

DDA4210/AIR6002 Advanced Machine Learning

Lecture 07 Non-Linear Dimensionality Reduction

Tongxin Li

School of Data Science, CUHK-Shenzhen


Spring 2024

Overview


- 1 Introduction
- 2 Locally Linear Embedding (LLE)
- 3 t-distributed stochastic neighbor embedding (t-SNE)
- 4 Autoencoder

- 1 Introduction
- 2 Locally Linear Embedding (LLE)
- 3 t-distributed stochastic neighbor embedding (t-SNE)
- 4 Autoencoder

Introduction: Dimensionality Reduction



Linear DR



Non-Linear DR

Popular linear DR algorithms

- PCA
- LDA (supervised)
- Multidimensional Scaling (MDS)

Introduction: Multidimensional Scaling (MDS)

- **Problem:** Given euclidean distances among points, recover the position of the points!

$$\mathbf{D} \in \mathbb{R}^{N \times N} \longrightarrow \mathbf{X} \in \mathbb{R}^{D \times N}$$

- **Example:** The road distances between 21 European cities (almost euclidean, but not quite) are as follows:

	Athens	Barcelona	Brussels	Calais	Cherbourg
Barcelona	3313				
Brussels	2963	1318			
Calais	3175	1326	204		
Cherbourg	3339	1294	583	460	
Cologne	2762	1498	206	409	785


Can we construct a map?

Introduction: Multidimensional Scaling (MDS)

Road distance


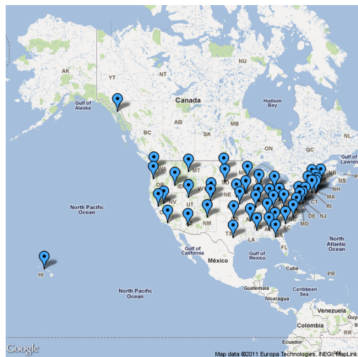
	Athens	Barcelona	Brussels	Calais	Cherbourg
Barcelona	3313				
Brussels	2963	1318			
Calais	3175	1326	204		
Cherbourg	3339	1294	583	460	
Cologne	2762	1498	206	409	785
	⋮				

Constructed 2-D map (data visualization)



Introduction: Multidimensional Scaling




Transform US city distances to city locations (2D)



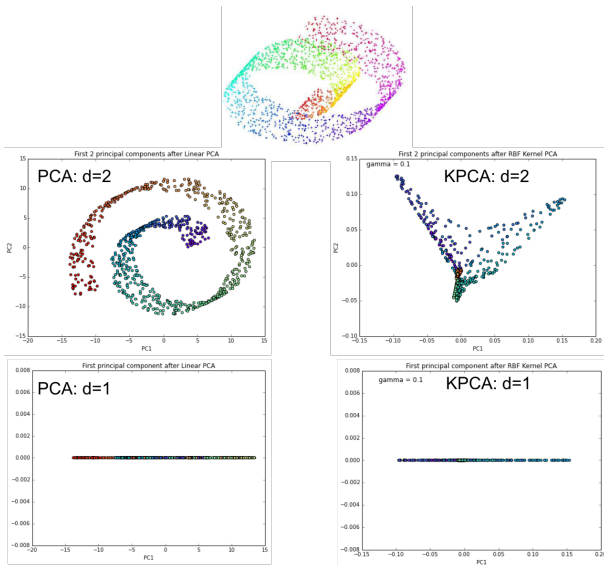
More about MDS

- https://www.stat.pitt.edu/sungkyu/course/2221Fall113/lec8_mds_combined.pdf
- Cox, M., Cox, T. (2008). Multidimensional Scaling. In: Handbook of Data Visualization. Springer Handbooks Comp.Statistics. Springer, Berlin, Heidelberg. <https://link.springer.com/content/pdf/10.1007/>

Introduction: A Toy Example of NLDR



Introduction: A Toy Example of NLDR



Introduction: NLDR

Important NLDR algorithms

- Kernel PCA¹
- Locally Linear Embedding²
- Isomap³
- Autoencoder⁴
- t-SNE⁵
- UMAP⁶

¹B. Scholkopf, A. Smola, and K.-R. Muller, Nonlinear component analysis as a kernel eigenvalue problem, *Neural Computation*, vol. 10, no. 5, pp. 1299-1319, 1998.

