# Data Mining Techniques 

CS 6220 - Section 3 - Fall 2016

## Lecture 13

Jan-Willem van de Meent (credit: David Lopez-Paz, David

Duvenaud, Laurens van der Maaten)


## Homework

- Homework 3 is out today (due 4 Nov)
- Homework 1 has been graded (we will grade Homework 2 a little faster)
- Regrading policy
- Step 1: E-mail TAs to resolve simple problems (e.g. code not running).
- Step 2: E-mail instructor to request regrading.
- We will regrade the entire problem set. The final grade can be lower than before.


## Review: PCA

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots \cdots & \mathbf{x}_{n} \\
\mid & & \mid
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Change of basis

$$
\begin{aligned}
\mathbf{z} & =\mathbf{U}^{\top} \mathbf{x} \\
\mathbf{z} & =\left(z_{1}, \ldots, z_{n}\right)^{\top} \\
z_{j} & =\mathbf{u}_{j}^{\top} \mathbf{x}
\end{aligned}
$$

Orthonormal Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} \cdot & \mid \\
\mid & \mathbf{u}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Inverse Change of basis

$$
\mathbf{x}=\mathbf{U} \mathbf{z}=\sum_{j=1}^{n} z_{j} \mathbf{u}_{j}
$$

## Review: PCA

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots & \cdots \\
\mid & & \mathbf{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigenvectors of Covariance

$$
\begin{gathered}
\mathbf{C}=\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j} \mathbf{x}_{j}^{\top}=\frac{1}{n} \mathbf{X} \mathbf{X}^{\top} \\
\mathbf{C u}_{j}=\lambda_{j} \mathbf{u}_{j}
\end{gathered}
$$

Orthonormal Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} & \cdot \\
\mid & \mathbf{u}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigen-decomposition

$$
\begin{gathered}
\boldsymbol{C}=\boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{\top} \\
\boldsymbol{\Lambda}=\left(\begin{array}{llll}
\lambda_{1} & & & \\
& \lambda_{2} & & \\
& & \cdots & \\
& & & \lambda_{n}
\end{array}\right)
\end{gathered}
$$

## Review: PCA

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots & \cdots \\
\mid & \mathbf{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigenvectors of Covariance

$$
\begin{gathered}
\mathbf{C}=\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j} \mathbf{x}_{j}^{\top}=\frac{1}{n} \mathbf{X} \mathbf{X}^{\top} \\
\mathbf{C u}_{j}=\lambda_{j} \mathbf{u}_{j}
\end{gathered}
$$

Orthonormal Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} \cdot \cdots & \mathbf{u}_{n} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigen-decomposition

$$
\begin{gathered}
\boldsymbol{C}=\boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{\top} \\
\boldsymbol{\Lambda}=\left(\begin{array}{llll}
\lambda_{1} & & & \\
& \lambda_{2} & & \\
& & \cdots & \\
& & & \lambda_{n}
\end{array}\right)
\end{gathered}
$$

Claim: Eigenvectors of a symmetric matrix are orthogonal

## $\rightarrow \cap \operatorname{Rin}^{\circ}$

For any real matrix $A$ and any vectors $\mathbf{x}$ and $\mathbf{y}$, we have

$$
\langle A \mathbf{x}, \mathbf{y}\rangle=\left\langle\mathbf{x}, A^{T} \mathbf{y}\right\rangle .
$$

Now assume that $A$ is symmetric, and $\mathbf{x}$ and $\mathbf{y}$ are eigenvectors of $A$ corresponding to distinct eigenvalues $\lambda$ and $\mu$. Then

$$
\lambda\langle\mathbf{x}, \mathbf{y}\rangle=\langle\lambda \mathbf{x}, \mathbf{y}\rangle=\langle A \mathbf{x}, \mathbf{y}\rangle=\left\langle\mathbf{x}, A^{T} \mathbf{y}\right\rangle=\langle\mathbf{x}, A \mathbf{y}\rangle=\langle\mathbf{x}, \mu \mathbf{y}\rangle=\mu\langle\mathbf{x}, \mathbf{y}\rangle .
$$

