

---

# Is an encoder within reach?

---

**Helene Hauschultz**

Department of Mathematics,  
Aarhus University  
hhauschultz@math.au.dk

**Rasmus Berg Palm**

rasmusbergpalm@gmail.com

**Pablo Moreno-Muñoz**

Technical University of  
Denmark, DTU Compute  
pabmo@dtu.dk

**Nicki Skafte Detlefsen**

Technical University of  
Denmark, DTU Compute  
nsde@dtu.dk

**Andrew Allan du Plessis**

Department of Mathematics,  
Aarhus University  
matadp@math.au.dk

**Søren Hauberg**

Technical University of  
Denmark, DTU Compute  
sohau@dtu.dk

## 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.

## 1 Encoders as projectors

A good *learned representation* has many desiderata (Bengio et al., 2013). The perhaps most elementary constraint placed over most learned representations is that a given observation  $\mathbf{x}$  should have a *unique* representation  $\mathbf{z}$ , at least in distribution. In practice this is ensured by letting the representation be given by the output of a function,  $\mathbf{z} = g(\mathbf{x})$ , often represented with a neural network. Having a unique representation for a given observation is important for downstream tasks, as otherwise the downstream loss function is complicated by having to be defined on sets (Zaheer et al., 2018). Uniqueness of representation is also important when a human investigator seeks information about the phenomenon underlying data, e.g. through visualizations, as uniqueness is a ubiquitous assumption.

The *autoencoder* (Rumelhart et al., 1986) 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*  $g_\psi : \mathbb{R}^D \rightarrow \mathbb{R}^d$  and a *decoder*  $f_\phi : \mathbb{R}^d \rightarrow \mathbb{R}^D$ , parametrized by  $\psi$  and  $\phi$ , respectively. These are trained by minimizing the *reconstruction error* of the training data  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ ,

$$\psi^*, \phi^* = \operatorname{argmin}_{\psi, \phi} \sum_{n=1}^N \|f(g(\mathbf{x}_n)) - \mathbf{x}_n\|^2. \quad (1)$$

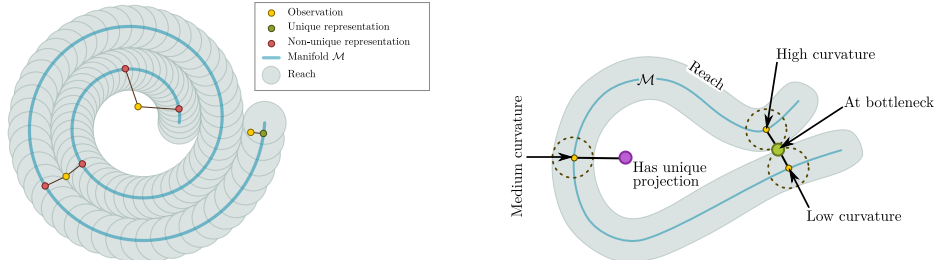


Figure 1: *Left*: The projection of a point (yellow) onto a nonlinear manifold can take unique (green) or multiple values (red) depending on the reach of the manifold. When training data is inside the reach, the encoder can match the projection resulting in more trustworthy representations. *Right*: 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.

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.

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

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

For any *nonlinear* choice of decoder  $f$ , this optimal encoder does *not* exist 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. 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 as good.

**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.

## 2 Reach and uniqueness of representation

Our starting question is *which observations  $\mathbf{x}$  have a unique representation  $\mathbf{z}$  for a given decoder  $f$ ?* To answer this, we first introduce the *reach* (Federer, 1959) of the manifold spanned by decoder  $f$ . This is a *global* scalar that quantifies how much 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.

### 2.1 Defining reach

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

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

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 that is 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 Federer (1959), the *reach* of  $\mathcal{M}$  provides us with an implicit way to understand which points are in and outside  $\text{Unp}(\mathcal{M})$ .

**Definition 2.1.** The *global reach* of a manifold  $\mathcal{M}$  is

$$\text{reach}(\mathcal{M}) = \inf_{\mathbf{x} \in \mathcal{M}} r_{\max}(\mathbf{x}), \quad (3)$$

where

$$r_{\max}(\mathbf{x}) = \sup\{r > 0 : B_r(\mathbf{x}) \subset \text{Unp}(\mathcal{M})\}. \quad (4)$$

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 2.1 does not immediately lend itself to computation. Fortunately, Federer (1959) provides a step in this direction, through the following result.

