Reading:
Vapnik 1998
Joachims 1999 (see class website)

Transductive SVMs

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

April 17th, 2006

Semi-supervised learning and discriminative models

- We have seen semin-supervised learning for generative models
  - EM

x is visible
y hidden

$$\max \log P(X) = \max \log \sum_y P(X, y)$$

- What can we do for discriminative models
  - Not regular EM

$P(Y|X) \rightarrow$ can't compute $P(X)$

    - 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?

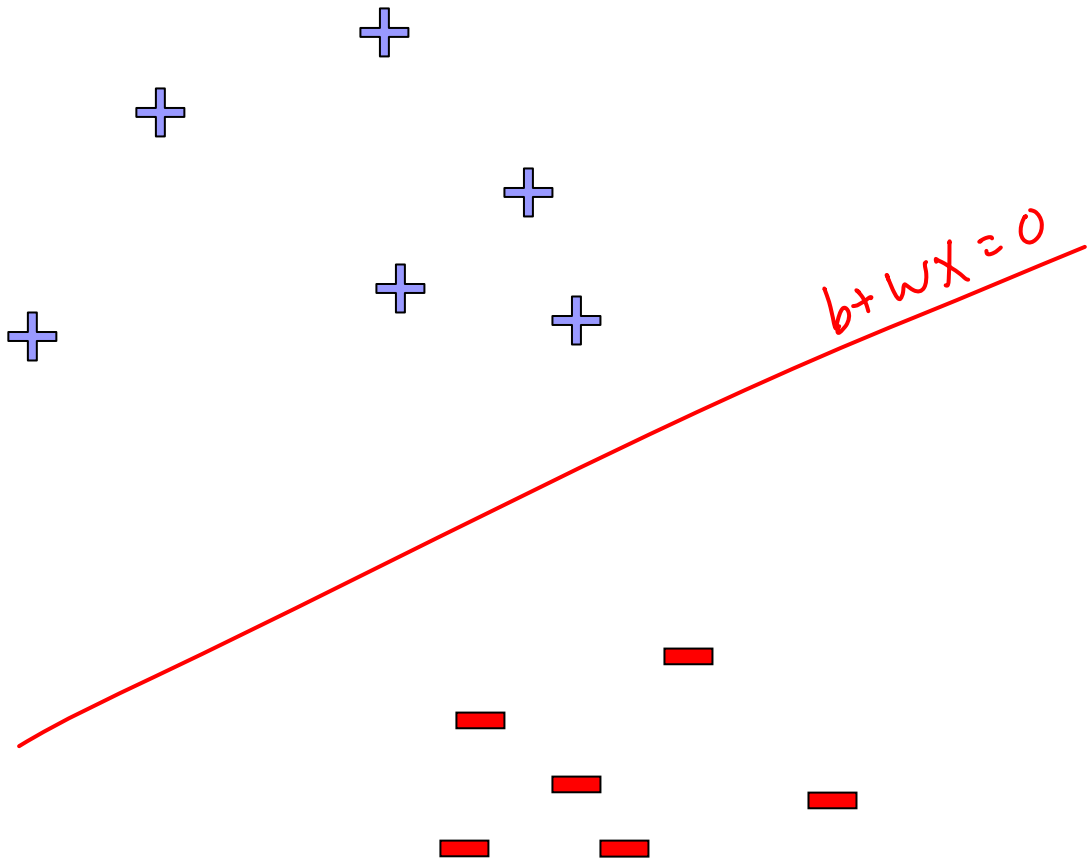

Data:

features

class

$$\left\langle x_1^{(1)}, \ldots, x_1^{(m)}, y_1 \right\rangle$$

$$\vdots$$

$$\left\langle x_n^{(1)}, \ldots, x_n^{(m)}, y_n \right\rangle$$

Example i:

$\left\langle x_i^{(1)}, \ldots, x_i^{(m)} \right\rangle$ — m features

$y_i \in \{-1, +1\}$ — class

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

Support vector machines (SVMs)

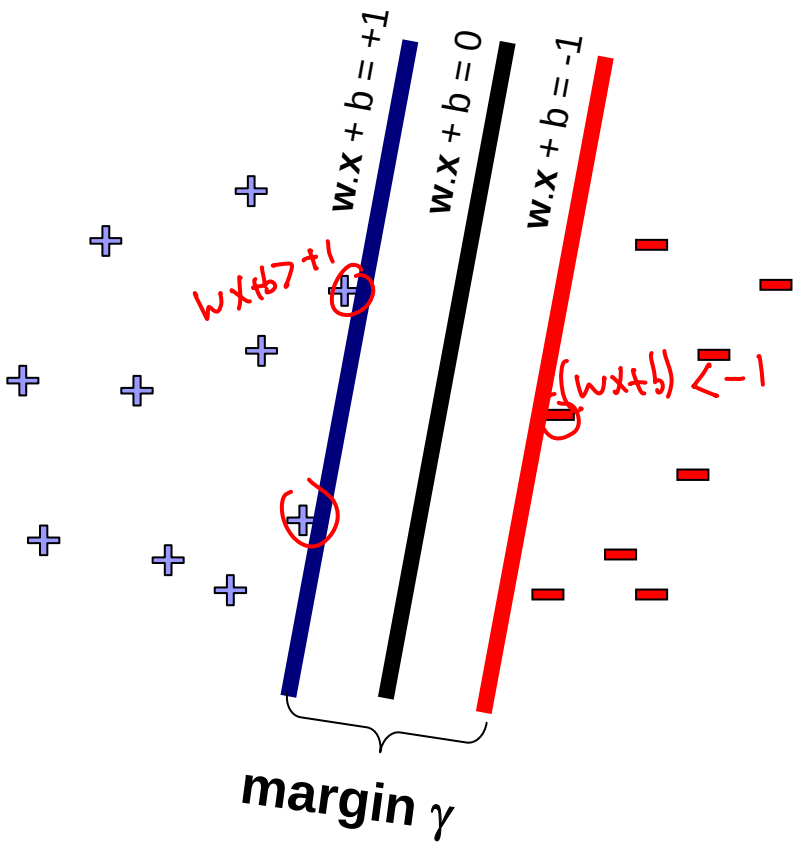

$$\text{minimize}_{\mathbf{w}} \quad \mathbf{w}.\mathbf{w}$$
$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1, \ \forall j$$

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

What if we have unlabeled data?