Therefore, $(\lambda-\mu)\langle\mathbf{x}, \mathbf{y}\rangle=0$. Since $\lambda-\mu \neq 0$, then $\langle\mathbf{x}, \mathbf{y}\rangle=0$, i.e., $\mathbf{x} \perp \mathbf{y}$.
Now find an orthonormal basis for each eigenspace; since the eigenspaces are mutually orthogonal, these vectors together give an orthonormal subset of $\mathbb{R}^{n}$. Finally, since symmetric matrices are diagonalizable, this set will be a basis (just count dimensions). The result you want now follows.
share cite improve this answer

## (from stack exchange)

## Review: PCA

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots & \cdots \\
\mid & \mathbf{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigenvectors of Covariance

$$
\begin{gathered}
\mathbf{C}=\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j} \mathbf{x}_{j}^{\top}=\frac{1}{n} \mathbf{X} \mathbf{X}^{\top} \\
\mathbf{C u}_{j}=\lambda_{j} \mathbf{u}_{j}
\end{gathered}
$$

Orthonormal Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} \cdot \cdots & \mathbf{u}_{n} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigen-decomposition

$$
\begin{gathered}
\boldsymbol{C}=\boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{\top} \\
\boldsymbol{\Lambda}=\left(\begin{array}{llll}
\lambda_{1} & & & \\
& \lambda_{2} & & \\
& & \cdots & \\
& & & \lambda_{n}
\end{array}\right)
\end{gathered}
$$

Claim: Eigenvectors of a symmetric matrix are orthogonal

## Review: PCA

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots & \cdots \\
\mid & \mathbf{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Eigenvectors of Covariance

$$
\begin{gathered}
\mathbf{C}=\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j} \mathbf{x}_{j}^{\top}=\frac{1}{n} \mathbf{X} \mathbf{X}^{\top} \\
\mathbf{C u}_{j}=\lambda_{j} \mathbf{u}_{j}
\end{gathered}
$$

Truncated Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} & \cdot \\
\mid & \mathbf{u}_{k} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times k}
$$

Truncated decomposition

$$
\begin{gathered}
\boldsymbol{C} \simeq \boldsymbol{U} \boldsymbol{\Lambda}^{(k)} \boldsymbol{U}^{\top} \\
\boldsymbol{\Lambda}^{(k)}=\left(\begin{array}{llll}
\lambda_{1} & & & \\
& \lambda_{2} & & \\
& & \cdots & \\
& & & \lambda_{k}
\end{array}\right)
\end{gathered}
$$

## Review: PCA

## Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots \cdots & \mathbf{x}_{n} \\
\mid & & \mid
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Projection / Encoding

$$
\mathbf{z}=\mathrm{U}^{\top} \mathbf{x}
$$

Truncated Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} & \cdot \\
\mid & \mathbf{u}_{k} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times k}
$$

Reconstruction / Decoding
$\tilde{\mathrm{x}}=\mathrm{Uz}$

## Review: PCA

## Top 2 components



Bottom 2 components



Data: three varieties of wheat: Kama, Rosa, Canadian
Attributes: Area, Perimeter, Compactness, Length of Kernel, Width of Kernel, Asymmetry Coefficient, Length of Groove

## PCA: Complexity

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathbf{x}_{1} & \cdots & \cdots \\
\mid & \mathbf{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Truncated Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} & \cdot \\
\mid & \mathbf{u}_{k} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times k}
$$

## Using eigen-value decomposition

- Computation of covariance $\mathbf{C}: O\left(n d^{2}\right)$
- Eigen-value decomposition: $O\left(d^{3}\right)$
- Total complexity: $O\left(n d^{2}+d^{3}\right)$


# PCA: Complexity 

Data

$$
\mathbf{X}=\left(\begin{array}{ccc}
\mid & \cdots & \mid \\
\mathrm{x}_{1} & \cdots & \cdots \\
\mid & \mathrm{x}_{n}
\end{array}\right) \in \mathbb{R}^{d \times n}
$$

Truncated Basis

$$
\mathbf{U}=\left(\begin{array}{cc}
\mid & \mid \\
\mathbf{u}_{1} & \cdot \\
\mid & \mathbf{u}_{k} \\
\mid & \mid
\end{array}\right) \in \mathbb{R}^{d \times k}
$$

Using singular-value decomposition

- Full decomposition: $O\left(\min \left\{n d^{2}, n^{2} d\right\}\right)$
- Rank-k decomposition: O(kdnlog(n)) (with power method)


