



Boosting



Recap

STAT/CSE 416: Machine Learning
Emily Fox
University of Washington
May 1, 2018

©2018 Emily Fox

Boosting question

"Can a set of weak learners be combined to create a stronger learner?" *Kearns and Valiant (1988)*



Yes! *Schapire (1990)*



Boosting



Amazing impact: • simple approach • widely used in industry • wins most Kaggle competitions

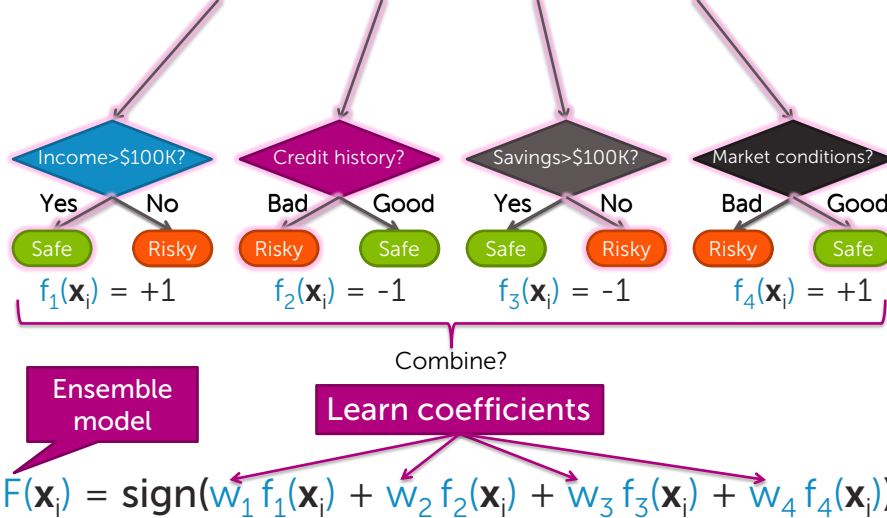
2

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Ensemble methods: Each classifier "votes" on prediction

$\mathbf{x}_i = (\text{Income}=\$120\text{K}, \text{Credit}=\text{Bad}, \text{Savings}=\$50\text{K}, \text{Market}=\text{Good})$

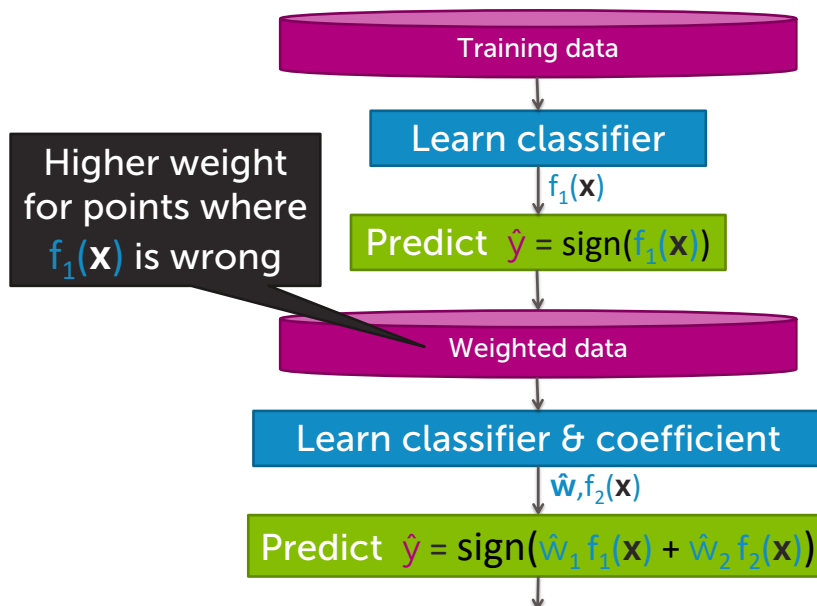


3

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Boosting = Greedy learning ensembles from data



4

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

AdaBoost: learning ensemble

[Freund & Schapire 1999]

- Start with same weight for all points: $\alpha_i = 1/N$
- For $t = 1, \dots, T$
 - Learn $f_t(\mathbf{x})$ with data weights α_i
 - Compute coefficient \hat{w}_t *Problem 1: How much do I trust f_t ?*
 - Recompute weights α_i *Problem 2: Weigh mistakes more*
- Final model predicts by:

$$\hat{y} = \text{sign} \left(\sum_{t=1}^T \hat{w}_t f_t(\mathbf{x}) \right)$$

5

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

AdaBoost: learning ensemble

- Start with same weight for all points: $\alpha_i = 1/N$
- For $t = 1, \dots, T$
 - - Learn $f_t(\mathbf{x})$ with data weights α_i
 - Compute coefficient \hat{w}_t
 - Recompute weights α_i
 - Normalize weights α_i
- Final model predicts by:

$$\hat{w}_t = \frac{1}{2} \ln \left(\frac{1 - \text{weighted_error}(f_t)}{\text{weighted_error}(f_t)} \right)$$

$$\alpha_i \leftarrow \begin{cases} \alpha_i e^{-\hat{w}_t}, & \text{if } f_t(\mathbf{x}_i) = y_i \\ \alpha_i e^{\hat{w}_t}, & \text{if } f_t(\mathbf{x}_i) \neq y_i \end{cases}$$

$$\hat{y} = \text{sign} \left(\sum_{t=1}^T \hat{w}_t f_t(\mathbf{x}) \right)$$

$$\alpha_i \leftarrow \frac{\alpha_i}{\sum_{j=1}^N \alpha_j}$$

6

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Learning from weighted data

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Learning from weighted data in general