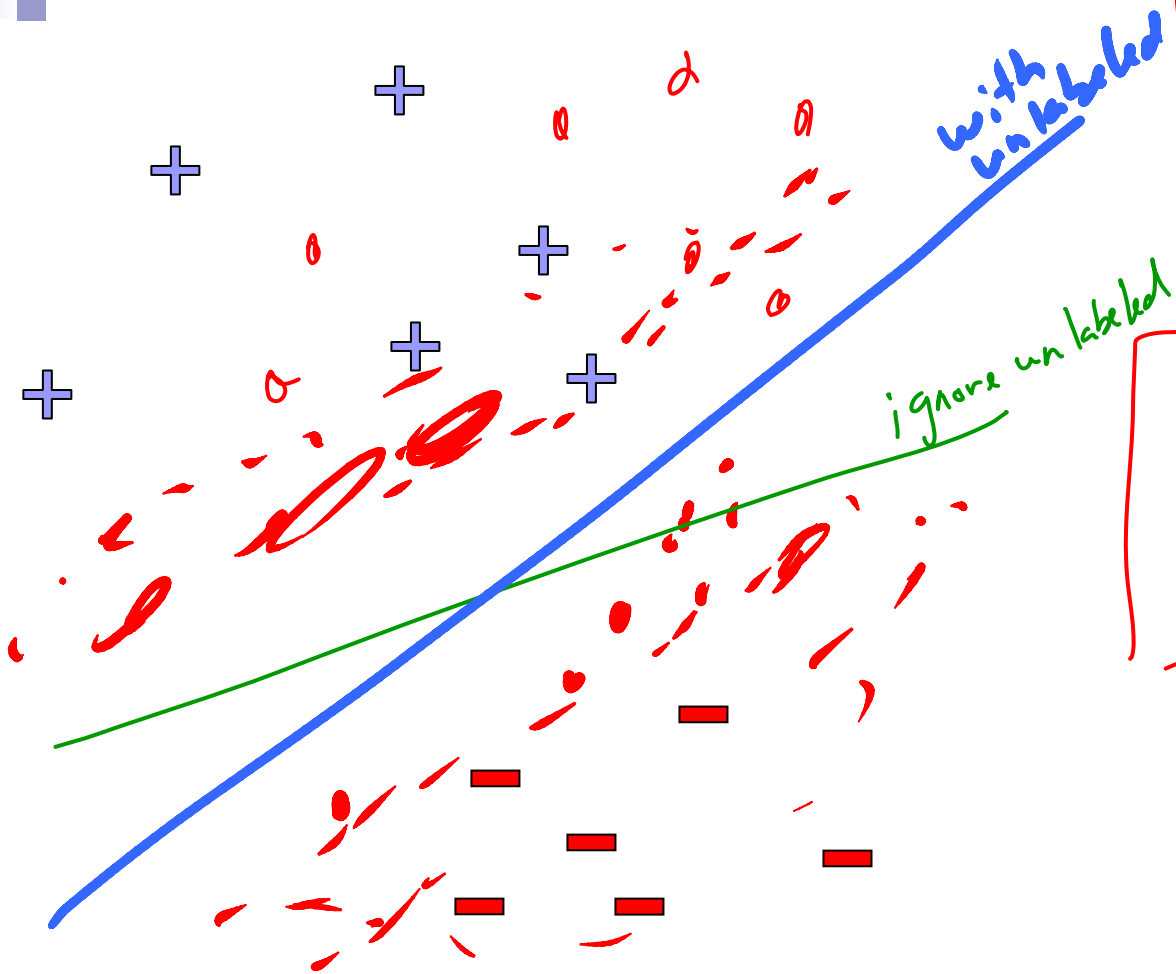

n_L Labeled Data:

$$\langle x_1^{(1)}, \ldots, x_1^{(m)}, y_1 \rangle$$
$$\vdots$$
$$\langle x_n^{(1)}, \ldots, x_n^{(m)}, y_{n_L} \rangle$$

Example i:

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

$y_i \in \{-1, +1\}$ — class

n_U Unlabeled Data:

$$\langle x_1^{(1)}, \ldots, x_1^{(m)}, ? \rangle$$
$$\vdots$$
$$\langle x_n^{(1)}, \ldots, x_{n_U}^{(m)}, ? \rangle$$

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

Transductive support vector machines (TSVMs)

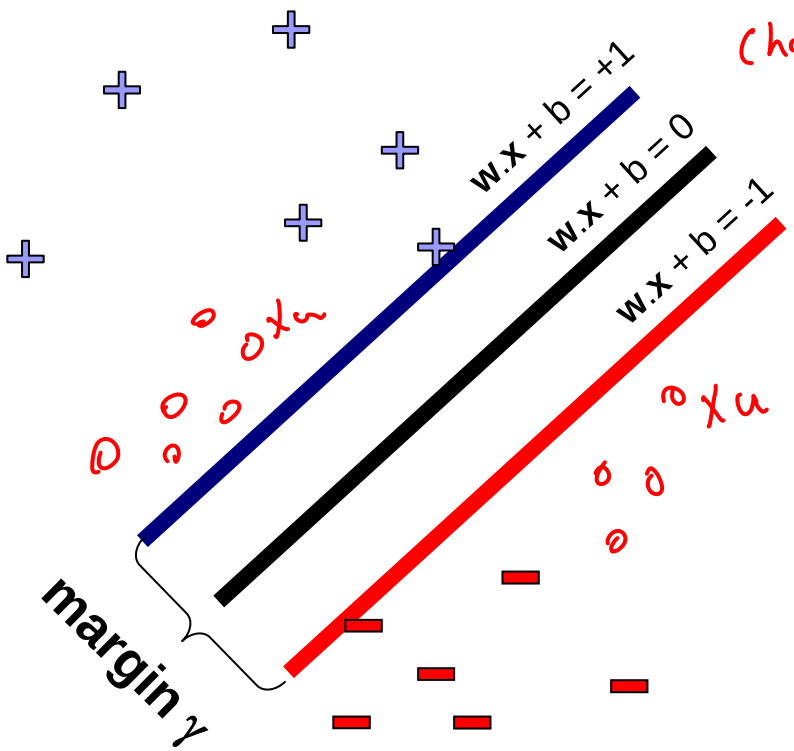

choose u

$$\text{minimize}_{\mathbf{w}} \quad \mathbf{w}.\mathbf{w}$$

choose $\{\hat{y}_1, \ldots, \hat{y}_{n+u}\}$

$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1, \quad \forall j \in L$$

$$(w.x_u + b)\,\hat{y}_u \geq 1 \quad \forall u \in U$$

$$\hat{y}_u \in \{+1, -1\} \quad \forall u \in U$$

Transductive support vector machines (TSVMs)

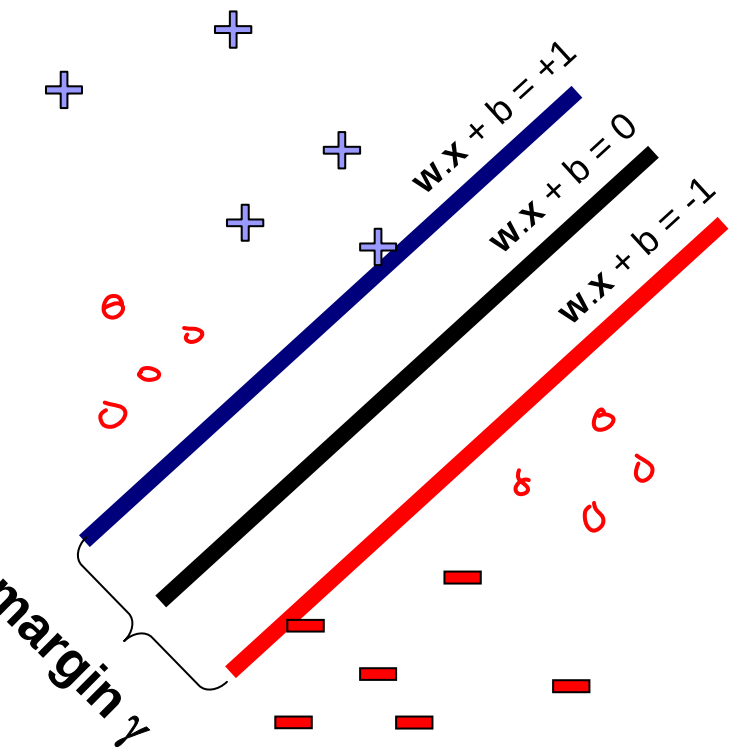

$$\text{minimize}_{\mathbf{w},\{\hat{y}_1,\ldots,\hat{y}_{n_U}\}} \quad \mathbf{w}.\mathbf{w}$$