## Singular Value Decomposition



Idea: Decompose a $\mathrm{d} x \mathrm{~d}$ matrix $M$ into

1. Change of basis $V$ (unitary matrix)
2. A scaling $\Sigma$
(diagonal matrix)
3. Change of basis $U$ (unitary matrix)

$$
\begin{aligned}
& M=\begin{array}{llll}
\mathbf{U} & \boldsymbol{\Sigma} & \mathbf{V}^{*}
\end{array} \\
& {\left[\begin{array}{rr}
1 & -1 \\
2 & 2
\end{array}\right]=\left[\begin{array}{cc}
0 & -i \\
1 & 0
\end{array}\right] \frac{1}{\sqrt{2}}\left[\begin{array}{ll}
2 & 0 \\
0 & 1
\end{array}\right] \frac{1}{\sqrt{2}}\left[\begin{array}{rr}
-1 & -1 \\
i & -i
\end{array}\right]}
\end{aligned}
$$

## Singular Value Decomposition



$$
M=U \cdot \Sigma \cdot V^{*}
$$

$$
\mathbf{X}=\mathbf{U}_{d \times d} \Sigma_{d \times n} \mathbf{V}_{n \times n}^{\top}
$$

Idea: Decompose the d x n matrix $X$ into

1. A $n \times n$ basis $V$ (unitary matrix)
2. Adxn matrix $\Sigma$
(diagonal projection)
3. A dxd basis $U$
(unitary matrix)

# Random Projections 



Borrowing from:<br>David Lopez-Paz<br>\& David Duvenaud

## Random Projections

Fast, efficient and $\mathcal{F}$ distance-preserving dimensionality reduction!


This result is formalized in the Johnson-Lindenstrauss Lemma

## Johnson-Lindenstrauss Lemma

For any $0<\epsilon<1 / 2$ and any integer $m>4$, let $k=\frac{20 \log m}{\epsilon^{2}}$. Then, for any set $V$ of $m$ points in $\mathbb{R}^{N} \exists f: \mathbb{R}^{N} \rightarrow \mathbb{R}^{k}$ s.t. $\forall \boldsymbol{u}, \boldsymbol{v} \in V$ :

$$
(1-\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^{2} \leq\|f(\boldsymbol{u})-f(\boldsymbol{v})\|^{2} \leq(1+\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^{2} .
$$

The proof is a great example of Erdös' probabilistic method (1947).


Paul Erdös 1913-1996


Joram Lindenstrauss 1936-2012


William B. Johnson 1944-

## Johnson-Lindenstrauss Lemma

For any $0<\epsilon<1 / 2$ and any integer $m>4$, let $k=\frac{20 \log m}{\epsilon^{2}}$. Then, for any set $V$ of $m$ points in $\mathbb{R}^{N} \exists f: \mathbb{R}^{N} \rightarrow \mathbb{R}^{k}$ s.t. $\forall \boldsymbol{u}, \boldsymbol{v} \in V$ :

$$
(1-\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^{2} \leq\|f(\boldsymbol{u})-f(\boldsymbol{v})\|^{2} \leq(1+\epsilon)\|\boldsymbol{u}-\boldsymbol{v}\|^{2} .
$$

Holds when $f$ is linear function with random coefficients

$$
f=\frac{1}{\sqrt{k}} \boldsymbol{A}, \boldsymbol{A} \in \mathbb{R}^{k \times N}, k<N \text { and } A_{i j} \sim \mathcal{N}(0,1)
$$

## Example: 20-newsgroups data

Data: 20-newsgroups, from 100.000 features to 300 (0.3\%)



## Example: 20-newsgroups data

## Data: 20-newsgroups, from 100.000 features to 1.000 (1\%)




## Example: 20-newsgroups data

Data: 20-newsgroups, from 100.000 features to $10.000(10 \%)$


Pairwise squared distances in original space


## Example: 20-newsgroups data

Data: 20-newsgroups, from 100.000 features to 10.000 (10\%)



Conclusion: RP preserves distances like PCA, but faster than PCA number of dimensions is vey large

# Stochastic Neighbor Embeddings 



Borrowing from:
Laurens van der Maaten
(Delft -> Facebook AI)

## Manifold Learning



Idea: Perform a non-linear dimensionality reduction in a manner that preserves proximity (but not distances)

## Manifold Learning



## Visualizing data using t-SNE

[PDF] jmlr.org
L Maaten, $\underline{G}$ Hinton - Journal of Machine Learning Research, 2008 - jmlr.org
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 ...
Cited by 1771 Related articles All 35 versions Cite Save