²Sam Roweis and Lawrence Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, v.290 no.5500 , Dec.22, 2000. pp.2323-2326

³J.B. Tenenbaum, V. de Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319-2323, 2000.

⁴Hinton, Geoffrey E. and Ruslan R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science* 313, no. 5786 (2006): 504-507.





⁵Van der Maaten, L.J.P. and Hinton, G.E. Visualizing Data Using t-SNE. *Journal of Machine Learning Research*. 2008, 9: 2579-2605.

⁶McInnes, Leland, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*,2018.

- 1 Introduction
- 2 Locally Linear Embedding (LLE)**
- 3 t-distributed stochastic neighbor embedding (t-SNE)
- 4 Autoencoder

Locally Linear Embedding (LLE)


- Based on a simple geometric intuition of local linearity
- Assume each sample and its neighbors lie on or close to a locally linear patch of the manifold
 - Manifold: low-dimensional surface embedded (nonlinearly) in high-dimensional space.
 - Examples of manifold



- LLE assumption: projection should preserve the neighborhood
 - Projected data point should have the same neighborhood as the original point
 - Locally linear representation should be preserved

LLE: Main Idea

- Neighborhood-preserved projection



$$\mathbf{x}_i \approx \sum_{j \in \mathcal{N}_i^k} w_{ij} \mathbf{x}_j$$

$$\sum_{j \in \mathcal{N}_i^k} w_{ij} = 1$$

- Input: D-dimensional data points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$, d, k ($k \geq d + 1$)
 - 1 For each data point \mathbf{x}_i , find its k nearest neighbors \mathcal{N}_i^k
 - 2 Find the locally linear reconstruction weights by solving

$$\mathbf{W} = \underset{\mathbf{W}}{\operatorname{argmin}} \sum_{i=1}^N \|\mathbf{x}_i - \sum_{j \in \mathcal{N}_i^k} w_{ij} \mathbf{x}_j\|^2, \quad \text{s.t. } \forall i \sum_{j \in \mathcal{N}_i^k} w_{ij} = 1$$

- 3 Use \mathbf{W} to compute the low-dimensional projections

$$\mathbf{Z} = \underset{\mathbf{Z}}{\operatorname{argmin}} \sum_{i=1}^N \|\mathbf{z}_i - \sum_{j \in \mathcal{N}_i^k} w_{ij} \mathbf{z}_j\|^2, \quad \text{s.t. } \sum_{i=1}^N \mathbf{z}_i = \mathbf{0}, \quad \frac{1}{N} \mathbf{Z} \mathbf{Z}^T = \mathbf{I}$$

- Output: Low-dimensional embedding $\mathbf{Z} \in \mathbf{R}^{d \times N}$

LLE: Algorithm

- Computation of \mathbf{W}
See the paper ⁷ if you are interested in it.
- Computation of \mathbf{Z}

$$\sum_{i=1}^N \|\mathbf{z}_i - \sum_{j \in \mathcal{N}_i^k} w_{ij} \mathbf{z}_j\|^2 = \|\mathbf{Z} - \mathbf{Z}\mathbf{W}\|_F^2 = \text{trace} \left(\mathbf{Z}(\mathbf{I} - \mathbf{W})(\mathbf{I} - \mathbf{W}^\top)\mathbf{Z}^\top \right)$$


Therefore, \mathbf{Z} should be composed of the d eigenvectors of $(\mathbf{I} - \mathbf{W})(\mathbf{I} - \mathbf{W}^\top)$ corresponding to the smallest d nonzero eigenvalues, i.e.,

$$\mathbf{Z} = [\mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_{d+1}]^\top$$

⁷Sam Roweis and Lawrence Saul. Nonlinear dimensionality reduction by locally linear embedding. Science, v.290 no.5500, Dec.22, 2000. pp.2323-2326


LLE: Applications

Toy examples



LLE: Applications

Face images




- 1 Introduction
- 2 Locally Linear Embedding (LLE)
- 3 t-distributed stochastic neighbor embedding (t-SNE)**
- 4 Autoencoder

t-SNE: Neighborhood Probability

- Given D high-dimensional data points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$, let

$$p_{j|i} = \begin{cases} \frac{\exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2/2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}_i - \mathbf{x}_k\|^2/2\sigma_i^2)}, & \text{if } i \neq j \\ 0, & \text{if } i = j \end{cases}$$