$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1, \;\; \forall j = 1,\ldots,n_L$$

$$\left(\mathbf{w}.\mathbf{x}_u + b\right) \hat{y}_u \geq 1, \;\; \forall u = 1,\ldots,n_U$$

bilinear constraint

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

integer constraint → hard!

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

- Not much, and
- A lot!!!

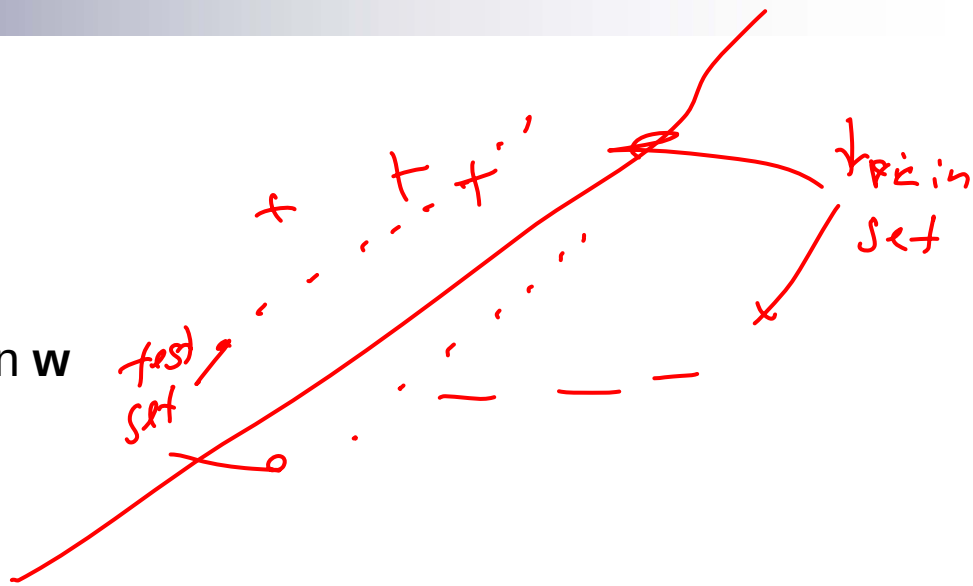

- Semi-supervised learning:
  - labeled and unlabeled data → learn **w**
  - use **w** on test data

- 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

$$\text{minimize}_{\mathbf{w},\{\hat{y}_1,\ldots,\hat{y}_{n_U}\}} \quad \mathbf{w}.\mathbf{w} + C\sum_j \xi_j + \hat{C}\sum_u \hat{\xi}_u$$

$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1 - \xi_j \quad \forall j = 1,\ldots,n_L$$

$$\left(\mathbf{w}.\mathbf{x}_u + b\right) \hat{y}_u \geq 1 - \hat{\xi}_u \quad \forall u = 1,\ldots,n_U$$

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

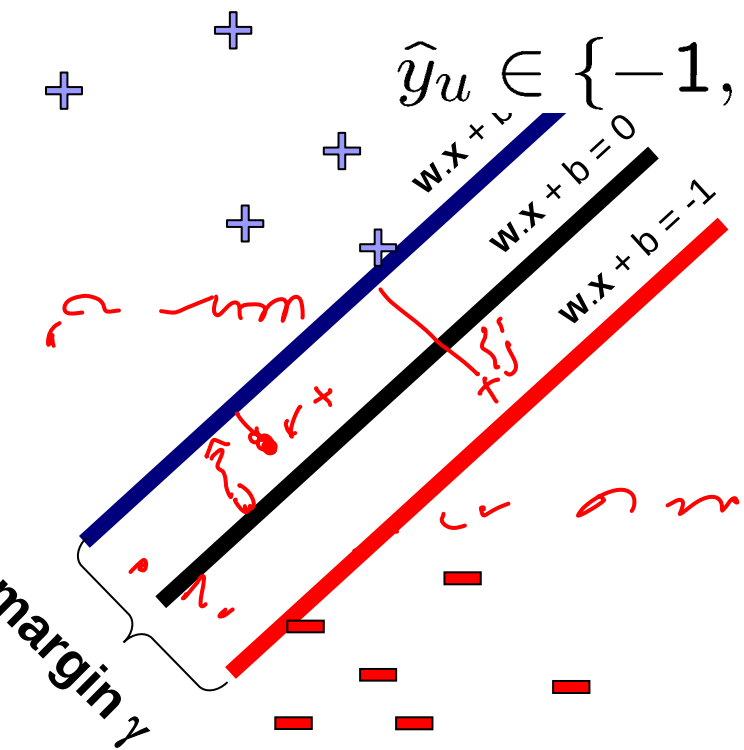

if $\hat{C} = 0$

learn only on labeled data

(ignore unlabeled data)

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

- Optimize $\mathbf{w}, \{\xi_1, ..., \xi_{n_L}\}, \{\hat{y}_1, ..., \hat{y}_{n_U}\}, \{\hat{\xi}_1, ..., \hat{\xi}_{n_U}\}$

  weights labeled slacks class of unlabeled unlabeled slacks

  minimize $\mathbf{w}.\mathbf{w} + C \sum_j \xi_j + \hat{C} \sum_u \hat{\xi}_u$

  $\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1 - \xi_j, \ \forall j = 1, ..., n_L$

  $\left(\mathbf{w}.\mathbf{x}_u + b\right) \hat{y}_u \geq 1 - \hat{\xi}_u, \ \forall u = 1, ..., n_u$

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

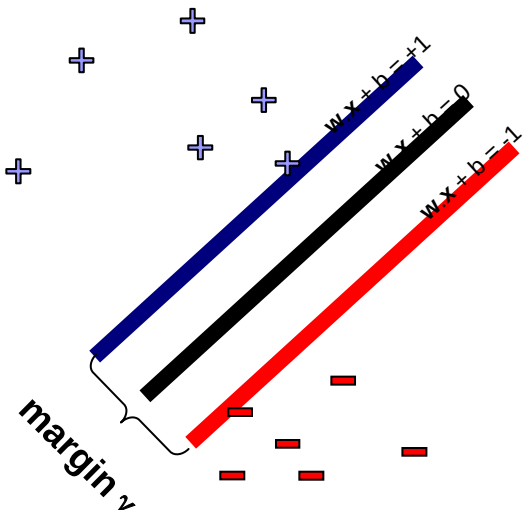

Learning Transductive SVMs is hard!

