

Lecture 9: Manifold learning

Nonlinear dimensionality reduction

Nils Olovsson

nils.olofsson@igp.uu.se

Introduction to data science and Python programming for medical research
Uppsala University

2025



UPPSALA
UNIVERSITET

Content

Introduction

Multidimensional scaling (MDS)

Isomap

Laplacian eigenmap

Locally linear embedding (LLE)

Comparison

Real world

tSNE

UMAP

Conclusion

Introduction




Introduction

Introduction: Problem

Dimensionality reduction with PCA is achieved by **projecting** data points on the first PC vectors.

This **embeds** the data in the **PCA coordinate system**.


PCA is a **linear method**.




Introduction: Data with nonlinear *shape*

Data points **can be** organized in **non-linear** ways.


The underlying **shape** or **embedding** created by the data points can be **curved** or **folded**.



Original 2D data.



Linear 3D
embedding.



Nonlinear 3D
embedding.





Introduction: PCA fails on non-linear data

PCA can only model **linearly** embedded data.

This means that the **data** has to be lying **on a flat plane**.

(But the *plane* can have more dimensions than two.)

It will **fail** to represent **non-linear** data.



Introduction: Manifolds in physics


In physics and mathematics a **manifold** defines a topological space embedded in a higher dimension.

A **2D manifold** is simply a **surface**.

Can be **flat (linear)** or **curved and folded (non-linear)**.



A crumpled paper is a 2D manifold in 3D space.



Gravity bends space.


Introduction: Locally flat

A **manifold** will look **flat** where you stand.

Manifold learning attempts to find this flat **embedding**.



Unfolding
3D surface
of earth on
2D plane



The earth is not flat. But to us on the surface it looks like it.

The unfolding distorts the surface, making the poles look much larger than they really are.

Introduction: Manifold analysis of biomedical data

Visualization of very high dimensional data, e.g. transcriptomics, proteomics etc.

Image **analysis** and **modeling**, e.g. respiratory motion.





Fig. 3. Schematic of group connections through simultaneous **manifold** embeddings. Motion fields from neighbouring and orthogonal slices can be embedded simultaneously using appropriate similarity kernels leading to aligned embeddings. Two close-up views of aligned **manifold** embeddings, originating from a dataset with 50 motion fields per slice position, are shown on the right hand side.

<https://doi.org/10.1038/s41467-019-13056-x>

<https://doi.org/10.1016/j.media.2016.06.005>

Introduction: Manifold learning methods

Find mapping such that points close in original data are close after mapping. (Locally flat.)

Will look at several methods that can be divided into two families.

MDS	}	Spectral methods
Isomap		
Eigenmap		
LLE		

t-SNE	}	Optimization (Gradient descent)
UMAP		

sklearn.manifold

Data embedding techniques.

User guide. See the [Manifold learning](#) section for further details.

Isomap	Isomap Embedding.
LocallyLinearEmbedding	Locally Linear Embedding.
MDS	Multidimensional scaling.
SpectralEmbedding	Spectral embedding for non-linear dimensionality reduction.
TSNE	T-distributed Stochastic Neighbor Embedding.
locally_linear_embedding	Perform a Locally Linear Embedding analysis on the data.
smacof	Compute multidimensional scaling using the SMACOF algorithm.
spectral_embedding	Project the sample on the first eigenvectors of the graph Laplacian.
trustworthiness	Indicate to what extent the local structure is retained.

<https://scikit-learn.org/stable/api/sklearn.manifold.html>

<https://umap-learn.readthedocs.io/en/latest/>

Introduction: Toy datasets

First set of methods illustrated using toy data.

Hello (Flat)

Should recover the 2D letters in *Hello*.

Hello (S-shaped)

Should recover the 2D letters in *Hello*.


Spiral

Should recover the 1D line going from blue to red.


Swiss roll

Should recover the 2D sheet going from blue to red.


"Hello"
Linear 2D embedded in 3D.




Spiral
Curved 1D embedded in 2D.



S-shaped "Hello"
Curved 2D embedded in 3D.



"Swiss roll"
Curved 2D embedded in 3D.



Multidimensional scaling (MDS)

Multidimensional scaling (MDS)


Multidimensional scaling (MDS): Distance matrix

Finds a **manifold embedding** while minimizing a *strain* between data points.


Fully connected **distance matrix**.

Every point is connected to all other points.


