

Co-Training for Semi-supervised learning (cont.)

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

April 23rd, 2007


©2005-2007 Carlos Guestrin

Exploiting redundant information in semi-supervised learning

- Want to predict Y from features \mathbf{X}
 - $f(\mathbf{X}) \mapsto Y$
 - have some labeled data \mathbf{L}
 - lots of unlabeled data \mathbf{U}
- Co-training assumption: \mathbf{X} is very expressive
 - $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2)$
 - can learn
 - $g_1(\mathbf{X}_1) \mapsto Y$
 - $g_2(\mathbf{X}_2) \mapsto Y$

Professor Faloutsos

my advisor



Christos Faloutsos

U.S. mail address:
Department of Computer Science
University of Maryland
College Park, MD 20742
(97-99: on leave at CMU)
Office: 3227 A.V. Williams Bldg.
Phone: (301) 405-2695
Fax: (301) 405-6707
Email: christos@cs.umd.edu

Current Position: Assoc. Professor of [Computer Science](#), (97-99: on leave at CMU)
Join Appointment: [Institute for Systems Research](#) (ISR).
Academic Degrees: Ph.D. and M.Sc. ([University of Toronto](#)), B.Sc. ([Nat. Tech. U. Athens](#))

Research Interests:

- Query by content in multimedia databases;
- Fractals for clustering and spatial access methods;
- Data mining.

can do a lot
with unlabeled
data, especially if $X_1 \perp X_2 | Y$

Co-Training Algorithm

[Blum & Mitchell '99]

(example of
the co-training
principle)

Given: labeled data L,

unlabeled data U

Loop:

Train g1 (^{x_1} hyperlink classifier) using L

Train g2 (^{x_2} page classifier) using L

Allow g1 to label p positive, n negative examps from U

Allow g2 to label p positive, n negative examps from U

~~Add~~ these self-labeled examples to L

move

Understanding Co-Training: A simple setting

- Suppose \mathbf{X}_1 and \mathbf{X}_2 are discrete

□ $|\mathbf{X}_1| = |\mathbf{X}_2| = N$

possible values

- No label noise

- Without unlabeled data, how hard is it to learn g_1 (or g_2)?

if X_1 is described by n binary features, $N = 2^n$

$|H| = 2^N$

#

training examples

hypothesis space

is dependent on N

1 $\{+, -\}$

$g_i \in H$

$\ln |H| = N \cdot \ln 2$

2 $\{+, -\}$


⋮

⋮

n $\{+, -\}$

Co-Training in simple setting – Iteration 0

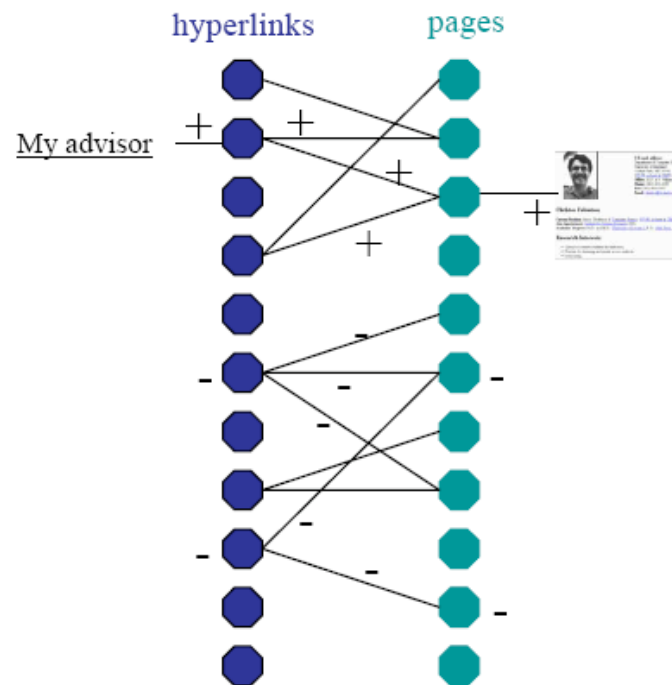
you get
a web page
with $x_1 = 12$
& $x_2 = 18$