- Optimize $\mathbf{w}, \{\xi_1, ..., \xi_{n_L}\}, \{\hat{y}_1, ..., \hat{y}_{n_U}\}, \{\hat{\xi}_1, ..., \hat{\xi}_{n_U}\}$

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

$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1 - \xi_j, \ \ \forall j = 1, ..., n_L$$

$$\left(\mathbf{w}.\mathbf{x}_u + b\right) \hat{y}_u \geq 1 - \hat{\xi}_u, \ \ \forall u = 1, ..., n_u$$

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

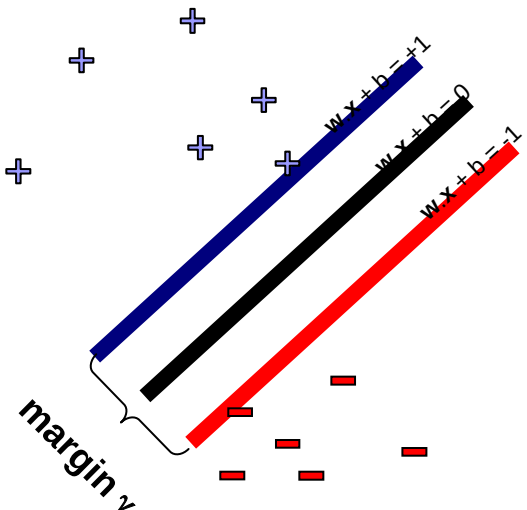

- Integer Program
  - NP-hard!!!
  - Well-studied solution algorithms, but will not scale up to very large problems 100 000 datapoints

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

minimize $\mathbf{w}.\mathbf{w} + C\sum_j \xi_j + \hat{C}\sum_u \hat{\xi}_u$

$$\left(\mathbf{w}.\mathbf{x}_j + b\right) y_j \geq 1 - \xi_j, \;\; \forall j = 1, ..., n_L$$

$$\left(\mathbf{w}.\mathbf{x}_u + b\right) \hat{y}_u \geq 1 - \hat{\xi}_u, \;\; \forall u = 1, ..., n_u$$

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

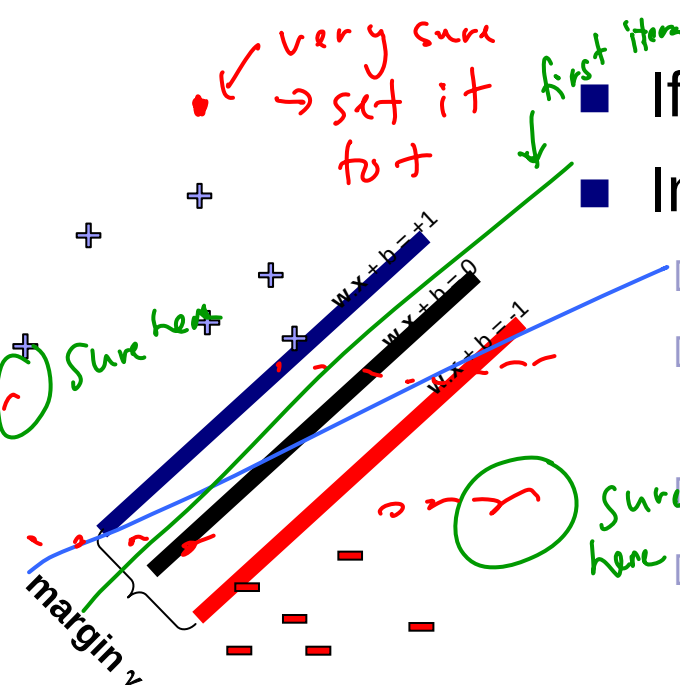

- If you set $\hat{C}$ to zero → ignore unlabeled data
- Intuition of algorithm:
  - start with small $\hat{C}$
  - add labels to some unlabeled data based on classifier prediction
  - slowly increase $\hat{C}$
  - keep on labeling unlabeled data and re-running classifier

if classifier very sure about x_u → label it

Some results classifying news articles – from [Joachims 99]

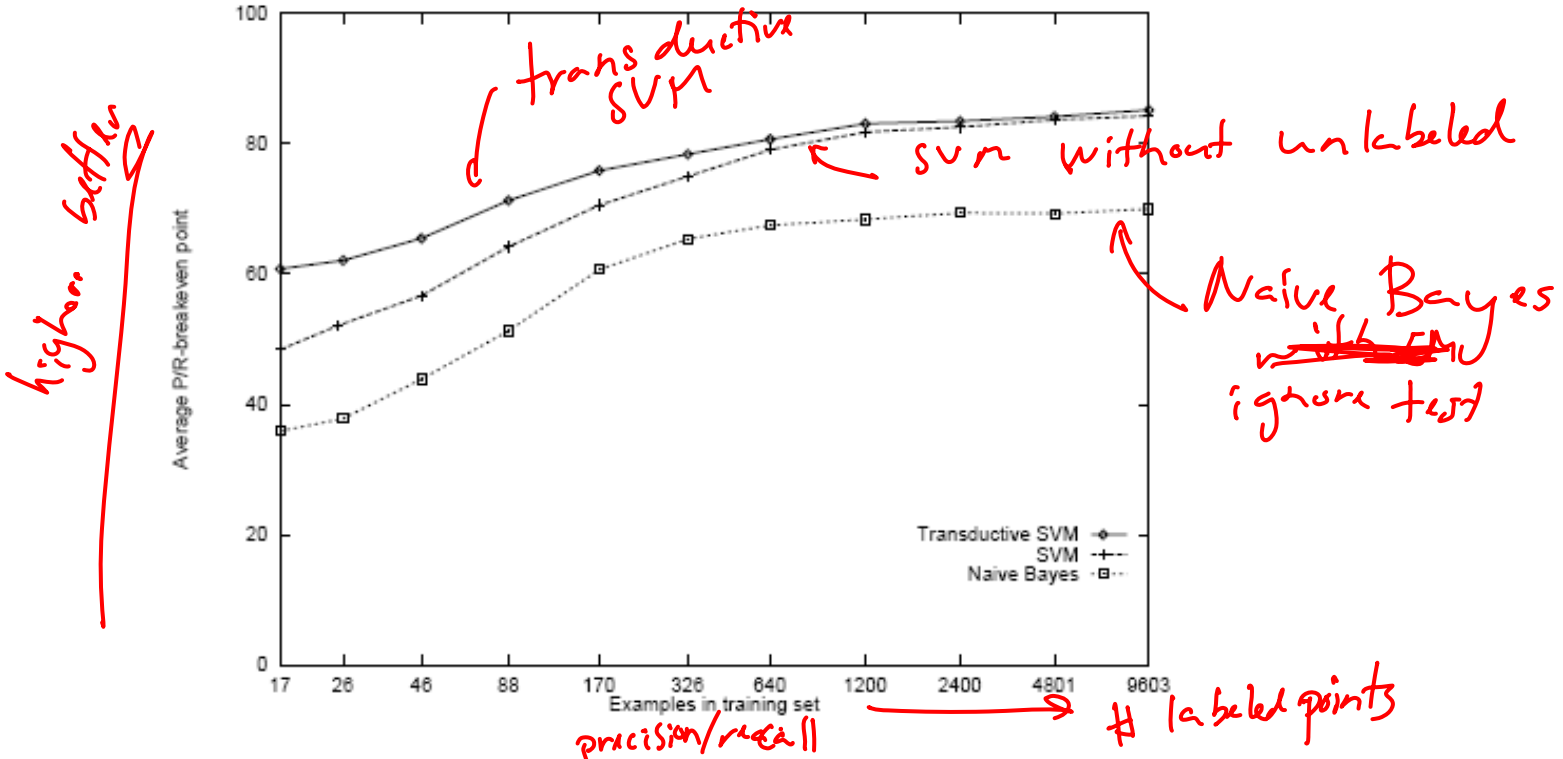

Figure 6: Average P/R-breakeven point on the Reuters dataset for different training set sizes and a test set size of 3,299.

