Distance Preserving Embeddings for General *n*-Dimensional Manifolds

Nakul Verma
UC San Diego
naverma@cs.ucsd.edu

Want to study

Manifold embeddings with ***provable*** guarantees.

Formally

Let X be a sample from an underlying n-dimensional manifold $M \subset \mathrm{R}^D$, and let $\mathcal{A}$ be an embedding of M from R^D to R^d.

Define **quality of embedding** of as $(1 \pm \epsilon)$-isometric, if

$$(1-\epsilon) \leq \frac{\mathrm{dist}(\mathcal{A}(p), \mathcal{A}(q))}{\mathrm{dist}(p,q)} \leq (1+\epsilon)$$

(we are interested in geodesic distances)

Questions:

- Can one come up with an embedding algorithm $\mathcal{A}$ That achieves $(1 \pm \epsilon)$-isometry for all points in M (not just X)?
- How much can one reduce the target dimension d and still have $(1 \pm \epsilon)$-isometry?
- What kinds of restrictions (if any) does one need on M and X ?

Previous work

Algorithms

Theorem (IsoMap): Let X be a sample from an n-dimensional manifold M that is isometric to some convex subset of R^n.
As $|X| \to \infty$, IsoMap can recover the n-dimensional parameterization of M (upto rigid transformations).

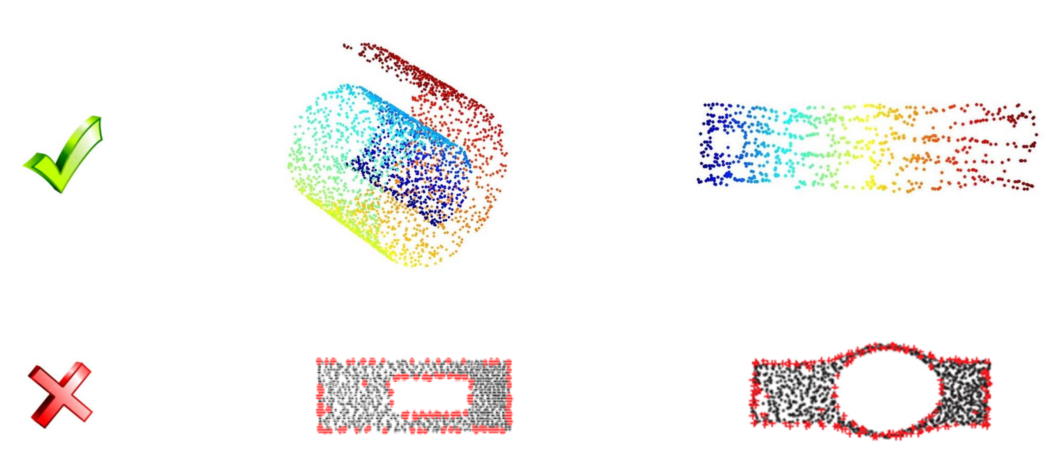

Theorem (random projections): For any $\epsilon > 0$. A random projection of n-dimensional manifold into a random subspace of dimension $d \geq \tilde{O}\left(\frac{n}{\epsilon^2}\right)$ is an $(1 \pm \epsilon)$-isometric embedding (with high probability).

✔ works large class of smooth (well-conditioned) manifolds!

✘ severe dependence on ϵ. If want all distances to be within 1%, d > 10,000!

Differential Geometry

Theorem (Nash'54): An n-dimensional manifold can be *isometrically* embedded in R^{2n+1}.

✔ no dependence on ϵ!

✘ uses manifold curvature tensor information, which is not accessible from samples.

Embedding Technique

Inspired by Nash's Theorem, we divide the embedding in two stages:

Embedding Stage

Map the given manifold M in a lower dimensional space, ***without*** having to worry about preserving distances.

This initial embedding distorts interpoint distances, but should not introduce any kinks, tears or discontinuities.

A good way is to apply a ***random projection***

Correction Stage

Fix the distorted distances (from the first stage) by applying ***local corrections*** to the embedded M.

Care needs to be taken so that the corrections don't interfere with each other.