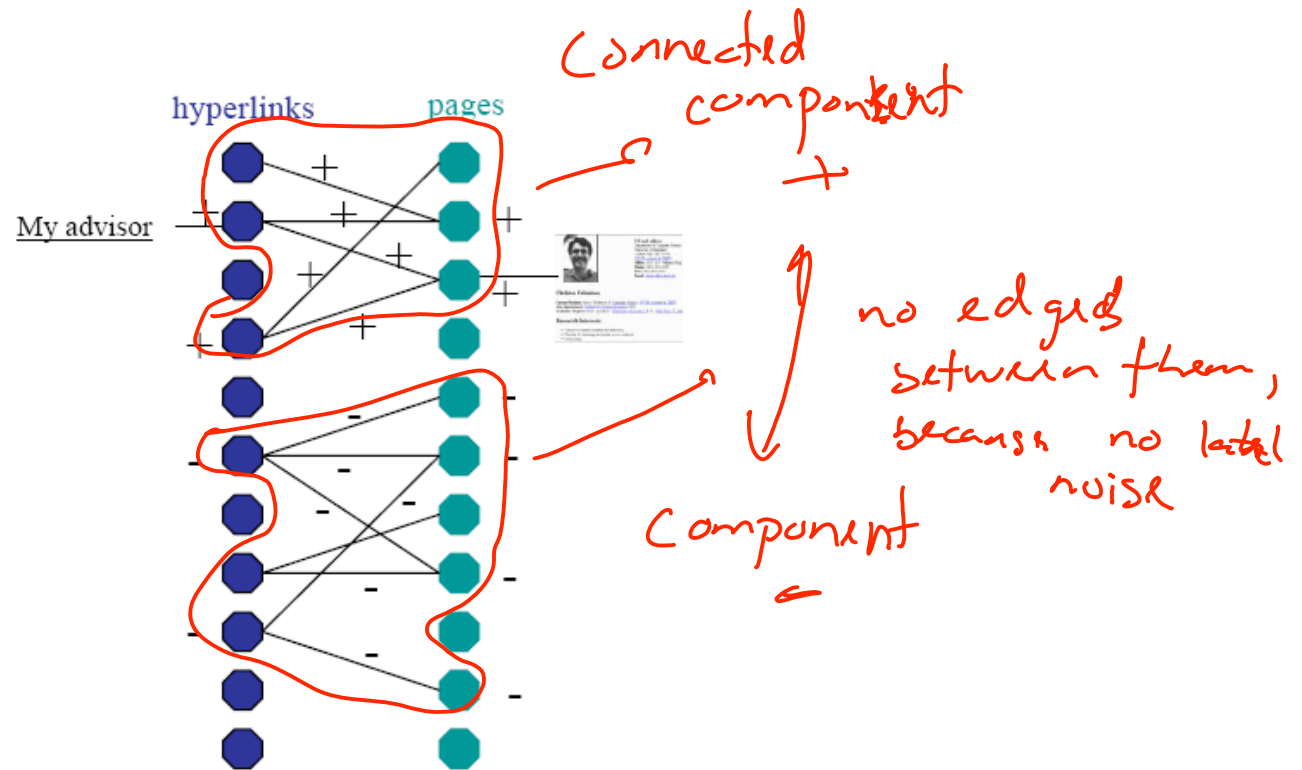
edge $x_1 = x_2$
to $x_2 = x_2$
means
 x_1 & x_2
co occurred
on a
web page

NO LABEL
NOISE

Co-Training in simple setting – Iteration 1



Co-Training in simple setting – after convergence



Co-Training in simple setting – Connected components

- Suppose infinite unlabeled data
 - Co-training must have at least one labeled example in each connected component of L+U graph

component g_j

- What's probability of making an error?

\exists connected component, where ~~no~~ no data was labeled

test point x

$E[\text{error}] = \sum_{g_j \in \text{components}} P(x \in g_j)$

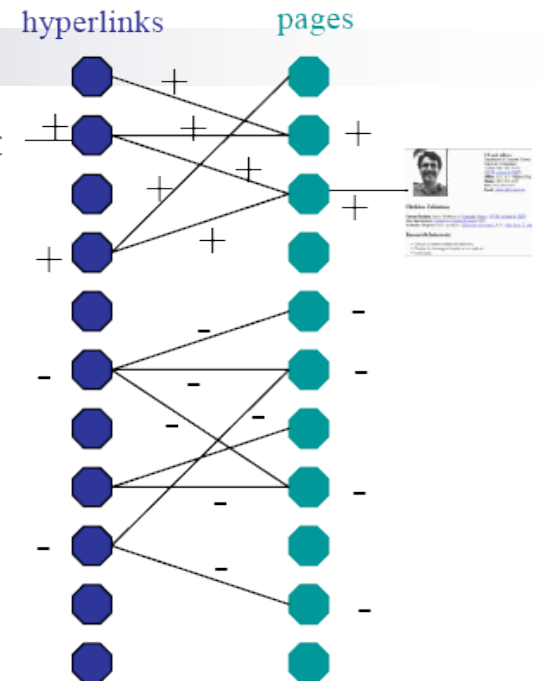
$(1 - P(x \in g_j))^m$

no training data

$$E[\text{error}] = \sum_j P(x \in g_j)(1 - P(x \in g_j))^m$$


- For k Connected components, how much labeled data?