- $p_{j|i}$ denotes the probability that \mathbf{x}_j is a neighbor of \mathbf{x}_i
- The parameter σ_i sets the size of the neighborhood of \mathbf{x}_i
- Set σ_i differently for each data point (according to "Perplexity")
- Results depend heavily on σ_i -it defines the neighborhoods we are trying to preserve




t-SNE: Neighborhood Probability

- Given D high-dimensional data points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$, let

$$p_{j|i} = \begin{cases} \frac{\exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2/2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}_i - \mathbf{x}_k\|^2/2\sigma_i^2)}, & \text{if } i \neq j \\ 0, & \text{if } i = j \end{cases}$$

- Let $p_{ij} = \frac{p_{j|i} + p_{i|j}}{2N}$
- Then $p_{ij} = p_{ji}$, $p_{ii} = 0$, $\sum_{i,j} p_{ij} = 1$





t-SNE: Perplexity (to determine σ_i) (optional)

- For each data point, define the perplexity:

$$\text{perp}(i) = 2^{H(p_{j|i})}, \quad H(p_{j|i}) = - \sum_{j=1}^N p_{j|i} \log p_{j|i}$$

- A low perplexity indicates the probability distribution is good at predicting the sample.
- Define the desired perplexity and set σ_i to get that (e.g. bisection)
- Values between 5-50 usually work well



t-SNE: Objective

- Learn a d -dimensional embedding $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ ($d < D$ and usually $d = 2$ or 3)⁸

$$q_{ij} = \frac{(1 + \|\mathbf{z}_i - \mathbf{z}_j\|^2)^{-1}}{\sum_k \sum_{l \neq k} (1 + \|\mathbf{z}_k - \mathbf{z}_l\|^2)^{-1}}$$

Such that $Q = [q_{ij}]_{N \times N}$ is close to $P = [p_{ij}]_{N \times N}$.

⁸Not $q_{ij} = \frac{\exp(-\|\mathbf{z}_i - \mathbf{z}_j\|^2)}{\sum_k \sum_{l \neq k} \exp(-\|\mathbf{z}_l - \mathbf{z}_k\|^2)}$, which corresponds to SNE. It has a crowding problem: in 2D or 3D, we do not have enough room to accommodate all neighbors when using Gaussian distribution.

t-SNE: Objective

- Learn a d -dimensional embedding $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ ($d < D$ and usually $d = 2$ or 3)⁸


$$q_{ij} = \frac{(1 + \|\mathbf{z}_i - \mathbf{z}_j\|^2)^{-1}}{\sum_k \sum_{l \neq k} (1 + \|\mathbf{z}_k - \mathbf{z}_l\|^2)^{-1}}$$

Such that $Q = [q_{ij}]_{N \times N}$ is close to $P = [p_{ij}]_{N \times N}$.

Student-t probability density

$$p(x) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi} \Gamma(\frac{\nu}{2})} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

- When $\nu = \infty$, $p(x)$ is Gaussian
- When $\nu = 1$, $p(x) \propto (1 + x^2)^{-1}$



⁸Not $q_{ij} = \frac{\exp(-\|\mathbf{z}_i - \mathbf{z}_j\|^2)}{\sum_k \sum_{l \neq k} \exp(-\|\mathbf{z}_l - \mathbf{z}_k\|^2)}$, which corresponds to SNE. It has a crowding problem: in 2D or 3D, we do not have enough room to accommodate all neighbors when using Gaussian distribution.

- Minimize the KL-divergence⁹

$$\mathcal{L} := \text{KL}(P \parallel Q) = \sum_{i \neq j} p_{ij} \log \frac{p_{ij}}{q_{ij}}$$


- Gradient based optimization:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{z}_i} = \sum_j (p_{ij} - q_{ij})(\mathbf{z}_i - \mathbf{z}_j)(1 + \|\mathbf{z}_i - \mathbf{z}_j\|^2)^{-1}, \quad i = 1, \dots, N$$

⁹A loss function to measure the difference between two distributions, similar to the cross entropy loss.





t-SNE: Applications

- Synthetic data



t-SNE: Applications

- MNIST handwritten digits (10 classes)




t-SNE: Applications


- COIL20 image data (20 objects with 72 different poses)




(a) Visualization by t-SNE.




