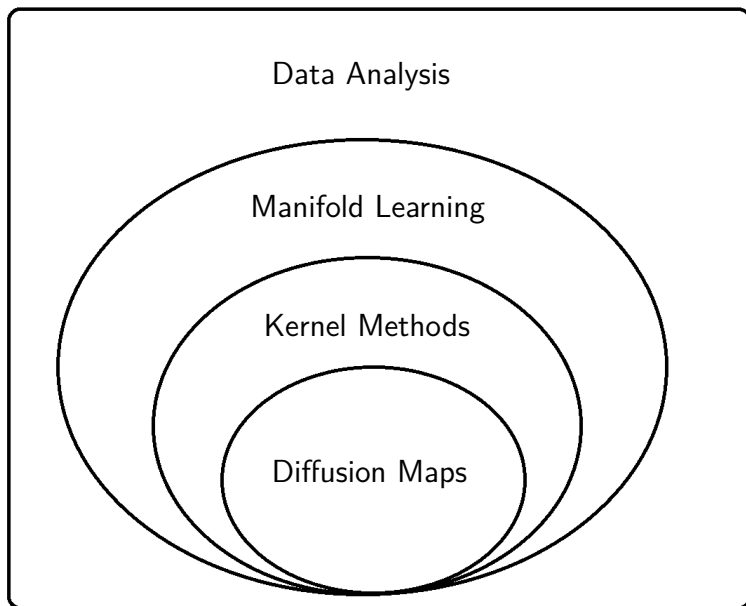


Diffusion Based Manifold Learning

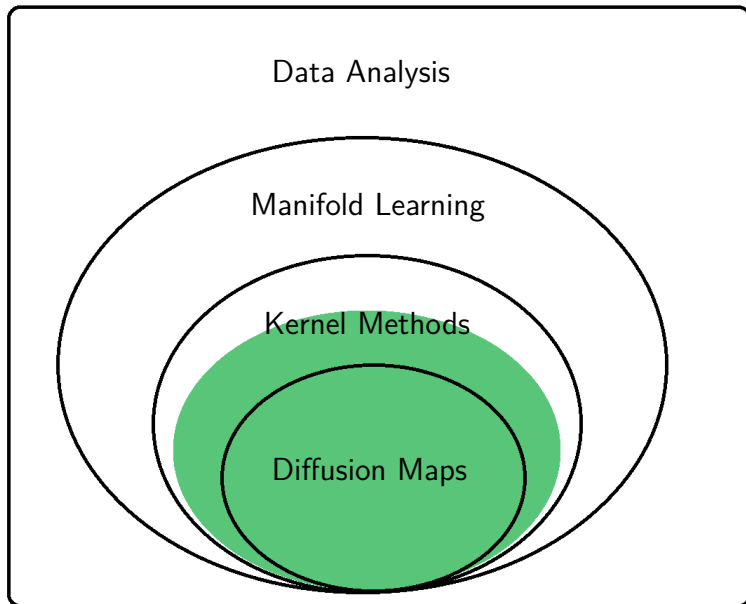
Matthew Hirn and Guy Wolf

October 23, 2013

Data Analysis Scope



Data Analysis Scope



Overview

- 1 Introduction
- 2 Manifold theory
- 3 Diffusion maps
- 4 Patch-to-Tensor Embedding
- 5 Diffusion maps for changing data

Overview

- 1 Introduction
 - Local distances
 - Lips manifold
 - Diffusion maps overview
- 2 Manifold theory
- 3 Diffusion maps
- 4 Patch-to-Tensor Embedding
- 5 Diffusion maps for changing data

Manifold learning

Why use manifolds?

availability of data

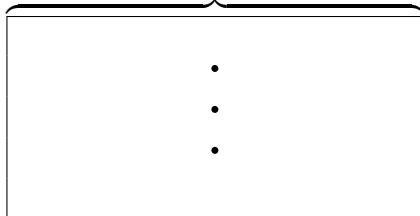


many observable parameters



high-dimensional recorded data

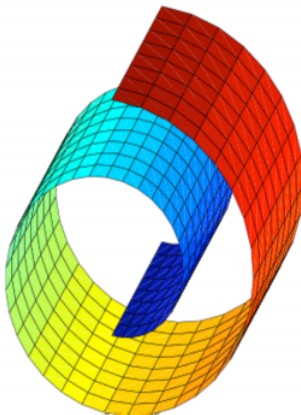
parameters = $O(10), O(100), \dots$



- Data contain dependencies & redundancies
- Observable space = non-linear mapping of few underlying factors
- Underlying locally low-dimensional geometry in high-dimensional ambient space

Manifold learning

The goal



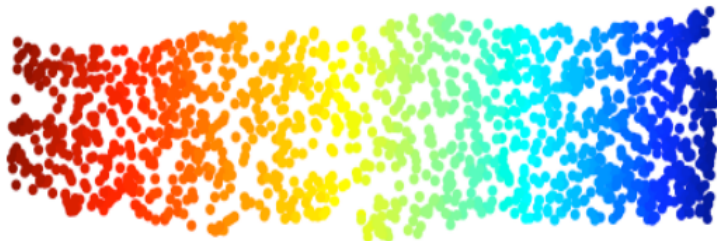
Manifold learning

The goal



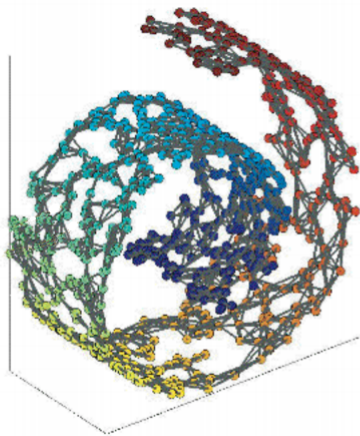
Manifold learning

The goal



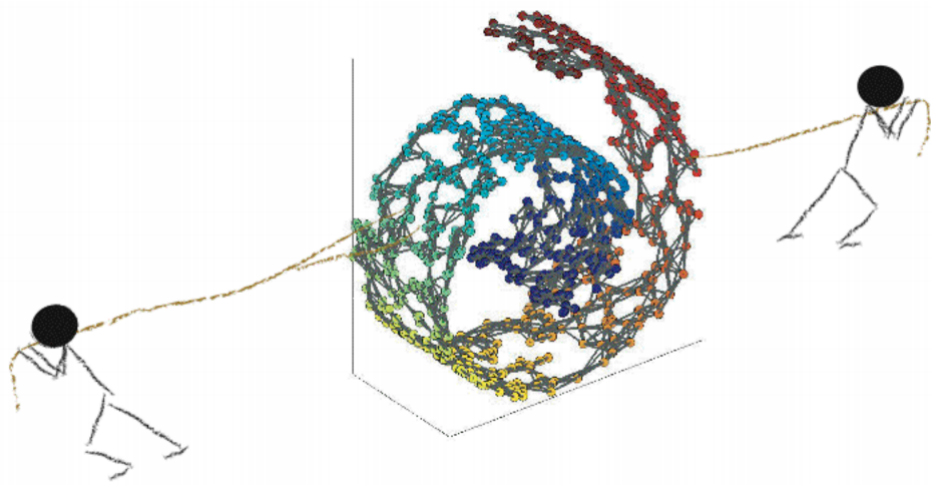
Manifold learning

Kernel methods



Manifold learning

Kernel methods



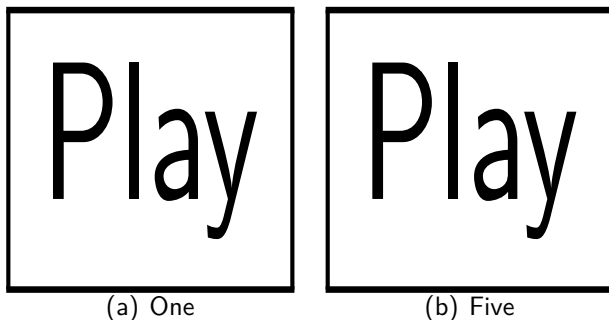


Figure : Recordings of someone saying “one” and “five.”¹

¹S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

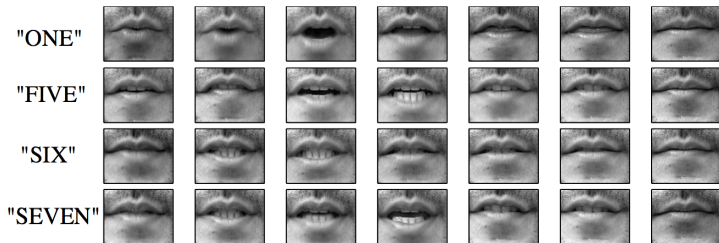


Figure : Typical frames for the words “one”, “five”, “six”, “seven.”²

²S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

The lips manifold

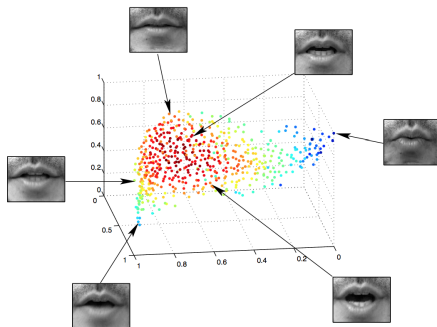
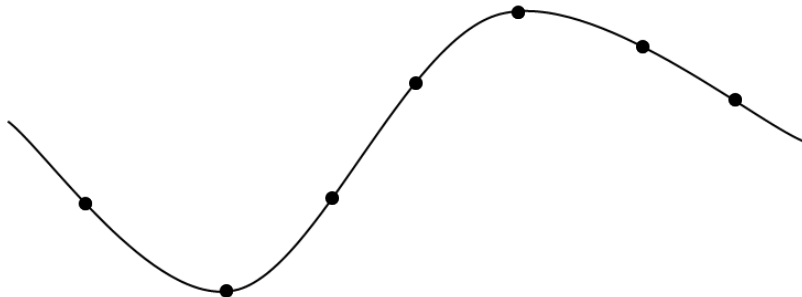


Figure : These manifold coordinates essentially capture two parameters: one controlling the opening of the mouth and the other measuring the portion of teeth that are visible.³

³S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

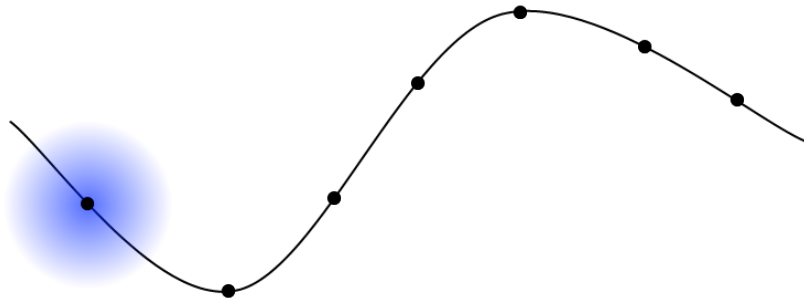
Diffusion maps overview

Gaussian kernel:



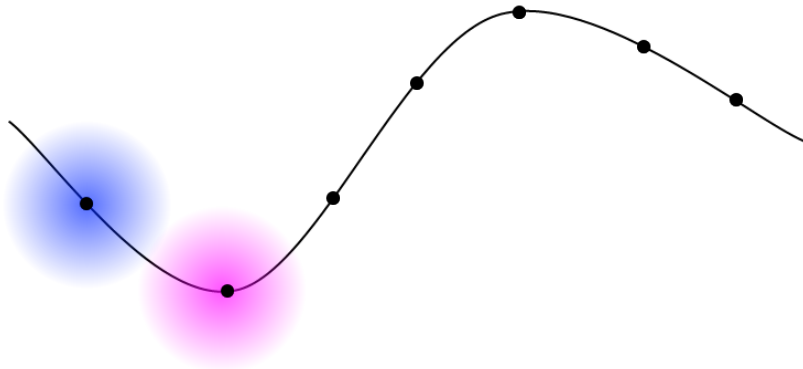
Diffusion maps overview

Gaussian kernel:



