Masking Schemes for Image Manifolds

Hamid Dadkhahi and Marco F. Duarte

Summary

- We consider the problem of *selecting an optimal mask for images modeled by a manifold*, i.e., choosing a subset of the dimensions of the image space that preserves the manifold structure present in the original data.
- Such masking implements a form of compressed sensing that reduces power consumption in emerging imaging sensor platforms.
- Our goal is for the manifold learned from masked images to resemble the manifold learned from full images as closely as possible.

Introduction

- High-dimensional data: data acquisition followed by dimensionality reduction is *inherently wasteful*.
- Increased flexibility in power consumption by allowing *pixel-level control of the acquisition process* via emerging imaging sensor architecture for embedded systems.
- Design *data-dependent image masking schemes* that reduce the number of pixels involved in acquisition while preserving the information of interest.

Manifold Learning and Linear Dimensionality Reduction:

- An ℓ-dimensional *manifold* $\mathcal{M} \subset \mathbb{R}^d$ is a set of data points $\mathcal{X} = \{x_1, x_2, \ldots, x_n\} \subset \mathbb{R}^d$ that have been generated according to an ℓ-dimensional parametric function.
- Goal: Given high-dimensional data set $\mathcal{X}$, find underlying parameterization of the manifold $\mathcal{M} : \mathbb{R}^\ell \to \mathbb{R}^d$.
- *Dimensionality reduction*: *embed* data $\mathcal{X}$ to low-dimensional space $\mathbb{R}^m$ ($m \ll d$) so that local geometry of $\mathcal{M}$ is preserved, i.e., distances in $\mathbb{R}^m$ correspond to parameter differences in $\mathbb{R}^\ell$.
- *Linear* dimensionality reduction: use a matrix projection $\Phi \in \mathbb{R}^{m \times d}$, e.g. principal component analysis (PCA), multidimensional scaling (MDS).
- PCA/MDS fail to preserve geometric structure of a *nonlinear manifold*.

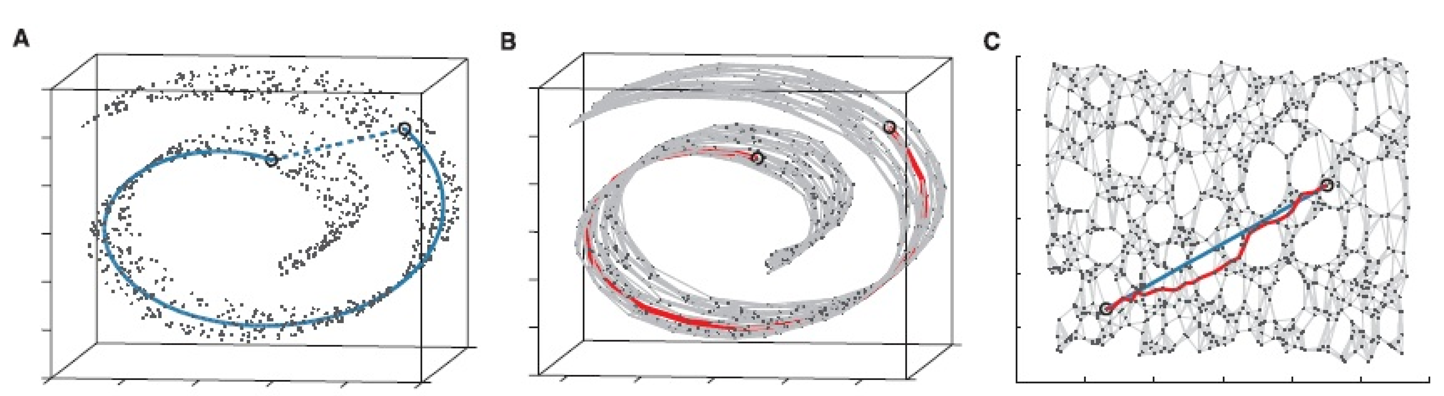

Nonlinear Manifolds and Manifold Learning:

- *Manifold learning methods* perform dimensionality reduction while preserving the local geometry of the manifold.
- *Isomap*: estimate *geodesic distances* (along the manifold, proportional to parameter difference) using shortest paths between data points across neighborhood graph; apply MDS to these distances [1].

Linear Dimensionality Reduction for Manifolds:

- *NuMax*: a data-dependent linear embedding obtained via convex optimization [2]; its goal is to preserve the norms (i.e., act as an isometry) of set of *secants* (pairwise differences between points in $\mathcal{X}$):

$$\mathcal{S} = \left\{ \frac{x_i - x_j}{\|x_i - x_j\|_2} : x_i, x_j \in \mathcal{M} \right\}.$$

- Optimization finds embedding Φ into a space of minimum dimension m such that for all $s \in \mathcal{S}$, $1 - \delta \le \|\Phi s\|_2^2 \le 1 + \delta$.
- If secant norms are preserved, then geodesic distance estimates are preserved as well.