Where g_j is the j th connected component of graph of L+U, m is number of labeled examples



How much unlabeled data?

Want to assure that connected components in the underlying distribution, G_D , are connected components in the observed sample, G_S



$O(\log(N)/\alpha)$ examples assure that with high probability, G_S has same connected components as G_D [Karger, 94]

N is size of G_D , α is min cut over all connected components of G_D

Co-Training theory

- Want to predict Y from features \mathbf{X}
 - $f(\mathbf{X})$ a Y
- Co-training assumption: \mathbf{X} is very expressive
 - $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2)$
 - want to learn $g_1(\mathbf{X}_1)$ a Y and $g_2(\mathbf{X}_2)$ a Y
- *Assumption:* $\exists g_1, g_2, \forall \mathbf{x} \ g_1(\mathbf{x}_1) = f(\mathbf{x}), g_2(\mathbf{x}_2) = f(\mathbf{x})$
- One co-training result [Blum & Mitchell '99]
 - If
 - $(\mathbf{X}_1 \perp \mathbf{X}_2 \mid Y)$
 - g_1 & g_2 are PAC learnable from noisy data (and thus f)
 - Then
 - f is PAC learnable from weak initial classifier plus unlabeled data

What you need to know about co-training

- Unlabeled data can help supervised learning (a lot) when there are (mostly) independent redundant features
- One theoretical result:
 - If $(\mathbf{X}_1 \perp \mathbf{X}_2 \mid Y)$ and g_1 & g_2 are PAC learnable from noisy data (and thus f)
 - Then f is PAC learnable from weak initial classifier plus unlabeled data
 - Disagreement between g_1 and g_2 provides bound on error of final classifier
- Applied in many real-world settings:
 - Semantic lexicon generation [Riloff, Jones 99] [Collins, Singer 99], [Jones 05]
 - Web page classification [Blum, Mitchell 99]
 - Word sense disambiguation [Yarowsky 95]
 - Speech recognition [de Sa, Ballard 98]
 - Visual classification of cars [Levin, Viola, Freund 03]



Transductive SVMs

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University


April 23rd, 2007

©2005-2007 Carlos Guestrin

Semi-supervised learning and discriminative models

- We have seen semi-supervised learning for generative models
 - EM
- What can we do for discriminative models
 - Not regular EM
 - we can't compute $P(x)$
 - But there are discriminative versions of EM
 - Co-Training!
 - Many other tricks... let's see an example

Linear classifiers – Which line is better?



Data:

$$\begin{aligned} &\langle x_1^{(1)}, \dots, x_1^{(m)}, y_1 \rangle \\ &\quad \vdots \\ &\langle x_n^{(1)}, \dots, x_n^{(m)}, y_n \rangle \end{aligned}$$


Example i:

$$\langle x_i^{(1)}, \dots, x_i^{(m)} \rangle \text{ — } m \text{ features}$$

$$y_i \in \{-1, +1\} \text{ — class}$$

$$\mathbf{w} \cdot \mathbf{x} = \sum_j w^{(j)} x^{(j)}$$


Support vector machines (SVMs)



$$\text{minimize}_w \quad w \cdot w$$
$$(w \cdot x_j + b) y_j \geq 1, \quad \forall j$$

- Solve efficiently by quadratic programming (QP)
 - Well-studied solution algorithms
- Hyperplane defined by support vectors

What if we have unlabeled data?



$$\mathbf{w} \cdot \mathbf{x} = \sum_j w^{(j)} x^{(j)}$$

n_L Labeled Data:

$$\begin{aligned} &\langle x_1^{(1)}, \dots, x_1^{(m)}, y_1 \rangle \\ &\vdots \\ &\langle x_{n_L}^{(1)}, \dots, x_{n_L}^{(m)}, y_{n_L} \rangle \end{aligned}$$