Often, learning from weighted data treats data point i as α_i replicates of that data point


Credit	Income	y	Weight α
A	\$130K	Safe	0.5
B	\$80K	Risky	1.5
C	\$110K	Risky	1.2
A	\$110K	Safe	0.8
A	\$90K	Safe	0.6
B	\$120K	Safe	0.7
C	\$30K	Risky	3
C	\$60K	Risky	2
B	\$95K	Safe	0.8
A	\$60K	Safe	0.7
A	\$98K	Safe	0.9

8

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Gradient ascent for logistic regression on *unweighted data*



init $\mathbf{w}^{(1)} = 0$ (or randomly, or smartly), $t=1$

while $\|\nabla \ell(\mathbf{w}^{(t)})\| > \epsilon$ ← small threshold
 Difference between truth and prediction

for $j=0, \dots, D$

$$\text{partial}[j] = \sum_{i=1}^N h_j(\mathbf{x}_i) (\mathbb{1}[y_i = +1] - P(y = +1 | \mathbf{x}_i, \mathbf{w}^{(t)}))$$

$$\mathbf{w}_j^{(t+1)} \leftarrow \mathbf{w}_j^{(t)} + \eta \text{partial}[j]$$

$t \leftarrow t + 1$

$$\mathbf{w}_j^{(t+1)} \leftarrow \mathbf{w}_j^{(t)} + \eta \sum_{i=1}^N h_j(\mathbf{x}_i) (\mathbb{1}(y_i = +1) - P(y = +1 | \mathbf{x}_i, \mathbf{w}^{(t)}))$$

$$\mathbb{1}(y_i = +1) = \begin{cases} 1 & \text{if } y_i = +1 \\ 0 & \text{otherwise} \end{cases}$$

9

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Modify the logistic regression gradient update for *weighted data*

$$\mathbf{w}_j^{(t+1)} \leftarrow \mathbf{w}_j^{(t)} + \eta \sum_{i=1}^N \alpha_i h_j(\mathbf{x}_i) (\mathbb{1}[y_i = +1] - P(y = +1 | \mathbf{x}_i, \mathbf{w}^{(t)}))$$

Sum over data points

Weigh each point by α_i

Resulting \mathbf{w}^* is "best" logistic classifier (with given features) on data with weights α_i

10

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How to learn the “best” decision stump?

Credit	Income	y	Weight α
A	\$130K	Safe	0.5
B	\$80K	Risky	1.5
C	\$110K	Risky	1.2
A	\$110K	Safe	0.8
A	\$90K	Safe	0.6
B	\$120K	Safe	0.7
C	\$30K	Risky	3
C	\$60K	Risky	2
B	\$95K	Safe	0.8
A	\$60K	Safe	0.7
A	\$98K	Safe	0.9

Goal:

Choose best **feature** (categorical input) or **feature/threshold pair** (real-valued input)

Questions:

1. For given feature or feature/threshold, how to determine classifier output using weighted data?
2. For that classifier, how to compute its accuracy on the training data?
3. For real-valued features, how to select the best threshold on weighted data?
4. Based on the above, select the best decision stump

11

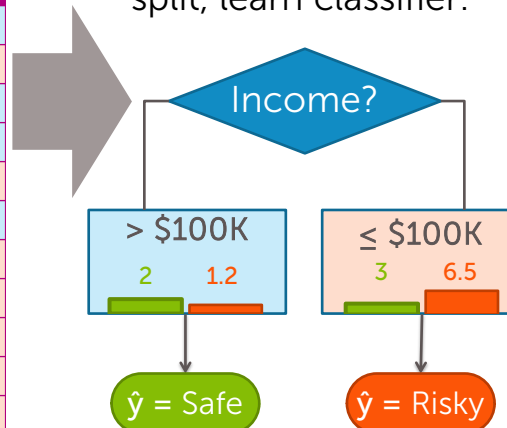
©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Learning a decision stump on weighted data

Credit	Income	y	Weight α
A	\$130K	Safe	0.5
B	\$80K	Risky	1.5
C	\$110K	Risky	1.2
A	\$110K	Safe	0.8
A	\$90K	Safe	0.6
B	\$120K	Safe	0.7
C	\$30K	Risky	3
C	\$60K	Risky	2
B	\$95K	Safe	0.8
A	\$60K	Safe	0.7
A	\$98K	Safe	0.9

For a given potential split, learn classifier:



12

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How to learn the “best” decision stump?