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

Recommended reading:
Bishop, Chapters 3.6, 8.6
Shlens PCA tutorial
Wall et al. 2003 (PCA applied to gene expression data)

Dimensionality reduction

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

April 24th, 2006

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

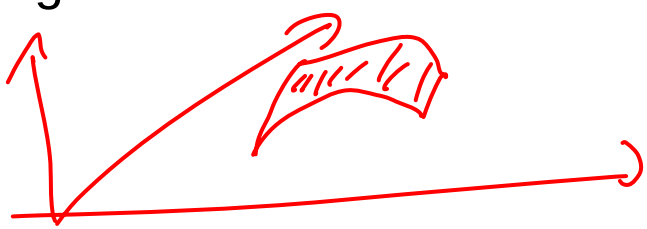

Feature selection

- Want to learn f:**X**↦Y for classification
  - **X**=<X_1,…,X_n>
  - 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 $F_t \cup \{X_j\}$
  - $F_{t+1} \leftarrow F_t \cup \{X_j\}$
  - 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 generate new features that are combinations of existing features

$X_{new} = 10X^2 + 17X^{17} - 28x^{-12}$

$X_{new} = g(X)$

- Let's see this in the unsupervised setting
  - just **X**, but no Y

ignore / don't have

Liner projection and reconstruction

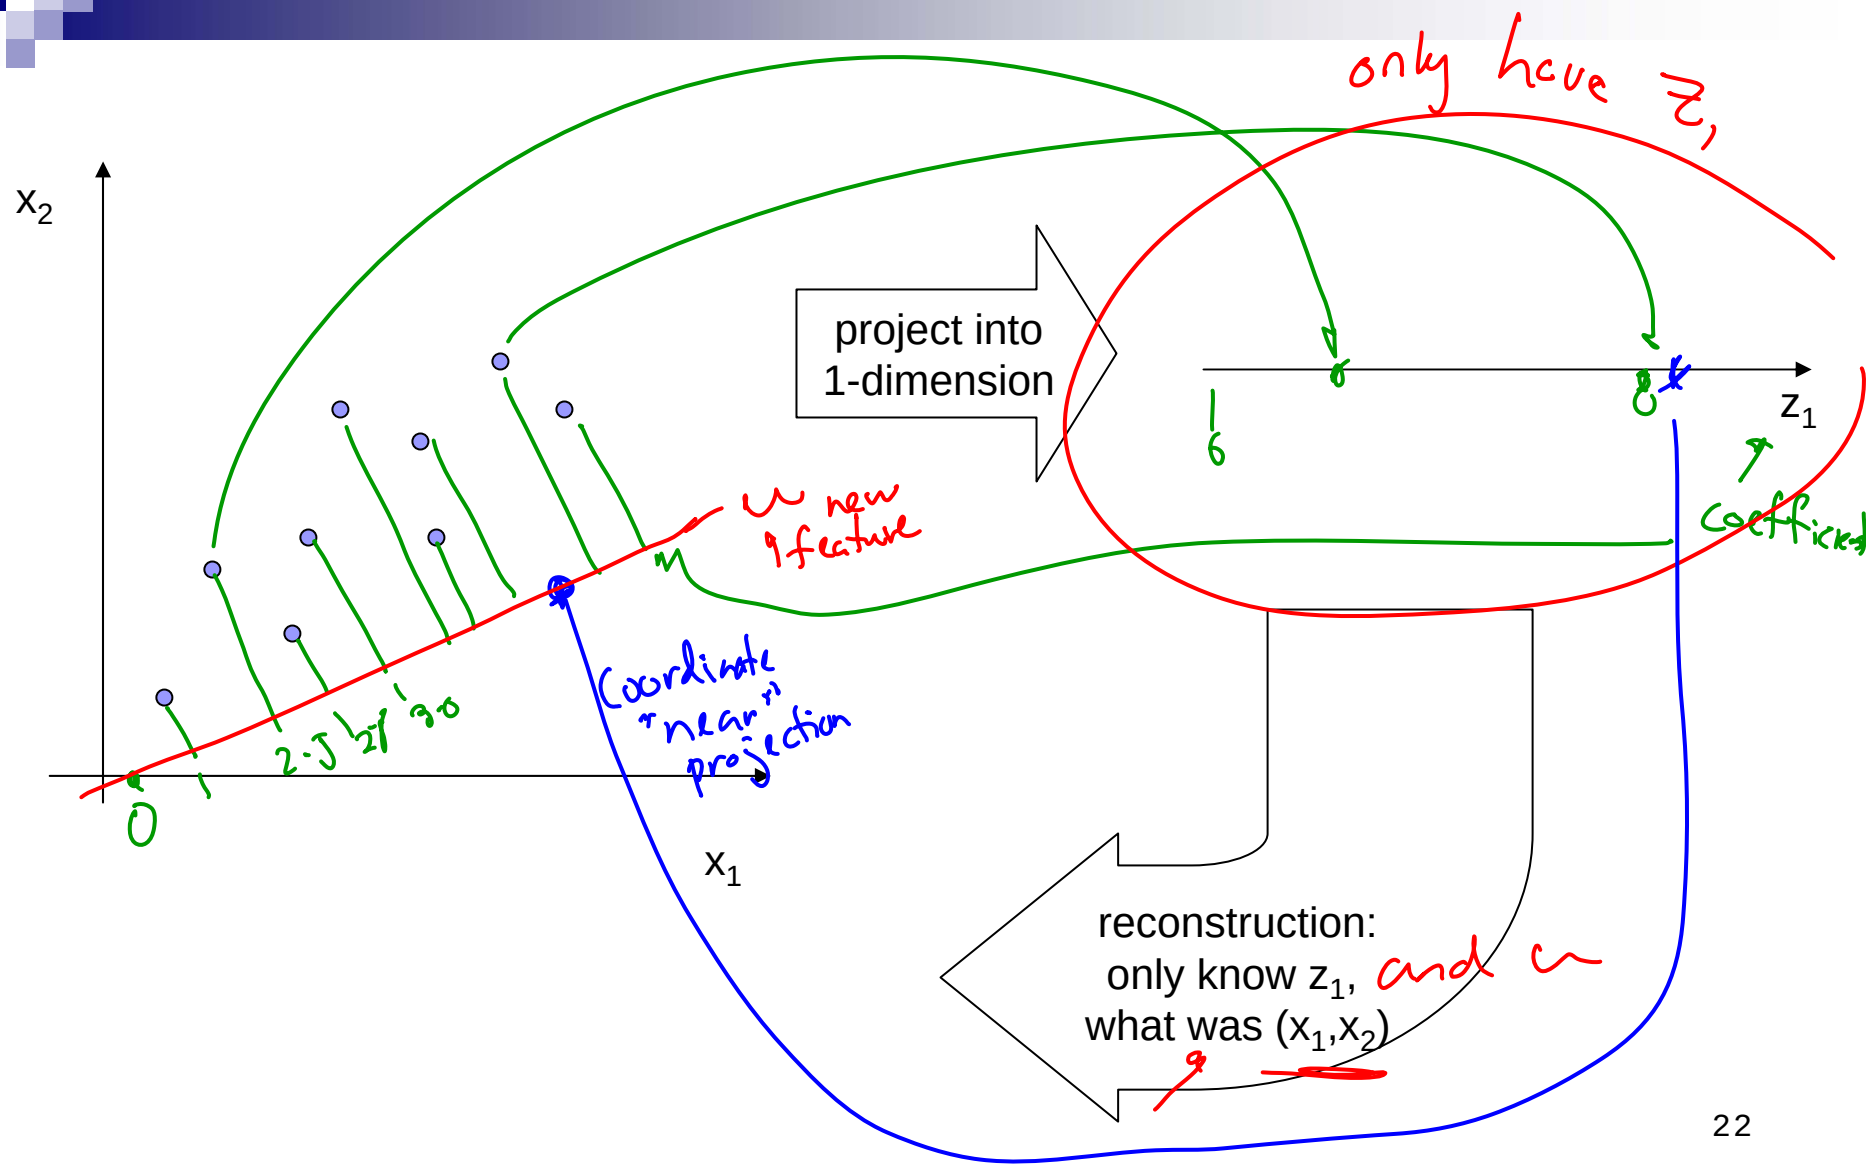

Principal component analysis – basic idea

- 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**: $\mathbf{x} = (x_1,\ldots,x_n)$
  - **select a basis** – set of basis vectors – $(\mathbf{u}_1,\ldots,\mathbf{u}_k)$
    - 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: $(z_1,\ldots,z_k)$, $z_i = (\mathbf{x}-\bar{\mathbf{x}})\cdot\mathbf{u}_i$
    - minimum squared error