Data points
(100 points)



Dense connections



Distance matrix
(100x100 matrix)



Multidimensional scaling (MDS): Eigenvalue problem

1. Distance matrix D
2. Double center the D .

$$B = -\frac{1}{2}CDC, \quad C = I - \frac{1}{N}J$$

3. Decompose into eigenvectors V and eigenvalues e
4. Embedding Y in first M dimensions is defined by

$$Y = \{\sqrt{e_1}V_1, \dots, \sqrt{e_M}V_M\}$$



```

1 def fit_transform(
2     self,
3     X: np.ndarray,
4 ) -> None:
5     # Pairwise distance between all points
6     D = sklearn.metrics.pairwise_distances(X)
7     D2 = np.pow(D, 2)
8
9     # Apply double centering
10    N = D.shape[0]
11    J = np.ones(D.shape, dtype=np.float64)
12    I = np.eye(N, dtype=np.float64)
13    C = I - (1.0/N) * J
14    B = -0.5 * C @ D2 @ C
15
16    # Eigendecomposition
17    e, V = np.linalg.eigh(B)
18    e = np.flip(e)
19    V = np.flip(V, axis=1)
20    V = V[:, 0:self._n_dimensions]
21    e = e[0:self._n_dimensions]
22    e = np.sqrt(e)
23
24    # Calculate low dimension embedding Y
25    Y = np.zeros(V.shape, dtype=np.float64)
26    for i in range(self._n_dimensions):
27        Y[:,i] = e[i] * V[:,i]
28    return Y

```

Multidimensional scaling (MDS): Example (linear)



```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import MDS
5 # Create data and project into 3D
6 N = 1000
7 X, y = make_hello(n_samples=N, rseed=0)
8 colorize = dict(c=X[:,0],
9                 cmap=plt.get_cmap('rainbow', 5))
10 X = random_projection_to_higher_dimension(X, 3,
11     rseed=0)
12 # Fit MDS, transform from 3D to 2D
13 mds = MDS(n_components=2)
14 X_mds = mds.fit_transform(X)
15 # Plot
16 fig = plt.figure(figsize=1.0 * plt.figaspect
17     (2.0/1.0))
18 gs = gridspec.GridSpec(nrows=2, ncols=1,
19     height_ratios=[1.5, 1])
20 # First subplot
21 ax0 = fig.add_subplot(gs[0], projection='3d')
22 ax1 = fig.add_subplot(gs[1])
23 ax0.scatter(X[:,0], X[:,1], X[:,2],
24     marker='o', **colorize)
25 ax1.scatter(X_mds[:,0], X_mds[:,1],
26     marker='o', **colorize)
27 #plt.show()
```



Multidimensional scaling (MDS): Example (S-shape)



Because MDS preserves all distances it **can't untangle non-linear data**.

Similarly to PCA it is limited to linear data which it can recover, disregarding orientation.



Multidimensional scaling (MDS): Example (roll)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import MDS
5 from sklearn.datasets import make_swiss_roll
6 # Create data and project into 3D
7 N = 2000
8 X, y = make_swiss_roll(n_samples=N,
9                        noise=0.1,
10                       random_state=0)
11 colorize = dict(c=y,
12                cmap=plt.get_cmap('turbo'))
13 # Fit MDS, transform from 3D to 2D
14 mds = MDS(n_components=2)
15 X_mds = mds.fit_transform(X)
16 # Plot
17 fig = plt.figure(figsize=1.0 * plt.figaspect
18                 (2.0/1.0))
19 gs = gridspec.GridSpec(nrows=2, ncols=1,
20                       height_ratios=[1.5, 1])
21 # First subplot
22 ax0 = fig.add_subplot(gs[0], projection='3d')
23 ax1 = fig.add_subplot(gs[1])
24 ax0.scatter(X[:,0], X[:,1], X[:,2],
25            marker='o', **colorize)
26 ax1.scatter(X_mds[:,0], X_mds[:,1],
27            marker='o', **colorize)
28 #plt.show()
```



Isomap

Isomap


Isomap: Nearest neighbor distance matrix

Finds a **manifold embedding** while minimizing a *strain* between data points same as for MDS.


Only **connect** points to their **nearest neighbors**.

We want the manifold to preserve **local structure**!


Data points
(100 points)



Connect to k=10
nearest neighbors



Distance matrix
(100x100 matrix)




Isomap: Geodesic distance

Distances between non-connected points are now computed by finding their shortest pairwise distances on the graph.

This approximates their **geodesic** distance.


A **geodesic** is the **shortest path** in a curved space, *i.e.* on a **manifold**.



<https://brickisland.net/ddg-web/>

Isomap: Geodesic matrix

The zeros in the distance matrix, D are filled in by traversing the connections to find the geodesic distances and create geodesic matrix G . (Floyd-Warshall algorithm)



$$D = \begin{pmatrix} 0.0 & 0.6 & 0.7 & 0.0 & 0.0 & 0.0 \\ 0.6 & 0.0 & 0.4 & 0.0 & 0.0 & 0.0 \\ 0.7 & 0.4 & 0.0 & 0.6 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.6 & 0.0 & 0.5 & 0.7 \\ 0.0 & 0.0 & 0.0 & 0.5 & 0.0 & 0.5 \\ 0.0 & 0.0 & 0.0 & 0.7 & 0.5 & 0.0 \end{pmatrix}$$