Example i:


$$\langle x_i^{(1)}, \dots, x_i^{(m)} \rangle \text{ — } m \text{ features}$$

$$y_i \in \{-1, +1\} \text{ — class}$$

n_U Unlabeled Data:

$$\begin{aligned} &\langle x_1^{(1)}, \dots, x_1^{(m)}, ? \rangle \\ &\vdots \\ &\langle x_{n_U}^{(1)}, \dots, x_{n_U}^{(m)}, ? \rangle \end{aligned}$$


Transductive support vector machines (TSVMs)



minimize_w $w \cdot w$

$$(w \cdot x_j + b) y_j \geq 1, \quad \forall j$$

Transductive support vector machines (TSVMs)



minimize $\mathbf{w}, \{\hat{y}_1, \dots, \hat{y}_{n_U}\}$ $\mathbf{w} \cdot \mathbf{w}$

$$(\mathbf{w} \cdot \mathbf{x}_j + b) y_j \geq 1, \quad \forall j = 1, \dots, n_L$$

$$(\mathbf{w} \cdot \mathbf{x}_u + b) \hat{y}_u \geq 1, \quad \forall u = 1, \dots, n_U$$

$$\hat{y}_u \in \{-1, +1\}, \quad \forall u = 1, \dots, n_U$$

What's the difference between transductive learning and semi-supervised learning?

- Not much, and
- A lot!!!
- Semi-supervised learning:
 - labeled and unlabeled data ! learn \mathbf{w}
 - use \mathbf{w} on test data
- Transductive learning
 - same algorithms for labeled and unlabeled data, but...
 - unlabeled data is test data!!!
- You are learning on the test data!!!
 - OK, because you never look at the labels of the test data
 - can get better classification
 - but be very very very very very very very very careful!!!
 - never use test data prediction accuracy to tune parameters, select kernels, etc.


Adding slack variables

$$\text{minimize}_{\mathbf{w}, \{\hat{y}_1, \dots, \hat{y}_{n_U}\}} \quad \mathbf{w} \cdot \mathbf{w}$$

$$(\mathbf{w} \cdot \mathbf{x}_j + b) y_j \geq 1 \quad \forall j = 1, \dots, n_L$$

$$(\mathbf{w} \cdot \mathbf{x}_u + b) \hat{y}_u \geq 1 \quad \forall u = 1, \dots, n_U$$

$$\hat{y}_u \in \{-1, +1\}, \quad \forall u = 1, \dots, n_U$$



Transductive SVMs – now with slack variables! [Vapnik 98]


■ Optimize $w, \{\xi_1, \dots, \xi_{n_L}\}, \{\hat{y}_1, \dots, \hat{y}_{n_U}\}, \{\hat{\xi}_1, \dots, \hat{\xi}_{n_U}\}$

$$\text{minimize } w \cdot w + C \sum_j \xi_j + \hat{C} \sum_u \hat{\xi}_u$$

$$(w \cdot x_j + b) y_j \geq 1 - \xi_j, \quad \forall j = 1, \dots, n_L$$

$$(w \cdot x_u + b) \hat{y}_u \geq 1 - \hat{\xi}_u, \quad \forall u = 1, \dots, n_u$$

$$\hat{y}_u \in \{-1, +1\}, \quad \forall u = 1, \dots, n_u$$



Learning Transductive SVMs is hard!

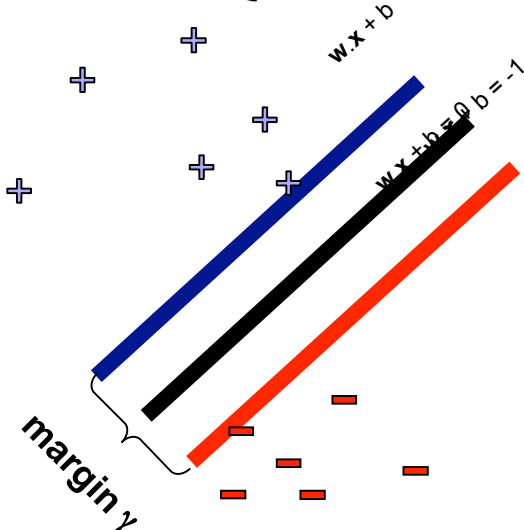
■ Optimize $w, \{\xi_1, \dots, \xi_{n_L}\}, \{\hat{y}_1, \dots, \hat{y}_{n_U}\}, \{\hat{\xi}_1, \dots, \hat{\xi}_{n_U}\}$