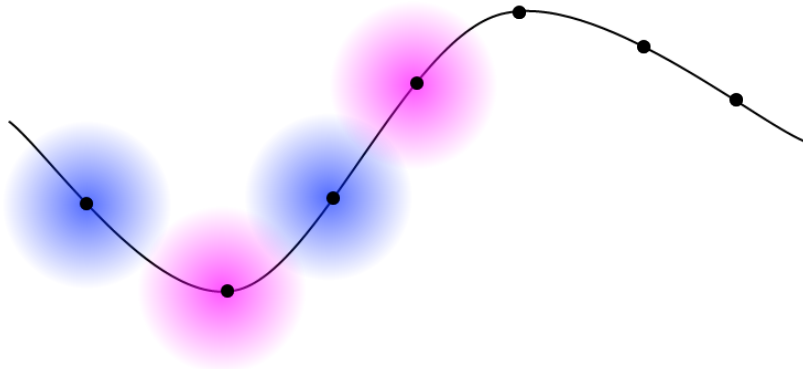
Diffusion maps overview

Gaussian kernel:



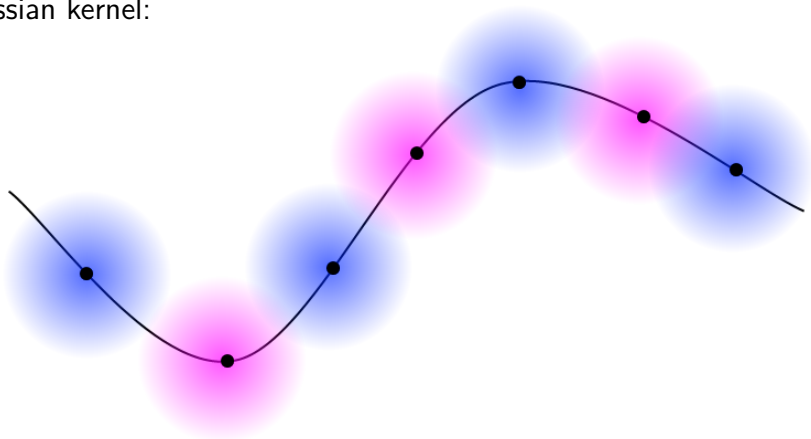
Diffusion maps overview

Gaussian kernel:



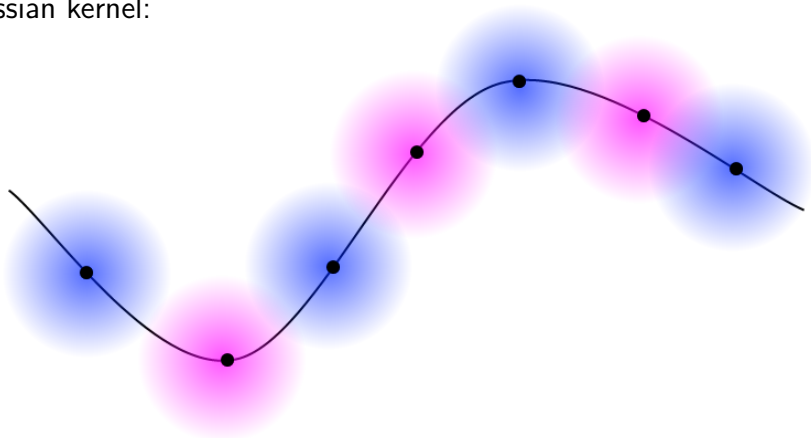
Diffusion maps overview

Gaussian kernel:



Diffusion maps overview

Gaussian kernel:

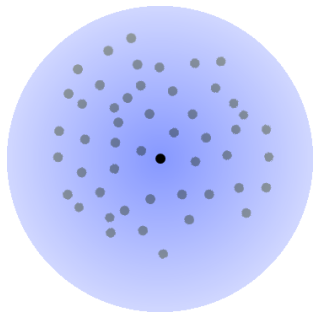


- Normalization \implies diffusion kernel
- Spectral analysis \implies map from $\mathcal{M} \subseteq \mathbb{R}^m$ to $\mathbb{R}^{\delta \ll m}$

Diffusion maps⁴ overview

Diffusion process & affinities

- Gaussian kernel:
 $k(x, y) \triangleq e^{-\frac{\|x-y\|}{\varepsilon}}$



- Degrees: $q(x) \triangleq \sum k(x, y)$
- Transition probabilities:

$$p(x, y) \triangleq \frac{k(x, y)}{q(x)}$$

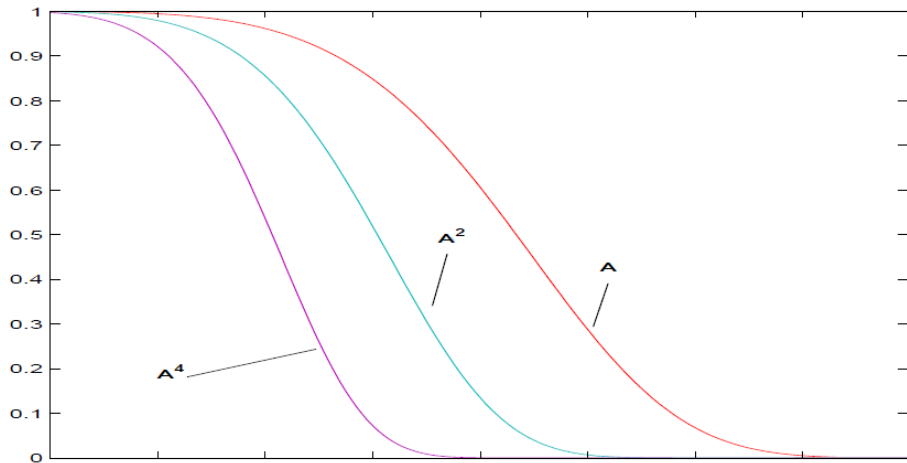
- Diffusion affinities:

$$\begin{aligned} a(x, y) &\triangleq \frac{k(x, y)}{\sqrt{q(x)}\sqrt{q(y)}} \\ &= q^{1/2}(x)p(x, y)q^{-1/2}(y) \end{aligned}$$

⁴R.R. Coifman and S. Lafon. “Diffusion Maps”. In: *Applied and Computational Harmonic Analysis* 21.1 (2006), pp. 5–30.

Diffusion maps overview

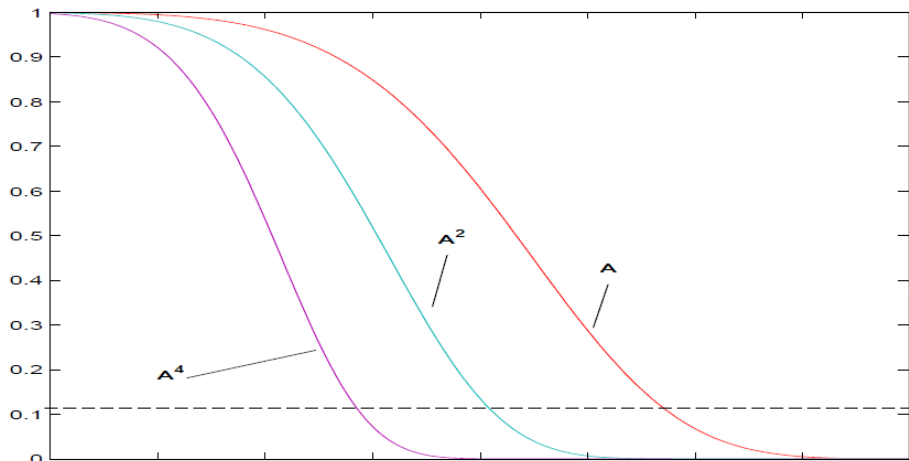
Spectral embedding



Spectrum (eigenvalues) of the diffusion affinity A and its powers

Diffusion maps overview

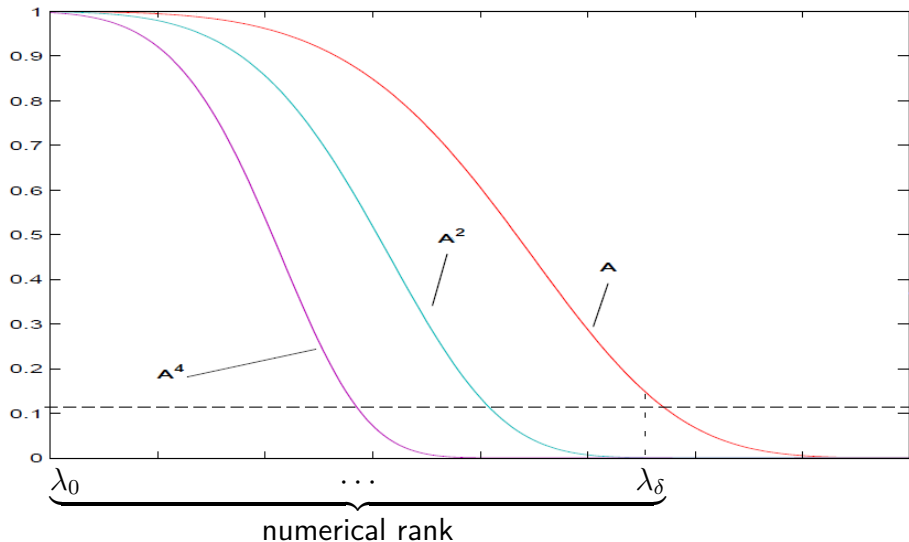
Spectral embedding



Spectrum (eigenvalues) of the diffusion affinity A and its powers

Diffusion maps overview

Spectral embedding



Diffusion maps overview

Spectral embedding

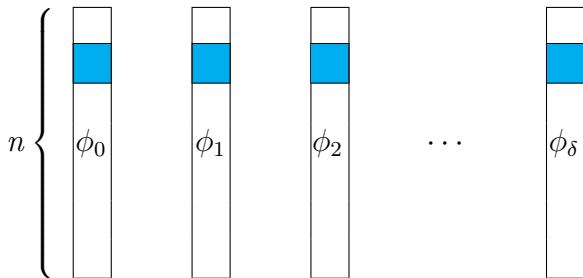
$$1 = \boxed{\lambda_0} \geq \boxed{\lambda_1} \geq \boxed{\lambda_2} \geq \dots \geq \boxed{\lambda_\delta} > 0$$

$$n \left\{ \begin{array}{c} \boxed{\phi_0} \\ \boxed{\phi_1} \\ \boxed{\phi_2} \\ \dots \\ \boxed{\phi_\delta} \end{array} \right.$$

Diffusion maps overview

Spectral embedding

$$1 = \boxed{\lambda_0} \geq \boxed{\lambda_1} \geq \boxed{\lambda_2} \geq \dots \geq \boxed{\lambda_\delta} > 0$$

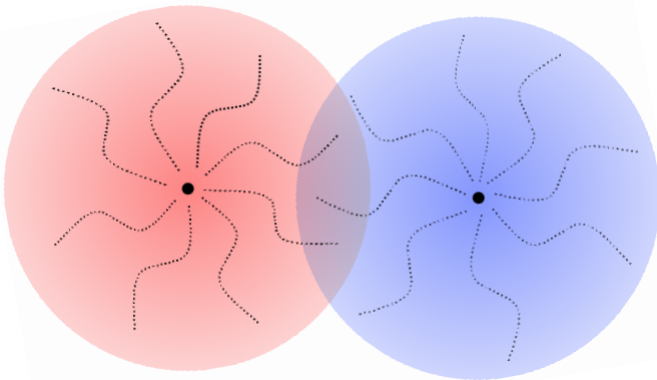


$$x \mapsto \Phi(x) \triangleq [\lambda_0 \phi_0(x), \lambda_1 \phi_1(x), \lambda_2 \phi_2(x), \dots, \lambda_\delta \phi_\delta(x)]^T$$