```
1 # For the zero entries in the distance matrix,  
2 # compute estimates of the geodesic distances  
3 # between pairs of points by computing  
4 # distances on the graph with weights  
5 # according to the Euclidean distances.  
6 G = scipy.sparse.csgraph.floyd_warshall(D)  
7 G2 = np.power(G, 2)  
8 G2 = np.maximum(G2, G2.T)
```

$$G_2 = \begin{pmatrix} 0.0 & 0.4 & 0.5 & 1.8 & 3.5 & 4.3 \\ 0.4 & 0.0 & 0.2 & 1.2 & 2.5 & 3.2 \\ 0.5 & 0.2 & 0.0 & 0.4 & 1.3 & 1.8 \\ 1.8 & 1.2 & 0.4 & 0.0 & 0.3 & 0.5 \\ 3.5 & 2.5 & 1.3 & 0.3 & 0.0 & 0.3 \\ 4.3 & 3.2 & 1.8 & 0.5 & 0.3 & 0.0 \end{pmatrix}$$


Isomap: Geodesic matrix vs MDS distance matrix

MDS:


Data points
(100 points)



Dense connections




Distance matrix
(100x100 matrix)




Isomap:


Data points
(100 points)



Connect to k=10
nearest neighbors



Geodesic distance matrix
(100x100 matrix)



Isomap: Eigenvalue problem


1. Sparse distance matrix D
2. Squared geodesic matrix G_2
3. Double center the G_2 .

$$B = -\frac{1}{2}CG_2C, \quad C = I - \frac{1}{N}J$$

4. Decompose into eigenvectors V and eigenvalues e
5. Embedding Y in first M dimensions is defined by



$$Y = \{\sqrt{e_1}V_1, \dots, \sqrt{e_M}V_M\}$$

Same eigenvalue problem as for MDS
but swap D for G_2 !





Isomap: Example (linear)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import Isomap
5 # Create data and project into 3D
6 N = 1000
7 X, y = make_hello(n_samples=N, rseed=0)
8 colorize = dict(c=X[:,0],
9                 cmap=plt.get_cmap('rainbow', 5))
10 X = random_projection_to_higher_dimension(X, 3,
11                                           rseed=0)
12 # Fit Isomap, transform from 3D to 2D
13 iso = Isomap(n_components=2,
14             n_neighbors=75,
15             eigen_solver='dense')
16 X_iso = iso.fit_transform(X)
17 # Plot
18 fig = plt.figure(figsize=1.0 * plt.figaspect
19                 (2.0/1.0))
20 gs = gridspec.GridSpec(nrows=2, ncols=1,
21                       height_ratios=[1.5, 1])
22 # First subplot
23 ax0 = fig.add_subplot(gs[0], projection='3d')
24 ax1 = fig.add_subplot(gs[1])
25 ax0.scatter(X[:,0], X[:,1], X[:,2],
26            marker='o', **colorize)
27 ax1.scatter(X_iso[:,0], X_iso[:,1],
28            marker='o', **colorize)
29 #plt.show()
```





Isomap: Example (S-shape)

Because Isomap preserves local distances
non-linear data can be untangled!



Isomap: Example (roll)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import Isomap
5 from sklearn.datasets import make_swiss_roll
6 # Create data and project into 3D
7 N = 2000
8 X, y = make_swiss_roll(n_samples=N,
9                       noise=0.1,
10                      random_state=0)
11
12 colorize = dict(c=y,
13               cmap=plt.get_cmap('turbo'))
14 # Fit Isomap, transform from 3D to 2D
15 iso = Isomap(n_components=2,
16             n_neighbors=10,
17             eigen_solver='dense')
18 X_iso = iso.fit_transform(X)
19 # Plot
20 fig = plt.figure(figsize=1.0 * plt.figaspect
21                 (2.0/1.0))
22 gs = gridspec.GridSpec(nrows=2, ncols=1,
23                       height_ratios=[1.5, 1])
24 # First subplot
25 ax0 = fig.add_subplot(gs[0], projection='3d')
26 ax1 = fig.add_subplot(gs[1])
27 ax0.scatter(X[:,0], X[:,1], X[:,2],
28            marker='o', **colorize)
29 ax1.scatter(X_iso[:,0], X_iso[:,1],
30            marker='o', **colorize)
31 #plt.show()
```



Laplacian eigenmap


Laplacian eigenmap

Laplacian eigenmap: Nearest neighbor adjacency matrix


Only connect points to their neighbors.
Discard distance, keep only
connections.

$\begin{cases} 1 : \text{Points are connected} \\ 0 : \text{Points are not connected} \end{cases}$


Data points
(100 points)



Connect to k=10
nearest neighbors




Adjacency matrix
(100x100 matrix)



Laplacian eigenmap: Adjacency matrix

Connect the k=2 nearest points.