$$\text{minimize} \quad w \cdot w + C \sum_j \xi_j + \hat{C} \sum_u \hat{\xi}_u$$

$$(\mathbf{w} \cdot \mathbf{x}_j + b) y_j \geq 1 - \xi_j, \quad \forall j = 1, \dots, n_L$$

$$(\mathbf{w} \cdot \mathbf{x}_u + b) \hat{y}_u \geq 1 - \hat{\xi}_u, \quad \forall u = 1, \dots, n_u$$

$$\hat{y}_u \in \{-1, +1\}, \quad \forall u = 1, \dots, n_u$$



■ Integer Program

- NP-hard!!!
- Well-studied solution algorithms, but will not scale up to very large problems


A (heuristic) learning algorithm for Transductive SVMs [Joachims 99]

$$\text{minimize} \quad \mathbf{w} \cdot \mathbf{w} + C \sum_j \xi_j + \hat{C} \sum_u \hat{\xi}_u$$

$$(\mathbf{w} \cdot \mathbf{x}_j + b) y_j \geq 1 - \xi_j, \quad \forall j = 1, \dots, n_L$$

$$(\mathbf{w} \cdot \mathbf{x}_u + b) \hat{y}_u \geq 1 - \hat{\xi}_u, \quad \forall u = 1, \dots, n_u$$

$$\hat{y}_u \in \{-1, +1\}, \quad \forall u = 1, \dots, n_u$$



- If you set \hat{C} to zero \rightarrow ignore unlabeled data
- Intuition of algorithm:
 - start with small \hat{C}
 - add labels to some unlabeled data based on classifier prediction
 - slowly increase \hat{C}
 - keep on labeling unlabeled data and re-running classifier

Some results classifying news articles – from [Joachims 99]




Figure 6: Average P/R-breakeven point on the Reuters dataset for different training set sizes and a test set size of 3,299.

What you need to know about transductive SVMs

- What is transductive v. semi-supervised learning
- Formulation for transductive SVM
 - can also be used for semi-supervised learning
- Optimization is hard!
 - Integer program
- There are simple heuristic solution methods that work well here



Dimensionality reduction

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

April 23rd, 2007

©2005-2007 Carlos Guestrin

Dimensionality reduction



- Input data may have thousands or millions of dimensions!
 - e.g., text data has
- **Dimensionality reduction:** represent data with fewer dimensions
 - easier learning – fewer parameters
 - visualization – hard to visualize more than 3D or 4D
 - discover “intrinsic dimensionality” of data
 - high dimensional data that is truly lower dimensional

Feature selection

- Want to learn $f: \mathbf{X} \mapsto Y$
 - $\mathbf{X} = \langle X_1, \dots, X_n \rangle$
 - but some features are more important than others
- **Approach:** select subset of features to be used by learning algorithm
 - **Score** each feature (or sets of features)
 - **Select** set of features with best score

Simple greedy **forward** feature selection algorithm

- Pick a dictionary of features
 - e.g., polynomials for linear regression
- Greedy heuristic:
 - Start from empty (or simple) set of features $F_0 = \emptyset$
 - Run learning algorithm for current set of features F_t
 - Obtain h_t
 - Select **next best feature** X_i
 - e.g., X_j that results in lowest cross-validation error learner when learning with $F_t \cup \{X_j\}$
 - $F_{t+1} \leftarrow F_t \cup \{X_i\}$
 - Recurse

Simple greedy **backward** feature selection algorithm

- Pick a dictionary of features
 - e.g., polynomials for linear regression
- Greedy heuristic:
 - Start from all features $F_0 = F$
 - Run learning algorithm for current set of features F_t
 - Obtain h_t
 - Select **next worst feature** X_i
 - e.g., X_j that results in lowest cross-validation error learner when learning with $F_t - \{X_j\}$
 - $F_{t+1} \leftarrow F_t - \{X_i\}$
 - Recurse

Impact of feature selection on classification of fMRI data [Pereira et al. '05]

Accuracy classifying
category of word read
by subject

#voxels	mean	subjects							
		233B	329B	332B	424B	474B	496B	77B	86B
50	0.735	0.783	0.817	0.55	0.783	0.75	0.8	0.65	0.75
100	0.742	0.767	0.8	0.533	0.817	0.85	0.783	0.6	0.783
200	0.737	0.783	0.783	0.517	0.817	0.883	0.75	0.583	0.783
300	0.75	0.8	0.817	0.567	0.833	0.883	0.75	0.583	0.767
400	0.742	0.8	0.783	0.583	0.85	0.833	0.75	0.583	0.75
800	0.735	0.833	0.817	0.567	0.833	0.833	0.7	0.55	0.75
1600	0.698	0.8	0.817	0.45	0.783	0.833	0.633	0.5	0.75
all (~2500)	0.638	0.767	0.767	0.25	0.75	0.833	0.567	0.433	0.733

Table 1: Average accuracy across all pairs of categories, restricting the procedure to use a certain number of voxels for each subject. The highlighted line corresponds to the best mean accuracy, obtained using 300 voxels.