Diffusion maps overview

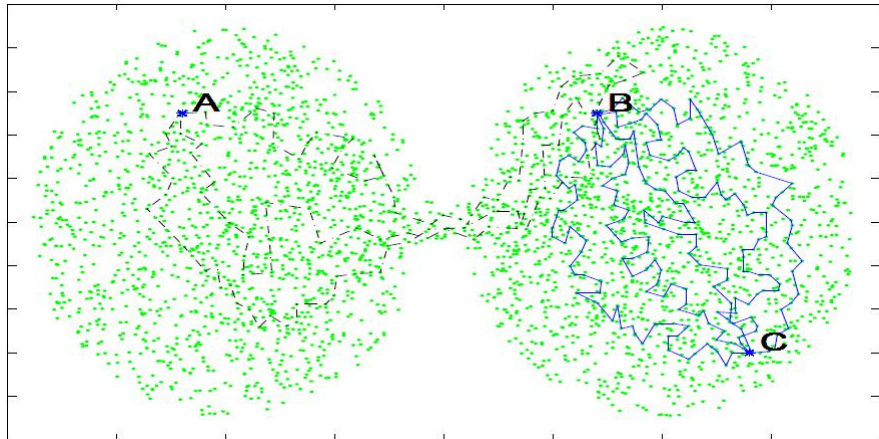
Embedded distances

$$\overbrace{\|\Phi(x) - \Phi(y)\|}^{\text{embedded distance}} = \overbrace{\|a(x, \cdot) - a(y, \cdot)\|}^{\text{diffusion distance}}$$



Diffusion maps overview

Diffusion distances



Overview

- 1 Introduction
- 2 Manifold theory
 - Heat equation
 - Embedding into Hilbert space
- 3 Diffusion maps
- 4 Patch-to-Tensor Embedding
- 5 Diffusion maps for changing data

Notation

- (\mathcal{M}, g) is a smooth, compact n -dimensional Riemannian manifold.
- Δ is the Laplace-Beltrami operator on \mathcal{M} .
- $L^2(\mathcal{M})$ is the space of functions $f : \mathcal{M} \rightarrow \mathbb{R}$ such that

$$\int_{\mathcal{M}} |f(x)|^2 dx < \infty.$$

- ℓ^2 is the space of real valued sequences $\{a_i\}_{i \geq 1}$ such that

$$\sum_{i=1}^{\infty} |a_i|^2 < \infty.$$

The heat equation

Definition

For a function $u(x, t)$ with domain and range $u : \mathcal{M} \times [0, T] \rightarrow \mathbb{R}$, the *heat equation* with initial solution $f(x)$ is

$$\begin{aligned} L(u) &\triangleq \frac{\partial u}{\partial t} - \Delta_x u = 0 \\ u(x, 0) &= f(x) \end{aligned}$$

- The fundamental solution of the heat equation is the heat kernel $K(t, x, y)$, which satisfies:
 - $L(K) = 0$ in (x, t) .
 - $\lim_{t \rightarrow 0} K(t, x, y) = \delta_y(x)$.
- The solution $u(x, t)$ is given by:

$$u(x, t) = (e^{-t\Delta} f)(x, t) \triangleq \int_{\mathcal{M}} K(t, x, y) f(y) dy.$$

Diffusion distance

Definition

The *diffusion distance* at time t is defined as:

$$D_t(x, y)^2 \triangleq \int_{\mathcal{M}} (K(t, x, z) - K(t, y, z))^2 dz.$$

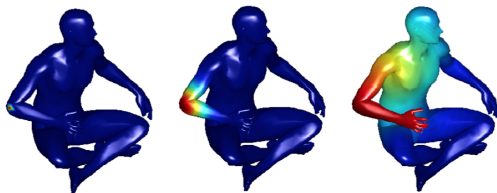


Figure : Heat kernel represents evolution of temperature.⁵

⁵H. Abdallah. Processus de Diffusion sur un Flot de Variétés Riemanniennes. PhD thesis, L'Université de Grenoble, 2010.

Spectral properties of the Laplacian

- The Laplacian Δ has a discrete sequence of eigenvalues

$$0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \cdots \nearrow +\infty.$$

- Let $\{\psi_i\}_{i \geq 0} \subset L^2(\mathcal{M})$ be a set of orthonormal eigenfunctions of Δ .
- One can write the heat kernel as:

$$K(t, x, y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \psi_i(x) \psi_i(y).$$

Embedding \mathcal{M} into ℓ^2

Definition

For any $t > 0$, define the map $\Psi_t : \mathcal{M} \rightarrow \ell^2$ as:

$$\Psi_t(x) \triangleq \sqrt{2}(4\pi)^{\frac{n}{2}} t^{\frac{n+2}{4}} \left\{ e^{\frac{-\lambda_i t}{2}} \psi_i(x) \right\}_{i \geq 1}$$

Theorem (Bérard, Besson, and Gallot, 1994⁶)

- 1 For all $t > 0$, the map Ψ_t is an embedding of \mathcal{M} into ℓ^2 .
- 2 The pulled-back metric Ψ_t^* is asymptotic to the metric g of \mathcal{M} when $t \rightarrow 0$.
- 3 Ψ_t^* (up to the normalization factor) is the diffusion distance.

⁶P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, volume 4, number 4, pages 373-398, 1994.

Overview

- 1 Introduction
- 2 Manifold theory
- 3 Diffusion maps
 - Kernels, graphs, and Markov chains
 - Diffusion embeddings
 - Statistics vs geometry
 - Discrete vs continuous
- 4 Patch-to-Tensor Embedding
- 5 Diffusion maps for changing data

Data, kernels, and graphs

- Our data set is X along with a measure μ that represents the distribution of the points.
- We assume that X lies on a low dimensional manifold \mathcal{M} .
- $k : X \times X \rightarrow \mathbb{R}$ is a symmetric, positive kernel.
- $k(x, y)$ measures the similarity between two points $x, y \in X$. The larger $k(x, y)$, the more similar.
- In high dimensions for points lying on a nonlinear manifold, we can only accurately measure local similarities, so k should behave accordingly. For example:

$$k(x, y) = k_\varepsilon(x, y) = e^{-\|x-y\|^2/\varepsilon}.$$

- X and k define a weighted graph $G = (X, k)$.

Starfish data



Figure : Data points sampled non-uniformly from a starfish shaped manifold.

Starfish weighted graph

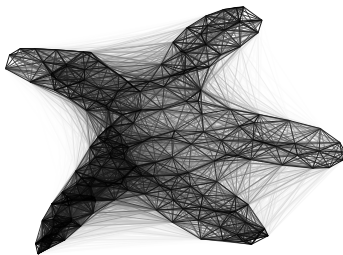


Figure : Graph with Gaussian weighted edges.

Markov chain

- To approximate heat diffusion on the underlying manifold, we construct a homogeneous Markov chain (random walk) on G .
- Define $d : X \rightarrow \mathbb{R}$ as:

$$d(x) \triangleq \int_X k(x, y) d\mu(y).$$

- Define the transition kernel $p : X \times X \rightarrow \mathbb{R}$ of the random walk as:

$$p(x, y) \triangleq \frac{k(x, y)}{d(x)}.$$

- If G is connected (which we will assume), then the Markov chain has a unique stationary distribution $\pi : X \rightarrow \mathbb{R}$ given by:

$$\pi(x) = \frac{d(x)}{\int_X d(y) d\mu(y)}.$$

Markov chain

- Define the corresponding diffusion operator $P : L^2(X) \rightarrow L^2(X)$ as:

$$(Pf)(x) \triangleq \int_X p(x, y)f(y) d\mu(y).$$

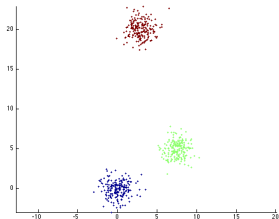
- Powers P^t of the diffusion operator run the random walk forward in time:

$$(P^t f)(x) = \int_X p_t(x, y)f(y) d\mu(y),$$

where $p_t(x, y)$ is the probability of transition from x to y in t steps.

- Running the random walk forward will integrate the local geometry and reveal the relevant geometric structures of X at different scales.

Markov chain on Gaussian data



(a) Data consisting of 3 Gaussian clouds

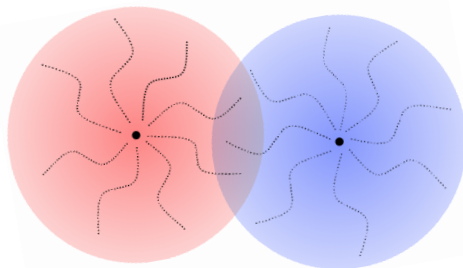
(b) Powers of P

Diffusion distance

Definition

The *diffusion distance* at time t is defined as:

$$D_t(x, y)^2 \triangleq \int_X (p_t(x, z) - p_t(y, z))^2 \frac{d\mu(z)}{\pi(z)}.$$



Diffusion maps

- The operator P has a discrete sequence of eigenvalues

$$1 = \lambda_0 > |\lambda_1| \geq |\lambda_2| \geq |\lambda_3| \cdots \searrow 0.$$

- Let $\{\psi_i\}_{i \geq 0} \subset L^2(X)$ be a set of eigenfunctions of P .

Definition

For any $t > 0$, define the *diffusion map* $\Psi_t : X \rightarrow \ell^2$ as:

$$\Psi_t(x) \triangleq \{\lambda_i^t \psi_i(x)\}_{i \geq 1}$$

Theorem (Coifman and Lafon, 2006⁷)

The diffusion map preserves the diffusion distance:

$$D_t(x, y) = \|\Psi_t(x) - \Psi_t(y)\|_{\ell^2}$$

⁷R. R. Coifman and S. Lafon. Diffusion Maps. *Applied and Computational Harmonic Analysis*, volume 21, pages 5-30, 2006.

Vehicle trajectory



(c) BTR-60

(d) BTR-60 driving in a circle

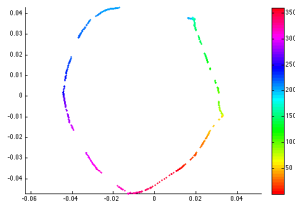
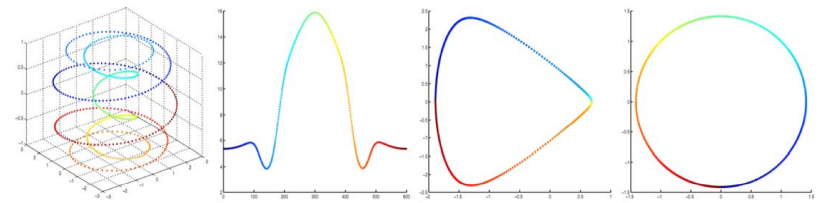


Figure : Diffusion map of the BTR video, restricted to the first two eigenvectors. The color corresponds to the angle on the circle at which the BTR is located.

Statistics vs geometry



From the left⁸:

- 1 A closed curve sampled non-uniformly.
- 2 Density of the samples.
- 3 Diffusion embedding of the curve in which the sampling skews the geometry.
- 4 A modified diffusion embedding of the curve in which the statistics have been separated from the geometry.

⁸R. R. Coifman and S. Lafon. Diffusion Maps. *Applied and Computational Harmonic Analysis*, volume 21, pages 5-30, 2006.

Starfish weighted graph

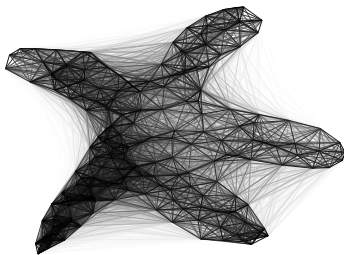


Figure : Graph with Gaussian weighted edges.

