

CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Homework and midterm recap

Assignment 1

Updated data!

(I knew that test error looked too low...)

Midterm on Tuesday!

← **Here it is!**
(with some Gaussian blurring)



CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Week 1 recap

Supervised versus unsupervised learning

Learning approaches attempt to **model data** in order to solve a problem

Unsupervised learning approaches find patterns/relationships/structure in data, but **are not** optimized to solve a particular predictive task

- E.g. PCA, community detection

Supervised learning aims to directly model the relationship between input and output variables, so that the output variables can be predicted accurately given the input

- E.g. linear regression, logistic regression

Linear regression

Linear regression assumes a predictor of the form

$$X\theta = y$$

matrix of features
(data)

unknowns
(which features are relevant)

vector of outputs
(labels)

The diagram illustrates the equation $X\theta = y$. Three green arrows point from descriptive text below to the terms in the equation: one from 'matrix of features (data)' to X , one from 'unknowns (which features are relevant)' to θ , and one from 'vector of outputs (labels)' to y .

(or $Ax = b$ if you prefer)

Mean-squared error (MSE)

$$\frac{1}{N} \|y - X\theta\|_2^2$$

$$= \frac{1}{N} \sum_{i=1}^N (y_i - X_i \cdot \theta)^2$$

Representing the month as a feature

How would you build a feature to represent the **month**?

- {"Jan": 1, "Feb": 2, "Mar": 3, "Apr": 4, "May": 5, "Jun": 6, ...}[mon]?
 - Jan = [1,0,0,0,0,0,0,0,0,0,0,0]
 - Feb = [0,1,0,0,0,0,0,0,0,0,0,0]
 - Nov = [0,0,0,0,0,0,0,0,0,0,1,0] (etc.)
 - Jan = [0,0,0,0,0,0,0,0,0,0,0,0]
 - Feb = [0,0,0,0,0,0,0,0,0,0,0,1]
 - Mar = [0,0,0,0,0,0,0,0,0,1,0] (etc.)

Any benefit of one vs. another?

Occam's razor

"Among competing hypotheses, the one with the fewest assumptions should be selected"



(image from personalspirituality.net)

Regularization

Regularization is the process of penalizing model complexity during training

$$\arg \min_{\theta} = \frac{1}{N} \|y - X\theta\|_2^2 + \lambda \|\theta\|_2^2$$

How much should we trade-off accuracy versus complexity?



Model selection

A **validation set** is constructed to “tune” the model’s parameters

- Training set: used to **optimize the model’s parameters**
- Test set: used to report how well we expect the model to perform on **unseen data**
- Validation set: used to **tune** any model parameters that are not directly optimized

Model selection

A few “theorems” about training, validation, and test sets

- The training error **increases** as lambda **increases**
- The validation and test error are at least as large as the training error (assuming infinitely large random partitions)
- The validation/test error will usually have a “sweet spot” between under- and over-fitting

CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Week 2

Classification

Pitch Black - Unrated Director's Cut R CC

PITCH BLACK ★★★★★ 777 IMDb 7.1/10

[Watch Trailer](#)

When their ship crash-lands on a remote planet, the marooned passengers soon learn that escaped convict Riddick (Vin Diesel) isn't the only thing they have to fear. Deadly creatures lurk in the shadows, waiting to attack in the dark, and the planet is rapidly plunging into the









[See More](#)

Starring: Vin Diesel, Radha Mitchell
Runtime: 1 hour, 53 minutes
Available to watch on [supported devices](#).

UNRATED

Will I **purchase**
this product?
(yes)

Shop for engagement rings on Google Sponsored ⓘ

 <p>French-Set Halo Diamond... \$1,990.00 Ritani</p>	 <p>18K White Gold Delicate... \$950.00 Brilliant Earth ★★★★★ (57)</p>	 <p>18K White Gold Fancy D... \$1,825.00 Brilliant Earth ★★★★★ (13)</p>	 <p>Chamise Diamond Eng... \$975.00 Brilliant Earth ★★★★★ (7)</p>
 <p>Vintage Cushion Halo... \$4,140.00</p>	 <p>Princess Cut Diamond Eng... \$1,906.82</p>	 <p>18K White Gold Hudson... \$975.00</p>	 <p>18K White Gold Harmon... \$1,675.00</p>

Will I **click on**
this ad?
(no)

Classification

What animal appears in this image?
(mandarin duck)



Classification

What are the **categories** of the item
being described?

(book, fiction, philosophical fiction)

From [Booklist](#)

Houellebecq's deeply philosophical novel is about an alienated young man searching for happiness in the computer age. Bored with the world and too weary to try to adapt to the foibles of friends and coworkers, he retreats into himself, descending into depression while attempting to analyze the passions of the people around him. Houellebecq uses his nameless narrator as a vehicle for extended exploration into the meanings and manifestations of love and desire in human interactions. Ironically, as the narrator attempts to define love in increasingly abstract terms, he becomes less and less capable of experiencing that which he is so desperate to understand. Intelligent and well written, the short novel is a thought-provoking inspection of a generation's confusion about all things sexual. Houellebecq captures precisely the cynical disillusionment of disaffected youth. *Bonnie Johnston --This text refers to an out of print or unavailable edition of this title.*

Linear regression

Linear regression assumes a predictor of the form

$$X\theta = y$$

matrix of features
(data)


unknowns
(which features are relevant)

vector of outputs
(labels)

The diagram shows the equation $X\theta = y$ centered at the top. Three green arrows point from descriptive text below to the terms in the equation: one from 'matrix of features (data)' to X , one from 'unknowns (which features are relevant)' to θ , and one from 'vector of outputs (labels)' to y .

Regression vs. classification

But how can we predict **binary** or **categorical** variables?