```

1 N = X.shape[0]
2 A0 = np.zeros((N, N), dtype=np.float64)
3 for i in range(N):
4     x_i = X[i, :]
5     distances = np.linalg.norm(X - x_i, axis
6                               =1)
7     indices = np.argsort(distances)
8     indices = indices[1:k + 1]
9     A0[i, indices] = 1
10 # Symmetrize the matrix
11 A = np.maximum(A0, A0.T)

```


$$D = \begin{pmatrix} 0.0 & 0.6 & 0.7 & \mathbf{1.3} & \mathbf{1.7} & \mathbf{2.0} \\ 0.6 & 0.0 & 0.4 & \mathbf{0.7} & \mathbf{1.0} & \mathbf{1.4} \\ \mathbf{0.7} & 0.4 & 0.0 & 0.6 & \mathbf{1.1} & \mathbf{1.3} \\ \mathbf{1.3} & \mathbf{0.7} & 0.6 & 0.0 & 0.5 & \mathbf{0.7} \\ \mathbf{1.7} & \mathbf{1.0} & \mathbf{1.1} & 0.5 & 0.0 & 0.5 \\ \mathbf{2.0} & \mathbf{1.4} & \mathbf{1.3} & 0.7 & 0.5 & 0.0 \end{pmatrix}$$

$$A_0 = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}$$

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ \mathbf{1} & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & \mathbf{1} \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}$$

Laplacian eigenmap: Degree matrix

Degree matrix H describes **how many connections** a data point have. The matrix is **diagonal**.



$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}$$

$$H = \begin{pmatrix} 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix}$$

Graph Laplacian operator is defines as $L = H - A$.


Normalized **graph Laplacian** for this eigenvalue problem defined as

$$L = H^{-\frac{1}{2}}(H - A)H^{-\frac{1}{2}}$$

Laplacian eigenmap: Eigenvalue problem



1. Distance matrix D
2. Degree matrix H .
3. Normalized graph Laplacian
$$L = H^{-\frac{1}{2}}(H - A)H^{-\frac{1}{2}}$$
4. Decompose into eigenvectors V and eigenvalues e
5. Ignore the first eigenvector since eigenvalue is always 0. Embedding Y in first M dimensions is defined by

$$Y = \{\sqrt{e_2}V_2, \dots, \sqrt{e_{M+1}}V_{M+1}\}$$





Laplacian eigenmap: Example (linear)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import SpectralEmbedding
5 # Create data and project into 3D
6 N = 1000
7 X, y = make_hello(n_samples=N, rseed=0)
8 colorize = dict(c=X[:,0],
9                 cmap=plt.get_cmap('rainbow', 5))
10 X = random_projection_to_higher_dimension(X, 3,
11                                           rseed=0)
12 # Fit Isomap, transform from 3D to 2D
13 lap = SpectralEmbedding(n_components=2,
14                         n_neighbors=75)
15 X_lap = lap.fit_transform(X)
16 # Plot
17 fig = plt.figure(figsize=1.0 * plt.figaspect
18                 (2.0/1.0))
19 gs = gridspec.GridSpec(nrows=2, ncols=1,
20                       height_ratios=[1.5, 1])
21 # First subplot
22 ax0 = fig.add_subplot(gs[0], projection='3d')
23 ax1 = fig.add_subplot(gs[1])
24 ax0.scatter(X[:,0], X[:,1], X[:,2],
25            marker='o', **colorize)
26 ax1.scatter(X_lap[:,0], X_lap[:,1],
27            marker='o', **colorize)
28 #plt.show()
```