Credit	Income	y	Weight α
A	\$130K	Safe	0.5
B	\$80K	Risky	1.5
C	\$110K	Risky	1.2
A	\$110K	Safe	0.8
A	\$90K	Safe	0.6
B	\$120K	Safe	0.7
C	\$30K	Risky	3
C	\$60K	Risky	2
B	\$95K	Safe	0.8
A	\$60K	Safe	0.7
A	\$98K	Safe	0.9

Goal:

Choose best **feature** (categorical input) or **feature/threshold pair** (real-valued input)

Questions:

1. For given feature or feature/threshold, how to determine classifier output using weighted data?
2. For that classifier, how to compute its accuracy on the training data?
3. For real-valued features, how to select the best threshold on weighted data?
4. Based on the above, select the best decision stump

13


©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Learning a decision stump on weighted data

Credit	Income	y	Weight α
A	\$130K	Safe	0.5
B	\$80K	Risky	1.5
C	\$110K	Risky	1.2
A	\$110K	Safe	0.8
A	\$90K	Safe	0.6
B	\$120K	Safe	0.7
C	\$30K	Risky	3
C	\$60K	Risky	2
B	\$95K	Safe	0.8
A	\$60K	Safe	0.7
A	\$98K	Safe	0.9

Compute weighted error:



Weighted error = $\frac{\text{total weight of mistakes}}{\text{total weight}}$


$$= \frac{1.2 + 0.6 + 0.8 + 0.7 + 0.9}{12.7} = \frac{4.2}{12.7} = 0.33$$

14

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How to learn the “best” decision stump?



Goal:

Choose best **feature** (categorical input) or **feature/threshold pair** (real-valued input)

Questions:

1. For given feature or feature/threshold, how to determine classifier output using weighted data?
2. For that classifier, how to compute its accuracy on the training data?
3. For real-valued features, how to select the best threshold on weighted data?
4. Based on the above, select the best decision stump

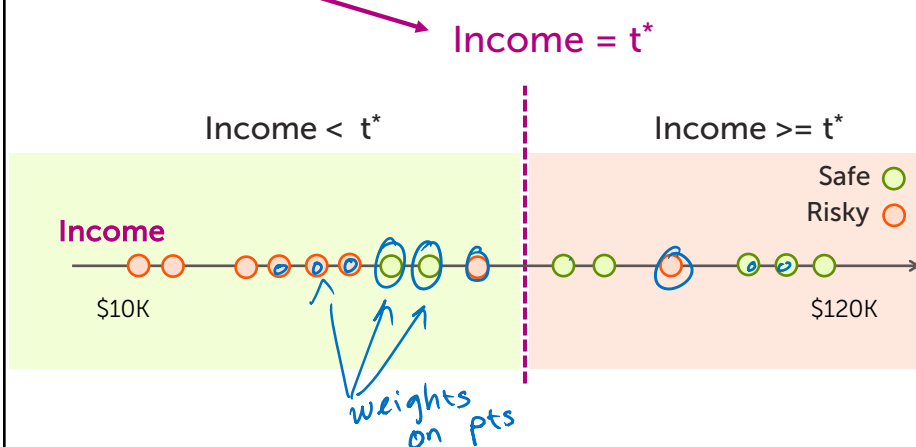
15

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Finding the best threshold split

Infinite possible values of t



16


©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Consider a threshold between points

Same **weighted error** for any threshold split between v_A and v_B

weighted error only changes when crossing over a (weighted) data pt



17


©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Only need to consider mid-points

Finite number of splits to consider

⇒ same alg. as before, but each considered threshold has a diff. weighted error.




18

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How to learn the “best” decision stump?



Goal:

Choose best **feature** (categorical input) or **feature/threshold pair** (real-valued input)

Questions:

1. For given feature or feature/threshold, how to determine classifier output using weighted data?
2. For that classifier, how to compute its accuracy on the training data?
3. For real-valued features, how to select the best threshold on weighted data?
4. Based on the above, select the best decision stump

19

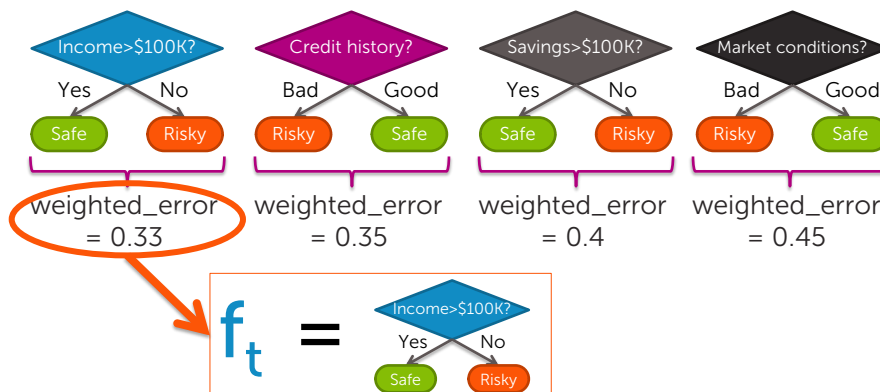
©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Finding best next decision stump $f_t(x)$

Which weak learner to choose?