**Theorem 2.2** (Federer (1959)). *Suppose  $\mathcal{M}$  is a manifold, then*

$$\text{reach}(\mathcal{M}) = \inf_{\substack{\mathbf{x}, \mathbf{y} \in \mathcal{M} \\ \mathbf{y} - \mathbf{x} \notin T_{\mathbf{x}}\mathcal{M}}} \frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\|P_{N_{\mathbf{x}}\mathcal{M}}(\mathbf{y} - \mathbf{x})\|}, \quad (5)$$

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

$$\mathcal{M}_r = \left\{ \mathbf{x} \mid \inf_{\mathbf{y} \in \mathcal{M}} \|\mathbf{y} - \mathbf{x}\| < \text{reach}(\mathcal{M}) \right\} \quad (6)$$

is a subset of  $\text{Unp}(\mathcal{M})$ . This implies that observations  $\mathbf{x}$  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. 1(right) illustrates this issue. Note how the global reach in the example is determined by the *bottleneck*<sup>1</sup> 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.

## 2.2 Pointwise normal reach

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

$$R(\mathbf{x}, \mathbf{y}) = \frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\|P_{N_{\mathbf{x}}\mathcal{M}}(\mathbf{y} - \mathbf{x})\|} \quad (7)$$

for  $\mathbf{x}, \mathbf{y}$  with  $\mathbf{y} - \mathbf{x} \notin T_{\mathbf{x}}\mathcal{M}$ , else we let  $R(\mathbf{x}, \mathbf{y}) = \infty$ . We then define the pointwise normal reach as the local infimum of eq. 5.

**Definition 2.3** (Pointwise normal reach). At a point  $\mathbf{x} \in \mathcal{M}$ , the pointwise normal reach is

$$r_N(\mathbf{x}) = \inf_{\mathbf{y} \in \mathcal{M}} R(\mathbf{x}, \mathbf{y}). \quad (8)$$

In theorem 2.4 below we prove that the local estimate  $r_N(\mathbf{x})$  describes how far we can move along a normal vector at  $\mathbf{x}$  and still stay within  $\text{Unp}(\mathcal{M})$ . This is useful as we know that  $\mathbf{x}$  will lie in the normal space of  $\mathcal{M}$  at  $\text{proj}_{\mathcal{M}}(\mathbf{x})$  (Federer (1959) Thm. 4.8).




Figure 2: Notation for the proof of theorem 2.4.

<sup>1</sup>Not to be confused with *bottleneck network architectures* or the *information bottleneck*.

**Theorem 2.4.** For all  $x \in \mathcal{M}$

$$B_{r_N(\mathbf{x})}(\mathbf{x}) \cap N_{\mathbf{x}}\mathcal{M} \subset \text{Unp}(\mathcal{M}), \quad (9)$$

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

*Proof.* Suppose for the sake of contradiction that there exists  $\mathbf{w} \in (B_{r_N(\mathbf{x})}(\mathbf{x}) \cap N_{\mathbf{x}}\mathcal{M}) \cap \text{Unp}(\mathcal{M})^c$ . That is, there exists  $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{M}$  such that

$$d(\mathbf{w}, \mathcal{M}) = \|\mathbf{y}_1 - \mathbf{w}\| = \|\mathbf{y}_2 - \mathbf{w}\|, \quad (10)$$

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

$$\|\mathbf{y} - \mathbf{w}\| \leq \|\mathbf{x} - \mathbf{w}\| < r_N(\mathbf{x}). \quad (11)$$

Now, let  $\theta_1$  denote the (acute) angle between  $T_{\mathbf{x}}\mathcal{M}$  and  $\mathbf{y} - \mathbf{x}$ , and let  $\theta_2$  denote the angle between  $\mathbf{y} - \mathbf{x}$  and  $\mathbf{w} - \mathbf{x}$ . The sum of  $\theta_1$  and  $\theta_2$  is a right angle, see Fig. 2. Let  $t$  be the distance from  $\mathbf{x}$  to  $\mathbf{y}$ . The altitude through the vertex  $\mathbf{w}$  divides  $\mathbf{y} - \mathbf{x}$  into two line segments. Denote the length of the segment from the foot of the altitude to  $\mathbf{x}$ ,  $t_1$ , and the length of the segment from the foot to  $\mathbf{y}$ ,  $t_2$ . Note,  $t_2$  will always be less or equal to  $t_1$ , as  $\|\mathbf{y} - \mathbf{w}\| \leq \|\mathbf{x} - \mathbf{w}\|$ .