Laplacian eigenmap: Example (roll)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import SpectralEmbedding
5 from sklearn.datasets import make_swiss_roll
6 # Create data and project into 3D
7 N = 2000
8 X, y = make_swiss_roll(n_samples=N,
9                       noise=0.1,
10                      random_state=0)
11 colorize = dict(c=y,
12                cmap=plt.get_cmap('turbo'))
13 # Fit Spec., transform from 3D to 2D
14 lap = SpectralEmbedding(n_components=2,
15                       n_neighbors=40)
16 X_lap = lap.fit_transform(X)
17 # Plot
18 fig = plt.figure(figsize=1.0 * plt.figaspect
19                 (2.0/1.0))
20 gs = gridspec.GridSpec(nrows=2, ncols=1,
21                       height_ratios=[1.5, 1])
22 # First subplot
23 ax0 = fig.add_subplot(gs[0], projection='3d')
24 ax1 = fig.add_subplot(gs[1])
25 ax0.scatter(X[:,0], X[:,1], X[:,2],
26            marker='o', **colorize)
27 ax1.scatter(X_lap[:,0], X_lap[:,1],
28            marker='o', **colorize)
29 #plt.show()
```



Locally linear embedding (LLE)


Locally linear embedding (LLE)

Locally linear embedding (LLE): Eigenvalue problem

Optimize a weight matrix W that reconstructs a data point from a weighted sum of its neighbors.



1. Weight matrix W
2. Reconstruction matrix
 $K = (I - W)^T(I - W)$
3. Decompose into eigenvectors V and eigenvalues e
4. Embedding Y in first M dimensions is defined by

$$Y = \{\sqrt{e_1} V_1, \dots, \sqrt{e_M} V_M\}$$





Locally linear embedding (LLE): Example (linear)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import LocallyLinearEmbedding
5 # Create data and project into 3D
6 N = 1000
7 X, y = make_hello(n_samples=N, rseed=0)
8 colorize = dict(c=X[:,0],
9                 cmap=plt.get_cmap('rainbow', 5))
10 X = random_projection_to_higher_dimension(X, 3,
11                                           rseed=0)
12 # Fit LLE, transform from 3D to 2D
13 lle = LocallyLinearEmbedding(n_components=2,
14                             n_neighbors=75,
15                             method='modified')
16 X_lle = lle.fit_transform(X)
17 # Plot
18 fig = plt.figure(figsize=1.0 * plt.figaspect
19                  (2.0/1.0))
20 gs = gridspec.GridSpec(nrows=2, ncols=1,
21                        height_ratios=[1.5, 1])
22 # First subplot
23 ax0 = fig.add_subplot(gs[0], projection='3d')
24 ax1 = fig.add_subplot(gs[1])
25 ax0.scatter(X[:,0], X[:,1], X[:,2],
26            marker='o', **colorize)
27 ax1.scatter(X_lle[:,0], X_lle[:,1],
28            marker='o', **colorize)
29 #plt.show()
```





Locally linear embedding (LLE): Example (S-shape)

Because LLE preserves local distances
non-linear data can be untangled!



Locally linear embedding (LLE): Example (roll)

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from matplotlib import gridspec
4 from sklearn.manifold import LocallyLinearEmbedding
5 from sklearn.datasets import make_swiss_roll
6 # Create data and project into 3D
7 N = 2000
8 X, y = make_swiss_roll(n_samples=N,
9                       noise=0.1,
10                      random_state=0)
11
12 colorize = dict(c=y,
13               cmap=plt.get_cmap('turbo'))
14 # Fit Spec., transform from 3D to 2D
15 lle = LocallyLinearEmbedding(n_components=2,
16                             n_neighbors=10,
17                             method='modified')
18 X_lle = lle.fit_transform(X)
19 # Plot
20 fig = plt.figure(figsize=1.0 * plt.figaspect
21                  (2.0/1.0))
22 gs = gridspec.GridSpec(nrows=2, ncols=1,
23                       height_ratios=[1.5, 1])
24 # First subplot
25 ax0 = fig.add_subplot(gs[0], projection='3d')
26 ax1 = fig.add_subplot(gs[1])
27 ax0.scatter(X[:,0], X[:,1], X[:,2],
28            marker='o', **colorize)
29 ax1.scatter(X_lle[:,0], X_lle[:,1],
30            marker='o', **colorize)
31 #plt.show()
```



Comparison

Comparison

Comparison: Eigenvalue problems

All methods so far have been based on **eigen decomposition** of some **connection matrix**.

Methods are **linear** but the **embeddings** are **non-linear**.

MDS

$$B = -\frac{1}{2}CDC$$

Isomap

$$B = -\frac{1}{2}CG_2C$$


Laplacian eigenmap

$$L = H^{-\frac{1}{2}}(H - A)H^{-\frac{1}{2}}$$

Locally linear embedding

$$K = (I - W)^T(I - W)$$

Comparison: Results



Comparison: Manifold learning vs PCA

Eigenvectors

PCA Embed data by projecting it on eigenvectors.

Manifolds Eigenvectors are the embedding.

Applying model

PCA Project data on eigenvectors.

Manifolds Data is the model. Find points close in data space and interpolate in embedded space.

Matrix size

PCA Covariance matrix size determines by dimensionality of data.

Manifolds Connection matrix size determined by number of data points.

(Matrix size matters for computational resources!

100 000 data points $\rightarrow$ 100 000 $\times$ 100 000 matrix requires 40 GB!)

```
1 import numpy as np
2 from sklearn.datasets import
  make_swiss_roll
3 from sklearn.metrics import
  pairwise_distances
4 X, _ = make_swiss_roll(n_samples=1000)
5 C = np.cov(X.T)
6 D = pairwise_distances(X)
7 print("Data dimensions:")
8 print(X.shape[1])
9 print("Number of data points:")
10 print(X.shape[0])
11 print("Covariance matrix size:")
12 print(C.shape)
13 print("Distance matrix size:")
14 print(D.shape)
```