A good way is to apply ***spirals***

Example:

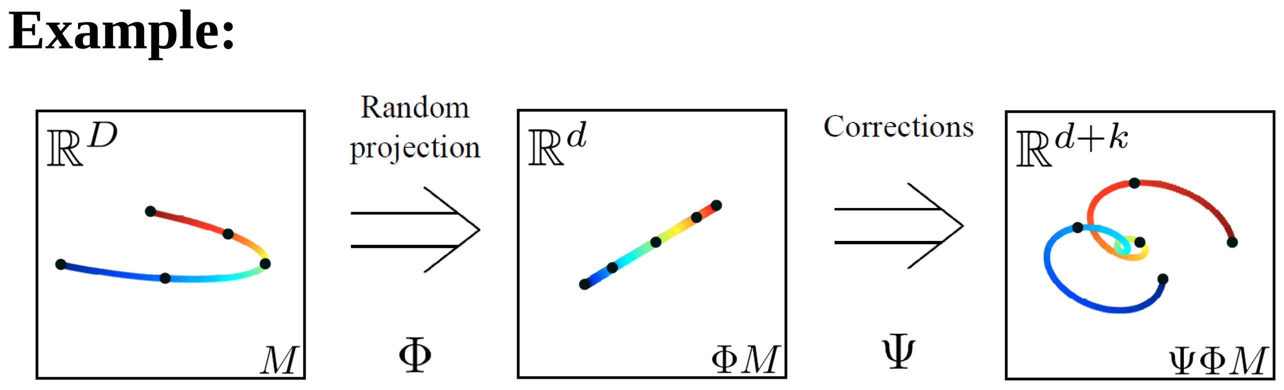

A low-dimensional manifold in some high dimensional space | Embedding of M by a random projection: tail-ends are distorted more than center. | Fix distances by applying spirals

Algorithm

Input: Sample X from M, local neighborhood size ρ.

Let Φ denote the initial random projection in $O(n)$ dim.

Preprocess:

- For each $x \in X$, let F_x be the local tangent space approximation using neighborhood size ρ.
- Let $U_x \Sigma_x V_x^\mathsf{T}$ be the SVD of ΦF_x.
- Estimate local correction around x as:

$$C_x := (\Sigma_x^{-2} - I)^{1/2} U_x^\mathsf{T}$$

Embedding: For any $p \in M$

- $t = \Phi p$.
- **for** every $x \in X$:
  - let $\Psi_{i-1}(t)$ be the embedding from previous iteration.
  - let η and ν be vectors normal to $\Psi_{i-1}(t)$.
  - let Λ_x be a localizing kernel.
  - apply correction
  $$\Psi_i = \Psi_{i-1} + \eta\sqrt{\Lambda_x}\sin(C_x t) + \nu\sqrt{\Lambda_x}\cos(C_x t)$$
- **return** $\Psi_{|X|}(t)$

Guarantee

Theorem: Fix an n-dimensional manifold $M \subset \mathrm{R}^D$. For any $\epsilon > 0$, let X be a $O(D/\epsilon)^n$-dense sample of M.

The above embedding procedure yields $(1 \pm \epsilon)$-isometric embedding of M in $O(n)$ dimensions.

Proof (sketch)

Length along any path γ on a manifold is given by:

$$\int \|\gamma'(s)\| ds$$

Length of a curve is infinitesimal sum of the length of tangent vectors along its path

To argue that our embedding preserves distances up to a factor of $(1 \pm \epsilon)$, it suffices to argue that the ***lengths of tangent vectors*** across the manifold are distorted no more than by the same factor.

How to analyze the effects of stages Φ and Ψ on tangent vectors?

Need to study the **derivative** or the **pushforward** map.