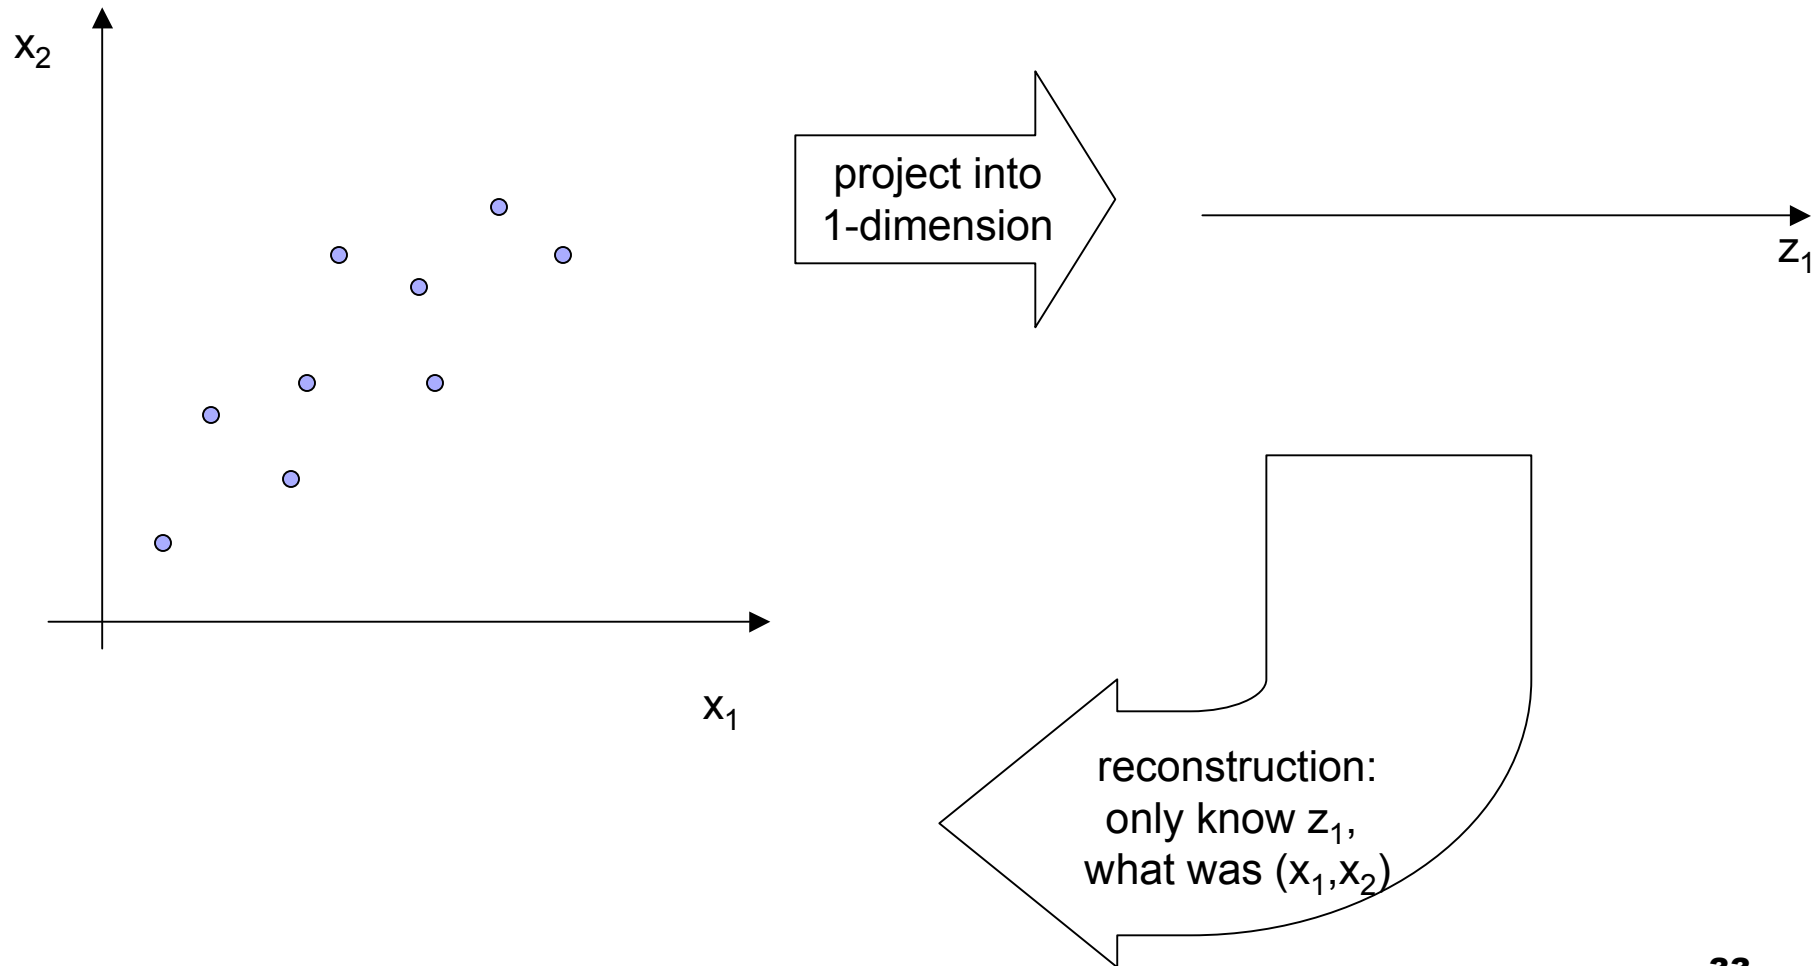
Voxels scored by p-value of regression to predict voxel value from the task