Modified Markov chain

- Let $k_\varepsilon(x, y) = e^{-\|x-y\|^2/\varepsilon}$, $x, y \in X = \mathcal{M}$.
- Approximate the density $q(x)dx = d\mu(x)$ via

$$q_\varepsilon(x) \triangleq \int_X k_\varepsilon(x, y) q(y) dy.$$

- Form a new kernel

$$k_\varepsilon^{(\alpha)}(x, y) \triangleq \frac{k_\varepsilon(x, y)}{q_\varepsilon(x)^\alpha q_\varepsilon(y)^\alpha}, \quad \alpha \in [0, 1]$$

- Normalize $k_\varepsilon^{(\alpha)}$ to be a Markov chain by defining

$$d_\varepsilon^{(\alpha)}(x) \triangleq \int_X k_\varepsilon^{(\alpha)}(x, y) q(y) dy,$$

and setting

$$p_\varepsilon^{(\alpha)}(x, y) \triangleq \frac{k_\varepsilon^{(\alpha)}(x, y)}{d_\varepsilon^{(\alpha)}(x)}.$$

Infinitesimal generator

- Define the diffusion operator $P_\varepsilon^{(\alpha)} : L^2(X) \rightarrow L^2(X)$,

$$(P_\varepsilon^{(\alpha)} f)(x) \triangleq \int_X p_\varepsilon^{(\alpha)}(x, y) f(y) q(y) dy.$$

Theorem (Coifman and Lafon, 2006⁹)

Let

$$L_\varepsilon^{(\alpha)} = \frac{I - P_\varepsilon^{(\alpha)}}{\varepsilon},$$

be the infinitesimal generator of the Markov chain. Then,

$$\lim_{\varepsilon \rightarrow 0} L_\varepsilon^{(\alpha)} f = \frac{\Delta(fq^{1-\alpha})}{q^{1-\alpha}} - \frac{\Delta(q^{1-\alpha})}{q^{1-\alpha}} f.$$

⁹R. R. Coifman and S. Lafon. Diffusion Maps. *Applied and Computational Harmonic Analysis*, volume 21, pages 5-30, 2006.

Corollaries

Two interesting cases:

- $\alpha = 0$. This was the original construction. In this case:

$$\lim_{\varepsilon \rightarrow 0} L_{\varepsilon}^{(0)} f = \frac{\Delta(fq)}{q} - \frac{\Delta(q)}{q} f,$$

and the density q fully influences the infinitesimal generator.

- $\alpha = 1$. In this case the density q is removed and we recover the Laplacian:

$$\lim_{\varepsilon \rightarrow 0} L_{\varepsilon}^{(1)} f = \Delta f.$$

This in turn implies that $P_{\varepsilon}^{(1)}$ approximates the heat kernel integral operator:

$$\lim_{\varepsilon \rightarrow 0} \left(P_{\varepsilon}^{(1)} \right)^{\frac{t}{\varepsilon}} = e^{-t\Delta}.$$

Head rotation data

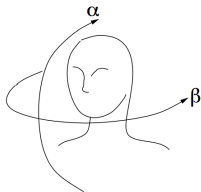


Figure : Each subject essentially moved their head along the two angles α and β . There was almost no tilting of the head.

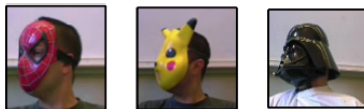


Figure : The three subjects each with their head rotated at a particular orientation.¹⁰

Head rotation data embeddings

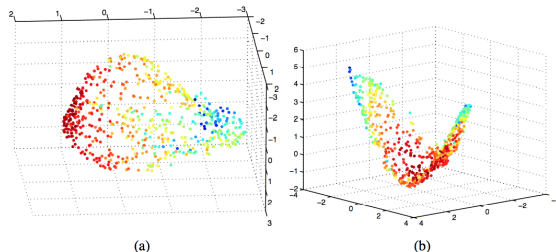


Figure : The embeddings of the YELLOW (a) and BLACK (b) sets in three diffusion coordinates without the density renormalization. The color encodes the density of the points. These embedded sets have very different shapes, and their alignment is impossible.¹¹

¹¹S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

Head rotation data embeddings

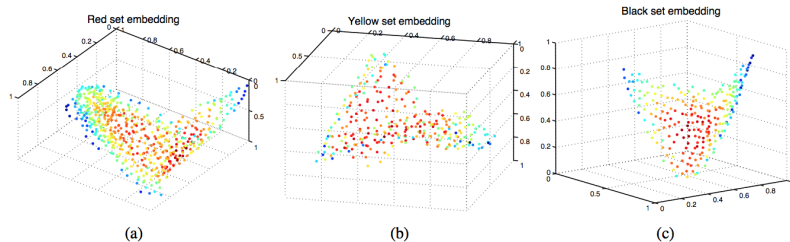


Figure : The embedding of each set in the first three diffusion coordinates with the density renormalization. All three sets share this butterfly shaped embedding.¹²

¹²S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

Head rotation data embeddings

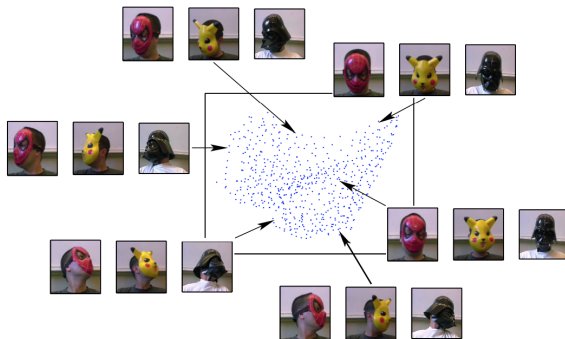


Figure : The embeddings of the YELLOW, RED, and BLACK sets after alignment.¹³

¹³S. Lafon, Y. Keller, and R. R. Coifman. Data Fusion and Multi-Cue Data Matching by Diffusion Maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 28, pages 1784-1797, 2006.

Discrete approximations

- Sample $\{x_1, \dots, x_m\} \subset X$ according to the density q .
- We have the following discrete approximations of the previous functions and operators:

$$\bar{q}_\varepsilon(x_i) \triangleq \sum_{j=1}^m k_\varepsilon(x_i, x_j)$$

$$\bar{k}_\varepsilon^{(\alpha)}(x_i, x_j) \triangleq \frac{k_\varepsilon(x_i, x_j)}{q_\varepsilon(x_i)^\alpha q_\varepsilon(x_j)^\alpha}$$

$$\bar{d}_\varepsilon^{(\alpha)}(x_i) \triangleq \sum_{j=1}^m \bar{k}_\varepsilon^{(\alpha)}(x_i, x_j)$$


$$\bar{p}_\varepsilon^{(\alpha)}(x_i, x_j) \triangleq \frac{\bar{k}_\varepsilon^{(\alpha)}(x_i, x_j)}{\bar{d}_\varepsilon^{(\alpha)}(x_i)}$$

$$(\bar{P}_\varepsilon^{(\alpha)} f)(x_i) \triangleq \sum_{j=1}^m \bar{p}_\varepsilon^{(\alpha)}(x_i, x_j) f(x_j)$$

Theorem (Singer, 2006¹⁴)

The convergence rate from the discrete to the continuous Markov chain is, with high probability,

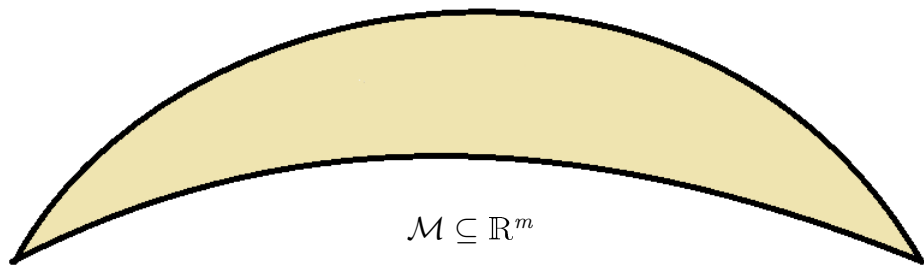
$$|(\overline{P}_\varepsilon^{(\alpha)} f)(x_i) - (P_\varepsilon^{(\alpha)} f)(x_i)| = O(m^{-\frac{1}{2}} \varepsilon^{-\frac{n}{4} - \frac{1}{2}}).$$

¹⁴A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, volume 21, pages 128-134, 2006. 

Overview

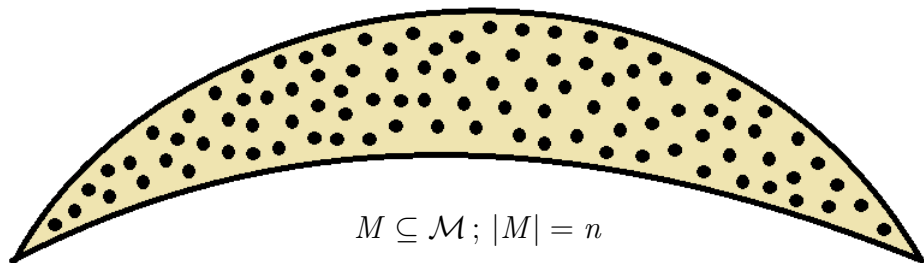
- 1 Introduction
- 2 Manifold theory
- 3 Diffusion maps
- 4 Patch-to-Tensor Embedding**
- 5 Diffusion maps for changing data

Patch-to-Tensor Embedding¹⁵



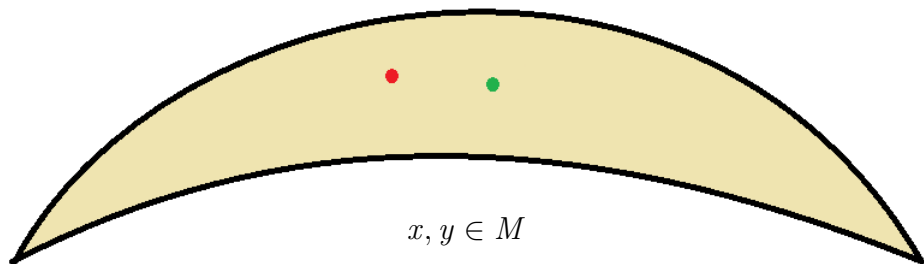
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



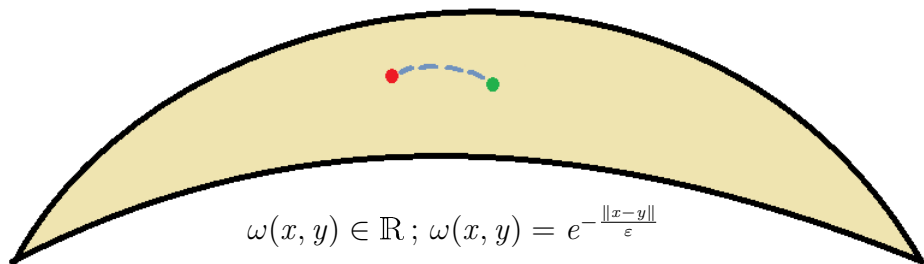
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



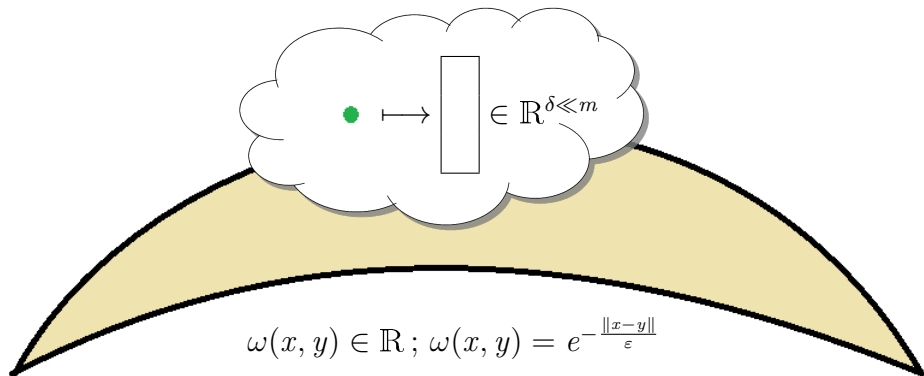
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



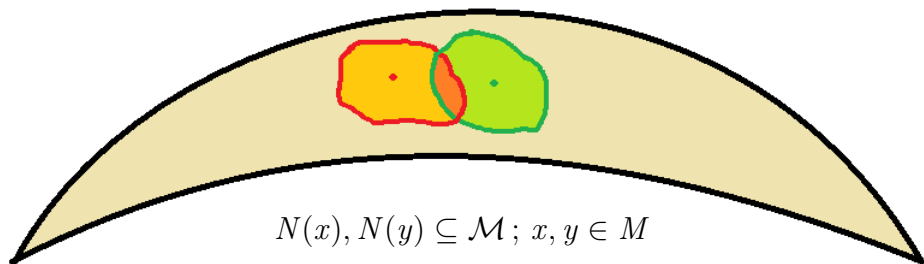
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