```
1 >> Data dimensions:
2 >> 3
3 >> Number of data points:
4 >> 1000
5 >> Covariance matrix size:
6 >> (3, 3)
7 >> Distance matrix size:
8 >> (1000, 1000)
```


Comparison: Problems

Isomap, Laplacian eigenmaps and LLE all need a parameter for **k = number of neighbors**. Can lead to **problems**.

Manifold shortcuts

Points that should not be connected are connected.


Solution: Decrease k.



Disjoint graph

The graph does not connect all points.
(Eigenvalue problem is ill posed.)


Solution: Increase k.




Real world

Real world


Real world: Single cell analysis




scRNA-Seq data analysis



ADT data analysis




- Differential expression



- tSNE




- PCA



Real world: Clustered data

Manifold learning used for visualization assumes **data** to be highly **clustered**, e.g. according to cell type.

Motivates specialized methods like **tSNE** and **UMAP**.




Real world: tSNE and UMAP

tSNE and UMAP are the two most commonly used manifold learning methods for visualizing data from cell assays.

Probabilistic interpretation of whether **points belong together** in **high** and **low dimensional** space.

Methods themselves are **nonlinear**, can't be solved by eigen decomposition. Must be **solved by** optimization, e.g. **gradient descent**.



Follow the gradient arrows to find values for parameters 1 and 2 that minimize the model error.

Introduction
oooooooo

MDS
oooooo

Isomap
oooooooo

Eigenmap
oooooo

LLE
ooooo

Comparison
ooooo

Real world
oooo

tSNE
●oooooooo

UMAP
oooooooo

Conclusion
ooo

tSNE

tSNE

tSNE: Overview

t-distributed stochastic neighborhood embedding (tSNE).

Described by four equations:

1. Distance in high dimensional data space.
2. Perplexity
3. Distance in low dimensional embedding.
4. Cost function to minimize.

Journal of Machine Learning Research 9 (2008) 2579-2605
Submitted 5/08; Revised 9/08; Published 11/08

Visualizing Data using t-SNE

Laurens van der Maaten
*TCC
 Tilburg University
 P.O. Box 90153, 5000 LE Tilburg, The Netherlands*
 LVDM.AATEN@GMAIL.COM

Geoffrey Hinton
*Department of Computer Science
 University of Toronto
 6 King's College Road, M5S 3G4 Toronto, ON, Canada*
 HINTON@CS.TORONTO.EDU

Editor: Yoshua Bengio

Abstract

We present a new technique called "t-SNE" that visualizes high-dimensional data by giving each datapoint a location in a two or three-dimensional map. The technique is a variation of Stochastic Neighbor Embedding (Hinton and Roweis, 2002) that is much easier to optimize, and produces significantly better visualizations by reducing the tendency to crowd points together in the center of the map. t-SNE is better than existing techniques at creating a single map that reveals structure at many different scales. This is particularly important for high-dimensional data that lie on several different, but related, low-dimensional manifolds, such as images of objects from multiple classes seen from multiple viewpoints. For visualizing the structure of very large data sets, we show how t-SNE can use random walks on neighborhood graphs to allow the implicit structure of all of the data to influence the way in which a subset of the data is displayed. We illustrate the performance of t-SNE on a wide variety of data sets and compare it with many other non-parametric visualization techniques, including Sammon mapping, Isomap, and Locally Linear Embedding. The visualizations produced by t-SNE are significantly better than those produced by the other techniques on almost all of the data sets.

Keywords: visualization, dimensionality reduction, manifold learning, embedding algorithms, multidimensional scaling

1. Introduction

Visualization of high-dimensional data is an important problem in many different domains, and deals with data of widely varying dimensionality. Cell nuclei that are relevant to breast cancer, for example, are described by approximately 30 variables (Street et al., 1993), whereas the pixel intensity vectors used to represent images or the word-count vectors used to represent documents typically have thousands of dimensions. Over the last few decades, a variety of techniques for the visualization of such high-dimensional data have been proposed, many of which are reviewed by de Oliveira and Levkowitz (2003). Important techniques include iconographic displays such as Chernoff faces (Chernoff, 1973), pixel-based techniques (Keim, 2000), and techniques that represent the dimensions in the data as vertices in a graph (Battista et al., 1994). Most of these techniques simply provide tools to display more than two data dimensions, and leave the interpretation of the

tSNE: (1) High dimensional distances

Probability of observing distances between points in high dimensional, data, space assumed to be **normally distributed**.
The probability is **normalized**.

Joint probability is **symmetrized** such that $p_{ij} = p_{ji}$.

Conditional probability:

$$p_{j|i} = \frac{\exp\left(\frac{-||x_i - x_j||^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-||x_i - x_k||^2}{2\sigma_i^2}\right)}$$

Joint probability:

$$p_{ij} = \frac{p_{i|j} + p_{j|i}}{2N}$$

tSNE: (2) Perplexity

The *perplexity* is a value provided by the user as method parameter that influences the solution.

It is then used "in reverse" to solve for a value of the standard deviations σ_i in equation 1.

Computed from the [entropy](#).

$$\text{perplexity} = 2^{-\sum_j (p_{j|i} \log_2(p_{j|i}))}$$

`perplexity : float, default=30.0`

The perplexity is related to the number of nearest neighbors that is used in other manifold learning algorithms. Larger datasets usually require a larger perplexity. Consider selecting a value between 5 and 50. Different values can result in significantly different results. The perplexity must be less than the number of samples.

tSNE: (3) Low dimensional distances

Probability of observing distances between points in low dimensional, embedding, space assumed to be **t-distributed**.

The probability is **normalized**.

(This is actually a t-distribution with one degree of freedom, also called Cauchy distribution.)

$$q_{ji} = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{i \neq j} (1 + \|y_i - y_j\|^2)^{-1}}$$

tSNE: (4) Cost function

The cost function is based on the **Kullback-Leibler divergence** between the probability distributions in high dimensional data space, P , and in low dimensional embedding, Q .

The embedding is optimized by minimizing the cost function.

This is done by calculating the gradient of KL with respect to each embedding position.


Minimized by gradient descent.

$$\text{KL}(P_i || Q_i) = \sum_i \sum_j p_{ij} \log \left(\frac{p_{ij}}{q_{ij}} \right)$$

tSNE: Example

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 from sklearn.manifold import TSNE
4
5 # Cancer associated fibroblast dataset and transform
6 X, y = load_CAF()
7 X = np.log(X + 1)
8
9 # Instantiate and apply TSNE
10 tsne = TSNE(n_components=2, random_state=0)
11 X_tsne = tsne.fit_transform(X)
12
13 # Plot result
14 ax = plt.figure(figsize=(3,3)).gca()
15 ax.scatter(X_tsne[:,0],
16           X_tsne[:,1],
17           marker='o',
18           c=y.astype(int),
19           edgecolor='k',
20           )
21 ax.set_xlabel('tSNE 1')
22 ax.set_ylabel('tSNE 2')
  
