is continuous, we can find a radius $s > 0$ such that $\frac{1}{\tau_y} \geq r + \frac{1}{2}a$, for all normal curvatures in y with $\|y - x\| < s$. In particular, we can choose $s \in (0, r + \frac{1}{2}a)$. For sufficiently large $k \in \mathbb{N}$ we have that

$$p_k, q_k \in \mathcal{M} \cap \bar{B}(x, s).$$

Then there is a geodesic from p_k to q_k, and due to the bound on the normal curvature, there is a bound on the curvature of the geodesic, thus it lies in the $(r+\frac{1}{2}a)$-spindle with endpoints p_k, q_k. However, for sufficiently large k, $d(z_k, \mathcal{M}) < (r+\frac{1}{2}a)$, hence the spindle sans the endpoints lies in the open ball $B(z_k, d(z_k, \mathcal{M}))$, which is a contradiction. □

Corollary A.17. *Suppose $\mathcal{M}$ is a smooth manifold, then*

$$r_N(x) = d(x, (N_x\mathcal{M} + x) \cap \overline{(\mathrm{Unp}(\mathcal{M}))^c}).$$

Proof. Lemma A.16 tells us that $r_N(x) \le d(x, (N_x\mathcal{M} + x) \cap \overline{(\mathrm{Unp}(\mathcal{M}))^c})$. Using lemma A.10 it is sufficient to show that

$$d(x, (N_x\mathcal{M} + x) \cap \overline{(\mathrm{Unp}(\mathcal{M}))c}) \le \inf_{v \in N_x\mathcal{M}} T_v.$$

For any $\varepsilon > 0$, there is a $v \in N_x\mathcal{M}$ such that $T_v < T + \varepsilon$. From Federer [2] (see Global reach, Theorem 1.3 6.) we know that $x + T_v v \notin \mathring{\mathrm{Unp}}(\mathcal{M})$, hence $x + T_v v \in \overline{\mathrm{Unp}(\mathcal{M})^c}$. Hence the claim is proved. □

Continuity of pointwise normal reach

We are now ready to show that the pointwise normal reach is continuous on smooth manifolds. The continuity of the pointwise normal reach implies that the value in y is close to the value in x if y is close to x. This is beneficial from a practical perspective, as we expect our model to embed some noise. That is, the learned representation of a data point $z \in \mathbb{R}^n$ might not be the actual nearest point on the manifold. We do, however, expect it to be close to the nearest point, and thus continuity implies that the calculated pointwise normal reach is close to the actual normal reach in the nearest point to z.

Theorem A.18. *Suppose $\mathcal{M}$ is a closed subset and smooth submanifold in $\mathbb{R}^n$ of dimension m, then $r_N : \mathcal{M} \to [0, \infty)$ is continuous.*

In order to prove the theorem we need the following result.

Lemma A.19. *Let the situation be as in Theorem A.18. For any $\varepsilon > 0$ there exists $z \in N_x\mathcal{M} + x$ such that $z \notin \mathrm{Unp}(\mathcal{M})$ and $\|z - x\| < r_N(x) + \varepsilon$.*

Proof (Rough sketch of proof). We know that there is a sequence $\{z_k\} \subset \mathrm{Unp}(\mathcal{M})^c$ such that $\lim_{k\to\infty}\|x - z_k\| = r_N(x)$ and $\lim_{k\to\infty} z_k = z \in N_x\mathcal{M} + x$. For each k we can find $y_k^1, y_k^2 \in \mathcal{M}$ such that

$$d(z_k, \mathcal{M}) = \|y_k^1 - z_k\| = \|y_k^2 - z_k\|.$$

As all these sequences are bounded, we can find appropriate convergent subsequences such that $\lim_{k\to\infty} y_k^1 = y^2$ and $\lim_{k\to\infty} y_k^2 = y^2$. If either y^1 or y^2 is not equal to x, then we are done. Otherwise $x = y^1 = y^2$. We can prove that $r_N(x)$ is a normal curvature of $\mathcal{M}$ in x, which proves the claim. □

Proof (Proof of Theorem). Let $x \in \mathcal{M}$ and suppose $\{y_k\}_{k=1}^{\infty}$ is a sequence in $\mathcal{M}$ which converges to x. We need to show that $\lim_{k\to\infty} r_N(y_k) = r_N(x)$. As $\mathcal{M}$ is a smooth submanifold, we can find a neighbourhood $U \ni x$ in $\mathcal{M}$ around x, and a smooth frame

$v_1, \ldots, v_{n-m}$ for the normal bundle in U. Hence for each $y \in U$, $v_1(y), \ldots v_{n-m}(y)$ is an orthonormal basis for $N_y\mathcal{M}$. We can assume $\{y_k\}_{k=1}^{\infty} \subset U$ by only considering the tail of the sequence if necessary. As r_N is invariant with regards to translations of $\mathcal{M}$ we can assume without loss of generality that $x = 0$.

First suppose there is $\delta > 0$ such that $r_N(y_k) < r_N(x) - \delta$ for all $k \in \mathbb{N}$. Hence for each $k \in \mathbb{N}$ there is a normal vector $\boldsymbol{n}_k \in N_{y_k}\mathcal{M}$ with $\|\boldsymbol{n}_k\| \le r_N(x) - \frac{\delta}{2}$ such that

$$\boldsymbol{n}_k + y_k \notin \operatorname{Unp}(\mathcal{M}).$$

For each $k \in \mathbb{N}$ we can find scalars $a_1(y_k), \ldots, a_{n-m}(y_k)$ such that $\boldsymbol{n}_k = a_1(y_k)v_1(y_k) + \cdots + a_{n-m}(y_k)v_{n-m}(y_k)$. For each $i = 1, \ldots, n-m$, $\{a_i(y_k)\}_{k=1}^{\infty}$ is a bounded sequence, hence we can find an appropriate subsequence, such that

$$\lim_{k\to\infty} a_i(y_k) = a_i,$$

for $i = 1, \ldots, n-m$.

Let $\boldsymbol{n} = a_1 v_1(x) + \cdots + a_{n-m}v_{n-m}(x)$, then

$$\begin{aligned}\|\boldsymbol{n}_k - \boldsymbol{n}\| &= \|\sum_{i=1}^{n-m} a_i(y_k)v_i(y_k) - a_i v_i(x)\| \\ &\le \sum_{i=1}^{n-m} |a_i(y_k) - a_i| + |a_i|\|v_i(x) - v_i(y_k)\|.\end{aligned}$$

As v_i is smooth, $v_i(y_k) \to v_i(x)$ as $k \to \infty$, hence $\boldsymbol{n}_k \to \boldsymbol{n}$ as $k \to \infty$. Thus it follows $\boldsymbol{n} \notin \mathring{\operatorname{Unp}}(\mathcal{M})$. However, $\|\boldsymbol{n}\| = \lim_{k\to\infty}\|\boldsymbol{n}_k\| \le r_N(x) - \delta < r_N(x)$, which is a contradiction due to Lemma A.16.

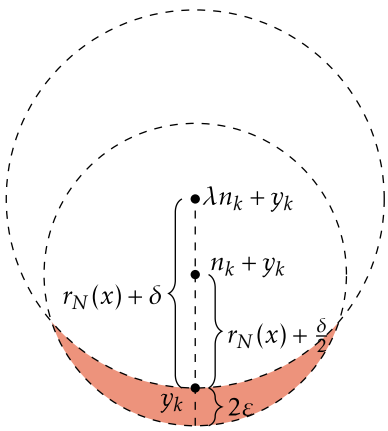

Figure 4: Illustration of the second part of the proof of Theorem A.18. We show that a point z lies outside the ball $B(\lambda n_k + y_k, r_N(x) + \delta)$ and inside the ball $B(n_k + y_k, r_N(x) + \frac{\delta}{2} + 2\varepsilon)$, which correspond to the red area in the figure. As $\varepsilon \to 0$, the red area becomes smaller until it disapears at $\varepsilon = 0$.

Next suppose there is a $\delta > 0$ such that $r_N(y_k) > r_N(x) + \delta$ for all $k \in \mathbb{N}$. Then there exists $\boldsymbol{n} \in N_x\mathcal{M}$ with $\|\boldsymbol{n}\| < r_N(x) + \frac{\delta}{2}$ such that we can find $z \in \mathcal{M}$, $z \neq x = 0$ such that $d(\boldsymbol{n}, \mathcal{M}) = \|\boldsymbol{n} - z\| \le r_N(x) + \frac{\delta}{2}$. We can write

$$\boldsymbol{n} = a_1 v_1(x) + \cdots + a_{n-m}v_{n-m}(x).$$

Let $\boldsymbol{n}_k = a_1 v_1(y_k) + \dots a_{n-m} v_{n-m}(y_k)$ and consider the distance from $\boldsymbol{n}_k + y_k$ to z

$$\|z - (\boldsymbol{n}_k + y_k)\| \le \|z - \boldsymbol{n} + \boldsymbol{n} - \boldsymbol{n}_k - y_k\| \le \|z - \boldsymbol{n}\| + \|\boldsymbol{n} - \boldsymbol{n}_k\| + \|y_k\|.$$

Let $\varepsilon > 0$, then we can find $K \in \mathbb{N}$ such that for $k \ge K$,

$$\|z - (\boldsymbol{n}_k + y_k)\| \le \|z - n\| + 2\varepsilon \le r_N(x) + \frac{\delta}{2} + 2\varepsilon. \tag{A.3}$$

Let $\lambda = \frac{r_N(x)+\delta}{\|\boldsymbol{n}_k\|}$, then $\|\lambda \boldsymbol{n}_k\| = r_N(x) + \delta < r_N(y_k)$, hence $\boldsymbol{n}_k + y_k \in \mathrm{Unp}(\mathcal{M})$, hence

$$\|z - (\lambda \boldsymbol{n}_k + y_k)\| > r_N(x) + \delta. \tag{A.4}$$

□

Hence, for every $k \ge K$ we get that z lies inside the $\bar{B}(\boldsymbol{n}_k + y_k, r_N(x) + \frac{\delta}{2} + 2\varepsilon)$ and outside the ball $B(\lambda \boldsymbol{n}_k + y_k, r_N(x) + \delta)$, see Figure 38. In the supplemental material (Section A.A) we prove in Proposition A.27 that for every $\varepsilon' > 0$ we can find $K \in N$ such that the norm of any point in $\bar{B}(\boldsymbol{n}_k + y_k, r_N(x) + \frac{\delta}{2} + 2\varepsilon) \cap B(\lambda \boldsymbol{n}_k + y_k, r_N(x) + \delta)^c$ is bounded by ε'. This implies that $z = x$ which is a contradiction.

To finish the proof, assume $y_k \to x$ is a sequence where $\{r_N(y_k)\}_{k=1}^{\infty}$ does not converge to $r_N(x)$. Then there must be a subsequence which satisfies either the first or the second situation. Hence, we have a contradiction.

A.4 Calculating uniqueness bounds for immersions

The motivation behind the pointwise normal reach is to obtain a computable local bound which ensures the existence of unique nearest points. The problem is that we, in general, have no easy way of computing general tangent cones. Instead, in this section, we want to utilise standard analytic tools to obtain tangent vectors.

Definition A.20. *Consider $f : \mathbb{R}^m \to \mathbb{R}^n$ an immersion and let $P_u = I - J_u(J_u^T J_u)^{-1} J_u^T$ for $u \in \mathbb{R}^m$, where J_u is the Jacobian in the point u. Then P_u is the projection matrix onto the orthogonal complement of* $\mathrm{span}(J_u)$.

For $u, v \in \mathbb{R}^m$ denote

$$\hat{R}(u,v) = \frac{\|f(v) - f(u)\|^2}{2\|P_u(f(v) - f(u))\|} \tag{A.5}$$

for $\|P_u(f(v) - f(u))\| \neq 0$, otherwise $\hat{R}(u,v) = \infty$. Then let

$$\hat{r}(u) = \inf_{v \in \mathbb{R}^m} \hat{R}(u,v). \tag{A.6}$$

Suppose $x = f(u)$, then it is clear that if $T_x\mathcal{M} = \mathrm{span}(J_u)$ then $d(y-x, T_x\mathcal{M}) = \|P_u(y-x)\|$. Hence, $\hat{r} : \mathbb{R}^m \to [0,\infty]$ is associated to the pointwise normal reach in the following way

Corollary A.21. *For $u \in \mathbb{R}^m$, $\hat{r}(u) = r_N(f(u))$ if* $\mathrm{span}(J_u) = T_{f(u)}\mathcal{M}$.

For an immersion f, its differential Df_u is injective, hence we know from the Inverse Function Theorem that f is locally diffeomorphic. Thus the image of f behaves locally as a manifold. As the image of f might have intersections, it is not necessarily globally a manifold. We thus characterise the points u in $\mathbb{R}^m$ regarding whether or not $f(\mathbb{R}^m)$ behaves as a manifold around $f(u)$.