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

In lemma 2.6 below we show that the pointwise normal reach bounds the reach. For this, we need theorem 4.8(7) from Federer (1959)

**Lemma 2.5** (Federer (1959)). Let  $\mathbf{x}, \mathbf{y}$  be points on  $\mathcal{M}$  with  $r_{\max}(\mathbf{x}) > 0$ , and let  $\mathbf{n}$  be a normal vector in  $N_{\mathbf{x}}\mathcal{M}$ , then

$$\langle \mathbf{n}, \mathbf{y} - \mathbf{x} \rangle \leq \frac{\|\mathbf{y} - \mathbf{x}\|^2 \|\mathbf{n}\|}{2r_{\max}(\mathbf{x})} \quad (12)$$

**Lemma 2.6.** For all  $\mathbf{x} \in \mathcal{M}$  we have that

$$\inf_{\mathbf{y} \in B_{2r_N(\mathbf{x})}(\mathbf{x}) \cap \mathcal{M}} r_N(\mathbf{y}) \leq r_{\max}(\mathbf{x}) \leq r_N(\mathbf{x}). \quad (13)$$

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

$$r_{\max}(\mathbf{x}) \leq \frac{\|\mathbf{x} - \mathbf{y}\|^2 \|\mathbf{n}\|}{2\|\mathbf{n}\| \|\mathbf{x} - \mathbf{y}\| \cos \theta}, \quad (14)$$

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

$$r_{\max}(\mathbf{x}) \leq \frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\|P_{N_{\mathbf{x}}\mathcal{M}}(\mathbf{y} - \mathbf{x})\|}, \quad (15)$$

proving the right inequality. Consider  $B = B_{r_N(\mathbf{x})}(\mathbf{x})$ . Suppose there exists  $\mathbf{w} \in B$  with  $\mathbf{w} \notin \text{Unp}(\mathcal{M})$ . Then  $\mathbf{w} \notin N_{\mathbf{x}}\mathcal{M}$ . Hence there exists  $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{M}$  such that  $d(\mathbf{w}, \mathcal{M}) = \|\mathbf{y}_1 - \mathbf{w}\| = \|\mathbf{y}_2 - \mathbf{w}\| < r_N(\mathbf{x})$ . From Federer (1959) theorem 4.8 we know that  $\mathbf{w} \in N_{\mathbf{y}_1}\mathcal{M}, N_{\mathbf{y}_2}\mathcal{M}$ . Combining this with lemma 2.4 gives that  $r_N(\mathbf{y}_1), r_N(\mathbf{y}_2) \leq d(\mathbf{w}, \mathcal{M})$  and that  $d(\mathbf{x}, \mathcal{M}) < \|\mathbf{w} - \mathbf{x}\|$ . We also have that  $\|\mathbf{y}_1 - \mathbf{x}\|, \|\mathbf{y}_2 - \mathbf{x}\| \leq 2r_N(\mathbf{x})$ . Combining these inequalities gives us that the distance from  $\mathbf{x}$  to any point not in  $\text{Unp}(\mathcal{M})$  is greater than  $\inf_{\mathbf{y} \in B_{2r_N(\mathbf{x})}(\mathbf{x}) \cap \mathcal{M}} r_N(\mathbf{y})$ , which implies that

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

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.

### 2.3 Estimating the pointwise normal reach

The definition of  $r_N$ , prompts us to minimize  $R(\mathbf{x}, \mathbf{y})$  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 2.3 with a minimization over the samples. Using that the projection matrix onto  $N_{\mathbf{x}}\mathcal{M}$  is given by  $P_{N_{\mathbf{x}}\mathcal{M}} = \mathbf{I} - \mathbf{J}(\mathbf{J}^\top\mathbf{J})^{-1}\mathbf{J}^\top$ , we get the following estimator

$$\hat{r}_N(\mathbf{x}) = \min_{\mathbf{y} \in \mathbf{S}} \frac{\|\mathbf{y} - \mathbf{x}\|^2}{2\|(\mathbf{I} - \mathbf{J}(\mathbf{J}^\top\mathbf{J})^{-1}\mathbf{J}^\top)(\mathbf{y} - \mathbf{x})\|}, \quad (16)$$

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

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 2.7.** *For any  $\mathbf{x}, \mathbf{y} \in \mathcal{M}$*

$$R(\mathbf{x}, \mathbf{y}) \geq \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|. \quad (17)$$

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

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. 16. The algorithm samples uniformly inside a ball centered on  $\mathbf{x}$  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  $\mathbf{x}$  onto  $\mathcal{M}$ .