```



tSNE: Problems

tSNE has some problems.

- ▶ Does not scale well with number of data points.
- ▶ Does not preserve global structure. Only inter-cluster distances are meaningful. (intra?)
- ▶ Works poorly directly on high dimensional data. PCA used to pre-process.
- ▶ Requires a lot of memory.
- ▶ Cost function is not convex. Random initialization can give different results.

UMAP

UMAP

UMAP: Overview

Uniform manifold approximation and projection (UMAP).

Alternative to tSNE that often performs better on real data.

Similarly described by four equations:

1. Distance in high dimensional data space.
2. k-nearest neighbors
3. Distance in low dimensional embedding.
4. Cost function to minimize.

arXiv:1802.03426v3 [stat.ML] 18 Sep 2020

UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction

Leland McInnes

Tutte Institute for Mathematics and Computing
leland.mcinnes@gmail.com

John Healy

Tutte Institute for Mathematics and Computing
jchealy@gmail.com

James Melville

jlmelville@gmail.com

September 21, 2020

Abstract

UMAP (Uniform Manifold Approximation and Projection) is a novel manifold learning technique for dimension reduction. UMAP is constructed from a theoretical framework based in Riemannian geometry and algebraic topology. The result is a practical scalable algorithm that is applicable to real world data. The UMAP algorithm is competitive with t-SNE for visualization quality, and arguably preserves more of the global structure with superior run time performance. Furthermore, UMAP has no computational restrictions on embedding dimension, making it viable as a general purpose dimension reduction technique for machine learning.

1 Introduction

Dimension reduction plays an important role in data science, being a fundamental technique in both visualisation and as pre-processing for machine

UMAP: (1) High dimensional distances

Probability of observing distances, $d()$, between points in high dimensional, data, space assumed to be **exponentially decaying**.

The probability is not normalized.

Joint probability is **symmetrized**, but differently to tSNE.

Conditional probability:

$$p_{j|i} = \exp\left(-\frac{d(x_i, x_j) - \rho_i}{\sigma_i}\right)$$

Joint probability:

$$p_{ij} = p_{i|j} + p_{j|i} - p_{i|j}p_{j|i}$$

UMAP: (2) k-nearest neighbors

The k nearest neighbors is a value provided by the user as method parameter that influences the solution.

It is then used "in reverse" to solve for a values of σ_i and ρ_i in equation 1.

$$k = 2 \sum_i p_{ij}$$

UMAP: (3) Low dimensional distances

Probability of observing distances between points in low dimensional, embedding, space is also similar to the one used in tSNE.

The parameters a and b are found by fitting the function for q_{ij} .

$$q_{ji} = \left(1 + a(y_i - y_j)^{2b}\right)^{-1}$$

$$q = \begin{cases} 1 & \text{if } (y_i - y_j) \leq \mu \\ \exp(-(y_i - y_j)) - \mu & \text{else} \end{cases}$$

UMAP: (4) Cost function

The cost function is based on the **binary cross entropy (BCE)** between the probability distributions in high dimensional data space, P , and in low dimensional embedding, Q .

The embedding is optimized by minimizing the cost function.