Definition A.22. *A point $u \in \mathbb{R}^m$ is called a manifold point if there is an $\varepsilon > 0$ and a neighbourhood $U \subset \mathbb{R}^m$, $u \in U$ such that*

$$f\big|_U : U \to \mathcal{M} \cap B(f(U), \varepsilon)$$

is a diffeomorphism. Otherwise, we call u a non-manifold point.

As Df_u is injective, there exists a neighbourhood $U \subset \mathbb{R}^m$, $u \in U$ such that

$$f\big|_U : U \to f(U) \tag{A.7}$$

is a diffeomorphism. Hence u is a manifold point if and only if there is an $\varepsilon > 0$ such that

$$\mathcal{M} \cap B(f(u), \varepsilon) \subset f(U).$$

Characterising continuity of $\hat{r}$

In this section we characterise the continuity of the $\hat{r}$ estimator. It turns out each point in $\mathbb{R}^m$ satisfies one of three situations.

Theorem A.23. *Suppose $f : \mathbb{R}^m \to \mathbb{R}^n$ is a smooth immersion, then the normal reach estimator $\hat{r} : \mathbb{R}^m \to [0, \infty]$ can be characterised in the following way*

1. *If $u \in \mathbb{R}^m$ is a manifold point, then $\hat{r}(u) = r_N(f(u))$ and $\hat{r}$ is continuous in u.*
2. *If $u \in \mathbb{R}^m$ is a non-manifold point and $\operatorname{span}(J_u) \neq T_{f(u)}\mathcal{M}$, then $\hat{r}(u) = 0$ and $\hat{r}$ is continuous in u.*
3. *If $u \in \mathbb{R}^m$ is a non-manifold point and $\operatorname{span}(J_u) = T_{f(u)}\mathcal{M}$, then $\hat{r}(u) = r_N(f(u))$ and we do not know a priori if $\hat{r}$ is continuous in u.*

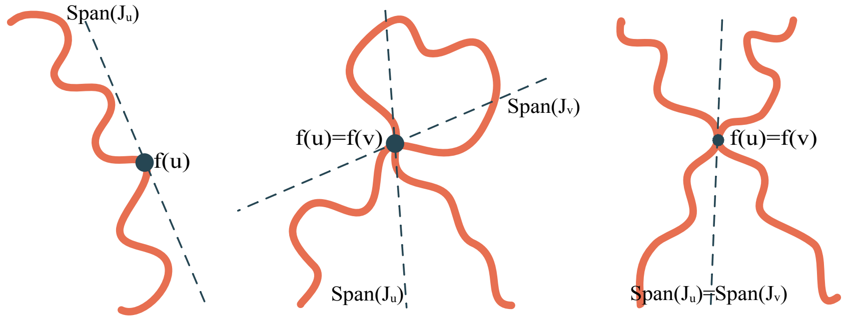

Figure 5: The left image shows the situation where u is a manifold point. The centre shows the situation where u is a non-manifold point, and the Jacobian does not span the tangent cone. The right image shows the situation where u is a non-manifold point, but the Jacobian does span the tangent cone.

Statement 1. follows from Lemma A.18. We will prove statement 2. in the subsequent Lemmas A.24 and A.25. The situation in the 3. statement is discussed in the following subsection.

Lemma A.24. *Suppose $u \in \mathbb{R}^m$ and* $\operatorname{span}(J_u) \neq T_{f(u)}\mathcal{M}$*, then* $\hat{r}(u) = 0$.

Proof. Let $x = f(u)$, and assume without loss of generality that $x = 0$. As $\operatorname{span}(J_u) \neq T_x\mathcal{M}$ there must be a $w \in T_x\mathcal{M}$, $\|w\| = 1$ and $w \notin \operatorname{span}(J_u)$. As $\operatorname{span}(J_u)$ is closed, $d(w, \operatorname{span} J_u) = \sin\theta$ for some angle θ such that $\sin\theta > 0$. As w is a tangent vector in $T_x\mathcal{M}$, it follows that for every $\varepsilon > 0$ there is a y in $\mathcal{M}$ such that

$$0 < \|y - x\| < \varepsilon \quad \text{and} \quad \|\frac{y - x}{\|y - x\|} - w\| < \varepsilon.$$

We observe that $\|\frac{y}{\|y\|} - w\| < \varepsilon$ if and only if $\|y - \|y\|w\| < \|y\|\varepsilon$. Suppose such an ε and y has been given, then

$$\begin{aligned}\|P_u(y)\| &= \|P_u(y - \|y\|w + \|y\|w)\| \\ &= \|P_u(\|y\|w) - (-1)P_u(y - \|y\|w)\| \\ &\geq \Big| \|P_u(\|y\|w)\| - \|P_u(y - \|y\|w)\| \Big|\end{aligned}$$

Notice P_u is a projection, so $\|P_u(y - \|y\|w)\| \leq \|y - \|y\|w\|$ and we have that $P_u(\|y\|w) = \|y\|\sin\theta$. Hence for $\varepsilon < \sin\theta$, we have that

$$\|P_u(y)\| > \|y\|\sin\theta - \|y\|\varepsilon.$$

Hence for sufficiently small ε and corresponding y we have that

$$\frac{\|y\|^2}{2\|P_u(y)\|} \leq \frac{\|y\|^2}{2\|y\|(\sin\theta - \varepsilon)} \leq \frac{\varepsilon}{2(\sin\theta - \varepsilon)}.$$

As $\varepsilon/2(\sin\theta - \varepsilon) \to 0$ as $\varepsilon \to 0$, it follows that we can find y and hence corresponding v with $y = f(v)$ such that $\hat{R}(u, v) < \varepsilon$ for any $0 < \varepsilon < \sin\theta$. □

Lemma A.25. *Suppose $u \in \mathbb{R}^m$ and that* $\operatorname{span}(J_u) \neq T_{f(u)}\mathcal{M}$*, then $\hat{r} : \mathbb{R}^m \to \mathbb{R}$ is continuous in u.*

Proof. To prove that $\hat{r}$ is continuous in u we need to prove that for every $\varepsilon > 0$ there is a $\delta > 0$ such that for $v \in R^m$ with $\|v - u\| < \delta$, then $\hat{r}(v) < \varepsilon$. Without loss of generality assume that $f(u) = 0$.

As $\operatorname{span}(J_u) \neq T_0\mathcal{M}$, there must be a $\boldsymbol{t} \in T_0\mathcal{M}$, $\|\boldsymbol{t}\| = 1$ with $\boldsymbol{t} \notin \operatorname{span}(J_u)$ such that $\|P_u(\boldsymbol{t})\| = \sin\theta > 0$. As $\boldsymbol{t}$ is a tangent vector, for each $\varepsilon > 0$ there is a $y \in \mathcal{M}$ with $0 < \|y\| < \varepsilon$ such that $\|\frac{y}{\|y\|} - \boldsymbol{t}\| < \varepsilon$.

Let $0 < \varepsilon < \sin\theta$ be given. For any $y \in \mathcal{M}$ with $y \neq 0$, we have that $\frac{\|P_u(y)\|}{\|y\|} = \|P_u(\frac{y}{\|y\|} - \boldsymbol{t} + \boldsymbol{t})\|$, thus

$$\left| \|P_u(\boldsymbol{t})\| - \|P_u\left(\frac{y}{\|y\|} - \boldsymbol{t}\right)\| \right| \leq \frac{\|P_u(y)\|}{\|y\|} \leq \|P_u\left(\frac{y}{\|y\|} - \boldsymbol{t}\right)\| + \|P_u(\boldsymbol{t})\|$$

As P_u is continuous and $P_u(\mathbf{0}) = 0$, we can find y_0 such that $\|P_u\left(\frac{y_0}{\|y_0\|} - \boldsymbol{t}\right)\| < \varepsilon$, thus for ε satisfying $0 < \varepsilon < \sin\theta$, we get that

$$\sin\theta - \varepsilon < \frac{\|P_u(y_0)\|}{\|y_0\|} < \sin\theta + \varepsilon. \tag{A.8}$$

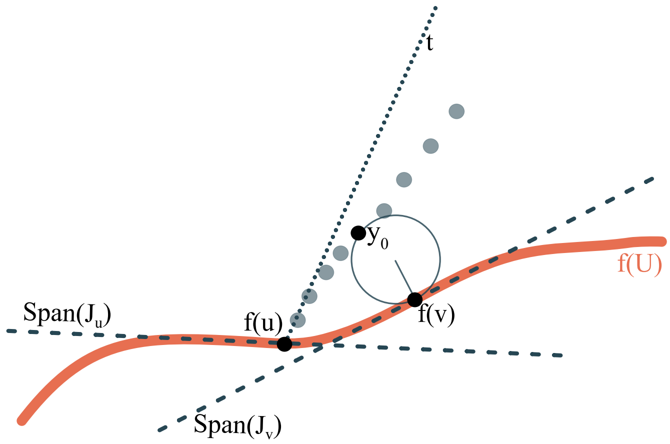

Figure 6: Illustration of the proof of Lemma A.25. If $\mathrm{span}(J_u) \neq T_{f(u)}\mathcal{M}$ there is a tangent vector which is not spanned by the Jacobian. In that case we can find points $y_0 \in \mathcal{M}$ for which the angle between $\boldsymbol{t}$ and the line from $f(u)$ to y_0 is arbitrarily close. As f is smooth, for $v \in \mathbb{R}^m$ close to u then the span of J_v is close to the span of J_u.

Next let $v \in \mathbb{R}^m$ and for ease of notation denote $x = f(v)$. We consider the fraction

$$\frac{\|y_0 - x\|^2}{2\|P_v(y_0 - x)\|} = \frac{\|y_0 - x\|^2}{2\|(P_u + P_v - P_u)(y_0 - x)\|} \leq \frac{\|y_0 - x\|^2}{2\left|\|P_u(y_0 - x)\| - \|(P_v - P_u)(y_0 - x)\|\right|}. \tag{A.9}$$

To start, we consider the denominator of the right hand side. We see that

$$\begin{aligned}\|P_u(y_0 - x)\| - \|(P_v - P_u)(y_0 - x)\| \geq& \|P_u(y_0 - x)\| - \|P_v - P_u\|\|y_0 - x\| \\ =& \|y_0 - x\|\left(\frac{\|P_u(y_0 - x)\|}{\|y_0 - x\|} - \|P_v - P_u\|\right)\end{aligned}$$

which if greater than 0 if and only if

$$\frac{\|P_u(y_0 - x)\|}{\|y_0 - x\|} - \|P_v - P_u\| > 0.$$

Let δ_0 be given such that $\|v\| < \delta_0$ implies $\|x\| = \|f(v)\| < \|y_0\|$. As f is smooth, $u \mapsto P_u = I - J_u\left(J_u^T J_u\right)^{-1} J_u^T$ is continuous. Thus we can find δ_1 such that $\|v - y\| < \delta_1$ implies $\|P_v - P_u\| < \varepsilon$. Let $\delta_2 = \min\{\delta_0, \delta_1\}$, then for $\|v - u\| < \delta_2$,

$$\begin{aligned}\frac{\|P_u(y_0 - x)\|}{\|y_0 - x\|} \geq& \frac{\|P_u(y_0)\| - \|P_u(x)\|}{\|y_0\| + \|x\|} \\ \geq& \frac{\|P_u(y_0)\| - \|P_u(x)\|}{2\|y_0\|} \\ =& \frac{\|P_u(y_0)\|}{2\|y_0\|} - \frac{\|P_u(x)\|}{2\|y_0\|}.\end{aligned}$$

As f is differentiable, we can find $\delta_3 > 0$ such that for v which satisfies $\|v - u\| < \delta_3$, we have that

$$f(v) = J_u(v - u) + \Phi(v - u)\|v - u\|,$$

where $\Phi : B(u,\delta_3) \to \mathbb{R}^n$ is an o-function, that is $\Phi(0) = 0$ and Φ is differentiable in 0. Thus it follows that

$$\|P_u(x)\| = \|P_u(f(v))\| = \|P_u\Big(J_u(v-u) + \Phi(v-u)\|v-y\|\Big)\| \leq \|v-u\|\|P_u(\Phi(v-u))\|.$$

This implies that we can find $\delta_4 < \delta_3$ such that for v which satisfies $\|v-u\| < \delta_4$,

$$\|P_u(x)\| \leq \|v-u\|\|P_u(\Phi(v-u))\| < \|y_0\|\varepsilon.$$

From the inequality (A.8) we know that $\frac{\|P_u(y_0)\|}{\|y_0\|} > \sin\theta - \varepsilon$. Combining that with the inequality above we get that

$$\frac{\|P_u(y_0-x)\|}{\|y_0-x\|} \geq \frac{\|P_u(y_0)\|}{2\|y_0\|} - \frac{\|P_u(x)\|}{2\|y_0\|} > \frac{\sin\theta-\varepsilon}{2} - \frac{\varepsilon}{2} = \frac{\sin\theta}{2} - \varepsilon.$$