Lower dimensional projections



- Rather than picking a subset of the features, we can new features that are combinations of existing features
- Let's see this in the unsupervised setting
 - just **X**, but no **Y**

Linear projection and reconstruction



Principal component analysis – basic idea

- Project n -dimensional data into k -dimensional space while preserving information:
 - e.g., project space of 10000 words into 3-dimensions
 - e.g., project 3-d into 2-d
- Choose projection with minimum reconstruction error

Linear projections, a review

- Project a point into a (lower dimensional) space:
 - **point:** $\mathbf{x} = (x_1, \dots, x_n)$
 - **select a basis** – set of basis vectors – $(\mathbf{u}_1, \dots, \mathbf{u}_k)$
 - we consider orthonormal basis:
 - $\mathbf{u}_i \cdot \mathbf{u}_i = 1$, and $\mathbf{u}_i \cdot \mathbf{u}_j = 0$ for $i \neq j$
 - **select a center** – $\bar{\mathbf{x}}$, defines offset of space
 - **best coordinates** in lower dimensional space defined by dot-products: (z_1, \dots, z_k) , $z_i = (\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{u}_i$
 - minimum squared error

PCA finds projection that minimizes reconstruction error


- Given m data points: $\mathbf{x}^i = (x_1^i, \dots, x_n^i)$, $i=1 \dots m$
- Will represent each point as a projection:

$$\square \hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j \quad \text{where: } \bar{\mathbf{x}} = \frac{1}{m} \sum_{i=1}^m \mathbf{x}^i \quad \text{and} \quad z_j^i = \mathbf{x}^i \cdot \mathbf{u}_j$$

- PCA:

- Given $k \ll n$, find $(\mathbf{u}_1, \dots, \mathbf{u}_k)$ minimizing reconstruction error:

$$error_k = \sum_{i=1}^m (\mathbf{x}^i - \hat{\mathbf{x}}^i)^2$$



Understanding the reconstruction error

- Note that \mathbf{x}^i can be represented exactly by n-dimensional projection:

$$\mathbf{x}^i = \bar{\mathbf{x}} + \sum_{j=1}^n z_j^i \mathbf{u}_j$$

- Rewriting error:

$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j \quad z_j^i = \mathbf{x}^i \cdot \mathbf{u}_j$$