Manifold Masking

- Emulate criteria used in linear/nonlinear embedding algorithms (Isomap, NuMax) to obtain *structure-preserving masking patterns for manifold-modeled data*.
- Seek *masking index set* $\Omega = \{\omega_1, \ldots, \omega_m\}$ that is a subset of the dimensions $[d] := \{1, 2, \ldots, d\}$ of $\mathcal{X} \subset \mathbb{R}^d$.
- Define masking linear operator $\Psi : x_i \mapsto \{x_i(j)\}_{j \in \Omega}$ corresponding to masking index set Ω.

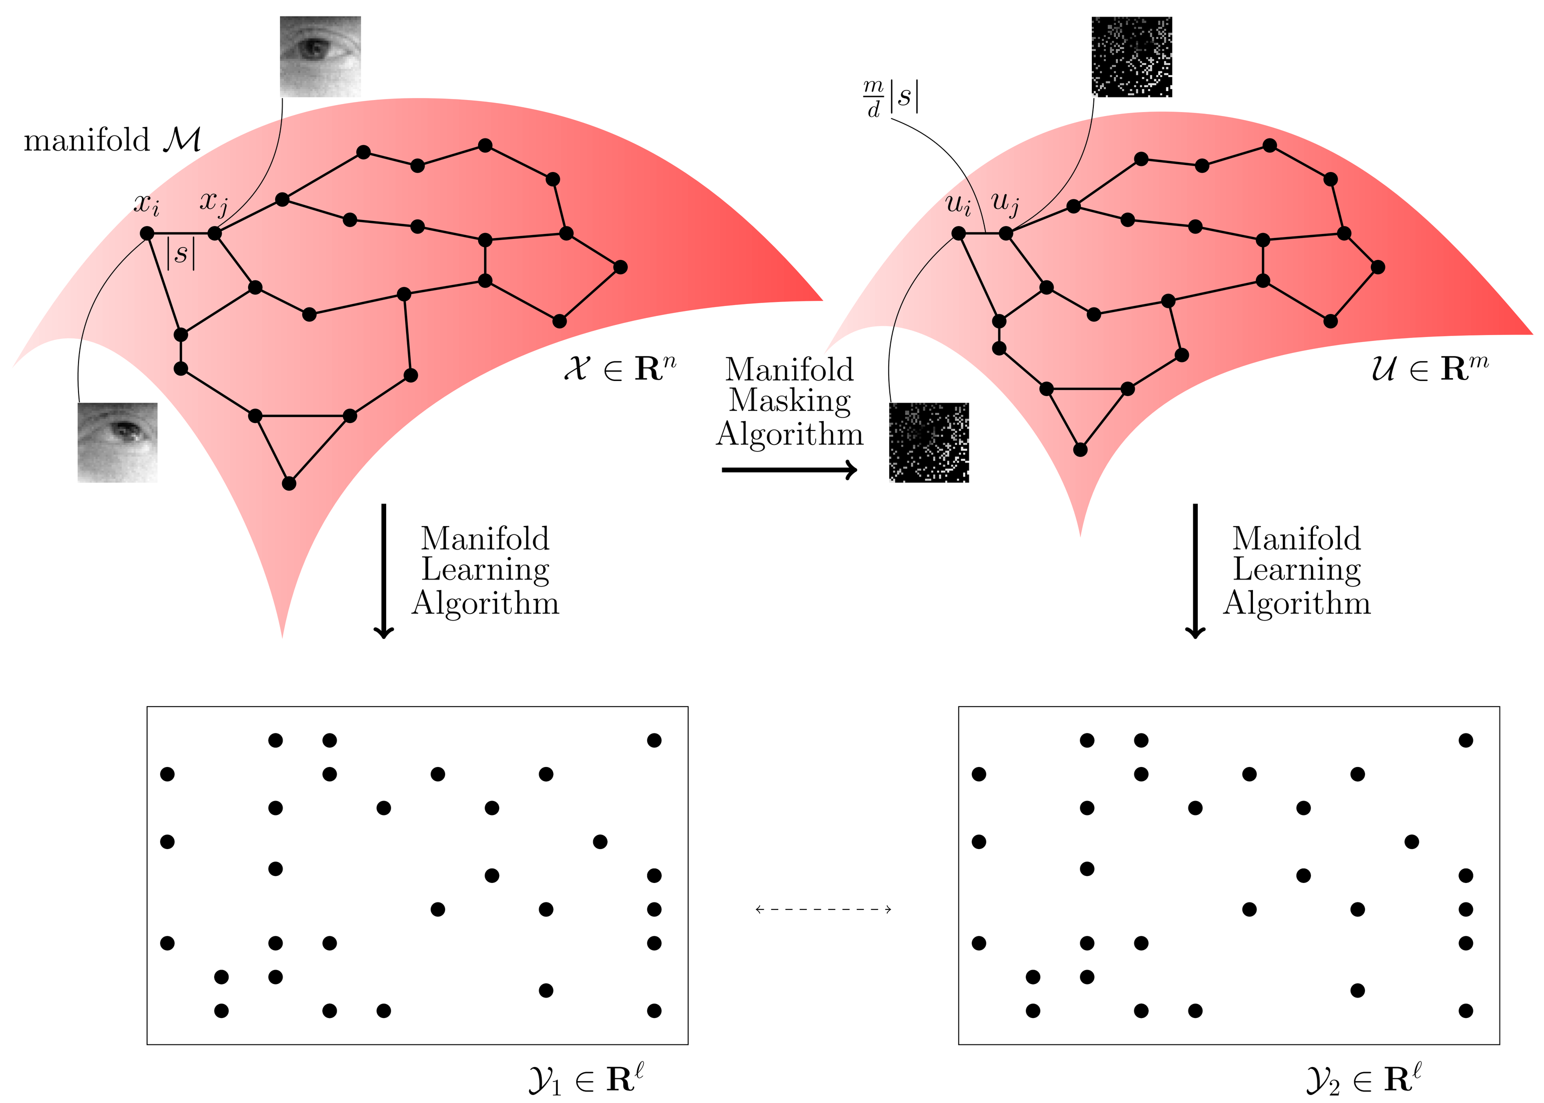

Optimization-Based Mask Selection:

- Minimize distortion incurred by secants with neighboring k data points:

$$\mathcal{S}_k = \left\{ \frac{x_i - x_j}{\|x_i - x_j\|_2} : i \in [n], x_j \in \mathcal{N}_k(x_i) \right\} \subseteq \mathcal{S}.$$

- Expectation of masked secant norms over masking index sets Ω drawn uniformly at random is $\mathbb{E}[\|\Psi s_i\|_2^2] = \frac{m}{d}$.
- Secants $\mathcal{S}_k$ inevitably subject to compaction factor of $\sqrt{\frac{m}{d}}$ in expectation by masking operator Ψ.
- Seek masking Ψ such that for all $s_i \in \mathcal{S}_k$, the squared norm of masked secants $\|\Psi s_i\|_2^2 \approx \frac{m}{d}$ as closely as possible.
- We have $\|\Psi s_i\|_2^2 = \sum_{j \in \Omega} s_i^2(j) = \sum_{j=1}^{d} s_i^2(j) z(j) = a_i^T z$, where $a_i = s_i^2$ entrywise and z is the d-dimensional indicator vector for index set Ω, i.e.,

$$z(j) = \begin{cases} 1 & \text{if} \quad j \in \Omega, \\ 0 & \text{otherwise.} \end{cases}$$

- Squared secants matrix A is an $|\mathcal{S}_k| \times d$ matrix defined by

$$A = [a_1 \;\; a_2 \;\; \cdots]^T = [s_1^2 \;\; s_2^2 \;\; \cdots]^T.$$