Without loss of generality we can assume $\varepsilon < \frac{\sin\theta}{4}$, Combining all our hard work, we let $\delta = \min\{\delta_2, \delta_4, \varepsilon\}$ and get that

$$\frac{\|P_u(y_0-x)\|}{\|y_0-x\|} - \|P_v - P_u\| > \frac{\sin\theta}{2} - 2\varepsilon > 0,$$

Applying the above to (A.9) we get that for v satisfying $\|u-v\| < \delta$

$$\begin{aligned}
\frac{\|y_0-x\|^2}{2\|P_v(y_0-x)\|} \leq & \frac{\|y_0-x\|^2}{2\left|\|P_u(y_0-x)\| - \|(P_v-P_u)(y_0-x)\|\right|} \\
= & \frac{\|y_0-x\|^2}{2\left(\|P_u(y_0-x)\| - \|(P_v-P_u)(y_0-x)\|\right)} \\
\leq & \frac{\|y_0-x\|^2}{2\left(\|P_u(y_0-x)\| - \|P_v-P_u\|\|y_0-x\|\right)} \\
= & \frac{\|y_0-x\|}{2\left(\frac{\|P_u(y_0-x)\|}{\|y_0-x\|} - \|P_v-P_u\|\right)} \\
< & \frac{\varepsilon}{2(\frac{\sin\theta}{2} - 2\varepsilon)} = \frac{\varepsilon}{\sin\theta - 4\varepsilon}.
\end{aligned}$$

As $\frac{\varepsilon}{\sin\theta-4\varepsilon} \to 0$ as $\varepsilon \to 0$, it is clear that for any $\varepsilon' > 0$ we can find δ such that $\|v-u\| < \delta$ implies

$$\hat{r}(v) \leq \frac{\|y_0-x\|^2}{2\|P_v(y_0-x)\|} < \frac{\varepsilon}{\sin\theta - 4\varepsilon} < \varepsilon'.$$ □

$\hat{r}(u)$ for a non-manifold point with $\mathrm{span}(J_u) = T_{f(u)}\mathcal{M}$

In this situation we cannot say if $\hat{r}$ is continuous or not. To understand why we first consider the following result..

Lemma A.26. *Suppose $u \in \mathbb{R}^m$ is a non-manifold point and let $x = f(u)$, then* $\mathrm{reach}(\mathcal{M}, x) = 0$.

Proof (Sketch of proof). In this proof we use the results in Appendix A.B. Suppose for contradiction that $\mathrm{reach}(\mathcal{M}, x) > 0$. Then it follows from Proposition A.40 that there is an $s > 0$ such that $\mathcal{M} \cap \bar{B}(x,s)$ has positive reach. It then follows from the arguments in Remark A.34 that $\text{span}(J_u) = T_x\mathcal{M}$.

Let $U \subset \mathbb{R}^m$ be a neighbourhood around u such that $f\big|_U : U \to f(U)$ is a diffeomorphism. From Proposition A.40 we get that we can restrict U such that $\text{reach}(f(U)) > r > 0$ for some $r > 0$. As u is a non-manifold point we know that there exists a sequence $\{y_k\}_{k\in\mathbb{N}}$ with $y_k \to x$ as $k \to \infty$, and $y_k \notin f(U)$ for all k. For sufficiently large k, we must have that $y_k \in \text{Unp}(f(U))$. We denote x_k the unique nearest point to y_k in $f(U)$. For sufficiently large k we get that both points x_k and y_k lie in $\mathcal{M} \cap \bar{B}(x,s)$. From results in the Global reach chapter, we know that there exists a geodesic between them which is contained in a spindle.

From arguments like Remark A.34 we get that $\text{span}(J_{u_k}) = T_{x_k}\mathcal{M}$, however the tangent vector given by the geodesic will not be spanned by J_{u_k}. □

The lemma implies that there is a sequence of points y_k in $\mathcal{M}$ such that $y_k \to x$ for $k \to \infty$, and $r_N(y_k) \to 0$, due to Proposition A.8. Then if $r_N(x) > 0$, $\hat{r}$ is not continuous in u. However, if we have that $r_N(x) = 0$, we do not have sufficient information to say that $\hat{r}$ is continuous.

A.A Appendix

Example A.11

The goal of this section is to prove that the pointwise normal reach is not continuous on the graph of h, where h is given by

$$h(x,y) = \begin{cases} 2\sqrt{r^2 - \left(\frac{y}{2}\right)^2} - \sqrt{r^2 - x^2} - r & \text{if } 0 \le |x| \le \frac{y}{2} \\ \sqrt{r^2 - (x-y)^2} - r & \text{if } \frac{y}{2} \le x \le y \\ \sqrt{r^2 - (x+y)^2} - r & \text{if } -y \le x \le -\frac{y}{2} \\ 0 & \text{otherwise} \end{cases}$$

for a given $r > 0$. Let $\mathcal{M}$ denote the graph of h. Then, $\mathcal{M}$ can be parametrised by

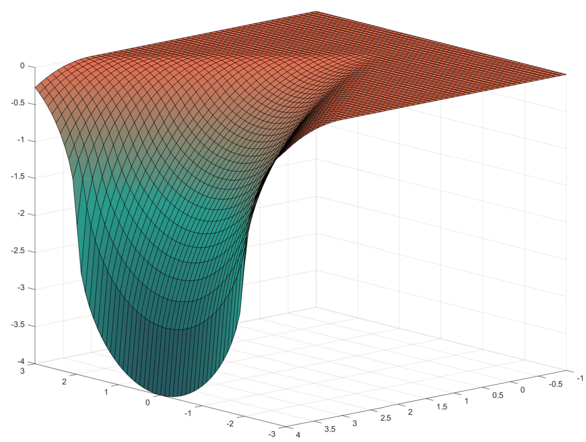

Figure 7: The graph of the function h.

$(x,y) \mapsto (x,y,h(x,y))$. We want to prove that r_N is not continuous in the point $(0,0,0)$.

Pointwise normal reach in $(0,0,0)$

To compute the pointwise normal reach, we start by finding the tangent plane in the point. Consider the curve on $\mathcal{M}$ given by $\gamma : t \mapsto (0,t,h(0,t))$. Then $\dot{\gamma}(t) = -\frac{y}{\sqrt{4r^2 - y^2}}$, and thus $\dot{\gamma}(0) = \mathbf{0}$. Hence, it follows that $T_{\mathbf{0}}\mathcal{M}$ is spanned by the x and y axis.

Using Definition A.3 the calculation for the pointwise normal reach is then the infimum of

$$R((0,0,0),(x,y,h(x,y))) = \frac{x^2 + y^2 + h(x,y)^2}{2|h(x,y)|}.$$

For ease of notation, denote $F(x,y) = R((0,0,0),(x,y,h(x,y)))$.

Suppose $\frac{y}{2} \leq |x| \leq y$, then the numerator of $F(x,y)$ becomes

$$\begin{aligned} x^2 + y^2 + h(x,y)^2 =& x^2 + y^2 + \left(\sqrt{(r^2 - (x \mp y)^2)} - r\right)^2 \\ =& \pm 2xy + 2r^2 - \sqrt{r^2 - (x \mp y)^2} \\ \geq& y^2. \end{aligned}$$

Next, considering the denominator, we get that

$$\begin{aligned} 2|h(x,y)| =& 2r - 2\sqrt{r^2 - (x \mp y)^2} \\ \leq& 2r - 2\sqrt{r^2 - \left(\frac{y}{2}\right)^2} = |h(0,y)|. \end{aligned}$$

Combining the two gives

$$F(x,y) \geq \frac{y^2}{|h(0,y)|}.$$

We claim that $\frac{y^2}{|h(0,y)|} \geq 2r$. This follows from the fact that $y^2 = y^2 + 4r^2 - 4r^2 = 4r^2 - \left(\sqrt{4r^2 - y^2}\right) \geq 4r^2 - 2r\sqrt{4r^2 - y^2}$.

Next we consider $F(x,y)$ for $|x| \leq \frac{y}{2}$. Again we want to show that $F(x,y) \geq 2r$. Consider that

$$\frac{x^2 + y^2 + h(x,y)^2}{2|h(x,y)|} \geq 2r$$

is equivalent to

$$x^2 + y^2 + h(x,y)^2 \geq 4r|h(x,y)|. \tag{A.10}$$

We start by considering the left hand side of $(A.10)$:

$$\begin{aligned} \text{'LHS'} =& x^2 + y^2 + \left(\sqrt{4r^2 - y^2} - \sqrt{r^2 - x^2} - r\right)^2 \\ =& 6r^2 - 2\sqrt{r^2 - x^2}\sqrt{4r^2 - y^2} - 2r\sqrt{4r^2 - y^2} + 2r\sqrt{r^2 - x^2}. \end{aligned}$$

The right hand side of the equation is:

$$\text{'RHS'} = 4r^2 + 4r\sqrt{r^2 - x^2} - 4r\sqrt{4r^2 - y^2}.$$

Combining the two sides we get that $(A.10)$ is equivalent to

$$\begin{aligned} & 2r^2 - 2\sqrt{r^2 - x^2}\sqrt{4r^2 - y^2} \geq 2r\sqrt{r^2 - x^2} - 2r\sqrt{4r^2 - y^2} \\ \Leftrightarrow \quad & 2r^2 + 2r\sqrt{4r^2 - y^2} \geq 2r\sqrt{r^2 - x^2} + 2\sqrt{r^2 - x^2}\sqrt{4r^2 - y^2} \\ \Leftrightarrow \quad & 2r(r + \sqrt{4r^2 - y^2}) \geq 2\sqrt{r^2 - x^2}(r + \sqrt{4r^2 - y^2}). \end{aligned}$$

Here, the last inequality holds as $r \geq \sqrt{r^2 - x^2}$ for all $|x| \leq \frac{y}{2}$, thus proving the claim.

Thus we have shown that $F(x,y) \geq 2r$ for all x, y. Next we consider

$$\lim_{y \to 0} F(0,y) = \lim_{y \to 0} \frac{y^2 + \left(\sqrt{4r^2 - y^2} - 2r\right)^2}{2\left(2r - \sqrt{4r^2 - y^2}\right)} = \lim_{y \to 0} \frac{y^2}{2\left(2r - \sqrt{4r^2 - y^2}\right)} + \left(2r - \sqrt{4r^2 - y^2}\right).$$

The second term of the sum goes to 0 as y goes to 0. Using L'Hôpital's rule we can show that the first term goes to $2r$. Thus $\lim_{y\to 0} F(x,y) = 2r$. This shows that $r_N((0,0,0)) = 2r$.

Pointwise normal reach in $(0,y,h(0,y))$

To prove that r_N is not continuous in $(0,0,0)$ we consider $r_N(0,y,h(0,y))$ as y goes to 0. For a given $y>0$ we consider the curve $\gamma : x \mapsto (x,y,h(x,y))$. We now want to compute the normal curvature at $(0,y,h(0,y))$ given by γ. We find the unit normal vector $\boldsymbol{n} = \sqrt{\frac{4r^2}{4r^2-y^2}}\left(0, \frac{y}{\sqrt{4r^2-y^2}}, 1\right)$, and $\ddot{\gamma}(0) = \left(0,0,\frac{1}{r}\right)$. Thus computing the normal curvature gives

$$\boldsymbol{n}\cdot\ddot{\gamma} = \frac{1}{r}\frac{2r}{\sqrt{2r^2-y^2}} = \frac{2}{\sqrt{4r^2-y^2}},$$

where the latter goes to $\frac{1}{r}$ as y goes to 0. This implies that for any $\varepsilon > 0$ we can find $\delta > 0$ such that for $y < \delta$, then the normal curvature at y generated by γ is larger than $\frac{1}{r+\varepsilon}$. This implies $r_N(0,y,h(0,y)) \le r+\varepsilon$. However as $r_N((0,0,0)) = 2r$, the pointwise normal reach is not continuous in $(0,0,0)$.

Supplement to continuity

Proposition A.27. *Let v be a unit vector in $\mathbb{R}^n$. Let $r_1, r_2 > 0$ with $r_1 < r_2$. Consider a ball B_2 with centre $r_2 v$ and radius r_2, and B_1 with centre $r_1 v$ and radius $r_1+\varepsilon$, for some $\varepsilon > 0$. For points which satisfies $z \in B_1 \cap B_2^c$, then $\|z\| \to 0$ as $\varepsilon \to 0$.*

Proof. Suppose $n = 2$. Consider the intersection points of the two boundary circle, that is the points which satisfy

$$x^2 + (y-r_2)^2 = r_2^2$$
$$x^2 + (y-r_1)^2 = (r_1+\varepsilon)^2$$

Solving these equations gives

$$\hat{x}^2 = 2\hat{y}r_2 - \hat{y}^2,$$
$$\hat{y} = \frac{2r_1\varepsilon + \varepsilon^2}{2(r_2-r_1)}.$$

Let $\alpha = \frac{2r_1\varepsilon+\varepsilon^2}{2(r_2-r_1)}$, then $B_1 \cap B_2^c \subset [-\sqrt{2\alpha r_2 - \alpha^2}, \sqrt{2\alpha r_2-\alpha^2}] \times [-\varepsilon, \alpha]$. However, $\alpha \to 0$ as $\varepsilon \to 0$, which proves the claim.