(linear) classification

We'll attempt to build **classifiers** that make decisions according to rules of the form

$$y_i = \begin{cases} 1 & \text{if } X_i \cdot \theta > 0 \\ 0 & \text{otherwise} \end{cases}$$

In week 2

1. Naïve Bayes

Assumes an **independence** relationship between the features and the class label and “learns” a simple model by counting

2. Logistic regression

Adapts the **regression** approaches we saw last week to binary problems

3. Support Vector Machines

Learns to classify items by finding a hyperplane that separates them

Naïve Bayes (2 slide summary)

$$(feature_i \perp\!\!\!\perp feature_j | label)$$



$$p(feature_i, feature_j | label)$$

=

$$p(feature_i | label)p(feature_j | label)$$

Double-counting: naïve Bayes vs Logistic Regression

Q: What would happen if we trained two regressors, and attempted to “naively” combine their parameters?


$$\text{no. of pages} = \alpha + \beta_1 \cdot \delta(\text{mentions wizards})$$

$$\text{no. of pages} = \alpha + \beta_2 \cdot \delta(\text{mentions witches})$$

$$\text{no. of pages} = \alpha + \beta_1 \cdot \delta(\text{mentions wizards}) + \beta_2 \cdot \delta(\text{mentions witches})$$

Logistic regression

A: sigmoid function: $\sigma(t) = \frac{1}{1+e^{-t}}$



Logistic regression

Training:

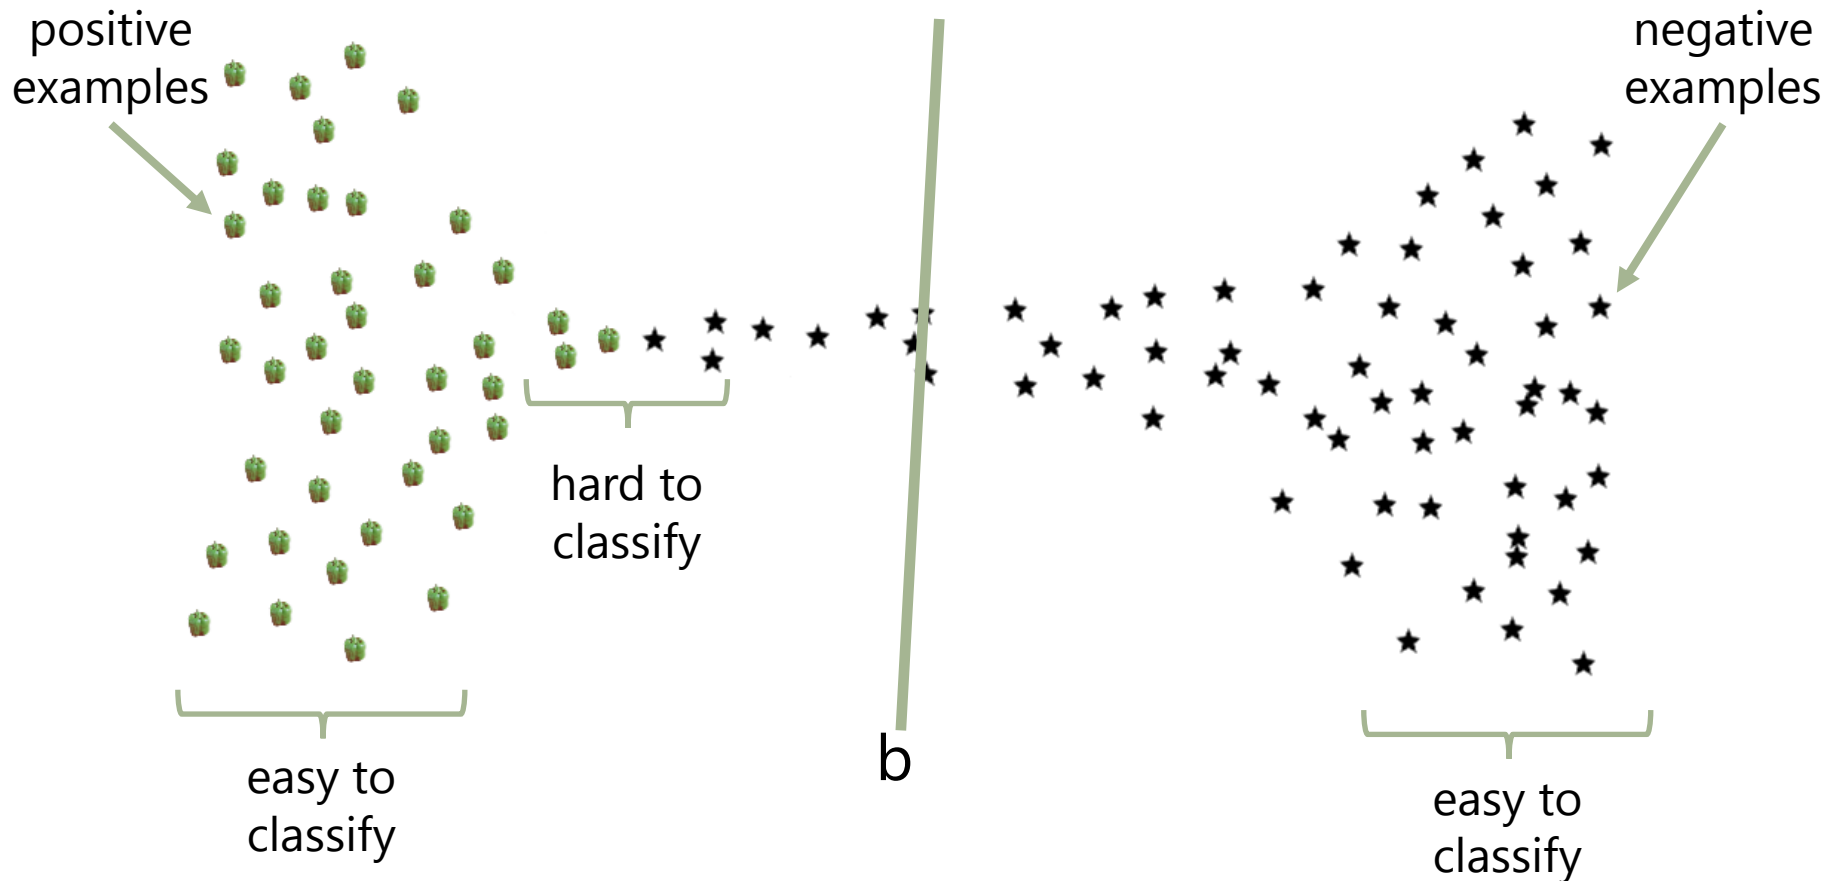
$X_i \cdot \theta$ should be maximized
when y_i is positive and
minimized when y_i is
negative

$$\arg \max_{\theta} \prod_i \delta(y_i = 1) p_{\theta}(y_i | X_i) + \delta(y_i = 0) (1 - p_{\theta}(y_i | X_i))$$