- Find optimal masking pattern by casting the following integer program:

$$z^* = \arg\min_z \left\| Az - \frac{m}{d} 1_{|\mathcal{S}_k|} \right\|_1 \quad \text{subject to} \quad 1_d^T z = m, z \in \{0,1\}^d, \qquad (1)$$

where 1_d denotes d-dimensional all-ones vector.
- Equality constraint $1_d^T z = m$: only m dimensions are to be retained in the masking process.
- Integer program (1) is computationally intractable (run only for 24 hours in experiments).

Manifold-Aware Pixel Selection (MAPS)

Inputs: normalized squared secants matrix A, number of dimensions m
Outputs: masking index set Ω
Initialize: $\Omega \leftarrow \{\}$
for $i = 1 \to m$ **do**
  $\bar{A}_\Omega \leftarrow A_\Omega \cdot 1_{|\Omega|}$ {compute current masked secant squared norms}
  $\omega_i \leftarrow \arg\min_{\omega \in \Omega^c} \|A_\omega + \bar{A}_\Omega - \frac{i}{d} 1_{|\mathcal{S}_k|}\|_1$ {minimize aggregate difference with $\mathbb{E}[\|\Psi s_i\|_2^2]$}
  $\Omega \leftarrow \Omega \cup \{\omega_i\}$ {add selected dimension to the masking index set}
end for

- Heuristic greedy algorithm that can find an approximate solution for (1) in drastically reduced time (seconds/minutes).
- MAPS iteratively selects elements of the masking index set Ω using the squared secants matrix A. At iteration i of the algorithm, MAPS finds a new dimension that, when added to the existing dimensions in Ω, causes the squared norm of the masked secants to match the expected value of $\frac{i}{d}$ as closely as possible on average.
- Computational complexity of MAPS is $\mathcal{O}(md|\mathcal{S}_k|) \approx \mathcal{O}(mdkn)$.