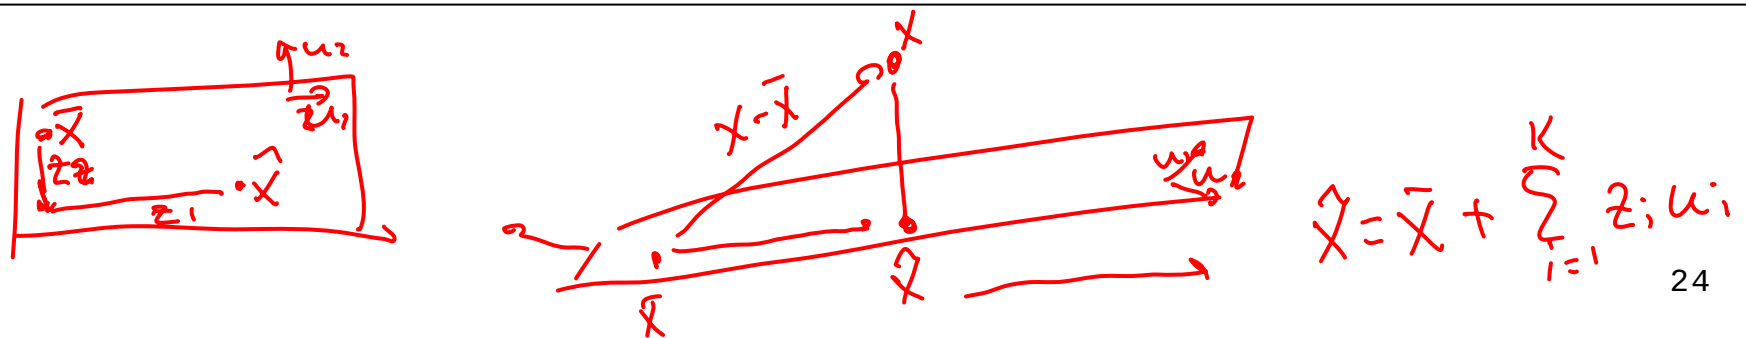

PCA finds projection that minimizes reconstruction error

- Given m data points: $\mathbf{x}^i = (x_1^i, \ldots, x_n^i)$, i=1…m
- Will represent each point as a projection:
  - $\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^{k} z_j^i \mathbf{u}_j$ where: $\bar{\mathbf{x}} = \frac{1}{m} \sum_{i=1}^{m} \mathbf{x}^i$ and $z_j^i = \mathbf{x}^i \cdot \mathbf{u}_j$
- PCA:
  - Given k≤n, find $(\mathbf{u}_1, \ldots, \mathbf{u}_k)$ minimizing reconstruction error:

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

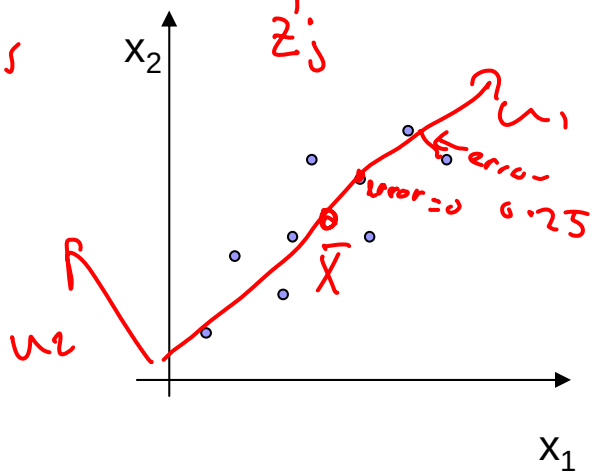

Understanding the reconstruction error

- 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 $k \le n$, find $(\mathbf{u}_1, \ldots, \mathbf{u}_k)$ minimizing reconstruction error:

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

Reconstruction error and covariance matrix

$a^2 = a^T a$

u_i is n by 1

$(a \cdot b)^2 = a^T b \, b^T a$

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

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

$$= \sum_{i=1}^{m} \sum_{j=k+1}^{n} [u_j^T (x^i - \bar{x}) (x^i - \bar{x})^T u_j]$$

$$= \sum_{j=k+1}^{n} \sum_{i=1}^{m} u_j^T (x^i - \bar{x}) (x^i - \bar{x})^T u_j$$

$$= \sum_{j=k+1}^{n} u_j^T \left[\sum_{i=1}^{m} (x^i - \bar{x}) (x^i - \bar{x})^T \right] u_j$$