 $\delta(\text{arg}) = 1$ if the argument is true, = 0 otherwise

Logistic regression

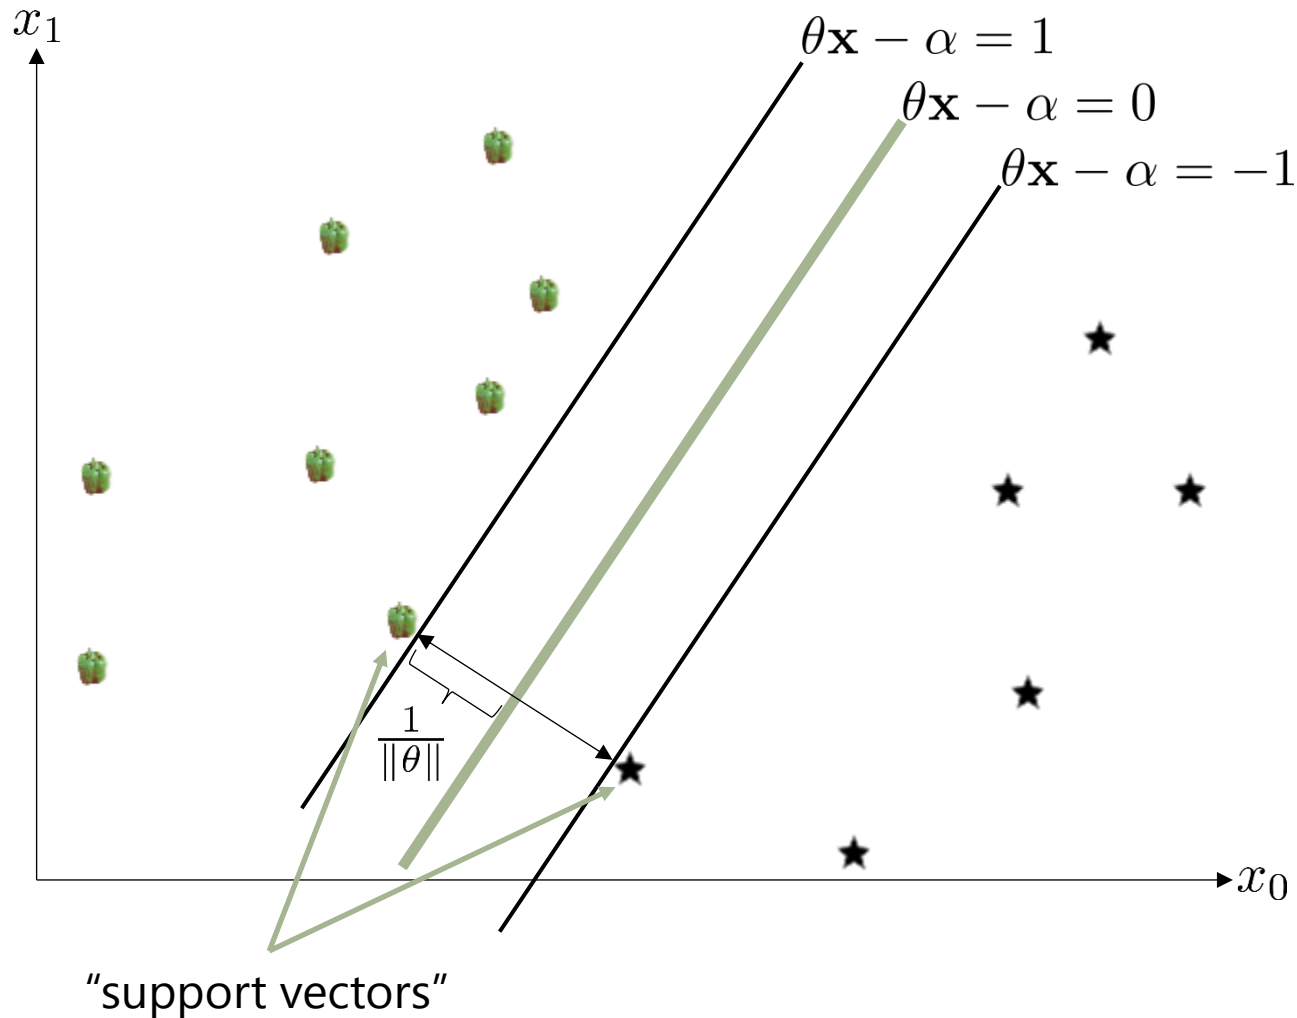
Q: Where would a logistic regressor place the decision boundary for these features?



Logistic regression

- Logistic regressors don't optimize the number of "mistakes"
- No special attention is paid to the "difficult" instances – every instance influences the model
- But "easy" instances can affect the model (and in a bad way!)
- How can we develop a classifier that optimizes the number of mislabeled examples?

Support Vector Machines



$$\arg \min_{\theta, \alpha} \frac{1}{2} \|\theta\|_2^2$$

such that

$$\forall_i y_i (\theta \cdot X_i - \alpha) \geq 1$$

Summary

The classifiers we've seen in Week 2 all attempt to make decisions by associating weights (θ) with features (x) and classifying according to

$$y_i = \begin{cases} 1 & \text{if } X_i \cdot \theta > 0 \\ 0 & \text{otherwise} \end{cases}$$

Summary

- **Naïve Bayes**

- Probabilistic model (fits $p(\text{label}|\text{data})$)
- Makes a conditional independence assumption of the form $(\text{feature}_i \perp\!\!\!\perp \text{feature}_j | \text{label})$ allowing us to define the model by computing $p(\text{feature}_i | \text{label})$ for each feature
- Simple to compute just by counting

- **Logistic Regression**

- Fixes the “double counting” problem present in naïve Bayes

- **SVMs**


- Non-probabilistic: optimizes the classification error rather than the likelihood

Which classifier is best?

1. When data are highly imbalanced

If there are far fewer positive examples than negative examples we may want to assign additional weight to negative instances (or vice versa)



e.g. will I purchase a product? If I purchase 0.00001% of products, then a classifier which just predicts "no" everywhere is 99.99999% accurate, but not very useful



Which classifier is best?