Consider splitting on each feature:



20

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Revising AdaBoost example: A visualization

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Boosted decision stumps

- Start same weight for all points: $\alpha_i = 1/N$
- For $t = 1, \dots, T$
 - Learn $f_t(\mathbf{x})$: pick decision stump with lowest weighted training error according to α_i
 - Compute coefficient \hat{w}_t
 - Recompute weights α_i
 - Normalize weights α_i
- Final model predicts by:

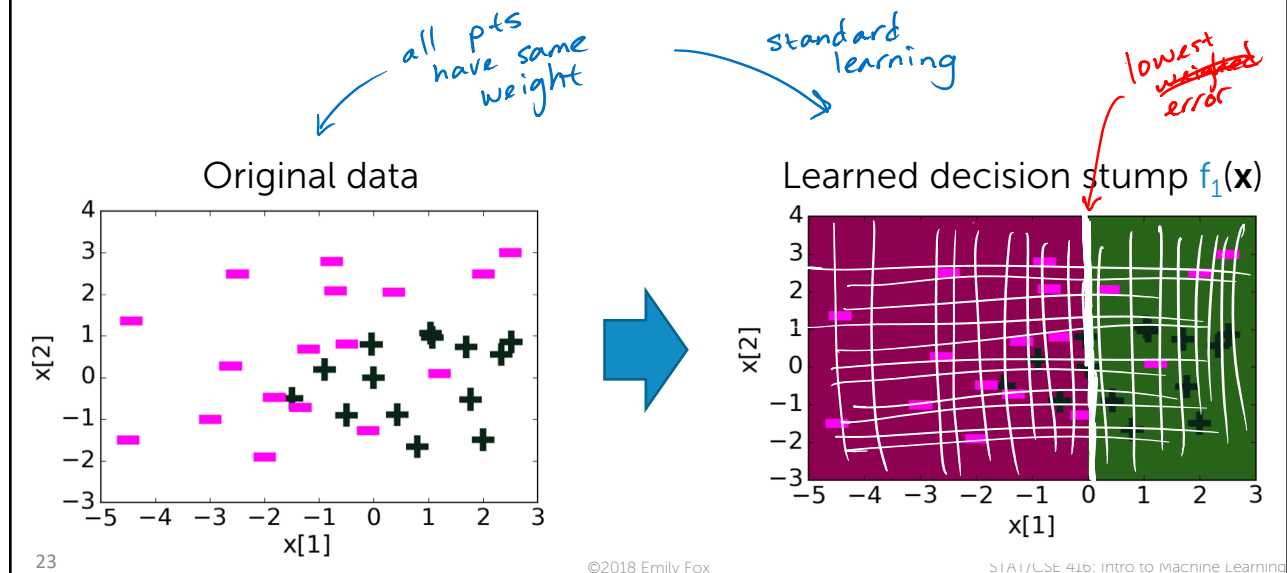
$$\hat{y} = \text{sign} \left(\sum_{t=1}^T \hat{w}_t f_t(\mathbf{x}) \right)$$

22

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

t=1: Just learn a classifier on original data

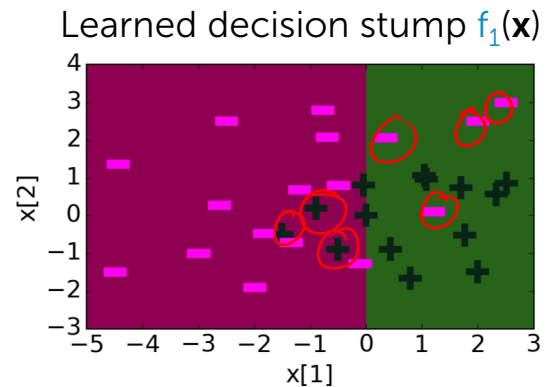


Determine weight (trust) of classifier


$$\hat{W}_t = \frac{1}{2} \ln \left(\frac{1 - \text{weighted_error}(f_t)}{\text{weighted_error}(f_t)} \right)$$

$$= \frac{1}{2} \ln \left(\frac{1 - \frac{7}{31}}{\frac{7}{31}} \right)$$


$$= 0.62$$




t=2: Learn classifier on weighted data



Weighted data: using α_1 chosen in previous iteration



Learned decision stump $f_2(x)$ on weighted data



©2018 Emily Fox

Ensemble becomes weighted sum of learned classifiers

Score(x) =


$$\hat{w}_1 \begin{bmatrix} f_1(x) \end{bmatrix} + \hat{w}_2 \begin{bmatrix} f_2(x) \end{bmatrix}$$

0.62 + 0.53

$$\hat{y} = \text{sign}(\underbrace{\hat{w}_1 f_1(x) + \hat{w}_2 f_2(x)}_{\text{score}(x)})$$

definite $\hat{y} = -1$

score space



uncertain


definite $\hat{y} = +1$

28

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Decision boundary of ensemble classifier after 30 iterations *(30 classifiers, $T=30$)*



*decision boundary
is crazy!*

*↓
probably
overfitting*

training_error = 0

29

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

AdaBoost: Revisiting convergence & overfitting

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

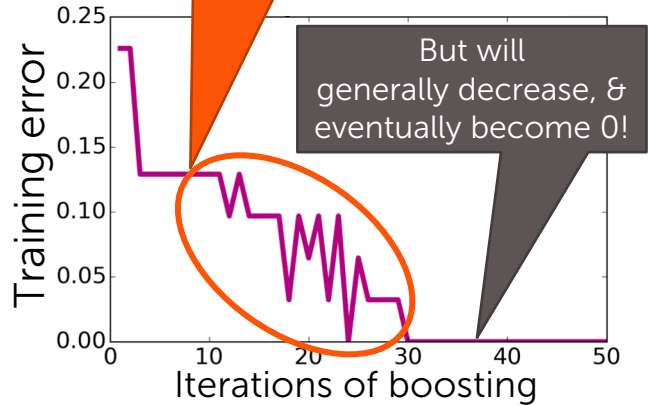
AdaBoost Theorem

Under some technical conditions...