$m \cdot \Sigma$ ← covariance matrix of data

$(n \times 1)(1 \times n) = (n \times n)$

$$(x^i - \bar{x}) (x^i - \bar{x})^T = \begin{bmatrix} \ddots & & \\ & (x_j^i - \bar{x}_j) \cdot (x_k^i - \bar{x}_k) & \\ & & \ddots \end{bmatrix} \leftarrow \sigma_{jk}$$

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

scalar ↓ ↓ n vector, z_j^i

$$(z_j^i u_j)^2$$

$$= z_j^i u_j^T u_j z_j^i$$

$u_j^T u_j = 1$ orthonormal

$$= z_j^i \cdot z_j^i$$

$$z_j^i = u_j \cdot (x^i - \bar{x})$$

Minimizing reconstruction error and eigen vectors

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

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

← consider ignored vectors in error

- Eigen vector:

u: $\Sigma u = \lambda u$ ⇒ $u^T \Sigma u = \lambda u^T u = \lambda$

eigen vector, eigen value

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

take eigen vectors / values of Σ

sort! $\lambda_i \geq \lambda_{i+1}$

u_1 λ_1

u_2 λ_2

⋮ ⋮

minimize $error_k$

$$\Rightarrow \min_u \sum_{j=k+1}^{n} u_j^T \Sigma u_j$$

$$= \min_u \sum_{j=k+1}^{n} \lambda_j$$

⇒ ignore vectors with low eigen vals

Basic PCA algoritm

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

PCA example

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