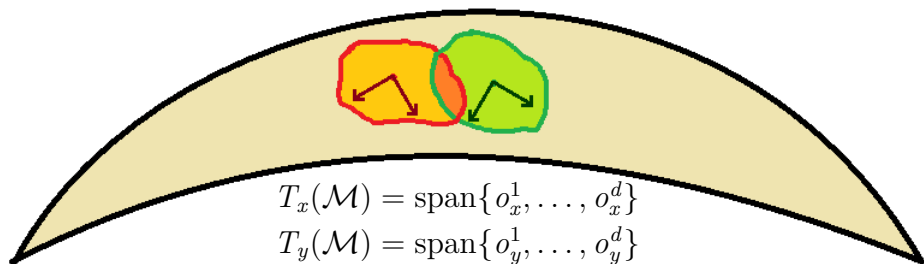
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



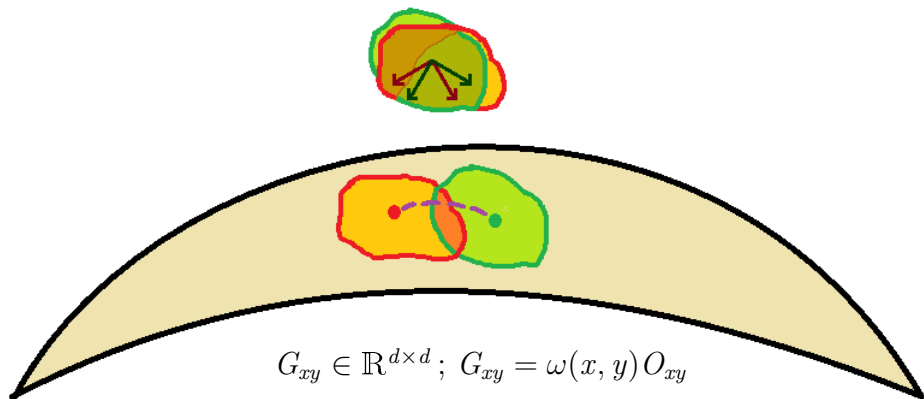
¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

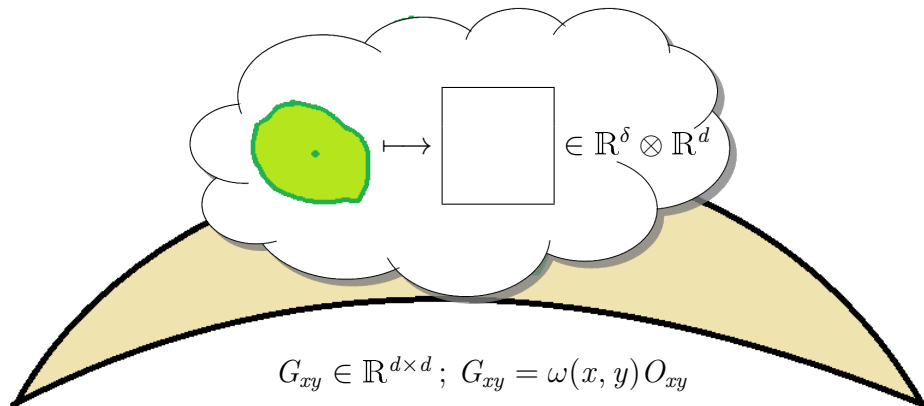
Patch-to-Tensor Embedding¹⁵



$$G_{xy} \in \mathbb{R}^{d \times d}; \quad G_{xy} = \omega(x, y) O_{xy}$$

¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Patch-to-Tensor Embedding¹⁵



¹⁵M. Salhov, GW, and A. Averbuch. “Patch-to-Tensor Embedding”. In: *Applied and Computational Harmonic Analysis* 33.2 (2012), pp. 182 –203.

Super-kernel

Vector field representation

A vector field can be represented in two equivalent ways:

$$\left\{ \begin{array}{c} \overbrace{[\text{---} \dots \text{---}]}^d \\ \vdots \\ \underbrace{[\text{---} \dots \text{---}]}_d \end{array} \right\}_n \equiv \left\{ \begin{array}{c} \left\{ \begin{array}{c} \vdots \\ \text{---} \end{array} \right\}_d \\ \vdots \\ \left\{ \begin{array}{c} \text{---} \\ \vdots \end{array} \right\}_d \end{array} \right\}_{nd}$$

Super-kernel

Definition (block matrix)

A super-kernel G is a block matrix / operator on vector-fields:

$$\left. \begin{matrix} n \times n \\ \left\{ \begin{matrix} d \times d \\ d \times d \end{matrix} \right\} \end{matrix} \right\} \left[\begin{array}{ccc} \boxed{} & \cdots & \boxed{} \\ \vdots & \boxed{G_{xy}} & \vdots \\ \boxed{} & \cdots & \boxed{} \end{array} \right] \left. \vphantom{\begin{matrix} n \times n \\ \left\{ \begin{matrix} d \times d \\ d \times d \end{matrix} \right\} \end{matrix}} \right\} nd \times nd$$

Super-kernel

Operator on vector-fields

The super-kernel as an operator on vector-fields:

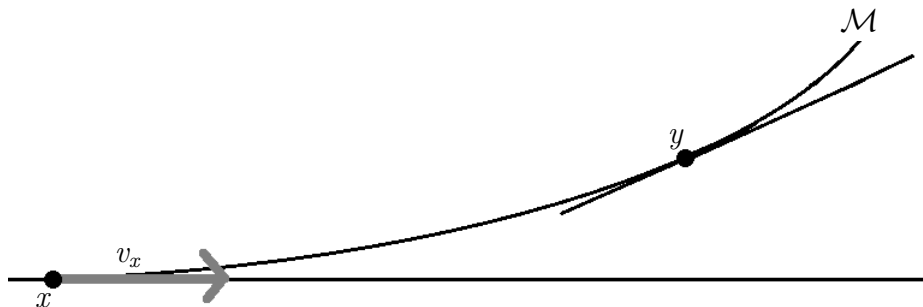
$$\vec{v}(x) \in T_x(\mathcal{M}) \xrightarrow{G} \sum G_{xy} \vec{v}(y) = G\vec{v}(x) \in T_x(\mathcal{M})$$

The diagram illustrates the super-kernel operator G as a sum over a diffusion affinity matrix O_{xy} weighted by a kernel $\omega(x,y)$. On the left, a vertical rectangle contains the expression $G\vec{v}(x)$. This is followed by an equals sign and a summation symbol \sum . Below the summation symbol is a small square containing ω with (x,y) below it. To the right of the summation is a large square labeled O_{xy} . Above this square is a horizontal brace labeled $d \times d$. Below the square is a horizontal brace labeled G_{xy} . To the right of the large square is another vertical rectangle containing $\vec{v}(x)$.

Diffusion super-kernel: $\omega(x, y) =$ the diffusion affinity $a(x, y)$

Linear-Projection Diffusion¹⁶

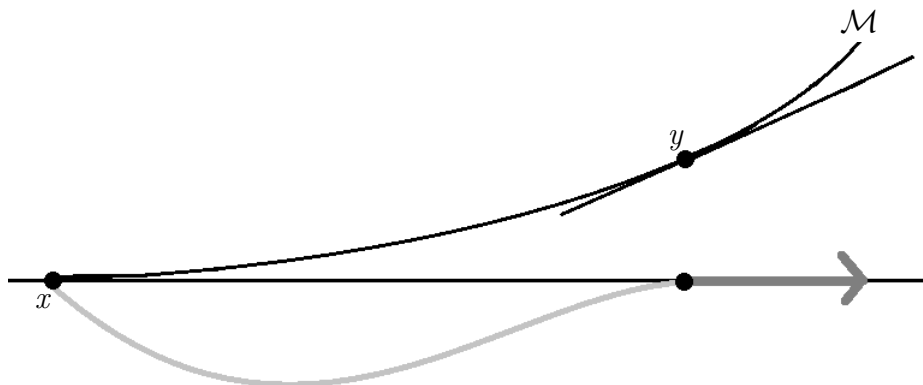
Diffusion step



¹⁶GW and A. Averbuch. “Linear-Projection Diffusion on Smooth Euclidean Submanifolds”. In: *Applied and Computational Harmonic Analysis* 34.1 (2013), pp. 1–14.

Linear-Projection Diffusion¹⁶

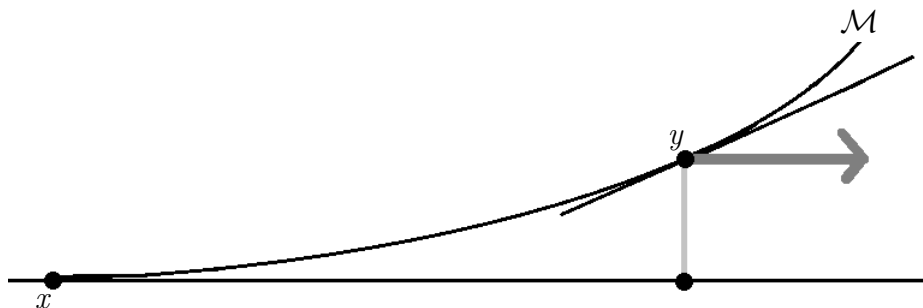
Diffusion step



¹⁶GW and A. Averbuch. “Linear-Projection Diffusion on Smooth Euclidean Submanifolds”. In: *Applied and Computational Harmonic Analysis* 34.1 (2013), pp. 1–14.

Linear-Projection Diffusion¹⁶

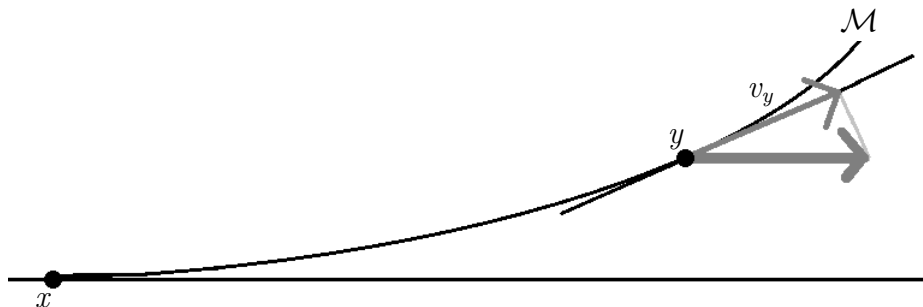
Diffusion step



¹⁶GW and A. Averbuch. “Linear-Projection Diffusion on Smooth Euclidean Submanifolds”. In: *Applied and Computational Harmonic Analysis* 34.1 (2013), pp. 1–14.

Linear-Projection Diffusion¹⁶

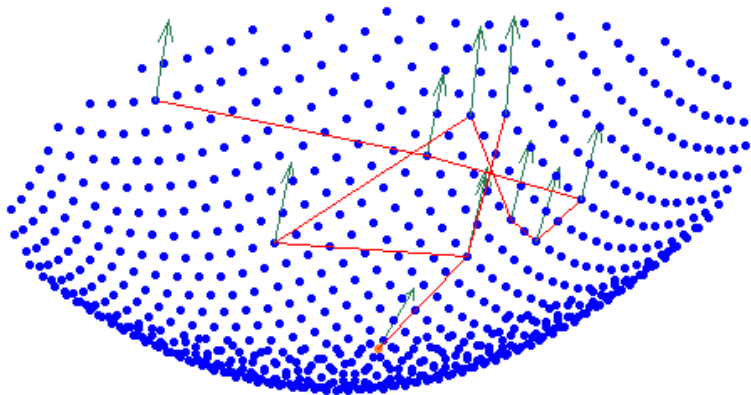
Diffusion step



¹⁶GW and A. Averbuch. “Linear-Projection Diffusion on Smooth Euclidean Submanifolds”. In: *Applied and Computational Harmonic Analysis* 34.1 (2013), pp. 1–14.