Simulation Results

- Compare (1) and MAPS with two baseline methods: *random masking* (select m out of d data dimensions uniformly at random) and *principal coordinate analysis (PCoA)*, (select indices of m dimensions with the highest variance across the dataset).
- *Eyeglasses* dataset: eye-tracking image captures via computational eyeglasses prototype that uses pixel-level imaging sensor array of [3].
- MAPS significantly outperforms random sampling and PCoA.
- For sufficiently large values of m, the performance of MAPS approaches or matches that of nonlinear/linear embedding algorithms on full images.

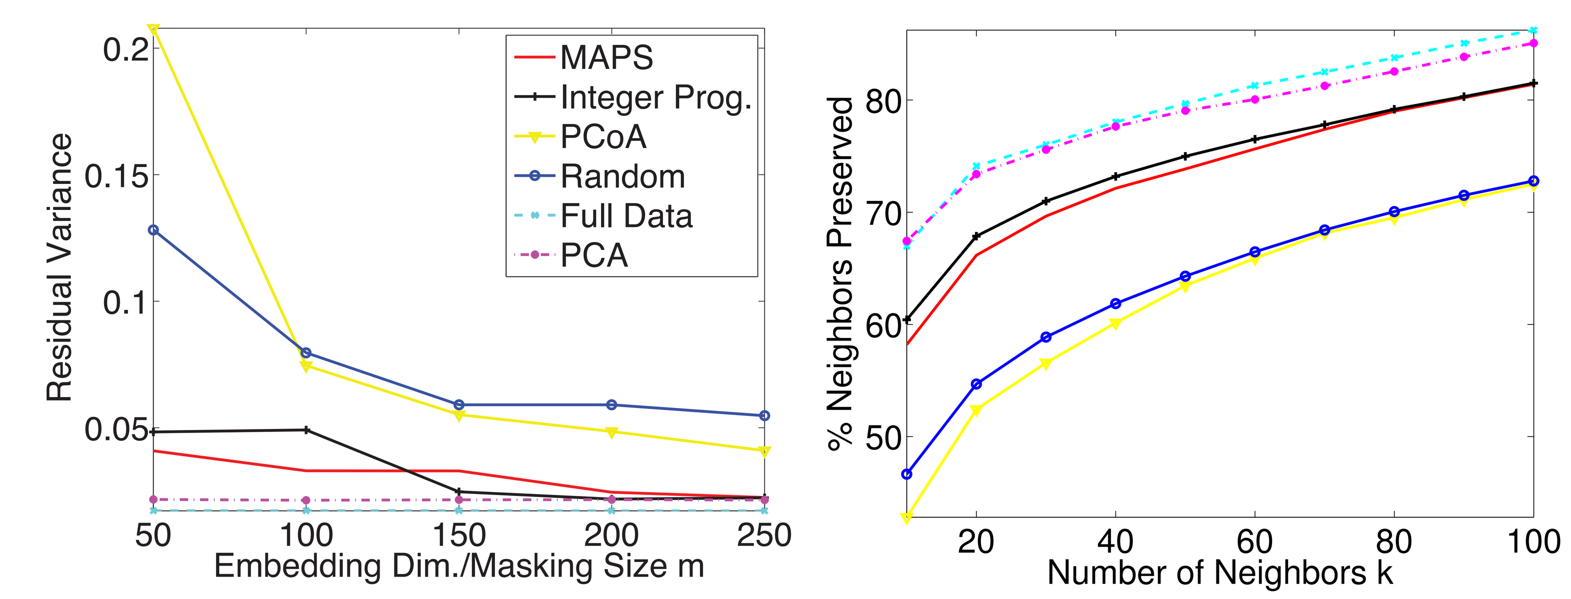

Performance comparison for linear embeddings (dashed lines) and masking algorithms (solid lines) with respect to full data. Residual variance as a function of m (left); percentage of preserved nearest neighbors for $m = 50$ (right).

Acknowledgments and References

This work was supported by NSF Grant IIS-1239341.

[1]. J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction", Science, vol. 290, no. 5500, pp. 2319-2323, 2000.
[2]. C. Hegde, A. C. Sankaranarayanan, and R.G. Baraniuk, "Near-isometric linear embeddings of manifolds", in SSP, August 2012.
[3]. A. Mayberry, P. Hu, B. Marlin, C. Salthouse, D. Ganesan, "iShadow: Design of a Wearable, Real-Time Mobile Gaze Tracker", in MobiSys 2014.

{hdadkhahi, mduarte}@ecs.umass.edu