Suppose $n > 2$, then for $z \in B_1 \cap B_2^c$, z and v lie in a plane P which passes through origo. Then the claim follows from the two dimensional case. □

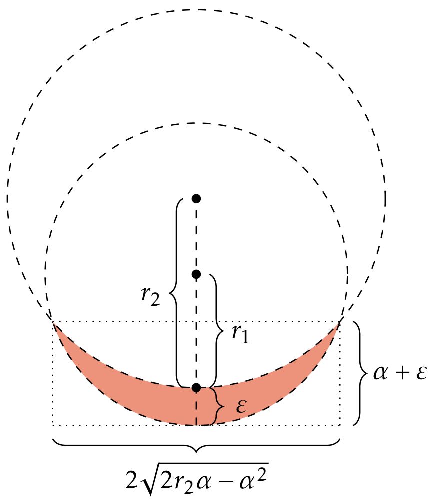

A.B Local reach for immersions

Dimension Theory

Definition A.28. *([5], III 1) Let A be a topological space. We define the topological dimension of A as follows.*

- *The empty set, and only the empty set, has dimension -1.*
- *A has dimension $\leq k$ at $x \in A$ if x has arbitrarily small neighbourhoods in A whose boundaries have dimension $\leq k-1$.*
- *A has dimension $\leq k$ if it has dimension $\leq k$ at every point in A. A has dimension k if it has dimension $\leq k$ but not dimension $\leq k-1$.*

It is immediate that the property of having dimension k is topological invariant.

In what follows in this section we assume that the topological spaces considered are metrisable and second-countable (which certainly holds for spaces which can be embedded in Euclidean space). We will make use of two classical results of dimension theory:

Theorem A.29. *([5], III 2) Let the topological space A be a countable union of closed subsets, each of dimension $\leq k$. Then the dimension of A is $\leq k$.*

Theorem A.30. *([5], IV 3 and Corollary 1, p. 46)*

1. *A subset of $\mathbb{R}^n$ has dimension n if and only if the subset has non-empty interior.*
2. *A subset of an n-dimensional manifold has dimension n if only if the subset has non-empty interior.*

These allow us to determine the dimension of an immersed manifold:

Proposition A.31. *Let $\mathcal{P}$ be a smooth k-dimensional manifold (paracompact and Hausdorff) and let $f : \mathcal{P} \to \mathbb{R}^n$ be a smooth immersion. Then $f(\mathcal{P})$ is of dimension n.*

Proof. For each $p \in P$ there exists an open neighbourhood U_p of p in $\mathcal{P}$ such that $f\big|_{U_p} : U_p \to f(U_p)$ is a smooth diffeomorphism. Choose now an open neighbourhood U_p' of p in $\mathcal{P}$ such that $\overline{U_p'} \subset U_p$. The collection $\left\{U_p' : p \in \mathcal{P}\right\}$ is an open covering of $\mathcal{P}$. Since $\mathcal{P}$ is paracompact there exists a countable refinement $\left\{U_\lambda' : \lambda \in \Lambda\right\}$ covering $\mathcal{P}$. It follows that $\bigcup_{\lambda\in\Lambda} f(U_\lambda') = f(\mathcal{P})$, and thus

$$\bigcup_{\lambda\in\Lambda} f(\overline{U_\lambda'}) = f(\mathcal{P}).$$

By Theorem A.30 (2), the closed set $\overline{U_\lambda'}$ has dimension k, since its interior U_λ' is non-empty. Thus, since $f\big|_{U_p} : U_p \to f(U_p)$ is a homeomorphism, $f(\overline{U_\lambda'})$ is also a closed set of dimension k. Thus by Theorem A.30, $f(\mathcal{P})$ is of dimension k. □

The topological dimension also plays a role in Federer's work on reach. Recall that we denote the tangent and normal cones at a in a subset A, T_aA and N_aA. We use the same notation to denote the tangent and normal space for a smooth manifold, as in this situation the two coincide.

Proposition A.32. *([2], 4.15 (4)) Let $A \subset \mathbb{R}^n$ have dimension k and positive reach, and let $a \in A$. Then the topological dimension of the normal cone satisfy*

$$\dim N_aA \geq n - k,$$

where if $\dim N_aA = n-k$ *then T_aA is a k-dimensional vector space.*

Corollary A.33. *Suppose, in the situation of Proposition A.32, that $T_aA \supset L$, where L is a vector space of dimension k. Then $T_aA = L$.*

Proof.

$$\begin{aligned} N_aA &= \{v \in \mathbb{R}^n : \langle v, u\rangle \leq 0 \text{ for all } u \in T_aA\} \\ &\subset \{v \in \mathbb{R}^n : \langle v, u\rangle \leq 0 \text{ for all } u \in L\} \\ &= L^\perp. \end{aligned}$$

Thus the topological dimension $\dim N_aA \leq \dim L^\perp = n-k$. Applying Proposition A.32, T_aA is a k-dimensional vector space. Thus $T_aA = L$. □

Remark A.34. Let $f(\mathcal{P})$ be as in Proposition A.31 , covered by $\left\{f(U_p) : p \in \mathcal{P}\right\}$ as in the proof of Proposition A.31. Since $f\big|_{U_p} : U_p \to f(U_p)$ is a smooth diffeomorphism, the tangent cone satisfies

$$T_{f(p)}f(U_p) = T(f(U_p))_{f(p)} = df_u(T(U_p)_p) = df_p(T_p\mathcal{P}),$$

and thus $T_{f(p)}f(M) \supset df_p(T_p\mathcal{P})$ for all $p \in \mathcal{P}$.

Suppose now that $f(\mathcal{P})$ has positive reach. Applying Corollary A.33 with L replaced by $df_p(T_p\mathcal{P})$ and A replaced by $f(\mathcal{P})$, we find that

$$T_{f(p)}f(\mathcal{P}) = df_p(T_p\mathcal{P})$$

for all $p \in \mathcal{P}$.

Transversality

Definition A.35. *Let $f : \mathcal{P} \to \mathcal{M}$, $g : \mathcal{Q} \to \mathcal{M}$ be two smooth maps between smooth manifolds. Then f, g are* transverse *if, for all $p \in \mathcal{P}$, $q \in \mathcal{Q}$, $m \in \mathcal{M}$ such that $f(p) = g(q) = m$,*

$$df_p(T_p\mathcal{P}) + dg_q(T_q\mathcal{Q}) = T_m\mathcal{M}.$$

Lemma A.36. *Let $f : \mathcal{P} \to \mathcal{M}$, $g : \mathcal{Q} \to \mathcal{M}$ be two smooth maps between smooth manifolds. Then f, g are transverse if and only if the product mapping*

$$f \times g : \mathcal{P} \times \mathcal{Q} \to \mathcal{M} \times \mathcal{M}$$

is transverse to the diagonal $\Delta_{\mathcal{M}} = \{(m,m) \in \mathcal{M} \times \mathcal{M} : m \in \mathcal{M}\}$.

Proof. This is left to the reader; it is often posed as an exercise in books on differential topology (see for example [4], Exercise 2.14(a), or [7], Problem 6-13(a)). □

Corollary A.37. *Let $f : \mathcal{P} \to \mathbb{R}^n$ be a smooth map, let $z \in \mathbb{R}^n$ and for any $s \in (0,\infty)$ let $G_s : S^{n-1} \to \mathbb{R}^n$ be given by $G_s(v) = v + sv$ for all $v \in S^{n-1}$ (the unit sphere centred at $\mathbf{0}$ in $\mathbb{R}^n$). Then f and G_s are transverse for almost all $s \in (0,\infty)$.*

Proof. Let $G : S^{n-1} \times (0,\infty) \to \mathbb{R}^n$ be defined by $G(v,s) = G_s(v)$ for all $v \in S^{n-1}$ and $s \in (0,\infty)$. It is immediate that G is a diffeomorphism onto its image $\mathbb{R}^n \setminus \{z\}$, and it follows that G is transverse to f. According to Lemma A.36 this is equivalent to $f \times G$ being transverse to the diagonal $\Delta_{\mathbb{R}^n}$ in $\mathbb{R}^n \times \mathbb{R}^n$.

Now $f \times G$ is a smooth parameterised family of maps $f \times G_s$. Applying the parameterised transversality theorem (e.g. [7], Theorem 6.35), $f \times G_s$ is transverse to $\Delta_{\mathbb{R}^n}$ for almost all $s \in (0,\infty)$. By Lemma A.36 then f and G_s are transverse for almost all $s \in (0,\infty)$.

We now introduce another of Federer's theorems on reach:

Theorem A.38. *([2], Theorem 4.10 (5)) Let A, B be closed subsets of R^n, with B compact. Let $C = A \cap B$ be non-empty. Suppose* $\operatorname{reach}(A,c), \operatorname{reach}(B,c) > r$ *for all $c \in C$. Suppose also that there exists no $c \in C$ and $v \in \mathbb{R}^n \setminus \{0\}$ such that $v \in N_cA$ and $-v \in N_cB$. Let*

$$\eta = \inf\left\{ \frac{\|w+v\|}{\|w\|+\|v\|} : c \in C, v \in N_cA, w \in N_cB, \|w\|+\|v\| > 0 \right\}.$$

Then $\eta \in (0,1]$ and C has reach $\geq \frac{1}{2}\eta r$.

We will need to deal with the case when A is not open:

Lemma A.39. *Let $A \subset \mathbb{R}^n$, and let $a \in A$ be such that* $\operatorname{reach}(A,a) > r > 0$. *Then*

$$A \cap \overline{B}(a,r) = \overline{A} \cap \overline{B}(a,r)$$

and $\operatorname{reach}(\overline{A}, r) > \frac{1}{2}r$.

Proof. We can find $r' > r$ such that $\text{reach}(A, a) > r'$. Since no point in $\overline{A}\backslash A$ has a nearest point in A, the intersection $(\overline{A}\backslash A) \cap \text{Unp}(A)$ is empty, and thus, since $\overline{B}(a, r') \subset \text{Unp}(A)$, we have that $(\overline{A}\backslash A) \cap \overline{B}(a, r')$ is empty. Furthermore, we have that $A \cap \overline{B}(a, r) = \overline{A} \cap \overline{B}(a, r)$.

Next suppose $z \in \mathbb{R}^n \setminus \text{Unp}(\overline{A})$, so there exists two points $a_1, a_2 \in \overline{A}$ such that $\|z - a_1\| = \|z - a_2\| = d(z, \overline{A}) = d(z, A)$. If $a_1, a_2 \in A$ then $z \notin \text{Unp}(A)$, and so $\|a - z\| > r'$. Otherwise, if one of a_1 or a_2, let us say a_1, is contained in $\overline{A} \setminus A$, so $\|a_1 - a\| > r'$. Since $\|a_1 - x\| = d(z, A)$, the distance $\|a_1 - z\| \leq \|z - a\|$. By the triangle inequality,

$$\|a_1 - a\| \leq \|a_1 - z\| + \|z - a\| \leq 2\|z - a\|,$$

thus $\|z - a\| \geq \frac{1}{2}\|a_1 - a\| > \frac{1}{2} r'$, completing the proof. □

Proposition A.40. *Let $f : \mathcal{P} \to \mathbb{R}^n$ be a smooth immersion and let $z \in f(\mathcal{P})$. Suppose that* $\text{reach}(f(\mathcal{P}), z) > 0$. *Then there exists $s > 0$ such that $f(\mathcal{P}) \cap \overline{B}(z, s)$ has positive reach.*

Proof. Let $r > 0$ be such that $\text{reach}(f(\mathcal{P}), z) > 4r$. Then according to Lemma A.39

$$f(\mathcal{P}) \cap \overline{B}(z, 4r) = \overline{f(\mathcal{P})} \cap \overline{B}(z, 4r)$$

and $\text{reach}(\overline{f(\mathcal{P})}, z) > 2r$. It follows that $\text{reach}(\overline{f(\mathcal{P})}, z') > r$ for all $z' \in \overline{B}(z, r)$.

According to Corollary A.37 we can find $s \in (0, r]$ such that G_s is not transverse to f. Let $B = \overline{B}(z, s)$. We claim that there exists no $c \in \overline{f(\mathcal{P})} \cap B$ and $v \in \mathbb{R}^n \setminus \{0\}$ such that $-v \in N_c \overline{f(\mathcal{P})}$ and $v \in N_c B$.

This is clear for c in the interior of B, since $N_c B = \{0\}$ at such points. Otherwise, let

$$c \in \overline{f(\mathcal{P})} \cap \partial B = f(\mathcal{P}) \cap \partial B = f(M) \cap G_s(S^{n-1}). \tag{A.11}$$

Thus for some $x \in \mathcal{P}$ and $w \in S^{n-1}$, $c = f(x) = G_s(w)$. Since f and G_s are transverse,

$$(dG_s)w(T(S^{n-1}_w)) + df_x(T_x\mathcal{P}) = \mathbb{R}^n.$$

$N_c B$ as a 1-dimensional half-space spanned by the unique outward pointing unit vector to ∂B – which is w.