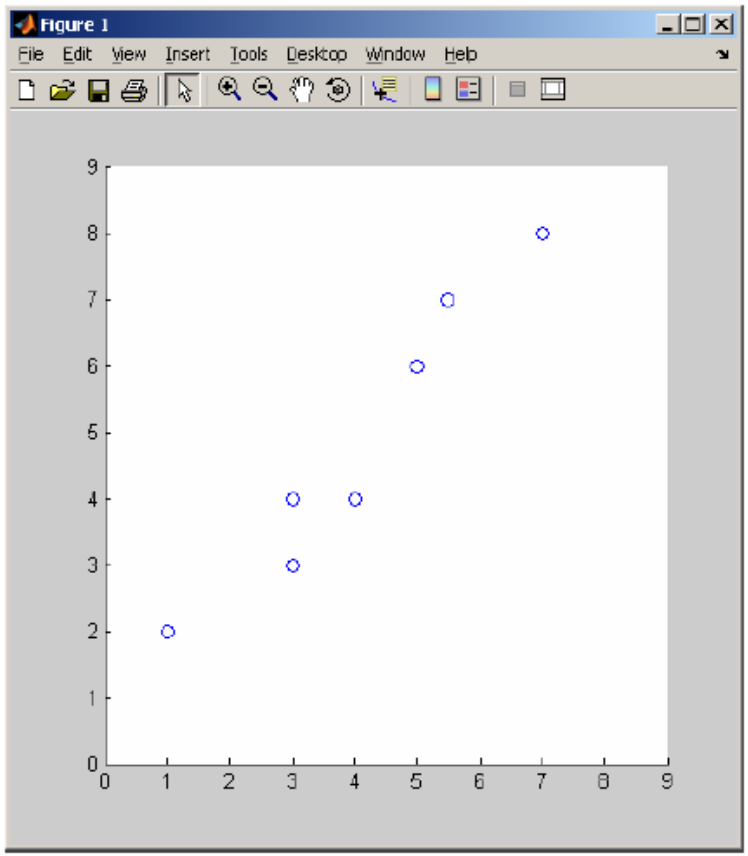

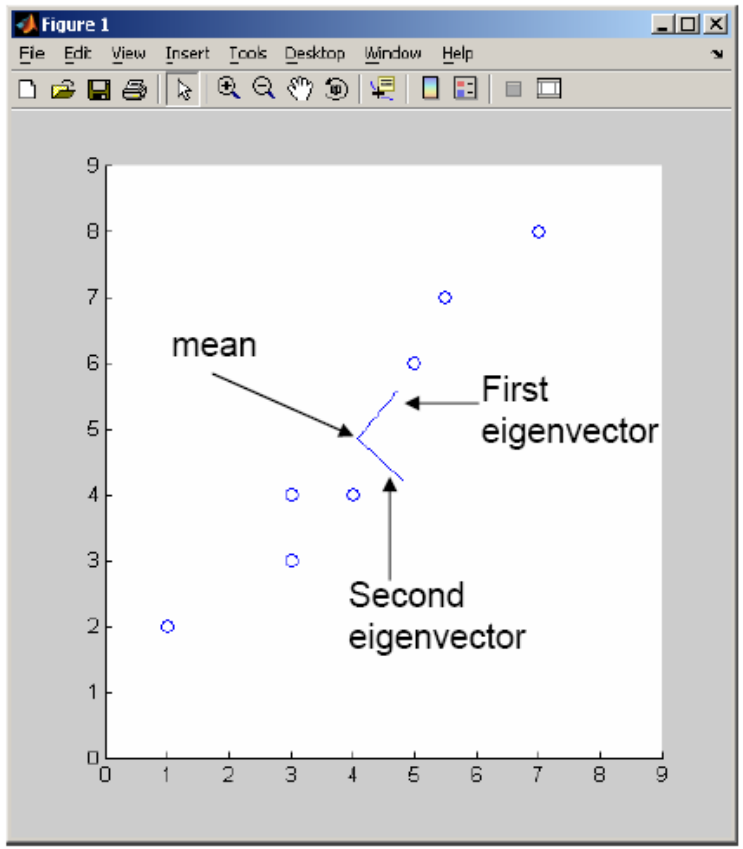

PCA example – reconstruction

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

only used first principal component

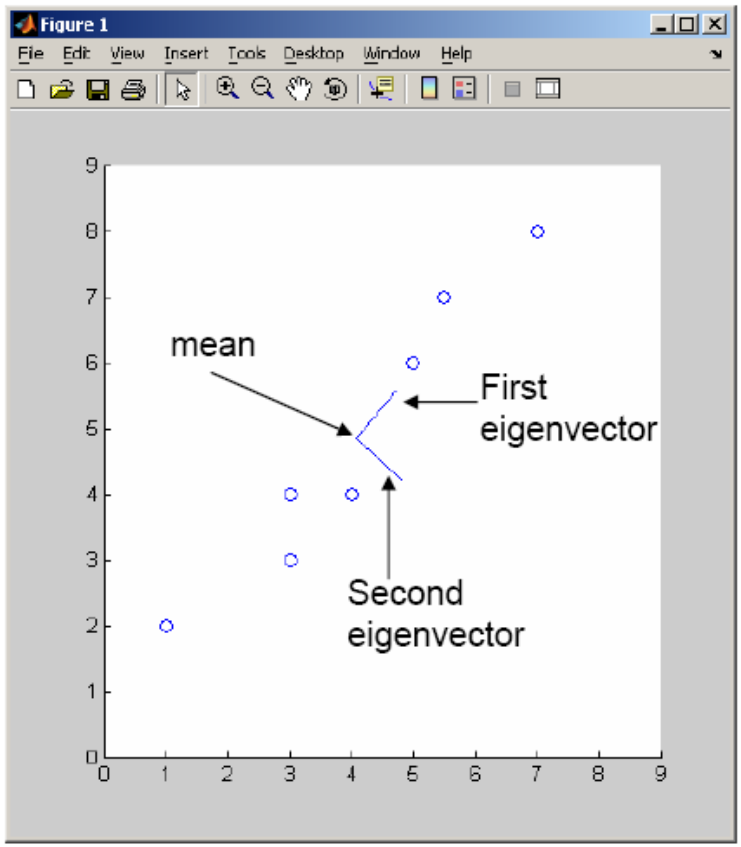

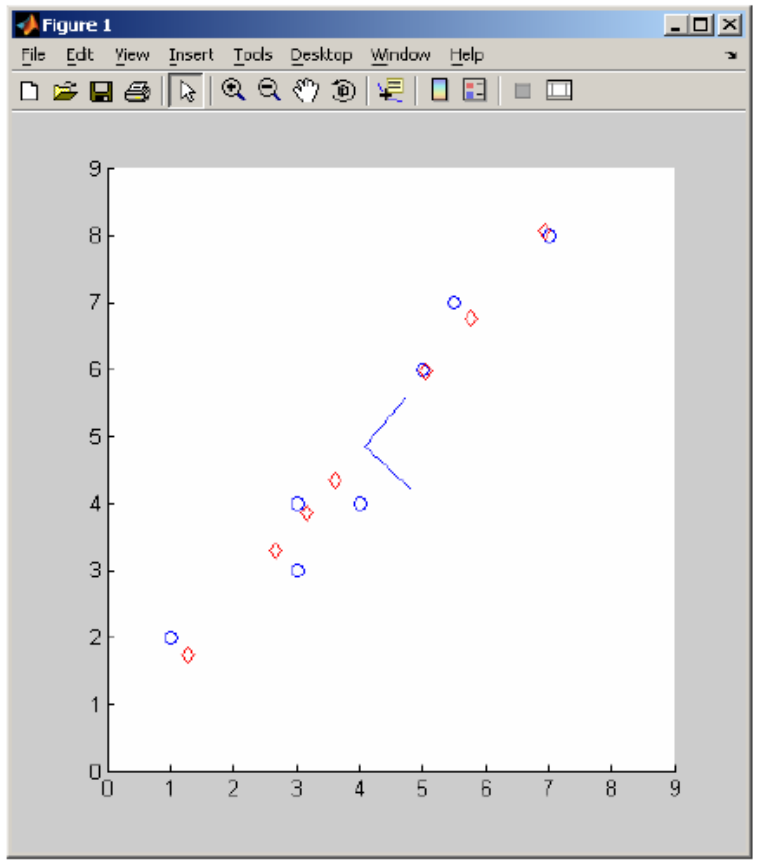

Eigenfaces [Turk, Pentland ’91]

- Input images:

- Principal components:

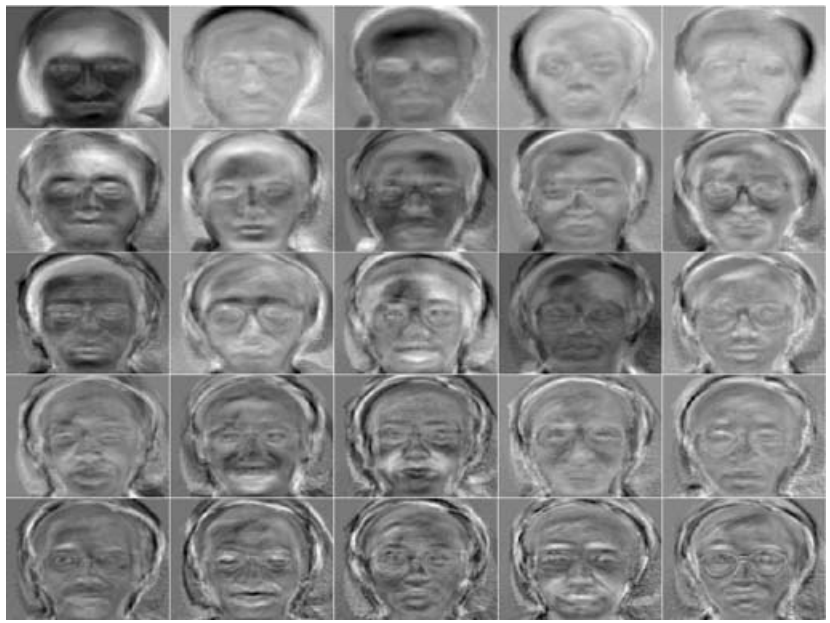

Eigenfaces reconstruction

- Each image corresponds to adding 8 principal components:

Relationship to Gaussians

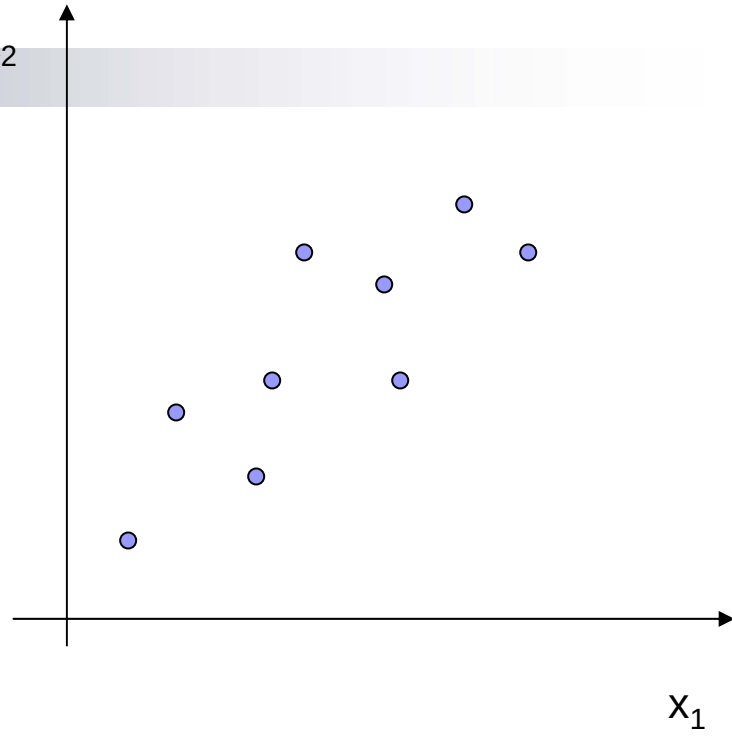

- PCA assumes data is Gaussian
  - $\mathbf{x} \sim N(\bar{\mathbf{x}};\Sigma)$
- Equivalent to weighted sum of simple Gaussians:

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

- 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 $\rightarrow$ $|\Sigma|$
  - 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 algoritm

- Start from m by n data matrix **X**
- **Recenter**: subtract mean from each row of **X**
  - $\mathbf{X}_c \leftarrow \mathbf{X} - \overline{\mathbf{X}}$
- 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:**X**$\mapsto$Y
  - **X**=<X_1,…,X_n>
  - but some features are more important than others

- **Approach**: Use PCA on **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