2. When mistakes are more costly in one direction

False positives are nuisances but false negatives are disastrous (or vice versa)

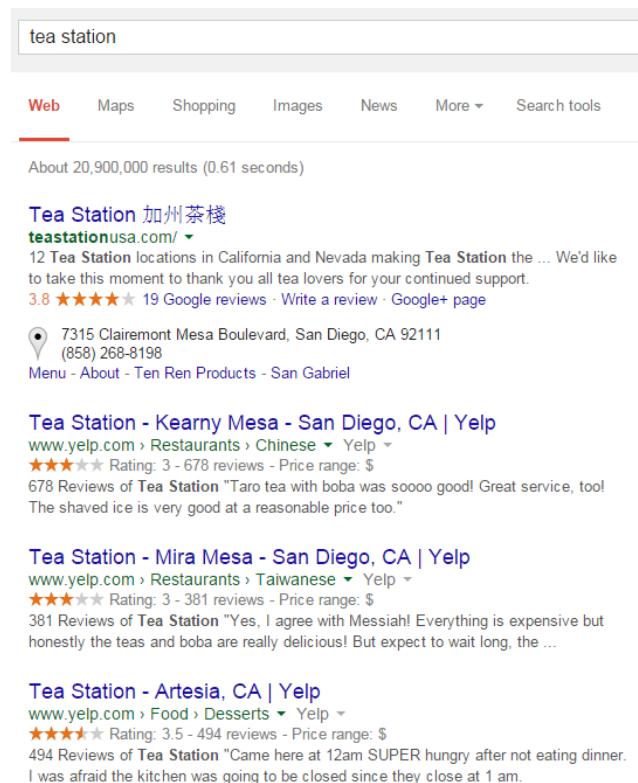


e.g. which of these bags contains a weapon?

Which classifier is best?

3. When we only care about the “most confident” predictions

e.g. does a relevant result appear among the first page of results?



tea station

Web Maps Shopping Images News More Search tools

About 20,900,000 results (0.61 seconds)

Tea Station 加州茶棧
teastationusa.com/ ▾
12 Tea Station locations in California and Nevada making Tea Station the ... We'd like to take this moment to thank you all tea lovers for your continued support.
3.8 ★★★★★ 19 Google reviews · Write a review · Google+ page


7315 Clairemont Mesa Boulevard, San Diego, CA 92111
(858) 268-8198
Menu · About · Ten Ren Products - San Gabriel

Tea Station - Kearny Mesa - San Diego, CA | Yelp
www.yelp.com › Restaurants › Chinese ▾ Yelp ▾
★★★★★ Rating: 3 - 678 reviews - Price range: \$
678 Reviews of Tea Station "Taro tea with boba was soooo good! Great service, too! The shaved ice is very good at a reasonable price too."


Tea Station - Mira Mesa - San Diego, CA | Yelp
www.yelp.com › Restaurants › Taiwanese ▾ Yelp ▾
★★★★★ Rating: 3 - 381 reviews - Price range: \$
381 Reviews of Tea Station "Yes, I agree with Messiah! Everything is expensive but honestly the teas and boba are really delicious! But expect to wait long, the ...

Tea Station - Artesia, CA | Yelp
www.yelp.com › Food › Desserts ▾ Yelp ▾
★★★★★ Rating: 3.5 - 494 reviews - Price range: \$
494 Reviews of Tea Station "Came here at 12am SUPER hungry after not eating dinner. I was afraid the kitchen was going to be closed since they close at 1 am.