---

#### Algorithm 1 Sampling-based reach estimator

---

```

radius  $\leftarrow$   $r_0$ 
reach  $\leftarrow$   $\infty$ 
for  $i \leftarrow 1, \dots, \text{num\_batches}$  do
  samples  $\leftarrow$  sample_ball( $\mathbf{x}$ , radius, batch_size)
  projected  $\leftarrow$  decode(encode(samples))
  reach  $\leftarrow$   $\min(\text{reach}, \text{reach\_est}(\mathbf{x}, \text{projected}))$ 
  radius  $\leftarrow$   $2 \cdot \text{reach}$ 
end for

```

---

### 2.4 Is a point within reach?

Suppose that a point  $\mathbf{x} \in \mathbb{R}^D$  is represented by a point on the manifold  $f(\mathbf{z})$ . From definition 2.1 we know that  $\mathbf{x}$  has a unique nearest point on the manifold if

$$\|\mathbf{x} - f(\mathbf{z})\| < r_{\max}(f(\mathbf{z})). \quad (18)$$

A point  $\mathbf{x}$  which does not satisfy this inequality risks not having a unique nearest point, and hence no unique representation. From lemma 2.6 we know that  $r_{\max}(f(\mathbf{z})) \leq r_N(f(\mathbf{z}))$ . So  $\mathbf{x}$  risks not having a unique nearest point if

$$\|\mathbf{x} - f(\mathbf{z})\| \geq r_N(f(\mathbf{z})) \geq r_{\max}(f(\mathbf{z})). \quad (19)$$

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

$$\hat{r}_N(f(\mathbf{z})) = \inf_{\substack{\mathbf{y} \in \mathcal{M} \cap B_{2\|\mathbf{x} - f(\mathbf{z})\|}(f(\mathbf{z})) \\ \mathbf{y} \neq f(\mathbf{z})}} R(f(\mathbf{z}), \mathbf{y}), \quad (20)$$

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

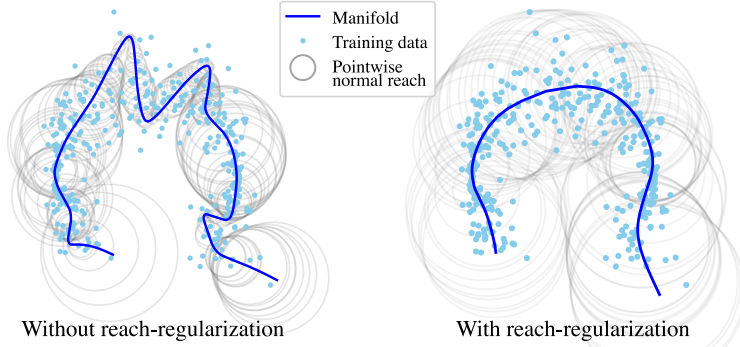


Figure 3: *Left*: An autoencoder trained on noisy points scattered along a circular arc. *Right*: The manifold spanned by the decoder of an autoencoder trained with reach regularization. In both panels, the gray circles illustrate the estimated pointwise normal reach at points along the autoencoder curve.

## 2.5 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

$$\mathcal{R}(\mathbf{x}) = \text{Softplus}(\|f(g(\mathbf{x})) - \mathbf{x}\| - \hat{r}_N(f(g(\mathbf{x}))). \quad (21)$$

The reach-regularized autoencoder then minimizes

$$\mathcal{L} = \sum_{n=1}^N \|f(g(\mathbf{x}_n)) - \mathbf{x}_n\|^2 + \lambda \sum_{n=1}^N \mathcal{R}(\mathbf{x}_n), \quad (22)$$

where we in practice use  $\lambda = 1$ . 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  $r_0 = 2\|f(g(\mathbf{x}_n)) - \mathbf{x}_n\|$ .

