

Outline

- Feature selection
- Single feature scoring criteria
- Search strategies
- Unsupervised dimension reduction using all features
- Principle Components Analysis
- Singular Value Decomposition
- Independent components analysis
- Supervised dimension reduction
- Fisher Linear Discriminant
- Hidden layers of Neural Networks


## Dimensionality Reduction

Why?

- Learning a target function from data where some features are irrelevant - reduce variance, improve accuracy
- Wish to visualize high dimensional data
- Sometimes have data whose "intrinsic" dimensionality is smaller than the number of features used to describe it recover intrinsic dimension


## Supervised Feature Selection

Problem: Wish to learn $f: X \rightarrow Y$, where $X=<X_{1}, \ldots X_{N}>$ But suspect not all $X_{i}$ are relevant

Approach: Preprocess data to select only a subset of the $X_{i}$

- Score each feature, or subsets of features - How?
- Search for useful subset of features to represent data - How?


## Scoring Individual Features $X_{i}$

Common scoring methods:

- Training or cross-validated accuracy of single-feature classifiers $f_{i}: X_{i} \rightarrow Y$
- Estimated mutual information between $X_{i}$ and $Y$ :
$\hat{I}\left(X_{i}, Y\right)=\sum_{k} \sum_{y} \hat{P}\left(X_{i}=k, Y=y\right) \log \frac{\hat{P}\left(X_{i}=k, Y=y\right)}{\hat{P}\left(X_{i}=k\right) \hat{P}(Y=y)}$
- $\chi^{2}$ statistic to measure independence between $X_{i}$ and $Y$
- Domain specific criteria
- Text: Score "stop" words ("the", "of", ...) as zero
- fMRI: Score voxel by T-test for activation versus rest condition
- ...


## Choosing Set of Features to learn $\mathrm{F}: \mathrm{X} \rightarrow \mathrm{Y}$

Common methods:

Forward1: Choose the n features with the highest scores

## Forward2:

- Choose single highest scoring feature $X_{k}$
- Rescore all features, conditioned on the set of already-selected features
- E.g., Score $\left(X_{i} \mid X_{k}\right)=I\left(X_{i}, Y \mid X_{k}\right)$
- E.g, Score $\left(X_{i} \mid X_{k}\right)=\operatorname{Accuracy}$ (predicting $Y$ from $X_{i}$ and $X_{k}$ )
- Repeat, calculating new scores on each iteration, conditioning on set of selected features


## Choosing Set of Features

Common methods:

Backward1: Start with all features, delete the n with lowest scores

Backward2: Start with all features, score each feature conditioned on assumption that all others are included. Then:

- Remove feature with the lowest (conditioned) score
- Rescore all features, conditioned on the new, reduced feature set
- Repeat



## Summary: Supervised Feature Selection

Approach: Preprocess data to select only a subset of the $X_{i}$

- Score each feature
- Mutual information, prediction accuracy, ...
- Find useful subset of features based on their scores
- Greedy addition of features to pool
- Greedy deletion of features from pool
- Considered independently, or in context of other selected features

Always do feature selection using training set only (not test set!)

- Often use nested cross-validation loop:
- Outer loop to get unbiased estimate of final classifier accuracy
- Inner loop to test the impact of selecting features

Unsupervised Dimensionality Reduction

Unsupervised mapping to lower dimension
Differs from feature selection in two ways:

- Instead of choosing subset of features, create new features (dimensions) defined as functions over all features
- Don't consider class labels, just the data points


## Principle Components Analysis

- Idea:
- Given data points in d-dimensional space, project into lower dimensional space while preserving as much information as possible
- E.g., find best planar approximation to 3D data
- E.g., find best planar approximation to $10^{4} \mathrm{D}$ data
- In particular, choose projection that minimizes the squared error in reconstructing original data

PCA: Find Projections to Minimize Reconstruction Error
Assume data is set of d-dimensional vectors, where nth vector is $\mathrm{x}^{n}=\left\langle x_{1}^{n} \ldots x_{d}^{n}\right\rangle$
We can represent these in terms of any d orthogonal basis vectors

$$
\mathbf{x}^{n}=\sum_{i=1}^{d} z_{i}^{n} \mathbf{u}_{i j} \quad \mathbf{u}_{i}^{T} \mathbf{u}_{j}=\delta_{i j}
$$