Evaluating classifiers

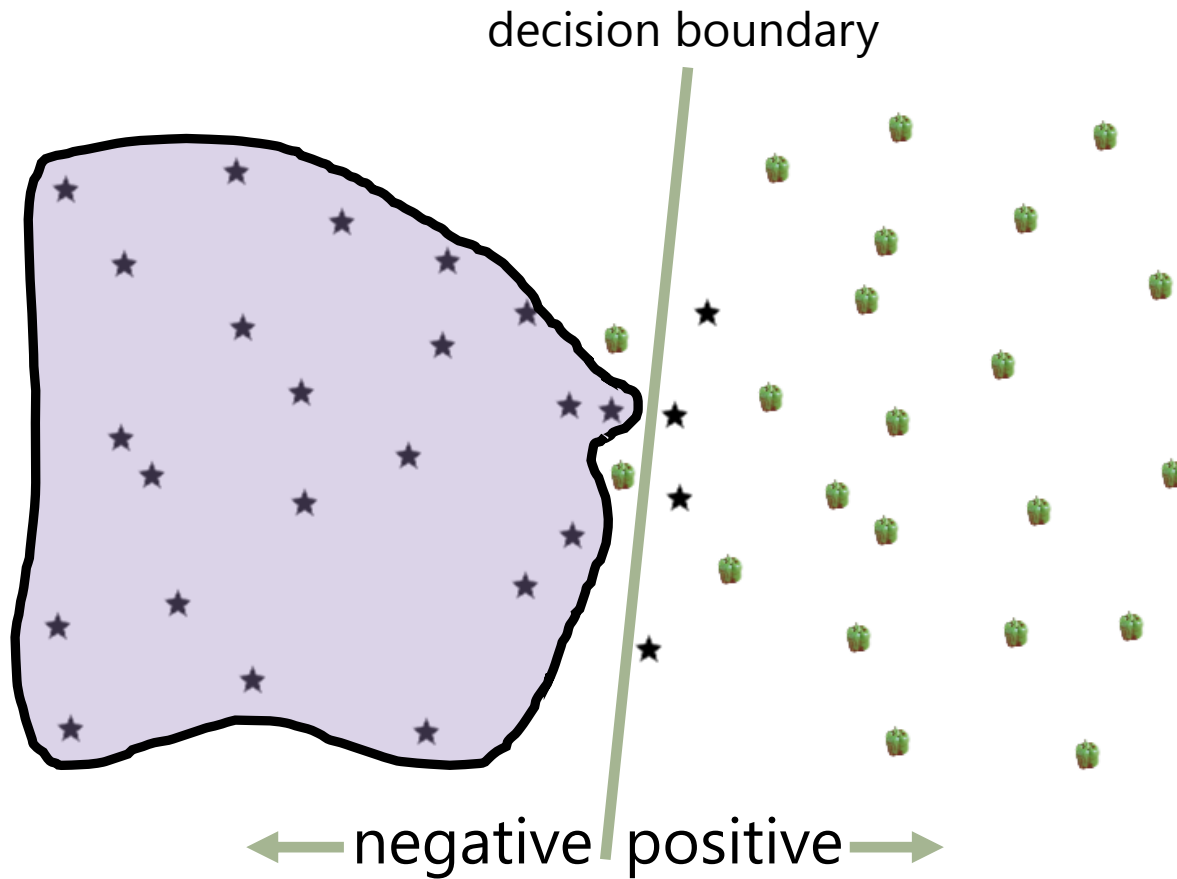


Evaluating classifiers



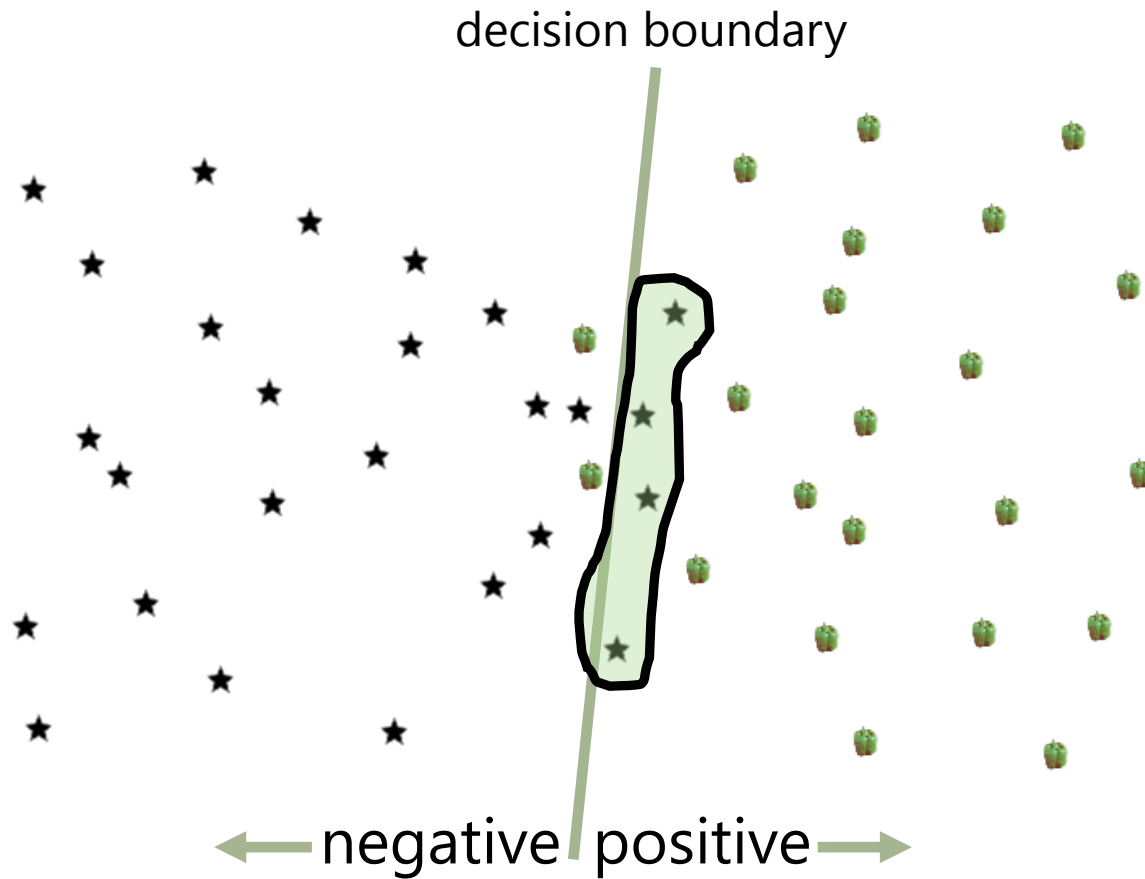
TP (true positive): Labeled as positive, predicted as positive

Evaluating classifiers




TN (true negative): Labeled as negative, predicted as negative

Evaluating classifiers



FP (false positive): Labeled as negative, predicted as positive

Evaluating classifiers



FN (false negative): Labeled as positive, predicted as negative

Evaluating classifiers

		Label	
		true	false
Prediction	true	true positive	false positive
	false	false negative	true negative

Classification accuracy = correct predictions / #predictions

$$= (TP + TN) / (TP + TN + FP + FN)$$

Error rate

= incorrect predictions / #predictions

$$= (FP + FN) / (TP + TN + FP + FN)$$

Week 2

- Linear classification – know what the different classifiers are and when you should use each of them. What are the advantages/disadvantages of each
- Know how to evaluate classifiers – what should you do when you care more about false positives than false negatives etc.

CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Week 3


Why dimensionality reduction?

Goal: take **high-dimensional** data, and describe it compactly using a small number of dimensions

Assumption: Data lies (approximately) on some **low-dimensional** manifold

(a few dimensions of opinions, a small number of topics, or a small number of communities)


Principal Component Analysis



Principal Component Analysis

Construct such vectors from 100,000 patches from real images and run PCA:

Color:




Principal Component Analysis

- We want to find a low-dimensional representation that best compresses or “summarizes” our data
- To do this we’d like to keep the dimensions with the highest variance (we proved this), and discard dimensions with lower variance. Essentially we’d like to capture the aspects of the data that are “hardest” to predict, while discard the parts that are “easy” to predict
- This can be done by taking the eigenvectors of the covariance matrix (we didn’t prove this, but it’s right there in the slides)