## PCA on MNIST Digits



## Swiss Roll



Euclidean distance is not always a good notion of proximity

## Non-linear Projection

High Dim


Low Dim


Bad projection: relative position to neighbors changes

## Non-linear Projection

High Dim


Low Dim


Intuition: Want to preserve local neighborhood

## Stochastic Neighbor Embedding

High Dim


Similarity in high dimension

$$
p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

Low Dim


Similarity in low dimension

$$
q_{j \mid i}=\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|y_{i}-y_{k}\right\|^{2}\right)}
$$

## Stochastic Neighbor Embedding

- Similarity of datapoints in High Dimension

$$
p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

- Similarity of datapoints in Low Dimension

$$
q_{j \mid i}=\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|y_{i}-y_{k}\right\|^{2}\right)}
$$

- Cost function

$$
C=\sum_{i} K L\left(P_{i} \| Q_{i}\right)=\sum_{i} \sum_{j} p_{j \mid i} \log \frac{p_{j \mid i}}{q_{j \mid i}}
$$

Idea: Optimize $y_{i}$ via gradient descent on C

## Stochastic Neighbor Embedding

- Similarity of datapoints in High Dimension

$$
p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

- Similarity of datapoints in Low Dimension

$$
q_{j \mid i}=\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|y_{i}-y_{k}\right\|^{2}\right)}
$$

- Cost function

$$
C=\sum_{i} K L\left(P_{i} \| Q_{i}\right)=\sum_{i} \sum_{j} p_{j \mid i} \log \frac{p_{j \mid i}}{q_{j \mid i}}
$$

Idea: Optimize $y_{i}$ via gradient descent on C

## Stochastic Neighbor Embedding

Gradient has a surprisingly simple form

$$
\frac{\partial C}{\partial y_{i}}=\sum_{j \neq i}\left(p_{j \mid i}-q_{j \mid i}+p_{i \mid j}-q_{i \mid j}\right)\left(y_{i}-y_{j}\right)
$$

The gradient update with momentum term is given by

$$
Y^{(t)}=Y^{(t-1)}+\eta \frac{\partial C}{\partial y_{i}}+\beta(t)\left(Y^{(t-1)}-Y^{(t-2)}\right)
$$

## Stochastic Neighbor Embedding

Gradient has a surprisingly simple form

$$
\frac{\partial C}{\partial y_{i}}=\sum_{j \neq i}\left(p_{j \mid i}-q_{j \mid i}+p_{i \mid j}-q_{i \mid j}\right)\left(y_{i}-y_{j}\right)
$$

The gradient update with momentum term is given by

$$
Y^{(t)}=Y^{(t-1)}+\eta \frac{\partial C}{\partial y_{i}}+\beta(t)\left(Y^{(t-1)}-Y^{(t-2)}\right)
$$

Problem: $p_{\mathrm{jli}}$ is not equal to $p_{\mathrm{ilj}}$

## Symmetric SNE

- Minimize a single KL divergence between a joint probability distribution

$$
C=K L(P \| Q)=\sum_{i} \sum_{j \neq i} p_{i j} \log \frac{p_{i j}}{q_{i j}}
$$

- The obvious way to redefine the pairwise similarities is

$$
\begin{aligned}
p_{i j} & =\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq l} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)} \\
q_{i j} & =\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|y_{l}-y_{k}\right\|^{2}\right)}
\end{aligned}
$$

## Symmetric SNE

- Minimize a single KL divergence between a joint probability distribution

$$
C=K L(P \| Q)=\sum_{i} \sum_{j \neq i} p_{i j} \log \frac{p_{i j}}{q_{i j}}
$$

- The obvious way to redefine the pairwise similarities is

$$
\begin{gathered}
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq l} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)} \\
q_{i j}=\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|y_{l}-y_{k}\right\|^{2}\right)}
\end{gathered}
$$

Problem: How should we choose $\sigma$ ?

## Choosing the bandwidth



$$
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)}
$$

Bad $\sigma$. Neighborhood is not local in manifold

## Choosing the bandwidth



$$
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|x_{I}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)}
$$

Good $\sigma$ : Neighborhood contains 5-50 points

## Choosing the bandwidth



$$
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)}
$$

Problem: optimal $\sigma$ may vary if density not uniform

## Choosing the bandwidth



$$
p_{i j}=\frac{p_{j \mid i}+p_{i \mid j}}{2 N} \quad p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

Solution: Define $\sigma_{i}$ per point.