Training error of boosted classifier $\rightarrow 0$ as $T \rightarrow \infty$

May oscillate a bit




31

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Decision trees on loan data



39% test error

Overfitting

8% training error

Boosted decision stumps on loan data



32% test error

Better fit & lower test error

28.5% training error

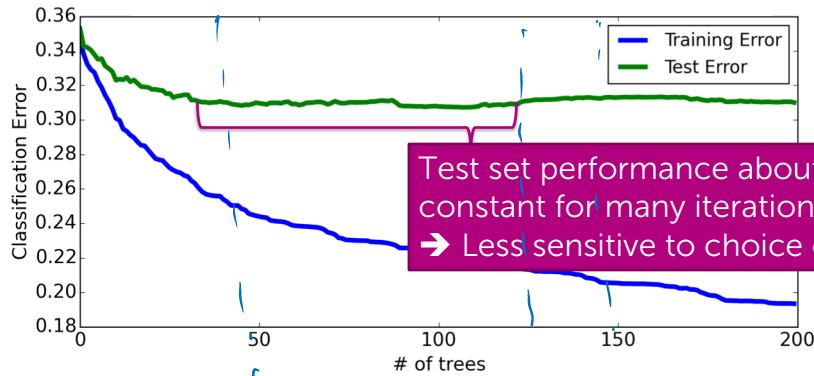
already, boosted decision stumps have better test error than best decision tree

32

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Boosting tends to be robust to overfitting




any of these values of T
would be fine

33

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

But boosting will eventually overfit,
so must choose max number of components T

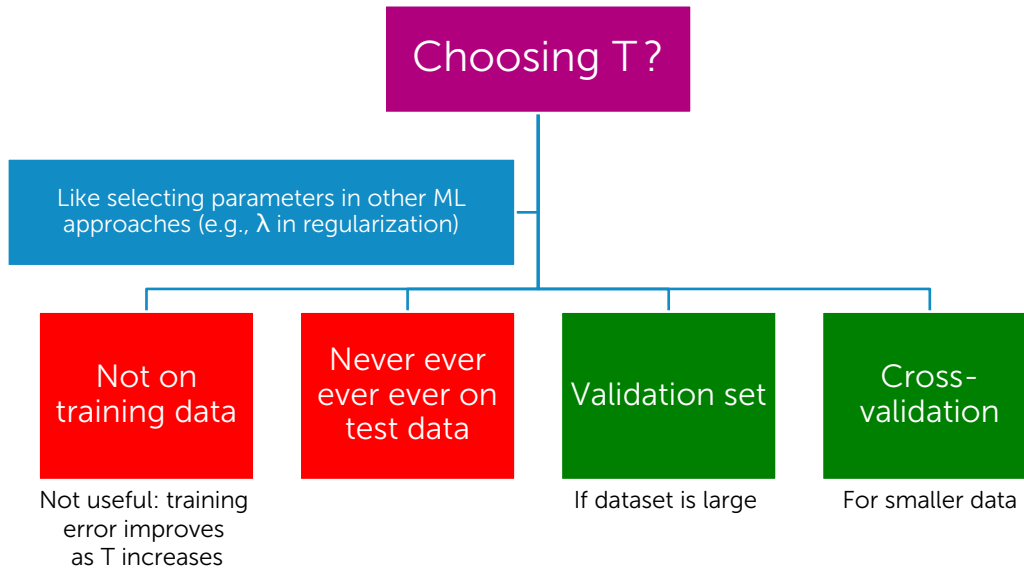


34

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How do we decide when to stop boosting?



35

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Summary of boosting

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Variants of boosting and related algorithms

There are hundreds of variants of boosting, most important:

Gradient boosting

- Like AdaBoost, but useful beyond basic classification

Many other approaches to learn ensembles, most important:

Random forests

- **Bagging**: Pick random subsets of the data
 - Learn a tree in each subset
 - Average predictions
- Simpler than boosting & easier to parallelize
- Typically higher error than boosting for same # of trees (# iterations T)

37

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Impact of boosting (*spoiler alert... HUGE IMPACT*)

Amongst most useful ML methods ever created

Extremely useful in computer vision

- Standard approach for face detection, for example (though deep learning is making serious inroads)

Used by **most winners** of ML competitions (Kaggle, KDD Cup,...)

- Malware classification, credit fraud detection, ads click through rate estimation, sales forecasting, ranking webpages for search, Higgs boson detection,...

Most deployed ML systems use model ensembles

- Coefficients chosen manually, with boosting, with bagging, or others

38

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What you can do now...