## 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.<sup>2</sup>

### 3.1 Analysing reach

#### 3.1.1 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 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 3(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.

#### 3.1.2 CelebA

To investigate the reach on a non-toy dataset, we train a deep autoencoder on the CelebA face dataset (Liu et al., 2015). The dataset consists of approximately 200 000 images of celebrity faces.

<sup>2</sup>The code is available at [https://github.com/HeleneHauschultz/is\\_an\\_encoder\\_within\\_reach](https://github.com/HeleneHauschultz/is_an_encoder_within_reach).



Figure 4: CelebA validation set reconstructions.

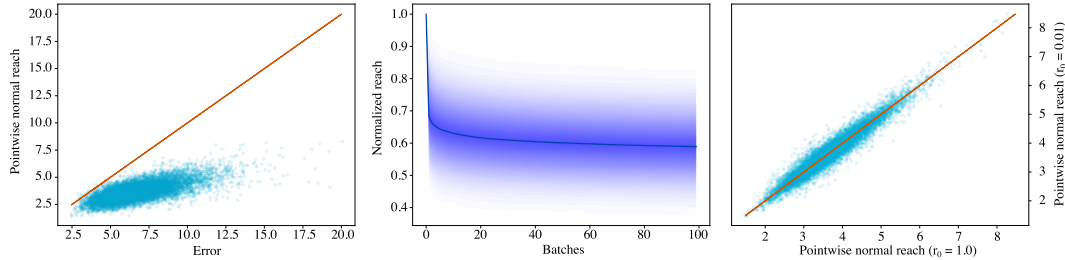


Figure 5: *Left*: Estimated reach for CelebA validation samples plotted against the L2 error. Samples below the diagonal red line does not have a unique encoder. *Center*: Normalized reach as a function of batches used to estimate the reach. The normalized reach is the estimated pointwise normal reach 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.

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. 4.

After training we estimate the reach of the validation set using the sampling based approach (Alg. 1). Fig. 5(left) plots the reconstruction error  $\|\mathbf{x} - f(\mathbf{z})\|$  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. 5(center) shows the average pointwise normal reach over the validation set as a function of the number of iterations in the sampling based estimator. 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. 5(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.

### 3.2 Reach regularization

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




Figure 6: 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. The red numbers are outside the reach, while green are within. Reach regularization smoothens the decoder to increase reach with minimal changes to both reconstructions and latent configuration.

### 3.2.1 Toy circle

Returning to the example from section 3.1.1, 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. 3(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.

### 3.2.2 MNIST

Next we train an autoencoder on 5000 randomly chosen images from the classes 2, 4 and 8 from MNIST (LeCun et al., 1998). We use a symmetric architecture reducing to two dimensional representation through a sequence of  $784 \rightarrow 500 \rightarrow 250 \rightarrow 150 \rightarrow 100 \rightarrow 50 \rightarrow 2$  linear layers with ELU non-linearities. We pretrain 5000 epochs without any regularization, and proceed with reach regularization enabled. Fig. 6(left) shows the percentage of points which lies within reach of the estimated manifold. We observe that reach regularization slightly increases the reconstruction error, 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 6(center) shows the latent representations given by the pretrained autoencoder without regularization, while Fig 6(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.

## 4 Related work

Representation learning is a foundational aspect of current machine learning, and the discussion paper by Bengio et al. (2013) is an excellent starting point. As is common, Bengio et al. (2013) 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* (Kingma and Welling, 2014; Rezende et al., 2014) amortizes the estimation of  $p(\mathbf{z}|\mathbf{x})$  such that it is parametrized by the output of a function. Alternatives relying on Monte Carlo estimates of  $p(\mathbf{z}|\mathbf{x})$  do allow for capturing non-uniqueness (Hoffman, 2017), 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. Fefferman et al. (2016) 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 Schonsheck et al. (2020), and the adaptive clustering of Besold and Spokoiny (2020). Both works rely on the reach as a tool of derivation. Similarly, Chae et al. (2021) 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 Aamari et al. (2019) which studies the convergence of an estimator of the global reach (5). 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.

## 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 implications 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 not ensured. This is indicative that the problem of uniqueness is not purely an academic question, but one of practical importance.

We provide a Monte Carlo 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.