PCA: given $M<d$. Find $\left\langle\mathbf{u}_{1} \ldots \mathbf{u}_{M}\right\rangle$
that minimizes $E_{M} \equiv \sum_{n=1}^{N}\left\|x^{n}-\hat{x}^{n}\right\|^{2}$
where $\hat{\mathrm{x}}^{n}=\overline{\mathrm{x}}+\sum_{i=1}^{M} z_{i}^{n} \mathbf{u}_{i}$
Mean

$$
\overline{\mathrm{x}}=\frac{1}{N} \sum_{n=1}^{N} \mathrm{x}^{n}
$$


$x_{1}$


Note we get zero error if $\mathrm{M}=\mathrm{d}$.
Therefore, $E_{M}=\sum_{i=M+1}^{d} \sum_{n=1}^{N}\left[\mathbf{u}_{i}^{T}\left(\mathrm{x}^{n}-\overline{\mathrm{x}}\right)\right]^{2}$
$=\sum_{i=M+1}^{d} \mathbf{u}_{i}^{T} \Sigma \mathbf{u}_{i}$
This minimized when $\boldsymbol{u}_{i}$ is eigenvector of $\Sigma$, i.e.,
when:

Covariance matrix: $\Sigma=\sum_{n}\left(\mathrm{x}^{n}-\overline{\mathrm{x}}\right)\left(\mathrm{x}^{n}-\overline{\mathrm{x}}\right)^{T}$

| PCA <br> Minimize $E_{M}=\sum_{i=M+1}^{d} \mathbf{u}_{i}^{T} \Sigma \mathbf{u}_{i}$ <br> $\rightarrow \quad \Sigma \mathbf{u}_{i}=\lambda_{i} \mathbf{u}_{i}{ }_{\text {EEigenvector of } \Sigma}$ <br> Eigenvalue $\rightarrow \quad E_{M}=\sum_{i=M+1}^{d} \lambda_{i}$ <br> PCA algorithm 1: <br> 1. $\mathrm{X} \leftarrow$ Create $\mathrm{N} \times \mathrm{d}$ data matrix, with one row vector $x^{n}$ per data point <br> 2. $\mathrm{X} \leftarrow$ subtract mean $\bar{x}$ from each row vector $x^{n}$ in X <br> 3. $\Sigma \leftarrow$ covariance matrix of $X$ <br> 4. Find eigenvectors and eigenvalues of $\Sigma$ <br> 5. PC 's $\leftarrow$ the M eigenvectors with largest eigenvalues |  |
| :---: | :---: |
|  |  |
|  |  |

PCA Example

$$
\hat{\mathbf{x}}^{n}=\overline{\mathbf{x}}+\sum_{i=1}^{M} z_{i}^{n} \mathbf{u}_{i}
$$





Very Nice When Initial Dimension Not Too Big
What if very large dimensional data?

- e.g., Images ( $d \geq 10^{\wedge} 4$ )

Problem:

- Covariance matrix $\Sigma$ is size ( $\mathrm{d} \times \mathrm{d}$ )
- $\mathrm{d}=10^{4} \rightarrow|\Sigma|=10^{8}$

Singular Value Decomposition (SVD) to the rescue!

- pretty efficient algs available, including Matlab SVD
- some implementations find just top N eigenvectors


## Singular Value Decomposition

To generate principle components:

- Subtract mean $\overline{\mathrm{x}}=\frac{1}{N} \sum_{n=1}^{N} \mathrm{x}^{n} \quad$ from each data point, to create zero-centered data
- Create matrix $X$ with one row vector per (zero centered) data point
- Solve SVD: $X=U S V^{T}$
- Output Principle components: columns of $V$ (= rows of $V^{T}$ ) - Eigenvectors in $V$ are sorted from largest to smallest eigenvalues
- $S$ is diagonal, with $s_{k}{ }^{2}$ giving eigenvalue for kth eigenvector

Independent Components Analysis

- PCA seeks directions $<Y_{1} \ldots Y_{M}>$ in feature space $X$ that minimize reconstruction error
- ICA seeks directions $<Y_{1} \ldots Y_{M}>$ that are most statistically independent. I.e., that minimize $I(Y)$, the mutual information between the $Y_{j}$ :

$$
I(Y)=\left[\sum_{j=1}^{J} H\left(Y_{j}\right)\right]-H(Y)
$$

Which maximizes their departure from Gaussianity!