Clearly, $df_x(T_x\mathcal{P}) \subset T_c f(\mathcal{P}) \subset T_c \overline{f(\mathcal{P})}$ so

$$c = \text{dual}\left(T_c \overline{f(\mathcal{P})}\right) \subset \text{dual}\,(df_x(T_x\mathcal{P})) = (df_x(T_x\mathcal{P}))^{\perp}.$$

Thus if there exists a non-zero vector $v \in N_c B$ with $-v \in N_c f(\mathcal{P})$ then $v \in (df_x(T_x\mathcal{P}))^{\perp}$. But v is a positive multiple of w, so $w \in (df_x(T_x\mathcal{P}))^{\perp}$. Dualising

$$df_x(T_x\mathcal{M}) \subset w^{\perp} = T(\partial B)_c = d(G_s)_w(T(S^{n-1})_w).$$

But this contradicts $(A.11)$, så no such v exists, and the claim is proved. It now follows directly from Theorem A.38 that $f(\mathcal{P}) \cap B = \overline{f(\mathcal{P})} \cap B$ is of positive reach. □

References

[1] Aamari, E., J. Kim, F. Chazal, B. Michel, A. Rinaldo and L. Wasserman (2019). Estimating the reach of a manifold. eng. *Electronic journal of statistics* 13(1), 1359–1399.

[2] Federer, H. (1959). Curvature Measures. eng. *Transactions of the American Mathematical Society* 93(3), 418–491.

[3] Fefferman, C., S. Mitter and H. Narayanan (2016). Testing the manifold hypothesis. eng. *Journal of the American Mathematical Society* 29(4), 983–1049.

[4] Hirsch, M. (1976). *Differential topology*. eng. Graduate texts in mathematics ; 33. New York: Springer.

[5] Hurewicz, W. and H. Wallman (1942). *Dimension Theory (PMS-4)*. eng. Princeton Mathematical Series. Princeton University Press.

[6] Kim, A.K.H. and H.H. Zhou (2015). Tight minimax rates for manifold estimation under Hausdorff loss. *Electronic Journal of Statistics* 9(1), 1562–1582.

[7] Lee, J.M (2013). *Introduction to smooth manifolds*. eng. 2. ed. Graduate texts in mathematics ; 218. New York: Springer.

[8] Niyogi, P., S. Smale and S. Weinberger (2008). Finding the Homology of Submanifolds with High Confidence from Random Samples. eng. *Discrete & Computational Geometry* 39(1-3), 419–441.

[9] Shao, H., A. Kumar and P.T. Fletcher (2018). "The Riemannian Geometry of Deep Generative Models". *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 428–4288.

[10] Svane, H.M (2019). "Reconstructing R-regular Objects from Trinary Digital Images". PhD dissertation. Aarhus Universitet, Institut for Matematik.

PAPER B

Is an encoder within reach?

Helene Hauschultz, Rasmus Berg Palm, Pablo Moreno-Muñoz, Nicki Skafte Detlefsen, Andrew du Plessis and Søren Hauberg

Abstract

The encoder network of an autoencoder is an approximation of the nearest point projection onto the manifold spanned by the decoder. A concern with this approximation is that, while the output of the encoder is always unique, the projection can possibly have infinitely many values. This implies that the latent representations learned by the autoencoder can be misleading. Borrowing from geometric measure theory, we introduce the idea of using the reach of the manifold spanned by the decoder to determine if an optimal encoder exists for a given dataset and decoder. We develop a local generalization of this reach and propose a numerical estimator thereof. We demonstrate that this allows us to determine which observations can be expected to have a unique, and thereby trustworthy, latent representation. As our local reach estimator is differentiable, we investigate its usage as a regularizer and show that this leads to learned manifolds for which projections are more often unique than without regularization.

B.1 Encoders as projectors

A good *learned representation* has many desiderata [2]. The perhaps most elementary constraint placed over most learned representations is that a given observation x should have a *unique* representation z, at least in distribution. In practice this is ensured by letting the representation be given by the output of a function, $z = g(x)$, often represented with a neural network.

The *autoencoder* [18] is an example where uniqueness of representation is explicitly enforced, even if its basic construction does not suggest unique representations. In the most elementary form, the autoencoder consists of an *encoder* $g_\psi : \mathbb{R}^D \to \mathbb{R}^d$ and a *decoder* $f_\phi : \mathbb{R}^d \to \mathbb{R}^D$, parametrized by ψ and ϕ, respectively. These are trained by minimizing the *reconstruction error* of the training data $\{x_1, \ldots, x_N\}$,

$$\psi^*, \phi^* = \underset{\psi,\phi}{\operatorname{argmin}} \sum_{n=1}^{N} \|f_\phi(g_\psi(x_n)) - x_n\|^2.$$

Here d is practically always smaller than D, such that the output of the encoder is a low-dimensional latent representation of high-dimensional data. The data is assumed to lie near a d-dimensional manifold $\mathcal{M}$ spanned by the decoder.

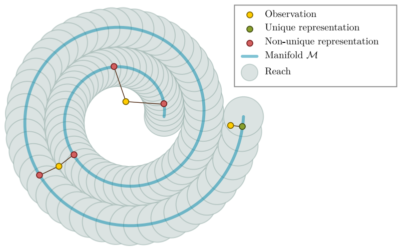

Figure 8: The projection of a point (yellow) onto a nonlinear manifold can take unique (green) or multiple values (red) depending on the reach of the manifold. When training data is inside the reach, the encoder can match the projection resulting in more trustworthy representations.

For a given decoder, we see that the optimal choice of the encoder is the projection onto $\mathcal{M}$, i.e.

$$g_{\text{optimal}}(x) = \text{proj}_{\mathcal{M}}(x) = \underset{z}{\operatorname{argmin}} \|x - f(z)\|^2.$$

For any *nonlinear* choice of decoder f, this optimal encoder does *not* exists everywhere. That is, multiple best choices of latent representation may exist for a given point, as the projection is not unique everywhere. As the learned encoder enforces a unique representation, it will choose arbitrarily among the potential representations (see Fig. 8). In this case, any analysis of the latent representations can be misleading, as it does not contain the information that another choice of representation would be equally good.

But does uniqueness of representations matter? Learned latent representations are used for a variety of tasks, most of which implicitly rely on the representations being unique. The simplest use case of learned representations is *visualization*, i.e. a scatter plot of the latent coordinates. Such plots are often used to form scientific hypotheses about the mechanics of the phenomena that generated the data, e.g. *protein evolution* [17, 6], or *identifying unexplored molecular structures* [19]. Scatter plots explicitly assume uniqueness of representations (one dot per observation), yet the non-uniqueness of projections (B.1) suggests that this assumption does not have mathematical backing.

Another common use case is *latent space statistics*. For example, it is common to perform *clustering* of high-dimensional data by finding low-dimensional representations, which are then clustered (either during training or post hoc; see e.g. recent surveys [14]). This may take the form of k-means-style latent clustering [9], which assumes that representation averages are well-defined. Another example is *Bayesian optimization* [15, 21] over the latent representations. This assumes the ability to fit stochastic processes to the latent representations. Both of these examples rely on the ability to perform statistical calculations with respect to the learned representations. Unfortunately, practically all statistical calculations rely on the assumption that observation representations are unique. For example, the average of a set of observations with non-unique representations is ill-defined; see e.g. the celebrated work of [4] for an excellent discussion of this issue.

The above examples of assuming unique representations are ever-present throughout the literature, yet the mathematical justification is lacking. We investigate methods for ensuring uniqueness, but one could alternatively fully embrace the lack of uniqueness. The latent representation of a single observation would in this case form a set rather than a vector. We do not investigate this direction but note that working with sets is feasible [22] albeit somewhat more complicated than vectorial representations.

In this paper we investigate the *reach* of the manifold $\mathcal{M}$ spanned by the decoder f. This concept, predominantly studied in geometric measure theory, informs us about regions of observation space where the projection onto $\mathcal{M}$ is unique, such that trustworthy unique representations exist. If training data resides inside this region we may have hope that a suitable encoder can be estimated, leading to trustworthy representations. The classic *reach* construction is global in nature, so we develop a local generalization that gives a more fine-grained estimate of the uniqueness of a specific representation. We provide a new local, numerical, estimator of this reach, which allows us to determine which observations can be expected to have unique representations, thereby allowing investigations of the latent space to disregard observations with non-unique representations. Empirically we find that in large autoencoders, practically all data is outside the reach and risk not having a unique representation. To counter this, we design a reach-based regularizer that penalizes decoders for which unique representations of given data do not exist. Empirically, this significantly improves the guaranteed uniqueness of representations with only a small penalty in reconstruction error.

B.2 Reach and uniqueness of representation

Our starting question is *which observations x have a unique representation z for a given decoder f?* To answer this, we first introduce the *reach* [7] of the manifold spanned by decoder f. This is a *global* scalar that quantifies how far points can deviate from the manifold while having a unique projection. Secondly, we contribute a generalization of this classic geometric construct to characterize the local uniqueness properties of the learned representation.

Defining reach

The nearest point projection $\text{proj}_{\mathcal{M}}$ (Eq. B.1) is a well-defined *function*[1] on all points for which there exists a unique nearest point. We denote this set

$$\text{Unp}(\mathcal{M}) = \{x \in \mathbb{R}^D : x \text{ has a unique nearest point in } \mathcal{M}\},$$

where $\mathcal{M} = f(\mathbb{R}^d)$ is the manifold spanned by mapping the entire latent space through the decoder. Observations that lie within $\text{Unp}(\mathcal{M})$ are certain to have a unique optimal representation, but there is no guarantee that the encoder will recover this. With the objective of characterizing the uniqueness of representation, the set $\text{Unp}(\mathcal{M})$ is a good starting point as here the encoder at least has a chance of finding a representation similar to that of a projection. However, for an arbitrary manifold $\mathcal{M}$ it is generally not possible to explicitly find the set $\text{Unp}(\mathcal{M})$. Introduced by [7], the *reach* of $\mathcal{M}$ provides us with an implicit way to understand which points are in and outside $\text{Unp}(\mathcal{M})$.

Definition B.1. *The* global reach *of a manifold $\mathcal{M}$ is*

$$\text{reach}(\mathcal{M}) = \inf_{x \in \mathcal{M}} r_{max}(x),$$

where

$$r_{max}(x) = \sup\{r > 0 : B_r(x) \subset \text{Unp}(\mathcal{M})\}.$$

Here B_r denotes the open ball of radius r.

Hence, $\text{reach}(\mathcal{M})$ is the greatest radius r such that any open r-ball centered on the manifold lies in $\text{Unp}(\mathcal{M})$. In the existing literature, the *global reach* is referred to as the *reach*; we emphasize the global nature of this quantity as we will later develop local counterparts.

Definition B.1 does not immediately lend itself to computation. Fortunately, [7] provides a step in this direction, through the following result.

Theorem B.2 ([7]). *Suppose $\mathcal{M}$ is a manifold, then*

$$\text{reach}(\mathcal{M}) = \inf_{\substack{x,y \in \mathcal{M} \\ y-x \notin T_x\mathcal{M}}} \frac{\|x-y\|^2}{2\|P_{N_x\mathcal{M}}(y-x)\|},$$

where $P_{N_x\mathcal{M}}$ is the orthogonal projection onto the normal space of $\mathcal{M}$ at x. If $y - x \in T_x\mathcal{M}$ for all pairs $x, y \in \mathcal{M}$ we let $\text{reach}(\mathcal{M}) = \infty$, as $\mathcal{M}$ will be flat and the projection is unique everywhere.

For our objective of understanding which observations have a unique representation, i.e. are inside $\text{Unp}(\mathcal{M})$, the global reach provides some information. Specifically, the set

$$\mathcal{M}_r = \left\{ x \,\middle|\, \inf_{y \in \mathcal{M}} \|y - x\| < \text{reach}(\mathcal{M}) \right\}$$

[1] We here stress that a function always returns a single output for a given input.

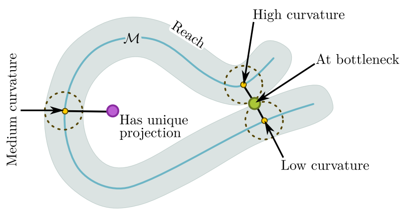

Figure 9: The global reach defines a region around the manifold $\mathcal{M}$ consisting of all points below a certain distance to $\mathcal{M}$. This captures both local manifold curvature as well as global shape.

is a subset of $\text{Unp}(\mathcal{M})$. This implies that observations x that are inside $\mathcal{M}_r$ will have a unique projection, such that we can expect the representation to be unique. The downside is that since $\text{reach}(\mathcal{M})$ is a global quantity, $\mathcal{M}_r$ is an overly restrictive small subset of $\text{Unp}(\mathcal{M})$. Fig. 9 illustrates this issue. Note how the global reach in the example is determined by the *bottleneck*[2] of the manifold. Even if this bottleneck only influences the uniqueness of projections of a single point, it determines the global reach of the entire manifold. This implies that many points exist outside the reach which nonetheless has a unique projection.