Clustering

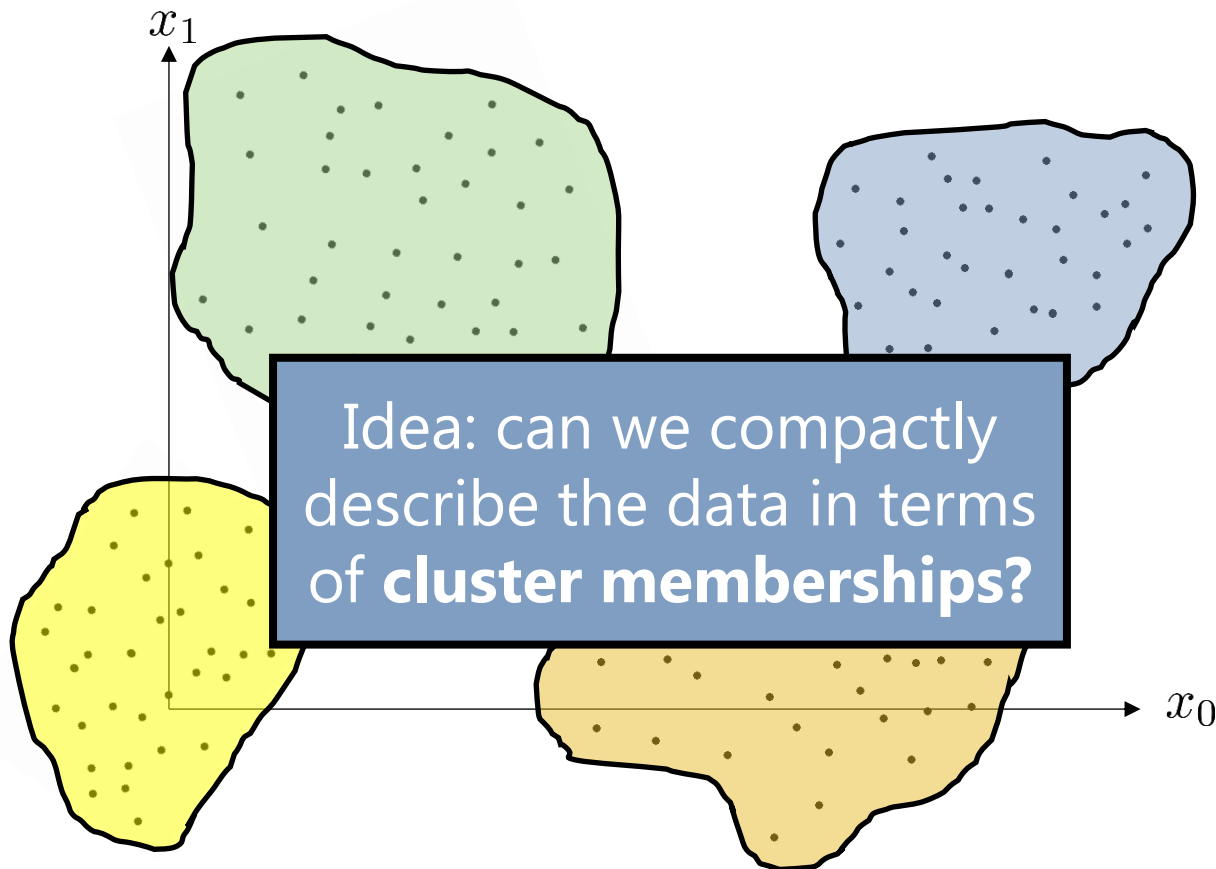
Q: What would PCA do with this data?

A: Not much, variance is about equal



Clustering

But: The data are highly **clustered**



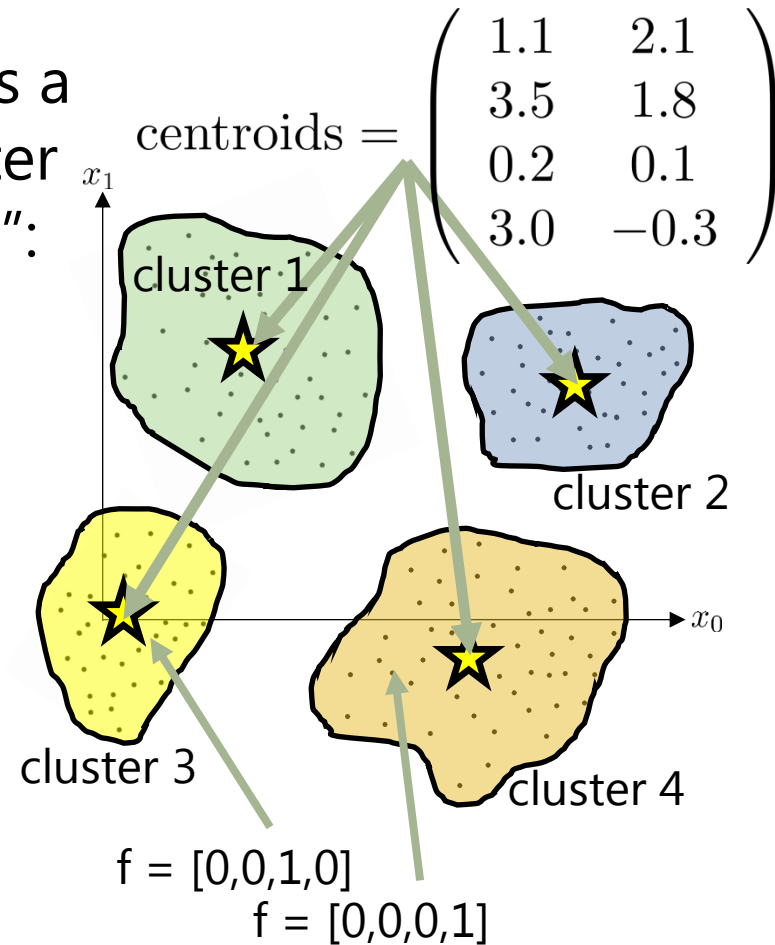
K-means Clustering

1. Input is still a matrix of features:

$$X = \begin{pmatrix} 5 & 3 & \dots & 1 \\ 4 & 2 & & 1 \\ 3 & 1 & & 3 \\ 2 & 2 & & 4 \\ 1 & 5 & & 2 \\ \vdots & & \ddots & \vdots \\ 1 & 2 & \dots & 1 \end{pmatrix}$$

3. From this we can describe each point in X by its cluster membership:

2. Output is a list of cluster "centroids":



$$Y = (1, 2, 4, 3, 4, 2, 4, 2, 2, 3, 3, 2, 1, 1, 3, \dots, 2)$$

K-means Clustering

Greedy algorithm:

1. Initialize C (e.g. at random)
2. Do
3. Assign each X_i to its nearest centroid.
4. Update each centroid to be the mean of points assigned to it
5. While (assignments change between iterations)


$$y_i = \arg \min_k \|X_i - C_k\|_2^2$$

$$C_k = \frac{\sum_i \delta(y_i=k) X_i}{\sum_i \delta(y_i=k)}$$

(also: reinitialize clusters at random should they become empty)

Hierarchical clustering

Q: What if our clusters are **hierarchical**?