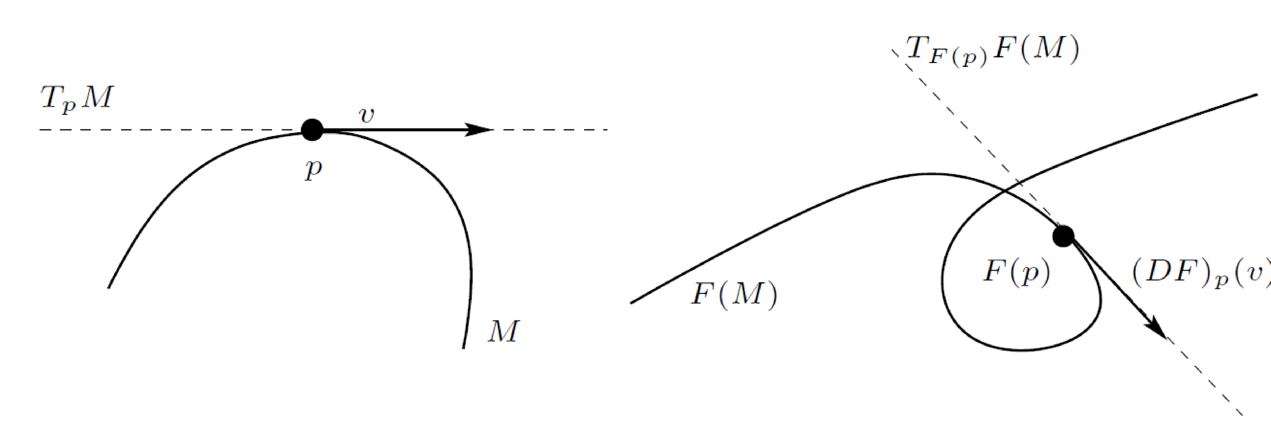

For any smooth function F that maps M to $F(M)$, there exists a derivative map DF that maps vectors tangent to M to vectors tangent to $F(M)$.

Effects of Φ

First argue that a random linear projection can shrink the lengths of tangent vectors by at most a constant amount (with high probability).

Effects of Ψ

For a fixed tangent vector, suppose its length was shrunk by a factor of L, how can we restore back its length?

Consider applying a spiral map:

$$\Psi : t \mapsto (t, \sin(Ct), \cos(Ct))$$

C is the correction size.

Tangent vector lengths change by the factor:

$$\|D\Psi\| = \|d\Psi/dt\| = \sqrt{1+C^2}$$

$\therefore$ setting $C = (L^{-2}-1)^{1/2}$ restores the lengths!

We localize the effect of a single (local) correction by applying a localizing kernel Λ

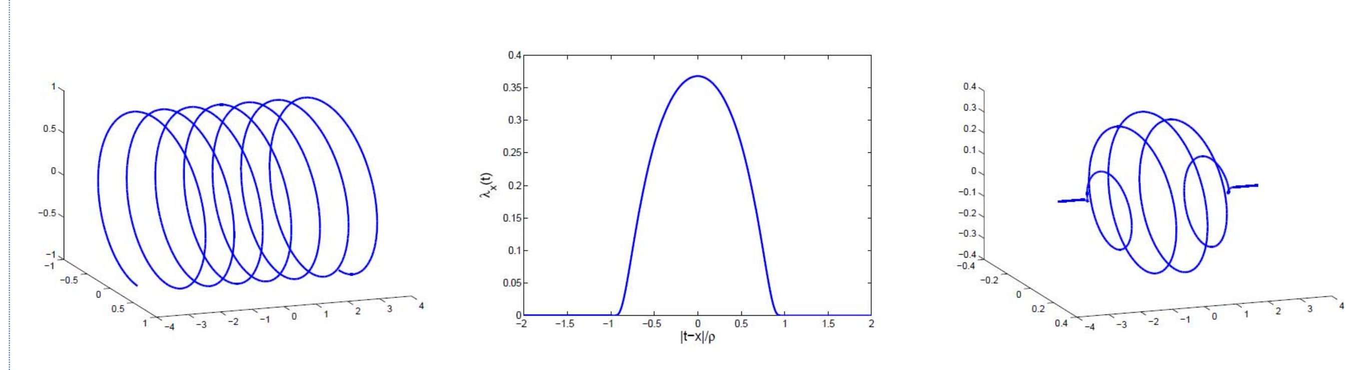

a spiral correction | localizing kernel | localized correction

We argue that the combined effect of different local corrections restores the lengths of all tangent vectors uniformly over the entire manifold.