Pointwise normal reach

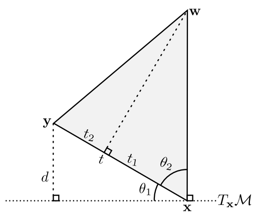

Figure 10: Notation for the proof of theorem B.4.

In order to get a more informative notion of reach, we develop a local version, which we, for reasons that will be clear, call the *pointwise normal reach*. For ease of notation denote

$$R(x,y) = \frac{\|x-y\|^2}{2\|P_{N_x\mathcal{M}}(y-x)\|}$$

for x, y with $y - x \notin T_x\mathcal{M}$, else we let $R(x,y) = \infty$. We then define the pointwise normal reach as the local infimum of eq. B.2.

Definition B.3 (Pointwise normal reach). *At a point $x \in \mathcal{M}$, the pointwise normal reach is*

$$r_N(x) = \inf_{y \in M} R(x,y).$$

[2]Not to be confused with *bottleneck network architectures* or the *information bottleneck*.

In theorem B.4 below we prove that the local estimate $r_N(x)$ describes how far we can move along a normal vector at x and still stay within $\mathrm{Unp}(\mathcal{M})$. This is useful as we know that x will lie in the normal space of $\mathcal{M}$ at $\mathrm{proj}_{\mathcal{M}}(x)$ ([7] Thm. 4.8).

Theorem B.4. *For all $x \in \mathcal{M}$*

$$B_{r_N(x)}(x) \cap N_x\mathcal{M} \subset \mathrm{Unp}(\mathcal{M}),$$

where $N_x\mathcal{M}$ denotes the normal space at x.

Proof. Suppose for the sake of contradiction that there exists $w \in \left(B_{r_N(x)}(x) \cap N_x\mathcal{M}\right) \cap \mathrm{Unp}(\mathcal{M})^c$. That is, there exists $y_1, y_2 \in \mathcal{M}$ such that

$$d(w, \mathcal{M}) = \|y_1 - w\| = \|y_2 - w\|,$$

where $d(w, \mathcal{M}) = \inf_{x \in \mathcal{M}} \|x - w\|$. In particular, we know there exists $y \in \mathcal{M}$ such that

$$\|y - w\| \leq \|x - w\| < r_N(x).$$

Now, let θ_1 denote the (acute) angle between $T_x\mathcal{M}$ and $y - x$, and let θ_2 denote the angle between $y - x$ and $w - x$. The sum of θ_1 and θ_2 is a right angle, see Fig. 10. Let t be the distance from x to y. The altitude through the vertex w divides $y - x$ into two line segments. Denote the length of the segment from the foot of the altitude to x, t_1, and the length of the segment from the foot to y, t_2. Note, t_2 will always be less or equal to t_1, as $\|y - w\| \leq \|x - w\|$.

By the definition of cosine, $\cos\theta_2 = \frac{t_1}{\|w-x\|} \geq \frac{t/2}{\|w-x\|}$. At the same time $\cos\theta_2 = \cos(\pi/2 - \theta_1) = \sin\theta_1 = \frac{d}{t}$, where $d = \|P_{w-x}(y-x)\|$, and as $w - x \in N_x\mathcal{M}$, $d \leq \|P_{N_x\mathcal{M}}(y - x)\|$. Thus, we have $\frac{t/2}{\|w-x\|} < \frac{d}{t} < \frac{\|P_{N_x\mathcal{M}}(y-x)\|}{t}$, implying $R(x, y) \leq \|w - x\|$, which contradicts $r_N(x) \leq R(x, y)$. □

In lemma B.6 below we show that the pointwise normal reach bounds the reach. For this, we need theorem 4.8(7) from [7]

Lemma B.5 ([7]). *Let x, y be points on $\mathcal{M}$ with $r_{max}(x) > 0$, and let $\mathbf{n}$ be a normal vector in $N_x\mathcal{M}$, then*

$$\langle \mathbf{n}, y - x \rangle \leq \frac{\|y - x\|^2 \|\mathbf{n}\|}{2 r_{max}(x)}$$

Lemma B.6. *For all $x \in \mathcal{M}$ we have that*

$$\inf_{y \in B_{2r_N(x)}(x) \cap \mathcal{M}} r_N(y) \leq r_{max}(x) \leq r_N(x).$$

Proof. Applying the result from Federer to the vector $\mathbf{n} = \frac{P_{N_x\mathcal{M}}(y-x)}{\|P_{N_x\mathcal{M}}(y-x)\|}$ gives

$$r_{\max}(x) \leq \frac{\|x - y\|^2 \|\mathbf{n}\|}{2\|\mathbf{n}\| \|x - y\| \cos\theta},$$

where θ is the angle between $x - y$ and $\mathbf{n}$. Hence, $\cos\theta = \frac{\|P_{N_x\mathcal{M}}(y-x)\|}{\|y-x\|}$. Thus, for all $y \in \mathcal{M}$

$$r_{\max}(x) \leq \frac{\|x - y\|^2}{2\|P_{N_x\mathcal{M}}(y - x)\|},$$

proving the right inequality. Consider $B = B_{r_N(x)}(x)$. Suppose there exists $w \in B$ with $w \notin \text{Unp}(\mathcal{M})$. Then $w \notin N_x\mathcal{M}$. Hence there exists $y_1, y_2 \in \mathcal{M}$ such that $d(w, \mathcal{M}) = \|y_1 - w\| = \|y_2 - w\| < r_N(x)$. From [7] theorem 4.8 we know that $w \in N_{y_1}\mathcal{M}, N_{y_2}\mathcal{M}$. Combining this with lemma B.4 gives that $r_N(y_1), r_N(y_2) \leq d(w, \mathcal{M})$ and that $d(x, \mathcal{M}) < \|w - x\|$. We also have that $\|y_1 - x\|, \|y_2 - x\| \leq 2r_N(x)$. Combining these inequalities gives us that the distance from x to any point not in Unp($\mathcal{M}$) is greater than $\inf_{y \in B_{2r_N(x)}(x) \cap \mathcal{M}} r_N(y)$, which implies that

$$\inf_{y \in B_{2r_N(x)}(x) \cap \mathcal{M}} r_N(y) \leq r_{\max}(x).$$

□

We presented the theoretical analysis under the assumption that $\mathcal{M} = f(\mathbb{R}^d)$ is a manifold. Although the theoretical results can be extended to arbitrary subsets of Euclidean space, the experimental setup requires the Jacobian to span the entire tangent space. This might not be the case if $\mathcal{M}$ has self-intersections. The theory can be extended to handle such self-intersections, but this significantly complicates the algorithmic development. See the appendix for a discussion.

Estimating the pointwise normal reach

The definition of r_N, prompts us to minimize $R(x, y)$ over all of $\mathcal{M}$, which is generally infeasible and approximations are in order. As a first step towards an estimator, assume that we are given a finite sample $\mathbf{S}$ of points on the manifold. We can then replace the infimum in definition B.3 with a minimization over the samples. Using that the projection matrix onto $N_x\mathcal{M}$ is given by $P_{N_x\mathcal{M}} = \mathbf{I} - \mathbf{J}(\mathbf{J}^\top\mathbf{J})^{-1}\mathbf{J}^\top$, we get the following estimator

$$\hat{r}_N(x) = \min_{y \in \mathbf{S}} \frac{\|y - x\|^2}{2\|(\mathbf{I} - \mathbf{J}(\mathbf{J}^\top\mathbf{J})^{-1}\mathbf{J}^T)(y - x)\|}, \tag{B.1}$$

where $\mathbf{J} \in \mathbb{R}^{D \times d}$ is the Jacobian matrix of f at x. Note that since we replace the infimum with a minimization over a finite set, we have that $\hat{r}_N(x) \geq r_N(x)$.

There are different choices of sampling sets $\mathbf{S}$. Given a trained autoencoder, a cheap way to obtain samples is to use the reconstructed training data as the sampling set. This will generally be sufficient if the training data is dense on the manifold, but this is rarely the case in high data dimensions. The following lemma provides us a way to restrict the area over which we must minimize.

Lemma B.7. *For any $x, y \in \mathcal{M}$*

$$R(x, y) \geq \frac{1}{2}\|x - y\|.$$

Proof. Recall that $y - x = P_{N_x\mathcal{M}}(y - x) + P_{T_x\mathcal{M}}(y - x)$, as $\mathbb{R}^D = T_x\mathcal{M} \oplus N_x\mathcal{M}$. Hence $\|y - x\| \geq \|P_{N_x\mathcal{M}}(y - x)\|$. The statement, thus, follows from the definition of R. □

Algorithm 1 Sampling-based reach estimator

```
radius ← r_0
reach ← ∞
for i ← 1,..., num_batches do
    samples ← sample_ball(x, radius, batch_size)
    projected ← decode(encode(samples))
    reach ← min(reach, reach_est(x, projected))
    radius ← 2· reach
end for
```

The lemma points towards a simple computational procedure for numerically estimating the pointwise normal reach, which is explicated in algorithm 1. Here `reach_est` refers to the application of eq. B.1. The algorithm samples uniformly inside a ball centered on x and repeatedly shrinks the radius of the ball as tighter estimates of the reach are recovered. We further use the autoencoding reconstruction as an approximation to the projection of x onto $\mathcal{M}$.

Is a point within reach?

Suppose that a point $x \in \mathbb{R}^D$ is represented by a point on the manifold $f(z)$. From definition B.1 we know that x has a unique nearest point on the manifold if

$$\|x - f(z)\| < r_{\max}(f(z)).$$

A point x which does not satisfy this inequality risks not having a unique nearest point, and hence no unique representation. From lemma B.6 we know that $r_{\max}(f(z)) \leq r_N(f(z))$. So x risks not having a unique nearest point if

$$\|x - f(z)\| \geq r_N(f(z)) \geq r_{\max}(f(z)).$$

We note that to show that $\|x - f(z)\| \geq r_N(f(z))$, it is enough to compute

$$\hat{r}_N(f(z)) = \inf_{\substack{y \in \mathcal{M} \cap B_{2\|x-f(z)\|}(f(z)) \\ y \neq f(z)}} R(f(z), y),$$

i.e. limit the search to a ball of radius $2\|x - f(z)\|$. Thus, when we only need to determine if a point is inside the pointwise normal reach, we can pick $r_0 = 2\|x - f(z)\|$ in Algorithm 1.

Notice that given any set of points on the manifold, the resulting estimation of r_N will always be larger than the true value. It means that any point which lies outside the estimated normal reach, will in fact lie outside the true normal reach. However, a point which lies inside the estimated normal reach, risks lying outside the true normal reach, and thus not having a unique projection.

Regularizing for reach

The autoencoder minimizes an l_2 error which is directly comparable to the pointwise normal reach. This suggests a regularizer that penalizes if the l_2 error is larger than the pointwise normal reach. In practice, we propose to use

$$\mathcal{R}(x) = \texttt{Softplus}\left(\|f(g(x)) - x\| - \hat{r}_N(f(g(x)))\right).$$

The reach-regularized decoder then minimizes

$$\mathcal{L} = \sum_{n=1}^{N} \|f(g(x_n)) - x_n\|^2 + \lambda \sum_{n=1}^{N} \mathcal{R}(x_n),$$

while we do not regularize the encoder. We also experimented with a ReLU activation instead of Softplus, but found the latter to yield more stable training. When estimating the pointwise normal reach, $\hat{r}_N$, we apply Algorithm 1 with an initial radius of $r_0 = 2\|f(g(x_n)) - x_n\|$.

B.3 Experiments

Having established a theory and algorithm for determining when a representation can be expected to be unique, we next investigate its use empirically. We first compute the pointwise local reach across a selection of models to see if it provides useful information. We then carry on to investigate the use of reach regularization. [3]

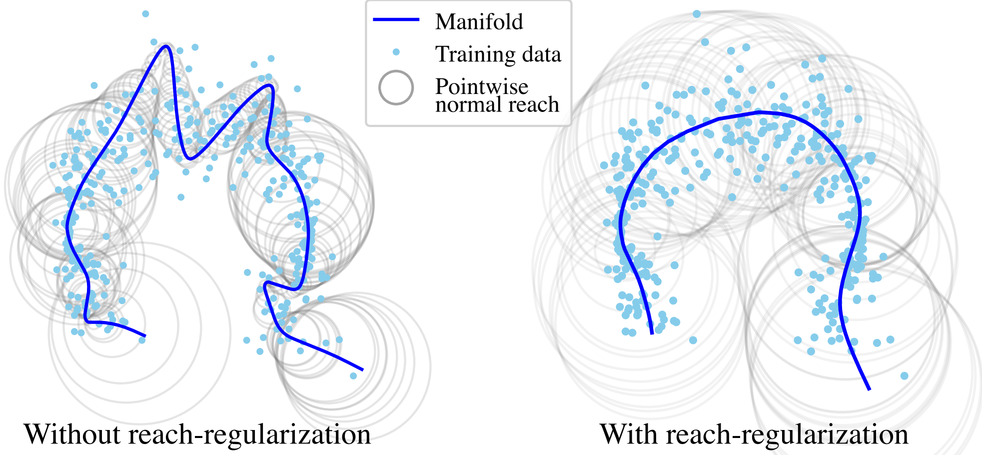

Figure 11: *Left:* An autoencoder trained on noisy points scattered along a circular arc. *Right:* The manifold spanned by the decoder of an autoencoder trained with reach regularization. In both panels, the gray circles illustrates the estimated pointwise normal reach at points along the autoencoder curve.