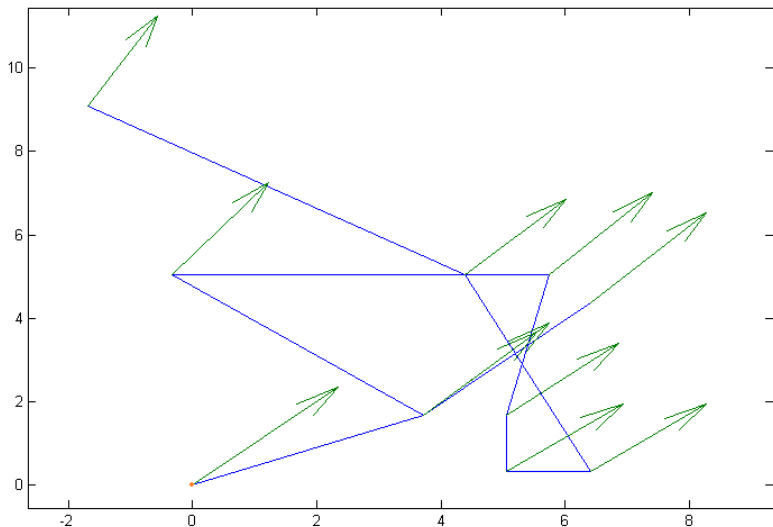
Linear-Projection Diffusion

Diffusion trajectory



Linear-Projection Diffusion

Projected trajectory



Linear-Projection Diffusion

Infinitesimal generator

Infinitesimal generator - a **partial derivative operator** that incorporates significant properties of a stochastic process

- Isotropic
- Anisotropic
- Y

Infinitesimal Generator:

$$\mathcal{L} = \lim_{\varepsilon \rightarrow 0} \frac{I - P_{\varepsilon}}{\varepsilon}$$

Linear-Projection Diffusion

Infinitesimal generator

Infinitesimal generator - a **partial derivative operator** that incorporates significant properties of a stochastic process

- Isotropic
- Anisotropic
- Vector

Heat Equation:

$$\frac{\partial u}{\partial t} = \alpha \nabla^2 u$$

Linear-Projection Diffusion

Infinitesimal generator

Infinitesimal generator - a **partial derivative operator** that incorporates significant properties of a stochastic process

- Isotropic Diffusion Maps: **Graph Laplacian**
- Anisotropic Diffusion Maps: **Laplace-Beltrami**
- Vector Diffusion Maps: **Connection Laplacian**

Linear-Projection Diffusion

Infinitesimal generator

Infinitesimal generator - a **partial derivative operator** that incorporates significant properties of a stochastic process

- Isotropic Diffusion Maps: **Graph Laplacian**
- Anisotropic Diffusion Maps: **Laplace-Beltrami**
- Vector Diffusion Maps: **Connection Laplacian**
- Linear-Projection Diffusion:
Vector-Laplacian of tangent projections

Linear-Projection Diffusion (LPD) Super-Kernel

Affinity blocks

Linear-Projection (LP) super-kernel consists of the blocks:

$$\underbrace{G_{xy}}_{d \times d} = \underbrace{\omega_{(x,y)}}_{(x,y)} \underbrace{O_x^T}_{d \times m} \underbrace{O_y}_{p \times u}$$

LPD super-kernel: $\omega(x, y) =$ the diffusion affinity $a(x, y)$

Linear-Projection Diffusion (LPD) Super-Kernel

Ambient coordinate independence

Proposition

The Linear-Projection super-kernel is independent of the coordinates of the ambient space.

Proof.

Consider a change of basis $O_x \mapsto BO_x$, where B is an orthogonal matrix:

$$\begin{aligned} G_{xy} = \omega(x, y)(BO_x)^T(BO_y) &= \omega(x, y)O_x^T B^T BO_y \\ &= \omega(x, y)O_x^T O_y \quad x, y \in \mathcal{M}. \end{aligned}$$

Therefore, the blocks G_{xy} remain the same in every basis of the ambient space. □

Linear-Projection Diffusion (LPD) Super-Kernel

Spectral properties

Ω = Scalar affinity kernel

G = Linear-Projection super-kernel

- Ω is positive (semi) definite $\Rightarrow G$ is positive (semi) definite
- Ω has operator norm $\|\Omega\| \Rightarrow G$ as operator norm $\|G\| \leq \|\Omega\|$

Spectrum of the LPD super-kernel

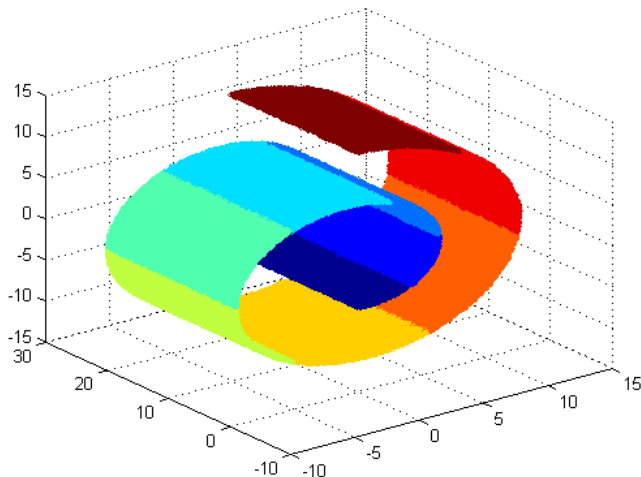
$$\Omega = A$$

$$\Downarrow$$

$$\|A\| = 1 \geq \lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots > 0$$

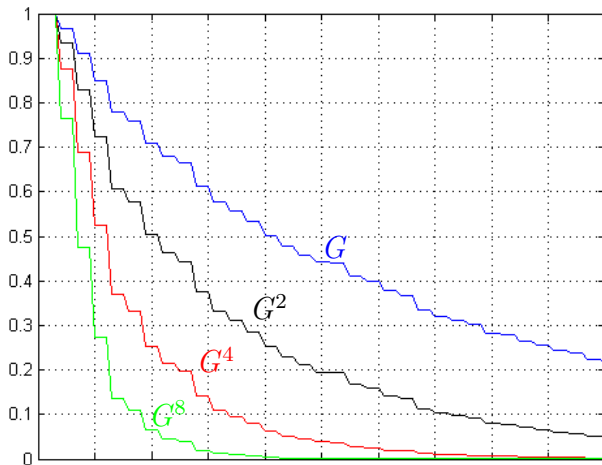
LPD-based Patch-to-Tensor Embedding

Spectral embedding



LPD-based Patch-to-Tensor Embedding

Spectral embedding

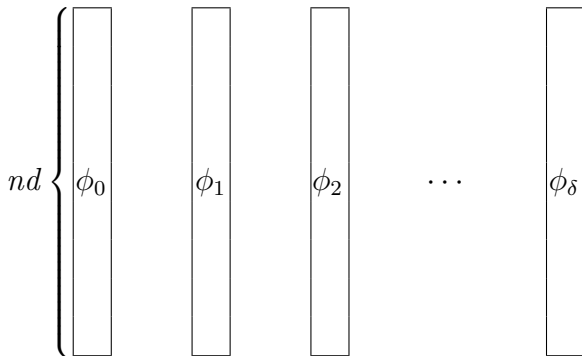


Spectrum (eigenvalues) of the LPD super-kernel G and its powers

LPD-based Patch-to-Tensor Embedding

Spectral embedding

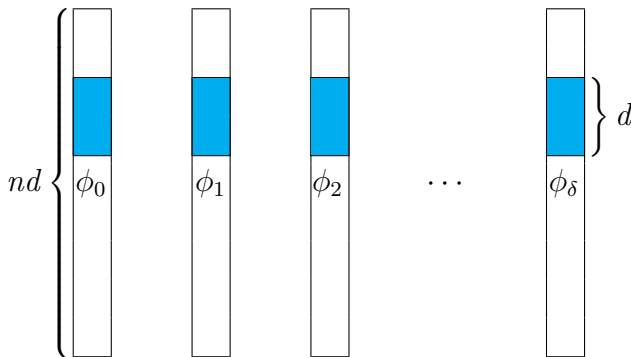
$$1 \geq \boxed{\lambda_0} \geq \boxed{\lambda_1} \geq \boxed{\lambda_2} \geq \dots \geq \boxed{\lambda_\delta} > 0$$



LPD-based Patch-to-Tensor Embedding

Spectral embedding

$$1 \geq \boxed{\lambda_0} \geq \boxed{\lambda_1} \geq \boxed{\lambda_2} \geq \dots \geq \boxed{\lambda_\delta} > 0$$

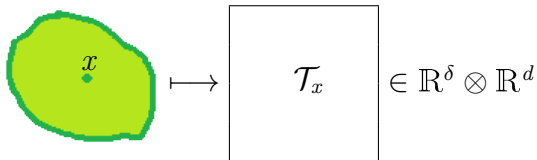


LPD-based Patch-to-Tensor Embedding

Embedded tensors

Embedded tensors

$$\mathcal{T}_x = \left(\begin{array}{c} \boxed{\lambda_1^t} \cdot \boxed{[- \vec{\phi}_1(x) -]} \\ \vdots \\ \boxed{\lambda_\delta^t} \cdot \boxed{[- \vec{\phi}_\delta(x) -]} \end{array} \right) \Bigg\} \delta \times d$$



LPD-based Patch-to-Tensor Embedding

Extended diffusion distances

PTE with the diffusion super-kernel

$$\overbrace{\|\mathcal{T}_x - \mathcal{T}_y\|_F^2}^{\text{embedded distances}} = \overbrace{\sum_{z \in M} \|a(x, z) O_{xz} - a(y, z) O_{yz}\|_F^2}^{\text{extended diffusion distances}}$$

PTE with the LPD super-kernel

$$\begin{aligned} \overbrace{\|\mathcal{T}_x - \mathcal{T}_y\|_F^2}^{\text{embedded distances}} &= \overbrace{\sum_{z \in M} \|a(x, z) O_x^T O_z - a(y, z) O_y^T O_z\|_F^2}^{\text{extended diffusion distances}} \\ &= \underbrace{\sum_{z \in M} \sum_{j=1}^d \|a(x, z) O_x^T o_z^j - a(y, z) O_y^T o_z^j\|^2}_{\text{LPD distances}} \end{aligned}$$

LPD-based Patch-to-Tensor Embedding

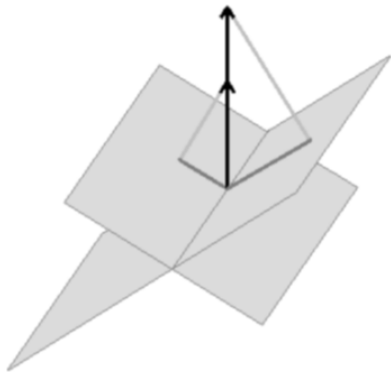
LPD distances

$$a(x, z) \quad a(y, z)$$

$$\|\mathcal{T}_x - \mathcal{T}_y\|_F^2 = \sum_{z \in M} \sum_{j=1}^d \|a(x, z) O_x^T o_z^j - a(y, z) O_y^T o_z^j\|^2$$

LPD-based Patch-to-Tensor Embedding

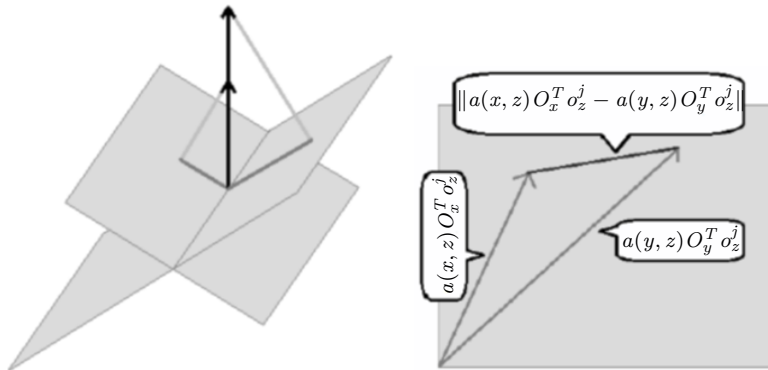
LPD distances



$$\|\mathcal{T}_x - \mathcal{T}_y\|_F^2 = \sum_{z \in M} \sum_{j=1}^d \|a(x, z) O_x^T o_z^j - a(y, z) O_y^T o_z^j\|^2$$