- Given $k \ll n$, find $(\mathbf{u}_1, \dots, \mathbf{u}_k)$

minimizing reconstruction error:

$$error_k = \sum_{i=1}^m (\mathbf{x}^i - \hat{\mathbf{x}}^i)^2$$

Reconstruction error and covariance matrix

$$error_k = \sum_{i=1}^m \sum_{j=k+1}^n [\mathbf{u}_j \cdot (\mathbf{x}^i - \bar{\mathbf{x}})]^2$$

$$\Sigma = \frac{1}{m} \sum_{i=1}^m (\mathbf{x}^i - \bar{\mathbf{x}})(\mathbf{x}^i - \bar{\mathbf{x}})^T$$

Minimizing reconstruction error and eigen vectors

- Minimizing reconstruction error equivalent to picking orthonormal basis $(\mathbf{u}_1, \dots, \mathbf{u}_n)$ minimizing:

$$error_k = \sum_{j=k+1}^n \mathbf{u}_j^T \Sigma \mathbf{u}_j$$



- Eigen vector:
- Minimizing reconstruction error equivalent to picking $(\mathbf{u}_{k+1}, \dots, \mathbf{u}_n)$ to be eigen vectors with smallest eigen values

Basic PCA algorithm

- Start from m by n data matrix \mathbf{X}
- **Recenter**: subtract mean from each row of \mathbf{X}
 - $\mathbf{X}_c \tilde{\mathbf{A}} \mathbf{X} - \bar{\mathbf{X}}$
- **Compute covariance matrix**:
 - $\Sigma \tilde{\mathbf{A}} \mathbf{X}_c^T \mathbf{X}_c$
- Find **eigen vectors and values** of Σ
- **Principal components**: k eigen vectors with highest eigen values

PCA example



$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j$$



PCA example – reconstruction

$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j$$

only used first principal component



Eigenfaces [Turk, Pentland '91]

■ Input images:



■ Principal components:



Eigenfaces reconstruction

- Each image corresponds to adding 8 principal components:



Relationship to Gaussians

- PCA assumes data is Gaussian

- $\mathbf{x} \sim N(\bar{\mathbf{x}}; \Sigma)$


- Equivalent to weighted sum of simple Gaussians:

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{j=1}^n z_j \mathbf{u}_j; \quad z_j \sim N(0; \sigma_j^2)$$

- Selecting top k principal components equivalent to lower dimensional Gaussian approximation:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \sum_{j=1}^k z_j \mathbf{u}_j + \varepsilon; \quad z_j \sim N(0; \sigma_j^2)$$

- $\varepsilon \sim N(0; \sigma^2)$, where σ^2 is defined by error_k



Scaling up



- Covariance matrix can be really big!
 - Σ is n by n
 - 10000 features ! $|\Sigma|$
 - finding eigenvectors is very slow...
- Use singular value decomposition (SVD)
 - finds to k eigenvectors
 - great implementations available, e.g., Matlab `svd`

SVD

- Write $\mathbf{X} = \mathbf{U} \mathbf{S} \mathbf{V}^T$
 - $\mathbf{X} \leftarrow$ data matrix, one row per datapoint
 - $\mathbf{U} \leftarrow$ weight matrix, one row per datapoint – coordinate of \mathbf{x}^i in eigenspace
 - $\mathbf{S} \leftarrow$ singular value matrix, diagonal matrix
 - in our setting each entry is eigenvalue λ_j
 - $\mathbf{V}^T \leftarrow$ singular vector matrix
 - in our setting each row is eigenvector \mathbf{v}_j

PCA using SVD algorithm

- Start from m by n data matrix \mathbf{X}
- **Recenter**: subtract mean from each row of \mathbf{X}
 - $\mathbf{X}_c \leftarrow \mathbf{X} - \bar{\mathbf{X}}$
- Call SVD algorithm on \mathbf{X}_c – ask for k singular vectors
- **Principal components**: k singular vectors with highest singular values (rows of \mathbf{V}^T)
 - **Coefficients** become:

Using PCA for dimensionality reduction in classification

- Want to learn $f: \mathbf{X} \mapsto Y$
 - $\mathbf{X} = \langle X_1, \dots, X_n \rangle$
 - but some features are more important than others
- **Approach:** Use PCA on \mathbf{X} to select a few important features

PCA for classification can lead to problems...

- Direction of maximum variation may be unrelated to “discriminative” directions:
- PCA often works very well, but sometimes must use more advanced methods
 - e.g., Fisher linear discriminant

What you need to know



- Dimensionality reduction
 - why and when it's important
- Simple feature selection
- Principal component analysis
 - minimizing reconstruction error
 - relationship to covariance matrix and eigenvectors
 - using SVD
 - problems with PCA