Analysing reach

Toy circle

We start our investigations with a simple toy example to get an intuitive understanding. We generate observations along a circular arc with added Gaussian noise

[3] The code is available at https://github.com/HeleneHauschultz/is_an_encoder_within_reach.

Figure 12: CelebA validation set reconstructions.

of varying magnitude. Specifically, we generate approximately 400 points as $z \mapsto t(\sin(z), -\cos(z)) + 1.5\cos(z)\epsilon$, where $\epsilon \sim \mathcal{N}(0,1)$. On this, we train an autoencoder with a one-dimensional latent space. The encoder and decoder both consist of linear layers, with three hidden layers with 128 nodes and with ELU non-linearities.

Figure 11(left) shows the data alongside the estimated manifold and its pointwise normal reach. We observe that the manifold spanned by the decoder has areas with small reach, where the manifold curves to fit the noisy data. The pointwise normal reach seems to well-reflect the curvature of the estimated manifold. The plot illustrates how some of the points end up further away from the manifold than the reach. For some of the points, this is not a problem, as they still have a unique projection onto the manifold. However, some of the points are equally close to different points on the manifold, such that their representation cannot be trusted.

CelebA

To investigate the reach on a non-toy dataset, we train a deep autoencoder on the CelebA face dataset [13]. The dataset consists of approximately 200 000 images of celebrity faces.

We train a symmetric encoder-decoder pair that maps the $64 \times 64 \times 3$ images to a 128 dimensional latent space, and back. The encoder consists of a single 2d convolution operation without stride followed by six convolution operations with stride 2, resulting in a $1 \times 1 \times C$ image. We use $C = 128$ channels for all convolutional operations, a filter size of 5 and Exponential Linear Unit (ELU) non-linearities. The decoder is symmetric, using transpose convolutions with stride 2 to upsample and ending with a convolution operation mapping to $64 \times 64 \times 3$. The model is trained for 1M gradient updates on the mean square error loss, with a batch size of 128, using the Adam optimizer with a learning rate of 10^{-4}. Example reconstructions on the validation set are provided in Fig. 12.

After training we estimate the reach of the validation set using the sampling based approach (Alg. 1). Fig. 13(left) plots the reconstruction error $\|x - f(z)\|$ versus the pointwise normal reach. We observe that almost all observations lie outside the pointwise normal reach, implying that we cannot guarantee a unique representation. This is a warning sign that our representations need not be trustworthy.

Next we analyze the empirical convergence properties of our estimator on the CelebA autoencoder. Fig. 13(center) shows the average pointwise normal reach over the validation set as a function of the number of iterations in the sampling based esti-

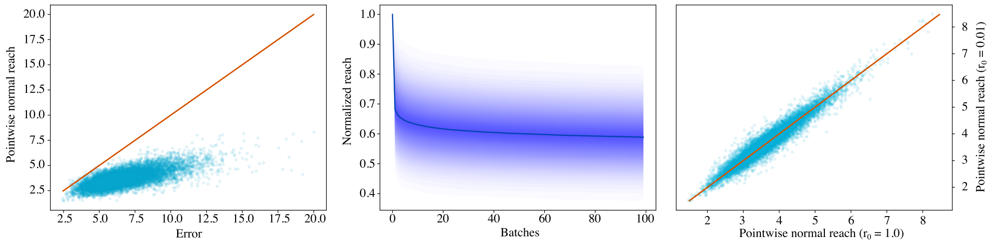

Figure 13: *Left:* Estimated reach for CelebA validation samples plotted against the L2 error. Samples below the diagonal red line does not have a unique encoder. *Center:* Normalized reach as a function of batches used to estimate the reach. The normalized reach is the estimated pointwise normal reach divided by the estimated pointwise normal reach after the first batch. The hyperball sampling reach estimator quickly converges. *Right:* Sensitivity analysis of the hyperball sampling reach estimator to the initial hyperball radius. The reach of CelebA validation samples are estimated with initial radii $r_0 = 1.0$ and $r_0 = 0.01$ respectively and their final reach after 100 batches are plotted against each other.

mator. We observe that the estimator converges after just a few iterations, suggesting that the estimator is practical.

The estimator relies on an initial radius for its search. Fig. 13(right) shows the estimated pointwise normal reach on the validation set, plotted for two different initial radii. We observe that the estimator converges to approximately the same value in both cases, suggesting that the method is not sensitive to this initial radius. However, initializing with a tight radius will allow for faster convergence.

Reach regularization

Having established that the pointwise normal reach provides a meaningful measure of uniqueness, we carry on to regularize accordingly.

Toy circle

Returning to the example from section B.3, we train an autoencoder of the same architecture with the reach regularization. We pretrain the network 100 epochs without regularization, and then 2000 iterations with reach regularization.

Fig. 11 (right) shows that reach regularization gives a significantly smoother manifold than without regularization (left panel). The gray circles on the plot indicate that almost all the points are now within the pointwise normal reach, and arguably the associated representations are now more trustworthy.

MNIST

Next we train an autoencoder on 5000 randomly chosen images from the classes 2, 4 and 8 from MNIST [12]. We use a symmetric architecture reducing to two dimensional representation through a sequence of $784 \to 500 \to 250 \to 150 \to 100 \to 50 \to 10$ linear layers with ELU non-linearities. We pretrain 5000 epochs without any regularization, and proceed with reach regularization enabled. Figure 14 (left) shows the percentage of points which lies within reach of the estimated manifold. We observe that reach regularization slightly increases the reconstruction error (see example reconstructions in fig. 15), as any regularization would, while significantly increasing the percentage of points that are known to have a unique representation. This suggests that reach regularization only minimally changes reconstructions while giving a significantly more smooth model, which is more reliable.

Figure 14 (center) shows the latent representations given by the pretrained autoencoder without regularization, while fig. 14 (right) shows the latent representations after an additional 200 epochs with reach regularization. The latent representations with corresponding data points outside reach, that is, where the reconstruction error is greater than the pointwise normal reach at the reconstructed point is plotted in red. The points inside reach are plotted in green. We observe that after regularization, significantly more points can be expected to be unique and thereby trustworthy. Note that the latent configuration is only changed slightly after reach regularization, which suggests that the expressive power of the model is largely unaffected by the reach regularization.

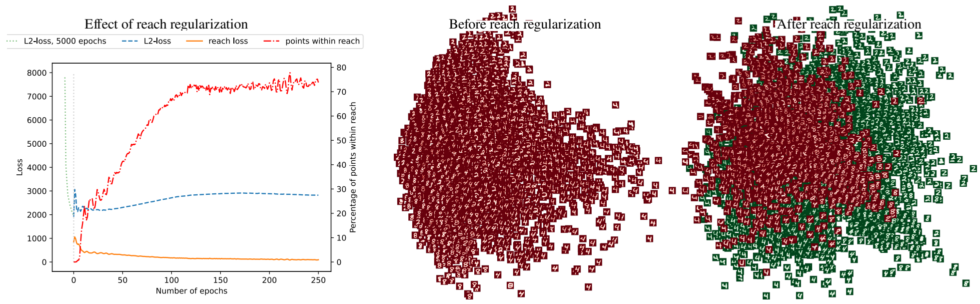

Figure 14: The effect of reach regularization on an MNIST model. *Left:* The plot shows that the percentage of points within reach increases, while the l_2-loss is nearly unchanged. We plot the loss curve from the initial 5000 epochs without regularization, to show how the l_2-loss behaves when regularizing. *Center & right:* Latent representations of the MNIST autoencoder before and after the reach regularization (visualized using PCA). The red numbers are outside the reach, while the green are within. Reach regularization smoothens the decoder to increase reach with minimal changes to both reconstructions and latent configuration.

Figure 15: Reconstructions of MNIST images before (top) and after (bottom) reach regularization. The reach regularization only minimally reduces reconstruction quality, while significantly improving upon representation uniqueness.

B.4 Related work

Representation learning is a foundational aspect of current machine learning, and the discussion paper by [2] is an excellent starting point. As is common, [2] defines a representation as the output of a function applied to an observation, implying that a representation is unique. In the specific context of autoencoders, we question this implicit assumption of uniqueness as many equally good representations may exist for a given observation. While only studied here for autoencoders, the issue applies more generally when representations span submanifolds of the observation space.

In principle, probabilistic models may place multimodal distributions over the representation of an observation in order to reflect lack of uniqueness. In practice, this rarely happens. For example, the highly influential *variational autoencoder* [11, 16] amortizes the estimation of $p(z|x)$ such that it is parametrized by the output of a function. Alternatives relying on Monte Carlo estimates of $p(z|x)$ do allow for capturing non-uniqueness [10], but this is rarely done in practical implementations. That Monte Carlo estimates provide state-of-the-art performance is perhaps indicative that coping with non-unique representations is important. Our approach, instead, aim to determine which observations can be expected to have a unique representation, which is arguably simpler than actually finding the multiple representations.

Our approach relies on the reach of the manifold spanned by the decoder. This quantity is traditionally studied in geometric measure theory as the reach is informative of many properties of a given manifold. For example, manifolds which satisfy that $\text{reach}(\mathcal{M}) > 0$ are $C^{1,1}$, i.e. the transition functions are differentiable with Lipschitz continuous derivatives. In machine learning, the reach is, however, a rarely used concept. [8] investigates if a manifold of a given reach can be fitted to observed data, and develops the associated statistical test. Further notable exceptions are the multichart autoencoder by [20], and the adaptive clustering of [3]. Both works rely on the reach as a tool of derivation. Similarly, [5] relies on the assumption of positive reach when deriving properties of deep generative models. These works all rely on the global reach, while we have introduced a local generalization.

The work closest to ours appears to be that of [1] which studies the convergence of an estimator of the global reach (B.2). This only provides limited insights into the uniqueness of a representation as the global reach only carries limited information about the local properties of the studied manifold. We therefore introduced the pointwise normal reach alongside an estimator thereof. This gives more precise information about which observations can be expected to have a unique representation.

B.5 Discussion

The overarching question driving this paper is *when can representations be expected to be unique?* Though commonly assumed, there is little mathematical reason to believe that the choice of optimal representation is generally unique. The theoretical implication of this is that enforcing uniqueness on non-unique representations leads to untrustworthy representations.

We provide a partial answer for the question in the context of autoencoders, through the introduction of the *pointwise normal reach*. This provides an upper bound for a radius centered around each point on the manifold spanned by the decoder, such that any observation within the ball has a unique representation. This bound can be directly compared to the reconstruction error of the autoencoding to determine if a given observation might not have a unique representation. This is a step towards a systematic quantification of the reliability and trustworthiness of learned representations.

Empirically, we generally find that most trained models do not ensure that representations are unique. For example, on CelebA we found that almost no observations were within reach, suggesting that uniqueness was far from ensured. This is indicative that the problem of uniqueness is not purely an academic question, but one of practical importance.

We provide a sampling estimator of the pointwise normal reach, which is guaranteed to upper bound the true pointwise normal reach. The estimator is easy to implement, with the main difficulty being the need to access the Jacobian of the decoder. This is readily accessible using forward-mode automatic differentiation, but it can be memory-demanding for large models.

It is easy to see that the sample-based pointwise normal reach estimator converges to the correct value in the limit of infinitely many samples. We, however, have no results on the rate of convergence. In practice we observe that the estimator converges in a few iterations for most models, suggesting the convergence is relatively fast. In practice, the estimator, however, remains computationally expensive.

While we can estimate the pointwise normal reach quite reliably even for large models within manageable time, the estimator is currently too expensive to use for regularization of large models. On small models, we observe significant improvements in the uniqueness properties of the representations at minimal cost in terms of reconstruction error. This is a promising result and indicative that it may be well worth using this form of regularization. While more work is needed to speed up the estimating of pointwise normal reach, our work does pave a path to follow.

Acknowledgements

This work was supported by research grants (15334, 42062) from VILLUM FONDEN. This project has also received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 757360). This work was funded in part by the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606). Helene Hauschultz is partly financed by Aarhus University Centre for Digitalisation, Big Data and Data Analytics (DIGIT).