- Identify notion ensemble classifiers
- Formalize ensembles as weighted combination of simpler classifiers
- Outline the boosting framework – sequentially learn classifiers on weighted data
- Describe the AdaBoost algorithm
 - Learn each classifier on weighted data
 - Compute coefficient of classifier
 - Recompute data weights
 - Normalize weights
- Implement AdaBoost to create an ensemble of decision stumps

39

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning



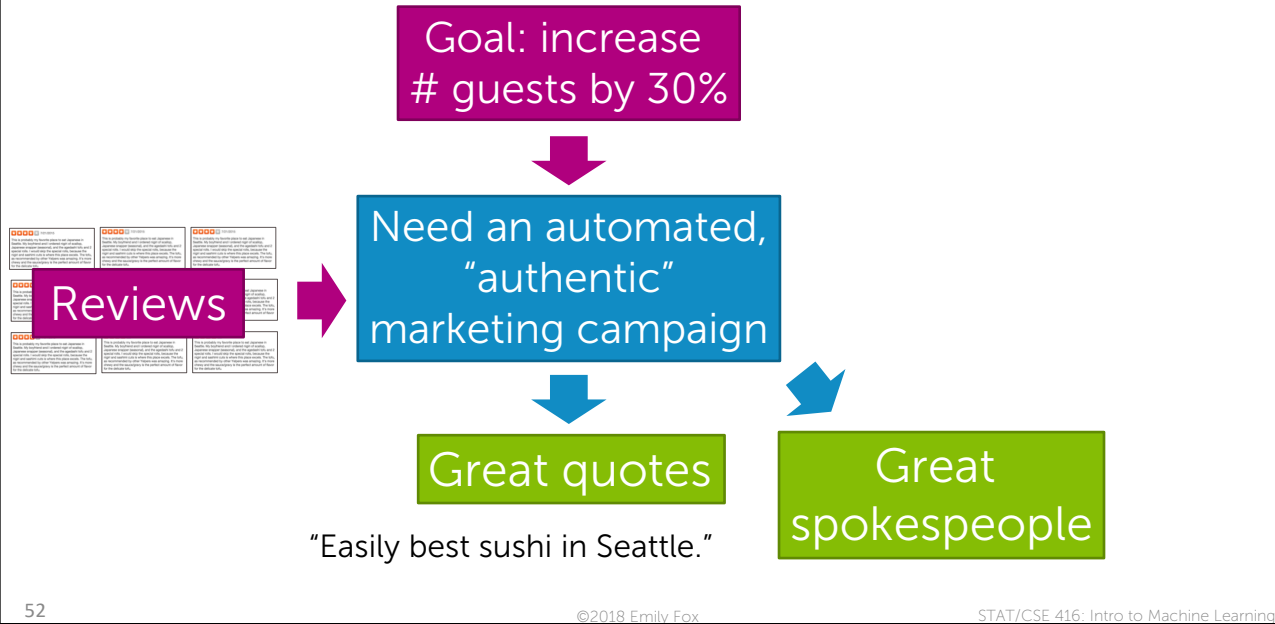
Evaluating classifiers:

Precision & Recall

STAT/CSE 416: Machine Learning
Emily Fox
University of Washington
May 1, 2018

©2018 Emily Fox

Using reviews to promote my restaurant



How do I find sentences with positive sentiment?


All reviews for my restaurant

★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.
★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.
★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.
★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.
★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.	★★★★ reviews This is probably the best place to eat Japanese food in the area. The food is delicious and the service is excellent. The atmosphere is also very nice. I highly recommend this place to anyone who is looking for a good meal in the area.



Intelligent restaurant review system

All reviews
for restaurant



Break all reviews
into sentences

The seaweed salad was just OK,
vegetable salad was just ordinary.

I like the interior decoration and
the blackboard menu on the wall.

All the sushi was delicious.

My wife tried their ramen and
it was pretty forgettable.

The sushi was amazing, and
the rice is just outstanding.

The service is somewhat hectic.

Easily best sushi in Seattle.

54

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Sentiment classifier

Input x_i : Easily best sushi in Seattle.



Sentence Sentiment
Classifier

Output: \hat{y}_i
Predicted
sentiment



Easily best sushi in Seattle.



55

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Use the sentiment classifier model!

Sentences from all reviews for my restaurant

The seaweed salad was just OK, vegetable salad was just ordinary.

I like the interior decoration and the blackboard menu on the wall.

All the sushi was delicious.

My wife tried their ramen and it was pretty forgettable.

The sushi was amazing, and the rice is just outstanding.

The service is somewhat hectic.

Easily best sushi in Seattle.



Classifier MODEL

Sentences predicted to be positive

$$\hat{y} = +1$$

Easily best sushi in Seattle.

I like the interior decoration and the blackboard menu on the wall.

All the sushi was delicious.

The sushi was amazing, and the rice is just outstanding.

Sentences predicted to be negative

$$\hat{y} = -1$$

The seaweed salad was just OK, vegetable salad was just ordinary.

My wife tried their ramen and it was pretty forgettable.

The service is somewhat hectic.

Show sentences with +1 prediction on website

56

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What does it mean for a classifier to be good?

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Previously, we asked the question:
“What is good accuracy?”

58

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

We explored accuracy of random classifier as baseline

- For binary classification:
 - Half the time, you’ll get it right! (on average)
 - classification error = 0.5
- For k classes, error = $1 - 1/k$
 - error = 0.666 for 3 classes, 0.75 for 4 classes,...

At the very, very, very least,
you should healthily beat random...
Otherwise, it’s (usually) pointless...

59

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

We explored the pitfalls of imbalanced problems:
Is 90% accuracy good? Depends ...

90% of sentences are negative!

90% accuracy by predicting every sentence is negative!!!

Amazing "performance" but not useful for me right now!