LPD-based Patch-to-Tensor Embedding

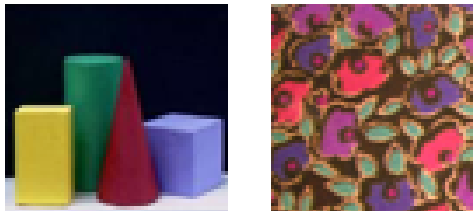
LPD distances



$$\|\mathcal{T}_x - \mathcal{T}_y\|_F^2 = \sum_{z \in M} \sum_{j=1}^d \|a(x, z) O_x^T o_z^j - a(y, z) O_y^T o_z^j\|^2$$

Image segmentation

Analysis algorithm



$H \times W$ RGB image $\longrightarrow (H \cdot W) \times 5$ dataset

- Patches of pixels $\xrightarrow{\text{PTE \& LPD}}$ embedded tensors
- Cluster the patches/tensors based on the embedded distances

Image segmentation

Achieved results

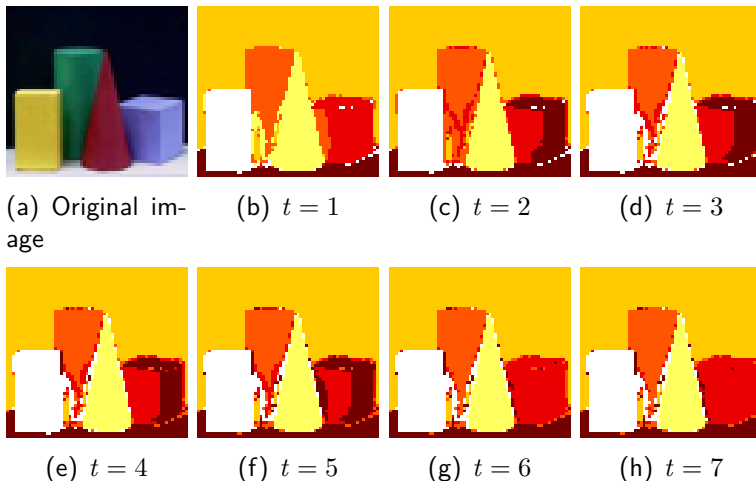


Figure : The PTE segmentation results shown at several diffusion times t .

Image segmentation

Additional results

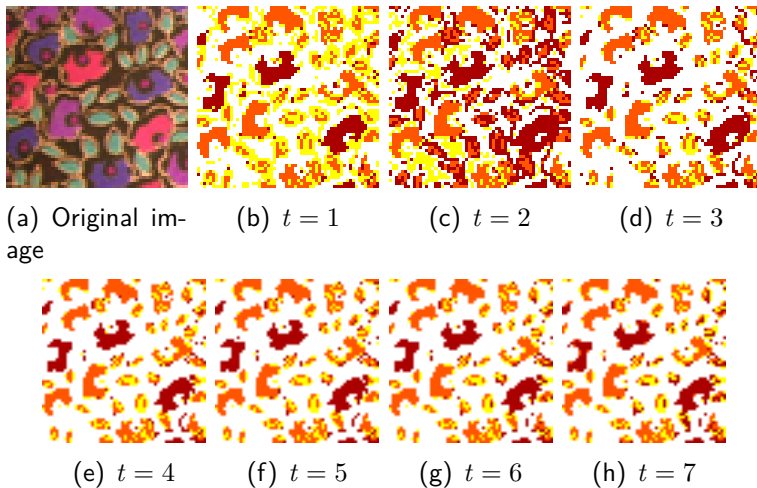


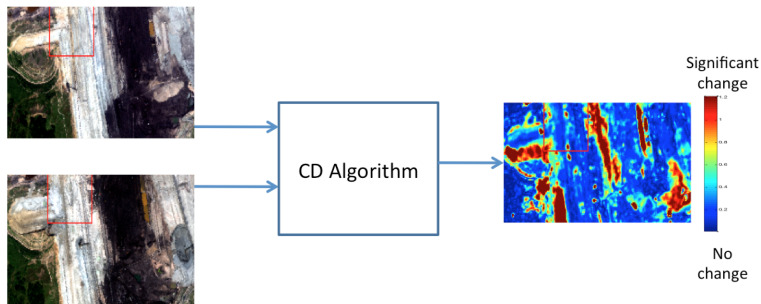
Figure : The PTE segmentation results shown at several diffusion times t .

Overview

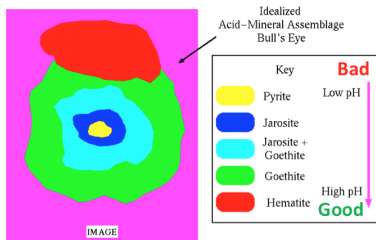
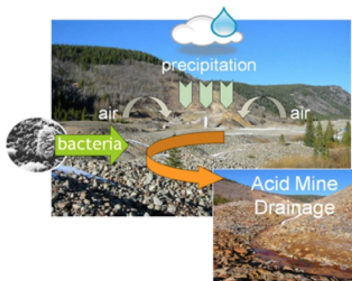
- 1 Introduction
- 2 Manifold theory
- 3 Diffusion maps
- 4 Patch-to-Tensor Embedding
- 5 Diffusion maps for changing data
 - Change detection
 - Time decoupled
 - Applications

Change detection

- Two images, different time acquisition.
- Where did the changes occur?



Mining hazards¹⁷



- Water acidity causes:
 - Vegetation stress and loss of biodiversity
 - Better transfer of heavy metals and other toxic components



¹⁷M. A. Sares, P. L. Hauff, D. C. Peters, D. W. Coulter, D. A. Bird, F. B. Henderson III, and E. C. Prosh. Characterizing Sources of Acid Rock Drainage and Resulting Water Quality Impacts Using Hyperspectral Remote Sensing - Examples from the Upper Arkansas River Basin, Colorado. *Advanced integration of geospatial technologies in mining reclamation*, December 7-9, 2004, Atlanta, Georgia, USA.

Problem description

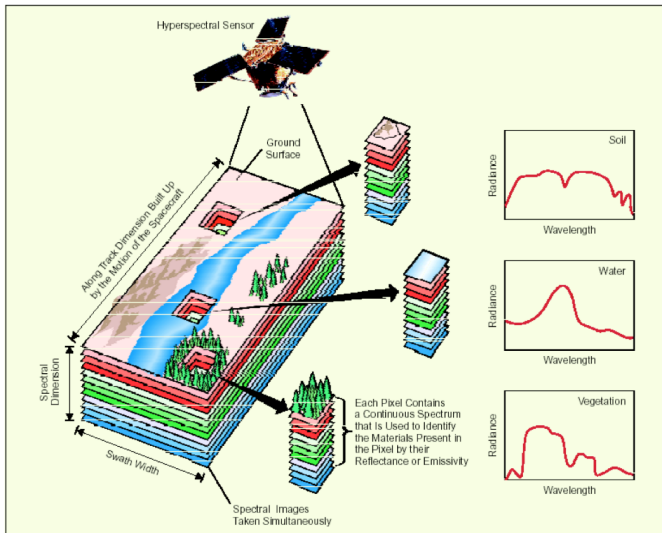
- Mining controlling authorities

- Field chemical measurements
- Area is too big

High cost of monitoring

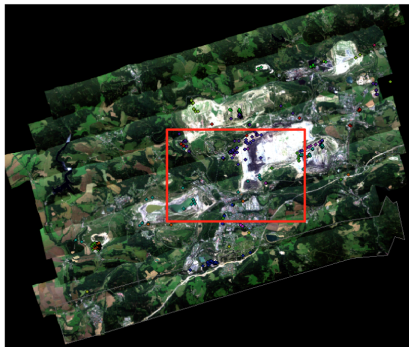


Solution: Imaging spectrometer



Sokolov coal mine area mosaic images

HyMap 2009



HyMap 2010



Sokolov coal mine

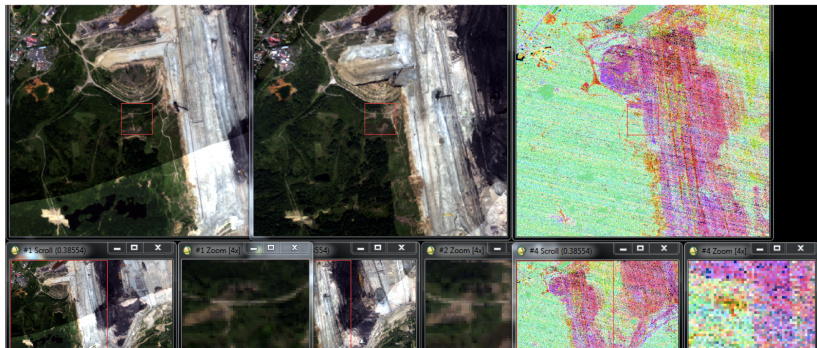
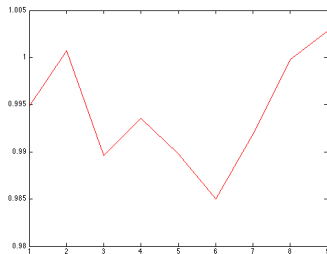
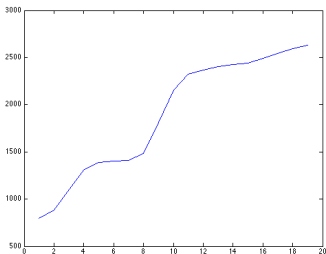


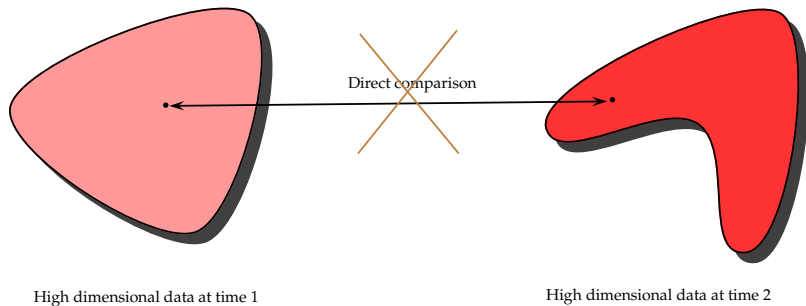
Figure : Sokolov coal mine in 2009 and 2010 (during the day, visible range, HyMap sensor), as well as 2011 (at night, thermal range, AHS sensor).

Sample spectral signatures

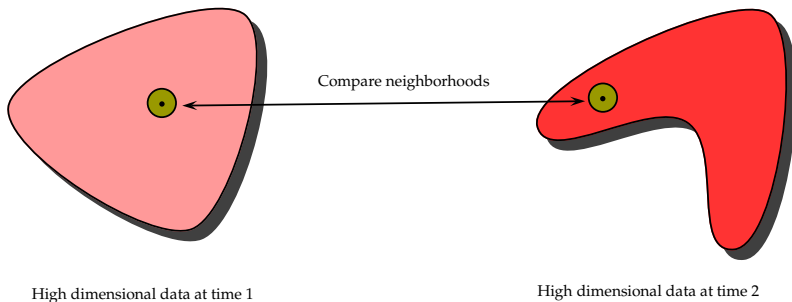


(a) Spectral signature of a vegetation pixel, 2009, during the day, visible range. (b) Spectral signature of the same pixel, 2011, at night, thermal range.

No direct comparison



Compare neighborhoods



- We are given several data sets $X_\tau = (X, d_\tau, \mu)$ indexed by $\tau \in \mathcal{I}$. Each data set is a metric measure space in which the metric depends on the parameter τ .
- We think of \mathcal{I} as the set of times or parameters for which we view/measure/observe (X, μ) .
- The changes in (X, μ) across \mathcal{I} are encoded in the metrics $\{d_\tau\}_{\tau \in \mathcal{I}}$.