References

[1] Aamari, E., J. Kim, F. Chazal, B. Michel, A. Rinaldo and L. Wasserman (2019). Estimating the reach of a manifold. eng. *Electronic journal of statistics* 13(1), 1359–1399.

[2] Bengio, Y., A. Courville and P. Vincent (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35(8), 1798–1828.

[3] Besold, F. and V. Spokoiny (2020). *Adaptive Manifold Clustering*. arXiv: 1912.04869 [math.ST].

[4] Billera, L.J., S.P Holmes and K. Vogtmann (2001). Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics* 27(4), 733–767.

[5] Chae, M., D. Kim, Y. Kim and L. Lin (2021). *A likelihood approach to nonparametric estimation of a singular distribution using deep generative models*. arXiv: 2105.04046 [stat.ML].

[6] Detlefsen, N.S., S. Hauberg and W. Boomsma (2022). Learning meaningful representations of protein sequences. *Nature Communications* 13(1), 1–12.

[7] Federer, H. (1959). Curvature Measures. eng. *Transactions of the American Mathematical Society* 93(3), 418–491.

[8] Fefferman, C., S. Mitter and H. Narayanan (2016). Testing the manifold hypothesis. eng. *Journal of the American Mathematical Society* 29(4), 983–1049.

[9] Hadipour, H., C. Liu, R. Davis, S.T. Cardona and P. Hu (2022). Deep clustering of small molecules at large-scale via variational autoencoder embedding and K-means. *BMC bioinformatics* 23(4), 1–22.

[10] Hoffman, M.D. (June 2017). "Learning Deep Latent Gaussian Models with Markov Chain Monte Carlo". *Proceedings of the 34th International Conference on Machine Learning*. Ed. by D. Precup and Y.W. Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, 1510–1519.

[11] Kingma, D.P and M. Welling (2014). "Auto-Encoding Variational Bayes". *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*.

[12] LeCun, Y., L. Bottou, Y. Bengio and P. Haffner (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11), 2278–2324.

[13] Liu, Z., P. Luo, X. Wang and X. Tang (Dec. 2015). "Deep Learning Face Attributes in the Wild". *Proceedings of International Conference on Computer Vision (ICCV)*.

[14] Min, E., X. Guo, Q. Liu, G. Zhang, J. Cui and J. Long (2018). A survey of clustering with deep learning: From the perspective of network architecture. *IEEE Access* 6, 39501–39514.

[15] Močkus, J. (1975). "On Bayesian methods for seeking the extremum". *Optimization techniques IFIP technical conference*. Springer, 400–404.

[16] Rezende, D.J., S. Mohamed and D. Wierstra (2014). "Stochastic Backpropagation and Approximate Inference in Deep Generative Models". *Proceedings of the 31st International Conference on Machine Learning (ICML)*.

[17] Riesselman, A.J., J.B. Ingraham and D.S. Marks (2018). Deep generative models of genetic variation capture the effects of mutations. *Nature methods* 15(10), 816–822.

[18] Rumelhart, D.E., G.E. Hinton and R.J. Williams (1986). Learning representations by back-propagating errors. *nature* 323(6088), 533–536.

[19] Sattarov, B., I.I. Baskin, D. Horvath, G. Marcou, E.J. Bjerrum and A. Varnek (2019). De novo molecular design by combining deep autoencoder recurrent neural networks with generative topographic mapping. *Journal of chemical information and modeling* 59(3), 1182–1196.

[20] Schonsheck, S., J. Chen and R. Lai (2020). *Chart Auto-Encoders for Manifold Structured Data*. arXiv: 1912.10094 [cs.LG].

[21] Stanton, S., W. Maddox, N. Gruver, P. Maffettone, E. Delaney, P. Greenside and A. G. Wilson (2022). Accelerating Bayesian Optimization for Biological Sequence Design with Denoising Autoencoders. *arXiv preprint arXiv:2203.12742*.

[22] Zaheer, M., S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov and A. Smola (2018). *Deep Sets*. arXiv: 1703.06114 [cs.LG].

B.A Appendix

Extending the pointwise normal reach to the non-manifold setting

[7] introduces reach for arbitrary subsets of Euclidean space. In this situation $T_x\mathcal{M}$ and $N_x\mathcal{M}$ denote the tangent- and normal cone.

Definition B.8. *Let $\mathcal{M} \subset \mathbb{R}^D$ denote an arbitrary subset and let $x \in \mathcal{M}$. Then $v \in \mathbb{R}^D$ is a tangent vector for $\mathcal{M}$ at x if either $v = 0$ or if for every $\varepsilon > 0$ exists $y \in \mathcal{M}$ with*

$$0 < \|y - x\| < \varepsilon \quad \textit{and} \quad \|\frac{y-x}{\|y-x\|} - \frac{v}{\|v\|}\| < \varepsilon. \tag{B.2}$$

Let $T_x\mathcal{M}$ denote the set of tangent vectors for $\mathcal{M}$ at x. A vector $w \in \mathbb{R}^D$ is a normal vector for $\mathcal{M}$ at x if

$$\langle w, v\rangle \leq 0 \textit{ for all } v \in T_x\mathcal{M}. \tag{B.3}$$

Let $N_x\mathcal{M}$ denote the set of all normal vectors for $\mathcal{M}$ at x.

We can extend theorem B.4 and Lemma B.6 to the general situation as defined by Federer. To extend Theorem B.4 it is sufficient to prove that for any $v \in N_x\mathcal{M}$ and $u \in \mathbb{R}^D$, $\|P_v(u)\| \leq d(u, T_x\mathcal{M})$.

Lemma B.9. *For any $v \in N_x\mathcal{M}$ and $u \in \mathbb{R}^D$ with $\langle v, u\rangle \geq 0$, $\|P_v(u)\| \leq d(u, T_x\mathcal{M})$.*

Proof. For a subset $A \subset \mathbb{R}^D$, $\text{dual}(A) = \{v \in \mathbb{R}^D : \langle a, v\rangle \leq 0 \text{ for all } a \in A\}$. First we prove that $d(u, \text{dual}(v)) = \|P_v(u)\|$. Note that we can write $u = u_v + u_{v^\perp}$, where $u_v = P_v(u)$ and $u_{v^\perp} \in v^\perp$. Then

$$\begin{aligned} d(u, \text{dual}(v)) &= \inf_{w \in \text{dual}(v)} \|u - w\| = \inf_{w \in \text{dual}(v)} \|u_v + u_{v^\perp} - w\| \\ &= \inf_{w \in \text{dual}(v)} \|u_v\| + \|u_{v^\perp} - w\| - 2\langle u_v, w\rangle. \end{aligned} \tag{B.4}$$

As $\langle u_v, w\rangle \leq 0$, it follows that the infimum is achieved when $w = u_{v^\perp}$. By the definition of the dual it follows that $\text{dual}(v) \supset \text{dual}(N_x\mathcal{M}) \supset T_x\mathcal{M}$. Hence $d(u, T_x\mathcal{M}) = \inf_{w \in T_x\mathcal{M}}\|u - w\| \geq \inf_{w \in \text{dual}(v)}\|u - w\| = \|P_u(v)\|$. □

To extend lemma B.6 note that if $r_{max}(x) > 0$, then $T_x\mathcal{M}$ is convex [7, Thm 4.8 (12)]. Let $y \in \mathcal{M}$. If $y - x \in T_x\mathcal{M}$ then $R(x, y) = \infty$. Otherwise, as $T_x\mathcal{M}$ is a convex cone, there exists $n \in N_x\mathcal{M}$ such that $\langle n, y - x\rangle \geq 0$. In that case $\langle n, x - y\rangle = \|P_n(y - x)\|$, so applying Lemma 2.5 gives the result.

Though the theory can be extended to general subspaces, the manifold assumption is important for the experimental setup. An important assumption for the estimator (B.1) is that the Jacobian spans the entire tangent space. If this is not the case, this estimator does not estimate the pointwise normal reach. The reason being, the length of the projection onto the orthogonal complement of the Jacobian is not necessarily the distance to the tangent space. It is clear that when we want to study the uniqueness of latent representations, if the decoder is not injective, it automatically has areas without unique representations. So if the decoder is not injective, we should already be wary about trusting the latent representations.

Reach estimation in increasing ambient dimension

In the following experiment we want to see the behavior of the reach estimator when the dimension in which the manifold is embedded increases. We consider the graph $(x,y) \mapsto U_n(x,y,x^2+y^2,0,\dots,0)$, where $U_n \in O(n)$ is an orthogonal matrix. That is, we embed the quadratic surface (x,y,x^2+y^2) isometrically into $\mathbb{R}^n$. We then estimate the pointwise normal reach in $\mathbf{0}$ with one iteration of Algorithm 1 with an initial radius of 5 and a sample size of 10. We estimate the pointwise normal reach 100 times in each dimension and take the average of these. The true value of the pointwise normal reach is $r_N(\mathbf{0}) = 0.5$. Figure B.A shows how the average overestimation of the

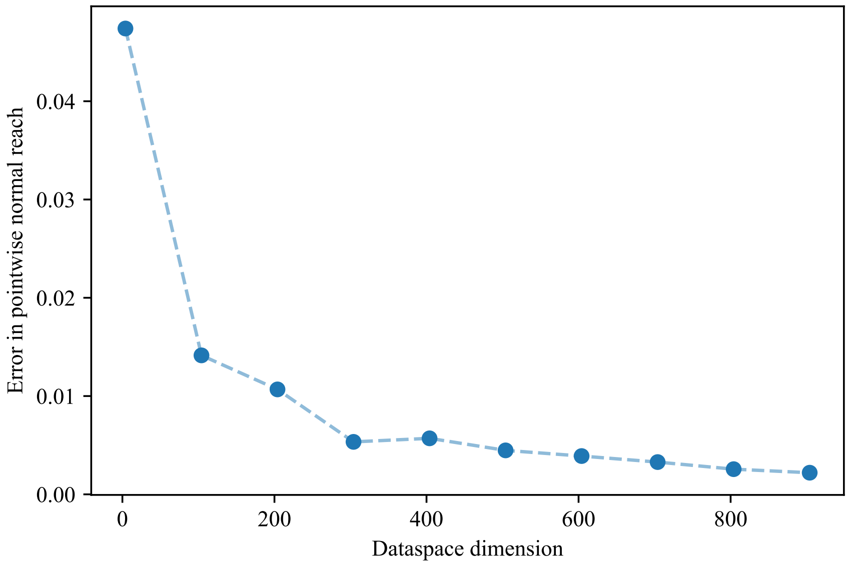

Figure 16: The plot shows how the average overestimation of the pointwise normal reach of an 2-dimensional quadratic surfaces isometrically embedded into a higher dimensional space goes down as the ambient dimension goes up.

pointwise normal reach goes down as the ambient dimension goes up.

Reconstruction error in test set during reach regularization

We extend the experiment from Section B.3 where we perform the reach regularization on an autoencoder trained on a subset of the MNIST data. At each iteration we compute the reconstruction error of a test set. We see that the test error is similar to the training error, suggesting that the model generalizes well to the data. This implies that a model having data points outside reach does not determine that the model does not generalize well to the data. Furthermore, reach regularization does necessarily impact the generalization of the model.

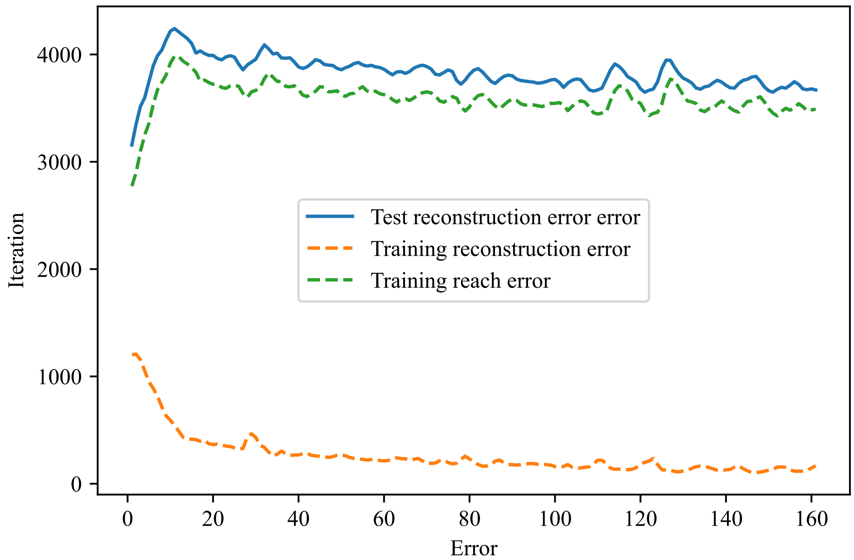

Figure 17: The plot shows 160 iterations of reach regularization of the autoencoder trained on the MNIST dataset, as in Section B.3. The blue line shows the average reconstruction error on the test set, the green line shows the average reconstruction error on the training set, and the orange line shows the reach loss.