60

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Automated marketing campaign cares about something else...

Website shows 10 sentences from recent reviews

PRECISION

Did I (mistakenly) show a negative sentence???

RECALL

Did I not show a (great) positive sentence???

Accuracy doesn't capture these issues well...

61

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Precision:
 Fraction of positive predictions that are
actually positive

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What fraction of the positive predictions are correct?

Sentences predicted
 to be positive: $\hat{y}_i = +1$

Easily best sushi in Seattle.	✓
The seaweed salad was just OK, vegetable salad was just ordinary.	✗
I like the interior decoration and the blackboard menu on the wall.	✓
The service is somewhat hectic.	✗
The sushi was amazing, and the rice is just outstanding.	✓
All the sushi was delicious.	✓

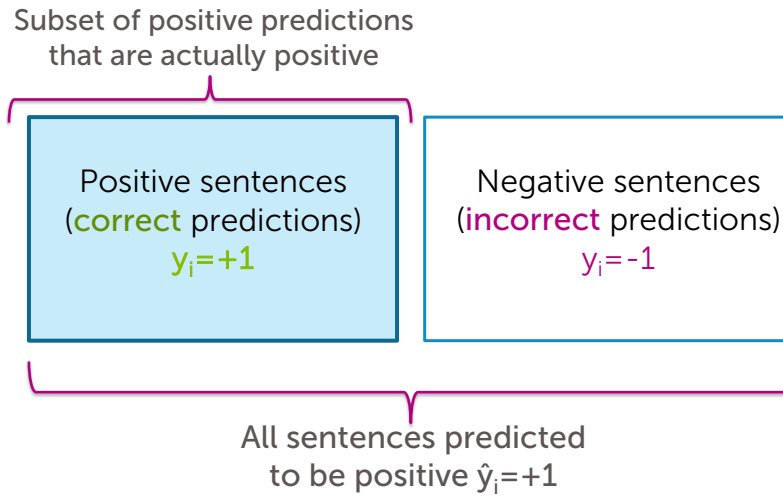
Only 4 out of 6
 sentences
 predicted to be
 positive are
 actually positive

63

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Precision: Fraction of positive predictions that are actually positive







64

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Types of error: *Review*





		Predicted label	
		 $\hat{y}_i = +1$	 $\hat{y}_i = -1$
True label	 $y_i = +1$	True Positive	False Negative
	 $y_i = -1$	False Positive	True Negative

65

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Confusion matrix for sentiment analysis

		Predicted sentiment	
		 $\hat{y}_i=+1$	 $\hat{y}_i=-1$
True sentiment	 $y_i=+1$	+1 sentence +1 prediction	+1 sentence -1 pred. <i>missed a sentence</i>
	 $y_i=-1$	-1 sentence +1 pred. <i>showed bad review on website</i>	-1 sentence -1 pred.

66

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning





Precision - Formula

Fraction of positive predictions that are correct

$$\text{precision} = \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false positives}}$$

- Best possible value : 1.0
- Worst possible value : 0.0

Focus on positive predictions

		Predicted label	
		 $\hat{y}_i=+1$	 $\hat{y}_i=-1$
True label	 $y_i=+1$	True Positive	False Negative
	 $y_i=-1$	False Positive	True Negative

67

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Example: Calculating precision

Sentences predicted to be positive: $\hat{y}_i = +1$

Easily best sushi in Seattle.	✓
The seaweed salad was just OK, vegetable salad was just ordinary.	✗
I like the interior decoration and the blackboard menu on the wall.	✓
The service is somewhat hectic.	✗
The sushi was amazing, and the rice is just outstanding.	✓
All the sushi was delicious.	✓

$$\text{precision} = \frac{4}{6} = \frac{2}{3}$$

4 correct

2 mistakes

68

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Why precision is important

Shown on website

Sentences predicted to be positive: $\hat{y}_i = +1$

Easily best sushi in Seattle.	✓
The seaweed salad was just OK, vegetable salad was just ordinary.	✗
I like the interior decoration and the blackboard menu on the wall.	✓
The service is somewhat hectic.	✗
The sushi was amazing, and the rice is just outstanding.	✓
All the sushi was delicious.	✓

2 negative sentences shown to potential customers... ☹️

High precision means positive predictions actually likely to be positive!

69

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Recall:
 Fraction of positive data
predicted to be positive

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Did I find all the
 positive sentences?


Sentences from
 all reviews
 for my restaurant

Classifier
 MODEL


Predicted positive $\hat{y}_i = +1$

Easily best sushi in Seattle. 

The seaweed salad was just OK,
 vegetable salad was just ordinary.

I like the interior decoration and
 the blackboard menu on the wall. 

The service is somewhat hectic.

The sushi was amazing, and
 the rice is just outstanding. 

All the sushi was delicious. 


Predicted negative $\hat{y}_i = -1$

The seaweed salad was just OK,
 vegetable salad was just ordinary.

My wife tried their ramen and
 it was delicious. 

The service is somewhat hectic.

My wife tried their ramen and
 it was pretty forgettable.

The service was perfect. 

 True positive
 sentences: $y_i = +1$

71

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What fraction of positive sentences were missed out?