## Independent Components Analysis

- ICA seeks to minimize $I(Y)$, the mutual information between the $Y_{j}$ :

$$
I(Y)=\left[\sum_{j=1}^{J} H\left(Y_{j}\right)\right]-H(Y)
$$

$$
\left|\begin{array}{c}
y_{1}(t) \\
\vdots \\
y_{m}(t)
\end{array}\right|=\mathrm{W}\left|\begin{array}{c}
x_{1}(t) \\
\vdots \\
x_{n}(t)
\end{array}\right|
$$

- Example: Blind source separation
- Original features $x_{i}(t)$ are microphones at a cocktail party
- Each receives sounds from multiple people speaking
- ICA outputs directions that correspond to individual speakers $y_{k}(t)$

ICA with independent spatial components


## 1. Fisher Linear Discriminant

- A method for projecting data into lower dimension to hopefully improve classification
- We'll consider 2-class case


Project data onto vector that connects class means?


| Fisher Linear Discriminant <br> Project data onto one dimension, to help classification $y=\mathbf{w}^{T} \mathbf{x}$ | Project data onto one dimension, to help classification $y=\mathbf{w}^{T} \mathbf{x}$ |
| :---: | :---: |
| Fisher Linear Discriminant : $\arg \max _{\mathrm{w}} \times \frac{\left(m_{2}-m_{1}\right)^{2}}{s_{1}^{2}+s_{2}^{2}}$ |  |
| is solved by : $\mathrm{w} \propto \mathrm{S}_{\mathbf{W}}{ }^{-1}\left(\mathrm{~m}_{2}-\mathrm{m}_{1}\right)$ |  |
| Where $\mathrm{S}_{\mathrm{w}}$ is sum of within-class covariances: |  |
| $\mathrm{s}_{\mathrm{W}} \equiv \sum_{n \in C_{1}}\left(\mathrm{x}^{n}-\mathrm{m}_{1}\right)\left(\mathrm{x}^{n}-\mathrm{m}_{1}\right)^{T}+\sum_{n \in C_{2}}\left(\mathrm{x}^{n}-\mathrm{m}_{2}\right)\left(\mathrm{x}^{n}-\mathrm{m}_{2}\right)^{T}$ |  |



## Summary: Fisher Linear Discriminant

- Choose $\mathrm{n}-1$ dimension projection for n -class classification problem
- Use within-class covariances to determine the projection
- Minimizes a different sum of squared error function



## 2. Hidden Layers in Neural Networks

When \# hidden units < \# inputs, hidden layer also performs dimensionality reduction.

Each synthesized dimension (each hidden unit) is logistic function of inputs

$$
h_{k}(\mathrm{x})=\frac{1}{1+\exp \left(w_{0}+\sum_{i=1}^{N} w_{i} x_{i}\right)}
$$

Hidden units defined by gradient descent to (locally) minimize squared output classification/regression error

$$
E=\sum_{n=1}^{N} \sum_{k}\left(\hat{y_{k}}\left(x^{n}\right)-y_{k}\left(x^{n}\right)\right)^{2}
$$

Also allow networks with multiple hidden layers
$\rightarrow$ highly nonlinear components (in contrast with linear subspace of Fisher LD, PCA)
subspace of Fisher LD, PCA)
家




Neural Nets for Face Recognition


## What you should know

- Feature selection
- Single feature scoring criteria
- Search strategies
- Common approaches: Greedy addition of features, or greedy deletion
- Unsupervised dimension reduction using all features
- Principle Components Analysis
- Minimize reconstruction error
- Singular Value Decomposition
- Efficient PCA
- Independent components analysis
- Supervised dimension reduction
- Fisher Linear Discriminant
- Project to $\mathrm{n}-1$ dimensions to discriminate n classes
- Hidden layers of Neural Networks
- Most flexible, local minima issues

Further Readings

- "Singular value decomposition and principal component analysis," Wall, M.E, Rechtsteiner, A., and L. Rocha, in A Practical Approach to Microarray Data Analysis (D.P. Berrar, W. Dubitzky, M. Granzow, eds.) Kluwer, Norwell, MA, 2003. pp. 91-109 LANL LA-UR-02-4001