(b) Visualization by Sammon mapping.



(c) Visualization by Isomap.




(d) Visualization by LLE.



- 1 Introduction
- 2 Locally Linear Embedding (LLE)
- 3 t-distributed stochastic neighbor embedding (t-SNE)
- 4 Autoencoder**


Autoencoder: Basic Structure

- Recall PCA: $\mathbf{Z} = \mathbf{U}^T \mathbf{X}$, $\hat{\mathbf{X}} = \mathbf{U} \mathbf{Z} = \mathbf{U} \mathbf{U}^T \mathbf{X}$




Autoencoder: Basic Structure

- Recall PCA: $\mathbf{Z} = \mathbf{U}^T \mathbf{X}$, $\hat{\mathbf{X}} = \mathbf{U}\mathbf{Z} = \mathbf{U}\mathbf{U}^T \mathbf{X}$



- Autoencoder: a neural network with output=input



Encoder/Decoder architecture

- Encoder: $f = \sigma(\mathbf{W}\mathbf{x})$
- Decoder: $g = \sigma(\mathbf{W}'\mathbf{h})$
- Hidden layer dimension < input dimension
- Predict the input by itself:
 $\mathbf{x} \approx g(f(\mathbf{x}))$

Autoencoder: Deep Model

- Stacked Autoencoders (SAE)
 - Use the middle layer as a representation
 - Out-of-sample extension: just feed new data into the encoder
 - Question: out-of-sample extension for PCA, LLE, and t-SNE?


Autoencoder: Deep Model

- Stacked Autoencoders (SAE)
 - Use the middle layer as a representation
 - Out-of-sample extension: just feed new data into the encoder
 - Question: out-of-sample extension for PCA, LLE, and t-SNE?
- Train the Autoencoder
 - For example, solve

$$\underset{\theta}{\text{minimize}} \quad \frac{1}{2} \sum_{i=1}^N \|\mathbf{x}_i - g(f(\mathbf{x}_i))\|^2$$



θ : network parameters

- Backpropagation
- Gradient-based optimization



* Autoencoder can be extended to convolutional neural networks—CAE

Autoencoder: Application-Data Compression

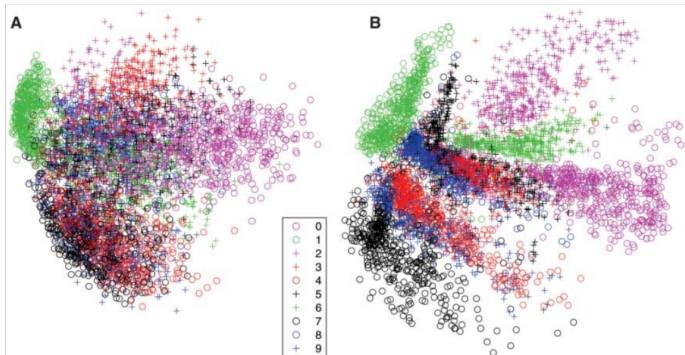


G. E. Hinton and R. R. Salakhutdinov. Science. 2006

Autoencoder: Application-Data Visualization

MNIST 2D visualization

Fig. 3. (A) The two-dimensional codes for 500 digits of each class produced by taking the first two principal components of all 60,000 training images. (B) The two-dimensional codes found by a 784-1000-500-250-2 autoencoder. For an alternative visualization, see (8).



G. E. Hinton and R. R. Salakhutdinov. Science. 2006

Generative Models

Limitation of Autoencoder

- Cannot generate meaningful data using the decoder

Generative models

- Variational Autoencoder (VAE) (Kingma and Welling. ICLR 2014)
- Generative Adversarial Network (GAN) (Goodfellow et al. arXiv 2014)

Generative Models

Limitation of Autoencoder


- Cannot generate meaningful data using the decoder

Generative models

- Variational Autoencoder (VAE) (Kingma and Welling. ICLR 2014)
- Generative Adversarial Network (GAN) (Goodfellow et al. arXiv 2014)

Examples of generated images (CelebA dataset)

Left: VAE. Right: GAN. Image from Pieters and Wiering 2018.



- Understand the basic ideas of LLE, t-SNE, and Autoencoder
- Know the limitations of LLE, t-SNE, and Autoencoder
- Be able to use t-SNE to visualize real data
- Be able to use Autoencoder to reduce the noise of real data