A: We'd like a representation that encodes that points have **some features** in common but not others


Hierarchical clustering

Hierarchical (agglomerative) clustering works by gradually fusing clusters whose points are closest together

```
Assign every point to its own cluster:  
Clusters = [[1],[2],[3],[4],[5],[6],..., [N]]  
While len(Clusters) > 1:  
    Compute the center of each cluster  
    Combine the two clusters with the nearest centers
```

1. Connected components

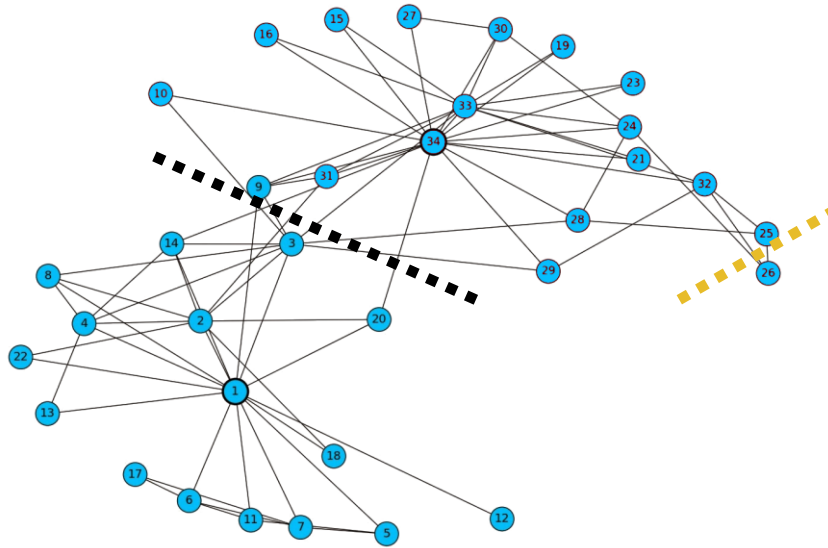
Define communities in terms of sets of nodes which are reachable from each other



- If a and b belong to a **strongly connected component** then there must be a path from $a \rightarrow b$ and a path from $b \rightarrow a$
- A **weakly connected component** is a set of nodes that **would be** strongly connected, if the graph were undirected

2. Graph cuts

What is the **Ratio Cut** cost of the following two cuts?



$$\text{Ratio Cut}(\text{---}) = \frac{1}{2} \left(\frac{3}{33} + \frac{3}{1} \right) = 1.54545$$

$$\text{Ratio Cut}(\text{- - -}) = \frac{1}{2} \left(\frac{9}{16} + \frac{9}{18} \right) = 0.53125$$

3. Clique percolation

- Clique percolation searches for “cliques” in the network of a certain size (K). Initially each of these cliques is considered to be its own community
- If two communities share a $(K-1)$ clique in common, they are merged into a single community
- This process repeats until no more communities can be merged

```
1. Given a clique size  $K$   
2. Initialize every  $K$ -clique as its own community  
3. While (two communities  $I$  and  $J$  have a  $(K-1)$ -clique in common):  
4.     Merge  $I$  and  $J$  into a single community
```

HW exercise: implement clique percolation on the FB ego network

Week 3

- Clustering & Community detection – understand the basics of the different algorithms
 - Given some features, know when to apply PCA vs. K-means vs. hierarchical clustering
 - Given some networks, know when to apply clique percolation vs. graph cuts vs. connected components

CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Midterm

1) Regression

Product Details

Genres	Science Fiction, Action, Horror
Director	David Twohy
Starring	Vin Diesel, Radha Mitchell
Supporting actors	Cole Hauser, Keith David, Lewis Fitz-Gerald, Claudia Black, Rhiana Gr Angela Moore, Peter Chiang, Ken Twoby
Studio	NBC Universal
MPAA rating	R (Restricted)
Captions and subtitles	English Details ▾
Rental rights	24 hour viewing period. Details ▾
Purchase rights	Stream instantly and download to 2 locations Details ▾
Format	Amazon Instant Video (streaming online video and digital download)


A. Phillips

Reviewer ranking: #17,230,554

90% helpful
votes received on reviews
(151 of 167)

ABOUT ME
Enjoy the reviews...

ACTIVITIES
Reviews (16)
Public Wish List (2)
Listmania Lists (2)
Tagged Items (1)



HipCzech


Aficionado
Male, from Texas

Profile Page

Member Since:	Jul 12, 2014	HipCzech was last seen:
Points:	175	Today at 12:19 AM
Beers:	108	
Places:	6	
Posts:	0	
Likes Received:	0	
Trading:	0% 0	

How can we use **features** such as product properties and user demographics to make predictions about **real-valued** outcomes (e.g. star ratings)?

How can we prevent our models from **overfitting** by favouring simpler models over more complex ones?



How can we assess our decision to optimize a particular error measure, like the MSE?

So far: Classification


Next we adapted these ideas to **binary** or **multiclass** outputs



What animal is in this image?




Will I **purchase** this product?



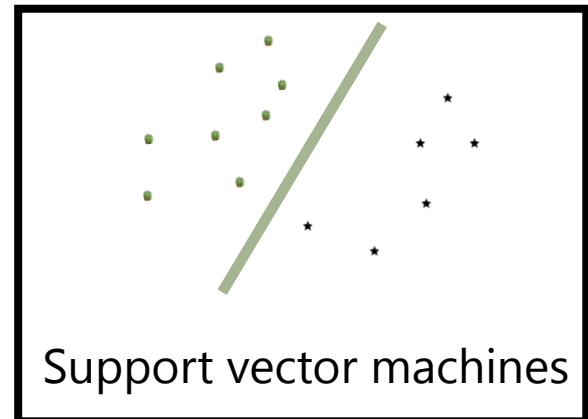
Will I **click on** this ad?