This is done by calculating the gradient of KL with respect to each embedding position.

Minimized by gradient descent.

$$\text{BCE}(P, Q) = \sum_i \sum_j \left(p_{ij} \log \left(\frac{p_{ij}}{q_{ij}} \right) + (1 - p_{ij}) \log \left(\frac{1 - p_{ij}}{1 - q_{ij}} \right) \right)$$

(Same cost function as in logistic regression. Can interpret it as probability of two data points belonging together.)

UMAP: Method compared with tSNE

Both methods are based on four similar steps.

tSNE

$$p_{j|i} = \frac{\exp\left(\frac{-||x_i - x_j||^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-||x_i - x_k||^2}{2\sigma_i^2}\right)}$$

$$\text{perplexity} = 2^{-\sum_j (p_{j|i} \log_2(p_{j|i}))}$$

$$q_{ji} = \frac{(1 + ||y_i - y_j||^2)^{-1}}{\sum_{i \neq j} (1 + ||y_i - y_j||^2)^{-1}}$$

$$\text{KL}(P_i || Q_i) = \sum_i \sum_j p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right)$$

UMAP

$$p_{j|i} = \exp\left(-\frac{d(x_i, x_j) - \rho_i}{\sigma_i}\right)$$

$$k = 2 \sum_i p_{ij}$$

$$q_{ji} = (1 + a(y_i - y_j)^{2b})^{-1}$$


$$q = \begin{cases} 1 & \text{if } (y_i - y_j) \leq \mu \\ \exp(-(y_i - y_j)) - \mu & \text{else} \end{cases}$$

$$\text{BCE}(P, Q) = \sum_i \sum_j \left(p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right) + (1 - p_{ij}) \log\left(\frac{1 - p_{ij}}{1 - q_{ij}}\right) \right)$$


UMAP: And tSNE method overview

Both UMAP and tSNE are based on the same principle.


1) Construct high-dimensional graph



2) Construct low-dimensional graph



3) Collapse the graphs together




Kullback-Leibler divergence


(Source: Nikolay Oskolkov)

UMAP: Example

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from umap import UMAP
4
5 # Cancer associated fibroblast dataset and transform
6 X, y = load_CAF()
7 X = np.log(X + 1)
8
9 # Instantiate and apply UMAP
10 umap = UMAP(n_components=2, random_state=0)
11 X_umap = umap.fit_transform(X)
12
13 # Plot result
14 ax = plt.figure(figsize=(3,3)).gca()
15 ax.scatter(X_umap[:,0],
16           X_umap[:,1],
17           marker='o',
18           c=y.astype(int),
19           edgecolor='k',
20           )
21 ax.set_xlabel('UMAP 1')
22 ax.set_ylabel('UMAP 2')
```



UMAP: Results compared with tSNE



Conclusion

Conclusion

Conclusion: Distorted or exaggerated relationships in the data

tSNE and UMAP can create or exaggerate clusters in the data.

Occasionally PCA might be a better choice because it preserves the fact that data points can have an interpretable position between clusters.

Biologists, stop putting UMAP plots in your papers

UMAP is a powerful tool for exploratory data analysis, but without a clear understanding of how it works, it can easily lead to confusion and misinterpretation.

AUTHOR

Rafael Irizarry

PUBLISHED

Dec. 23, 2024

<https://nikolay-oskolkov.medium.com/is-umap-accurate-fad1b3f14fb5>

<https://archive.is/4t0yA>

<https://simplystatistics.org/posts/2024-12-23-biologists-stop-including-umap-plots-in-your-papers/>

MATHEMATICAL STATISTICS AND MACHINE LEARNING FOR LIFE SCIENCES


Is UMAP accurate?

Addressing some fair and unfair criticism



Nikolay Oskolkov · [Subscribe](#)

20 min read · Feb 22, 2025




UMAP in single cell biology vs. UMAP in population genomics. Image by author

Conclusion: Example

Methods for **non-linear dimensionality reduction** in two families.

Isomap and LLE can be used when PCA seemingly fails.

Single cell data often visualized using tSNE or UMAP.



sklearn.manifold	
Data embedding techniques.	
User guide. See the Manifold learning section for further details.	
Isomap	Isomap Embedding.
LocallyLinearEmbedding	Locally Linear Embedding.
MDS	Multidimensional scaling.
SpectralEmbedding	Spectral embedding for non-linear dimensionality reduction.
TSNE	T-distributed Stochastic Neighbor Embedding.
locally_linear_embedding	Perform a Locally Linear Embedding analysis on the data.
smacof	Compute multidimensional scaling using the SMACOF algorithm.
spectral_embedding	Project the sample on the first eigenvectors of the graph Laplacian.
trustworthiness	Indicate to what extent the local structure is retained.

<https://scikit-learn.org/stable/api/sklearn.manifold.html>

<https://umap-learn.readthedocs.io/en/latest/>