Markov chains and diffusion distance

- For each $\tau \in \mathcal{I}$ we construct a homogenous Markov chain $P^{(\tau)}$ along with the corresponding transition kernel $p^{(\tau)}$.
- Let $\pi^{(\tau)}$ be the stationary distribution of $P^{(\tau)}$.
- Let $\{\lambda_i^{(\tau)}\}_{i \geq 0}$ denote the eigenvalues of $P^{(\tau)}$, and let $\{\psi_i^{(\tau)}\}_{i \geq 0}$ be a set of corresponding eigenfunctions.

Definition

The *diffusion* distance at time t between $x_\tau \in X_\tau$ and $y_\sigma \in X_\sigma$ is defined as:

$$D_t(x_\tau, y_\sigma)^2 \triangleq \int_X \left(\frac{p_t^{(\tau)}(x, z)}{\sqrt{\pi^{(\tau)}(z)}} - \frac{p_t^{(\sigma)}(y, z)}{\sqrt{\pi^{(\sigma)}(z)}} \right)^2 d\mu(z).$$

Diffusion maps

- For each $\tau \in \mathcal{I}$ we have the diffusion map

$$\Psi_t^{(\tau)}(x) \triangleq \left\{ \left(\lambda_i^{(\tau)} \right)^t \psi_i^{(\tau)}(x) \right\}_{i \geq 1}.$$

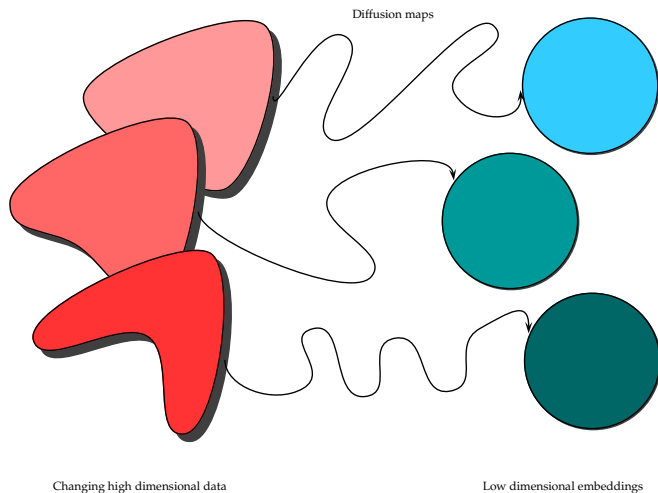
- The standard diffusion maps theory tells us that

$$D_t(x_\tau, y_\tau) = \left\| \Psi_t^{(\tau)}(x) - \Psi_t^{(\tau)}(y) \right\|_{\ell^2}.$$

- On the other hand, if $\tau \neq \sigma$, then

$$D_t(x_\tau, y_\sigma) \neq \left\| \Psi_t^{(\tau)}(x) - \Psi_t^{(\sigma)}(y) \right\|_{\ell^2}.$$

Multiple low dimensional embeddings



Symmetric diffusion operators

- Symmetrize $p^{(\tau)}$ by conjugating with the squareroot of the density function $d^{(\tau)}$,

$$a^{(\tau)}(x, y) \triangleq \sqrt{d^{(\tau)}(x)} \cdot p^{(\tau)}(x, y) \cdot \frac{1}{\sqrt{d^{(\tau)}(y)}}.$$

- Define the corresponding integral operator as:

$$(A^{(\tau)}f)(x) \triangleq \int_X a^{(\tau)}(x, y)f(y) d\mu(y).$$

- The eigenvalues of $A^{(\tau)}$ and $P^{(\tau)}$ are the same. Let $\{\varphi_i^{(\tau)}\}_{i \geq 0}$ denote the eigenfunctions of $A^{(\tau)}$.

Aligning the embeddings

- Let $a = \{a_i\}_{i \geq 0} \in \ell^2$.
- For $\tau, \sigma \in \mathcal{I}$, define the map $O_{\tau \rightarrow \sigma} : \ell^2 \rightarrow \ell^2$ as:

$$O_{\tau \rightarrow \sigma}(a) \triangleq \left\{ \sum_{j \geq 0} a_j \cdot \langle \varphi_j^{(\tau)}, \varphi_i^{(\sigma)} \rangle \right\}_{i \geq 0}.$$

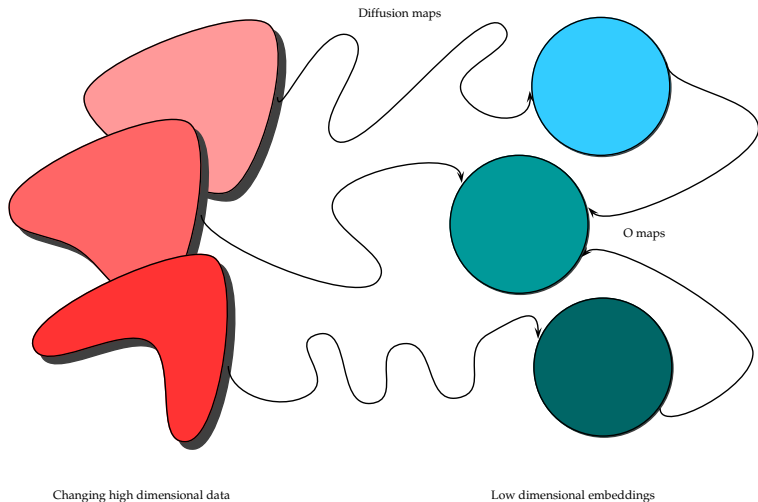
Theorem (Coifman and H., 2013¹⁸)

Fix $\gamma \in \mathcal{I}$. Then, for any $\tau, \sigma \in \mathcal{I}$, $t > 0$, and $x_\tau \in X_\tau$, $y_\sigma \in X_\sigma$,

$$D_t(x_\tau, y_\sigma) = \left\| O_{\tau \rightarrow \gamma} \Psi_t^{(\tau)}(x) - O_{\sigma \rightarrow \gamma} \Psi_t^{(\sigma)}(y) \right\|_{\ell^2}.$$

¹⁸R. R. Coifman and M. H. Diffusion maps for changing data. *Applied and Computational Harmonic Analysis*, In press.

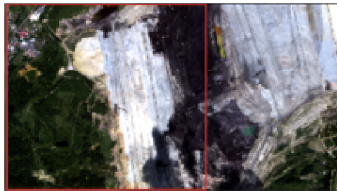
Single embedding



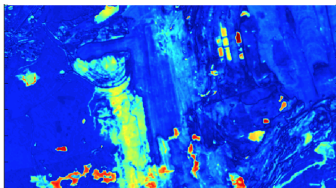
Change detection using diffusion maps



(c) Sokolov Mine 2009 / Day / Visible / HyMap sensor

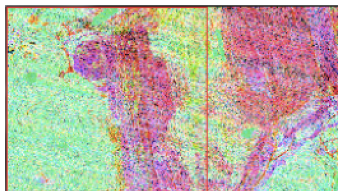


(d) Sokolov Mine 2011 / Day / Visible / AHS sensor

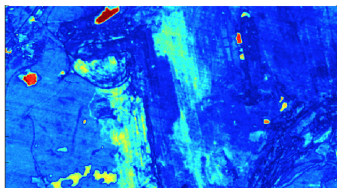


(e) Change map using diffusion maps.

Change detection using diffusion maps

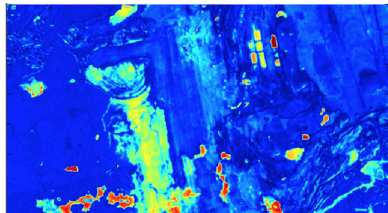


(f) Sokolov Mine 2009 / Day / Visible / HyMap sensor (g) Sokolov Mine 2011 / Night / Thermal / AHS sensor

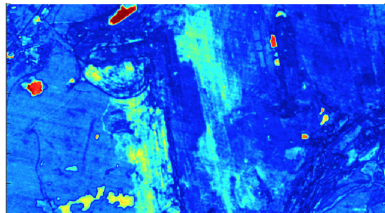


(h) Change map using diffusion maps

Comparison of change maps



(i) 2009 / Day / Visible / HyMap sensor vs. 2011 / Day / Visible / AHS sensor

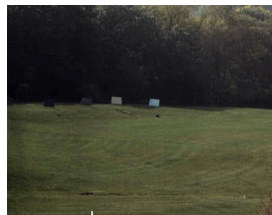


(j) 2009 / Day / Visible / HyMap sensor vs. 2011 / Night / Thermal / AHS sensor

Change detection via diffusion maps as $t \rightarrow \infty$



(k) September



(l) October

(m) $D^{(t)}(x_{\text{Sep}}, x_{\text{Oct}})$ as $t \rightarrow \infty$

The Standard Map

- The standard map is an area preserving chaotic map from a square with side length 2π onto itself.
- It is a discrete difference equation defined as:

$$\begin{aligned}p_{i+1} &= p_i + K \cdot \sin(\theta_i) \\ \theta_{i+1} &= \theta_i + p_{i+1}\end{aligned}$$

- $K = 0$: the map is linear, and only periodic and quasi-periodic orbits are possible.
- $K > 0$: the map is nonlinear, and certain initial conditions give chaotic dynamics. The larger K , the more chaotic.
- Used to model what can happen in certain plasma driven fusion experiments.

Ergodic quotient

- Let $x = (x_0, x_1, x_2, \dots)$ be a discrete trajectory on a manifold \mathcal{M} with $\dim(\mathcal{M}) = n$.
- Test the trajectory x against Fourier modes:

$$\phi_x(\omega) = \sum_{j=0}^{\infty} e^{2\pi i \omega \cdot x_j}, \quad \omega \in \mathbb{Z}^n.$$

- The distance between two trajectories x and y is then defined as:

$$d(x, y) = \sum_{\omega \in \mathbb{Z}^n} \frac{|\phi_x(\omega) - \phi_y(\omega)|^2}{(1 + |\omega|^2)^s},$$

where $s = (n + 1)/2$.

Diffusion Embedding of the Standard Map

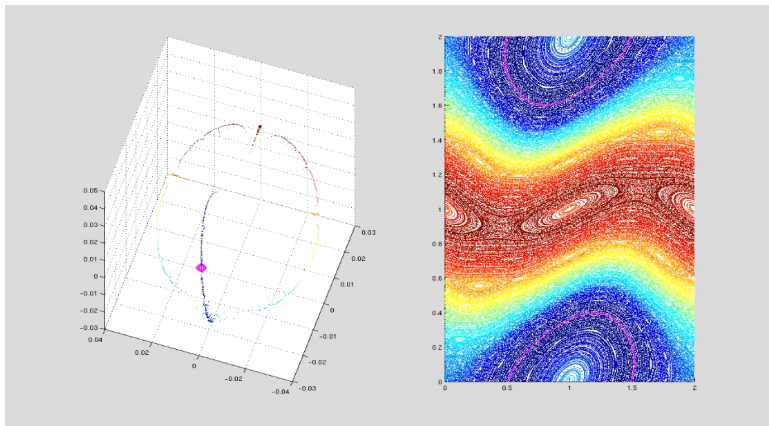


Figure : Standard map on right; diffusion embedding on left. The orbit of a particular color on the right corresponds to the embedded point of the same color on the left.

Diffusion embedding as K changes

Figure : Standard map on right; diffusion embedding on left.

Common embedding

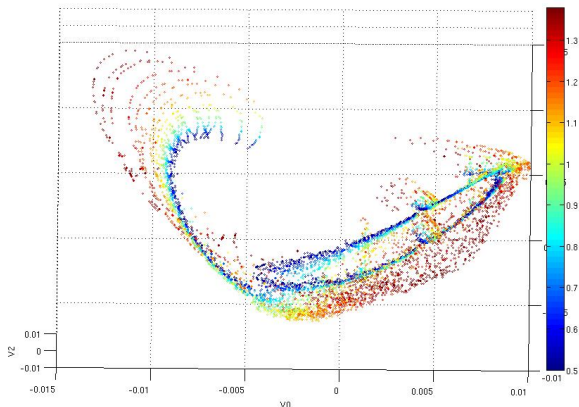


Figure : Diffusion embeddings of the standard map for various K all mapped into the same Euclidean space. The color of the points corresponds to the value of K (see color bar on right).

The End