## Choosing the bandwidth



$$
p_{i j}=\frac{p_{j \mid i}+p_{i \mid j}}{2 N} \quad p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

Solution: Define $\sigma_{i}$ per point.

## Choosing the bandwidth



$$
p_{i j}=\frac{p_{j \mid i}+p_{i \mid j}}{2 N} \quad p_{j \mid i}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma_{i}^{2}\right)}{\sum_{k \neq i} \exp \left(-\left\|x_{i}-x_{k}\right\|^{2} / 2 \sigma_{i}^{2}\right)}
$$

Solution: Define $\sigma_{i}$ per point.

## Choosing the bandwidth



$$
\operatorname{Perp}\left(\mathbf{p}_{j \mid i}\right)=\exp H\left(\mathbf{p}_{j \mid i}\right)=\exp ^{-\sum_{j} \mathbf{p}_{j \mid i} \log \mathbf{p}_{j \mid i}}
$$

Set $\sigma_{i}$ to ensure constant perplexity

## t-SNE: SNE with a t-Distribution



Similarity in High Dimension

$$
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)}
$$

Similarity in Low Dimension

$$
q_{i j}=\frac{\left(1+\left\|y_{i}-y_{j}\right\|^{2}\right)^{-1}}{\sum_{k \neq l}\left(1+\left\|y_{k}-y_{l}\right\|^{2}\right)^{-1}}
$$

Gradient

$$
\frac{\partial C}{\partial y_{i}}=4 \sum_{j \neq i}\left(p_{i j}-q_{i j}\right)\left(1+\left\|y_{i}-y_{j}\right\|^{2}\right)^{-1}\left(y_{i}-y_{j}\right)
$$

## t-SNE: SNE with a t-Distribution



Similarity in High Dimension

$$
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)}
$$

Similarity in Low Dimension

$$
q_{i j}=\frac{\left(1+\left\|y_{i}-y_{j}\right\|^{2}\right)^{-1}}{\sum_{k \neq l}\left(1+\left\|y_{k}-y_{l}\right\|^{2}\right)^{-1}}
$$

Gradient

$$
\frac{\partial C}{\partial y_{i}}=4 \sum_{j \neq i}\left(p_{i j}-q_{i j}\right)\left(1+\left\|y_{i}-y_{j}\right\|^{2}\right)^{-1}\left(y_{i}-y_{j}\right)
$$

## Symmetric SNE

- Minimize a single KL divergence between a joint probability distribution

$$
C=K L(P \| Q)=\sum_{i} \sum_{j \neq i} p_{i j} \log \frac{p_{i j}}{q_{i j}}
$$

- The obvious way to redefine the pairwise similarities is

$$
\begin{gathered}
p_{i j}=\frac{\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)}{\sum_{k \neq l} \exp \left(-\left\|x_{l}-x_{k}\right\|^{2} / 2 \sigma^{2}\right)} \\
q_{i j}=\frac{\exp \left(-\left\|y_{i}-y_{j}\right\|^{2}\right)}{\sum_{k \neq 1} \exp \left(-\left\|y_{l}-y_{k}\right\|^{2}\right)}
\end{gathered}
$$

Problem: How should we choose $\sigma$ ?

## PCA on MNIST Digits



## t-SNE on MNIST Digits



# t-SNE on MNIST Digits 



## t-SNE on Olivetti Faces



## t-SNE on Olivetti Faces



## t-SNE on Olivetti Faces



## t-SNE on ImageNet



## t-SNE on ImageNet



## t-SNE on ImageNet



## Next lecture: Recommender Systems

|  | $\frac{5 m s}{x i m s}$ |  |  |  | 3 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 |  |  | 4 | 5 | 2.94* |
| $2$ | 5 |  | 4 |  |  | 1 |
|  |  |  | 5 |  | 2 | 2.48* |
|  |  | 1 |  | 5 |  | 4 |
|  |  |  | 4 |  |  | 2 |
|  | 4 | 5 |  | 1 |  | 1.12* |
| $\operatorname{sim}(\mathrm{i}, \mathrm{j})$ | -1 | 1 | 0.86 | 1 | NA |  |