## Acknowledgments and Disclosure of Funding

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

- E. Aamari, J. Kim, F. Chazal, B. Michel, A. Rinaldo, and L. Wasserman. Estimating the reach of a manifold, 2019.
- Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- F. Besold and V. Spokoiny. Adaptive manifold clustering, 2020.
- M. Chae, D. Kim, Y. Kim, and L. Lin. A likelihood approach to nonparametric estimation of a singular distribution using deep generative models, 2021.
- H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. ISSN 0002-9947.
- C. Fefferman, S. Mitter, and H. Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016. ISSN 0894-0347.
- M. D. Hoffman. Learning deep latent Gaussian models with Markov chain Monte Carlo. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1510–1519. PMLR, 06–11 Aug 2017.
- D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, 2014.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- S. Schonsheck, J. Chen, and R. Lai. Chart auto-encoders for manifold structured data, 2020.
- M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. Smola. Deep sets, 2018.

## A Appendix

### A.1 Extending the pointwise normal reach to the non-manifold setting

Federer (1959) introduces reach for arbitrary subsets of Euclidean space. In this situation  $T_{\mathbf{x}}\mathcal{M}$  and  $N_{\mathbf{x}}\mathcal{M}$  denote the tangent- and normal cone.

**Definition A.1.** 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  $\mathbf{x}$  if either  $v = 0$  or if for every  $\varepsilon > 0$  exists  $\mathbf{y} \in \mathcal{M}$  with

$$0 < \|\mathbf{y} - \mathbf{x}\| < \varepsilon \quad \text{and} \quad \left\| \frac{\mathbf{y} - \mathbf{x}}{\|\mathbf{y} - \mathbf{x}\|} - \frac{v}{\|v\|} \right\| < \varepsilon. \quad (23)$$

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

$$\langle w, v \rangle \leq 0 \text{ for all } v \in T_{\mathbf{x}}\mathcal{M}. \quad (24)$$

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

We can extend theorem 2.4 and Lemma 2.6 to the general situation as defined by Federer. To extend Theorem 2.4 it is sufficient to prove that for any  $v \in N_{\mathbf{x}}\mathcal{M}$  and  $u \in \mathbb{R}^D$ ,  $\|P_v(u)\| \leq d(u, T_{\mathbf{x}}\mathcal{M})$ .

**Lemma A.2.** For any  $v \in N_{\mathbf{x}}\mathcal{M}$  and  $u \in \mathbb{R}^D$  with  $\langle v, u \rangle \geq 0$ ,  $\|P_v(u)\| \leq d(u, T_{\mathbf{x}}\mathcal{M})$ .

*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

$$\begin{aligned} d(u, \text{dual}(v)) &= \inf_{w \in \text{dual}(v)} \|u - w\| = \inf_{w \in \text{dual}(v)} \|u_v + u_{v^\perp} - w\| \\ &= \inf_{w \in \text{dual}(v)} \|u_v\| + \|u_{v^\perp} - w\| - 2\langle u_v, w \rangle. \end{aligned} \quad (25)$$

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  $\text{dual}(v) \supset \text{dual}(N_{\mathbf{x}}\mathcal{M}) \supset T_{\mathbf{x}}\mathcal{M}$ . Hence  $d(u, T_{\mathbf{x}}\mathcal{M}) = \inf_{w \in T_{\mathbf{x}}\mathcal{M}} \|u - w\| \geq \inf_{w \in \text{dual}(v)} \|u - w\| = \|P_u(v)\|$ .  $\square$

To extend lemma 2.6 note that if  $r_{max}(x) > 0$ , then  $T_{\mathbf{x}}\mathcal{M}$  is convex (Federer, 1959, Thm 4.8 (12)). Let  $\mathbf{y} \in \mathcal{M}$ . If  $\mathbf{y} - \mathbf{x} \in T_{\mathbf{x}}\mathcal{M}$  then  $R(\mathbf{x}, \mathbf{y}) = \infty$ . Otherwise, as  $T_{\mathbf{x}}\mathcal{M}$  is a convex cone, there exists  $n \in N_{\mathbf{x}}\mathcal{M}$  such that  $\langle n, \mathbf{y} - \mathbf{x} \rangle \geq 0$ . In that case  $\langle n, \mathbf{x} - \mathbf{y} \rangle = \|P_n(\mathbf{y} - \mathbf{x})\|$ , 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 (16) 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.