Combining features using naïve Bayes models


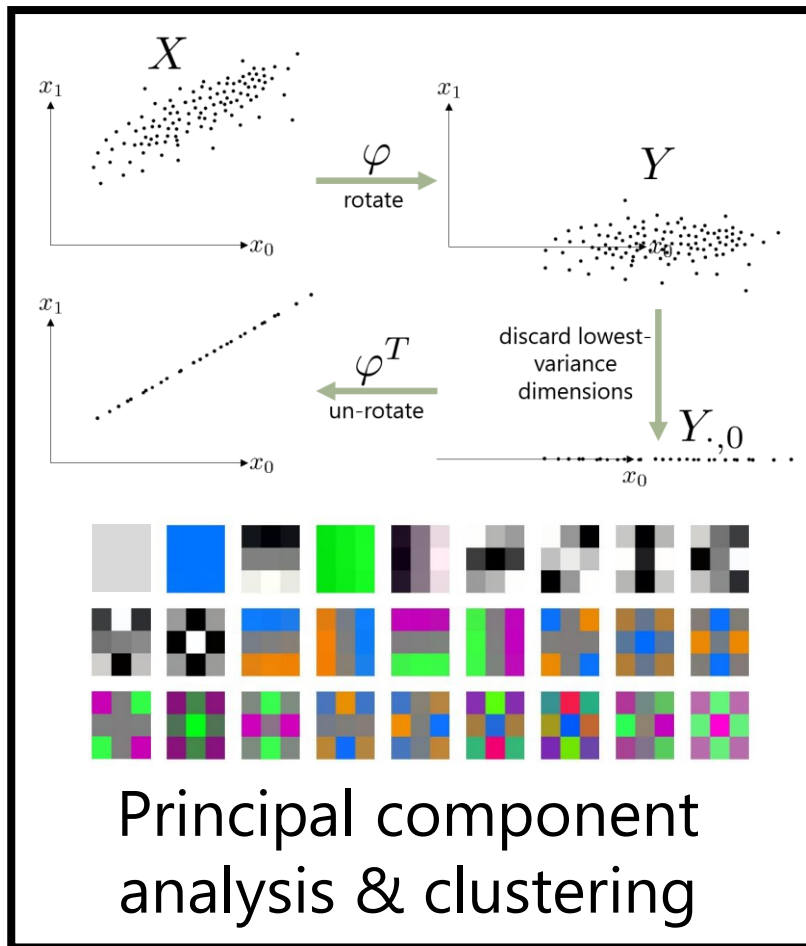


Logistic regression



Support vector machines

Week 3: Dimensionality Reduction



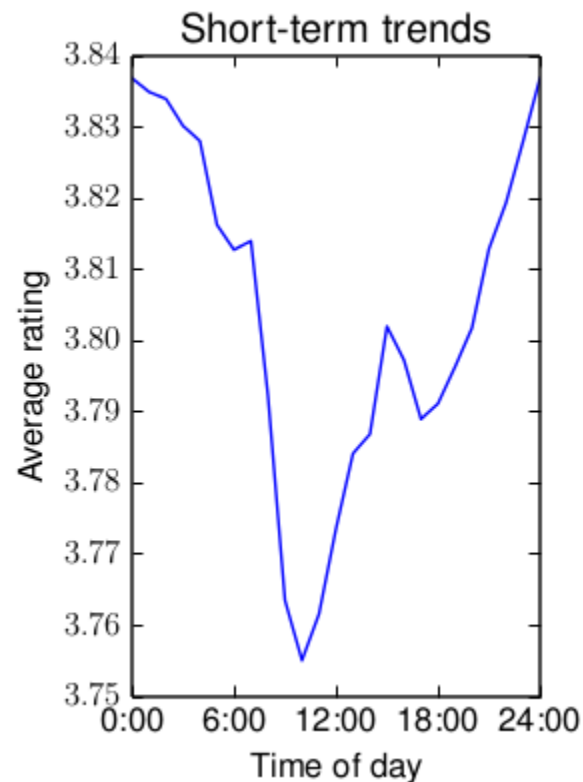
CSE 190 – Lecture 10

Data Mining and Predictive Analytics

Misc. questions

Representing the **day** as a feature

How would you build a feature to represent the **day**?



Homework 2 – “5-d mean”?

Is just the mean for each of the 5 dimensions...

K-means

```
centroids = [[0,0,0,0,1],[0,0,0,1,0]]
print centroids

assignment = defaultdict(set)
while (True):
    old_assignment = dict(assignment)
    assignment = defaultdict(set)
    for i in range(len(X)):
        minJ = None
        minD = None
        for j in range(len(centroids)):
            d = sum([(x-y)*(x-y) for x,y in zip(X[i], centroids[j])])
            if minJ == None or d < minD:
                minJ = j
                minD = d
        assignment[minJ].add(i)
    if dict(assignment) == old_assignment: break
    centroids = []
    for i in assignment.keys():
        centroids.append(sum([numpy.array(X[j]) for j in assignment[i]]) / len(assignment[i]))

distances = []
for k in assignment.keys():
    ck = centroids[k]
    distances += [sum([(x-y)*(x-y) for x,y in zip(X[i], ck)]) for i in assignment[k]]

print sum(distances)
```

Clique percolation

```
def cliquePerc(cliques3, cliques4):
    communities = dict([(str(c),c) for c in cliques4])
    while mergel(communities, cliques3):
        pass
    return communities

def mergel(communities, cliques3):
    coms = communities.keys()
    for x in range(len(coms)):
        for y in range(x+1, len(coms)):
            i = coms[x]
            j = coms[y]
            inter = set(communities[i]).intersection(set(communities[j]))
            for a in inter:
                for b in inter:
                    for c in inter:
                        if (a,b,c) in cliques3:
                            newcommunity = set(communities[i]).union(set(communities[j]))
                            print "merging", communities[i], "and", communities[j], "into", newcommunity
                            del communities[i]
                            del communities[j]
                            communities[i+'+'+j] = newcommunity
                            return True
    return False

print cliquePerc(cliques3, cliques4).values()
```

Interpretation of linear models

- Suppose we have a linear regression model to predict college GPA
 - One of the features of this model encodes whether a student owns a car
 - The fitted model looks like:

$$y = \dots - 0.4[\text{owns a car}] + \dots$$

Conclusion: "The GPA of the average student *who owns a car* is 0.4 lower than that of the average student"

Q: is this conclusion reasonable?

Interpretation of linear models

I tried rounding my ratings to integers but
it didn't help?

Q: Why not?