Predicted positive $\hat{y}_i = +1$

- Easily best sushi in Seattle.
- The seaweed salad was just OK, vegetable salad was just ordinary.
- I like the interior decoration and the blackboard menu on the wall.
- The service is somewhat hectic.
- The sushi was amazing, and the rice is just outstanding.
- All the sushi was delicious.

Predicted negative $\hat{y}_i = -1$

- The seaweed salad was just OK, vegetable salad was just ordinary.
- My wife tried their ramen and it was delicious.
- The service is somewhat hectic.
- My wife tried their ramen and it was pretty forgettable.
- The service was perfect.

Found 4 positive sentences

Model could not find 2 sentences that were actually positive

Missed 2 positive sentences

72
©2018 Emily Fox
STAT/CSE 416: Intro to Machine Learning

Recall: Fraction of positive data predicted to be positive

All positive data points $y = +1$

Sentences predicted to be positive
(**correct** predictions)
 $\hat{y}_i = +1$

Sentences predicted to be negative
(**incorrect** predictions)
 $\hat{y}_i = -1$

Subset of positive data points correctly identified

73
©2018 Emily Fox
STAT/CSE 416: Intro to Machine Learning

Recall - Formula

Fraction of positive data points correctly classified

$$\text{Recall} = \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false negatives}}$$

- Best possible value : 1.0
- Worst possible value : 0.0

focus on positive examples

		Predicted label	
		$\hat{y}_i = +1$	$\hat{y}_i = -1$
True label	$y_i = +1$	True Positive	False Negative
	$y_i = -1$	False Positive	True Negative

74

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Why is recall important?

Predicted positive $\hat{y}_i = +1$

- Easily best sushi in Seattle.
- The seaweed salad was just OK, vegetable salad was just ordinary.
- I like the interior decoration and the blackboard menu on the wall.
- The service is somewhat hectic.
- The sushi was amazing, and the rice is just outstanding.
- All the sushi was delicious.

Predicted negative $\hat{y}_i = -1$

- The seaweed salad was just OK, vegetable salad was just ordinary.
- My wife tried their ramen and it was delicious.
- The service is somewhat hectic.
- My wife tried their ramen and it was pretty forgettable.
- The service was perfect.

Want to show positive sentences on website

2 positive sentences not shown to potential customers... 😞

High recall means positive data points are very likely to be discovered!

75

©2018 Emily Fox

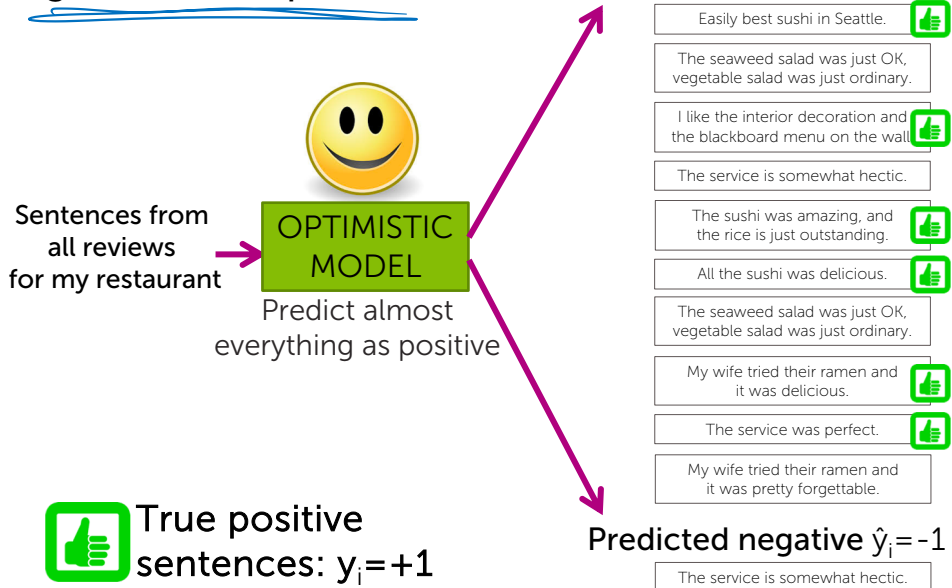
STAT/CSE 416: Intro to Machine Learning

Precision-recall extremes

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Optimistic model: High recall, low precision



77

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Pessimistic model: High precision, low recall

Sentences from all reviews for my restaurant → **PESSIMISTIC MODEL**
Predict positive only when **very** sure

Predicted positive $\hat{y}_i = +1$

- Easily best sushi in Seattle.
- The sushi was amazing, and the rice is just outstanding.

Predicted negative $\hat{y}_i = -1$

- I like the interior decoration and the blackboard menu on the wall.
- The service is somewhat hectic.
- The seaweed salad was just OK, vegetable salad was just ordinary.
- All the sushi was delicious.
- The seaweed salad was just OK, vegetable salad was just ordinary.
- My wife tried their ramen and it was delicious.
- The service was perfect.
- My wife tried their ramen and it was pretty forgettable.
- The service is somewhat hectic.

True positive sentences: $y_i = +1$

78 ©2018 Emily Fox STAT/CSE 416: Intro to Machine Learning

Balancing precision & recall

PESSIMISTIC MODEL (sad face icon)
Finds few positive sentences, but includes no false positives

Want to find many positive sentences, but minimize risk of incorrect predictions!!

OPTIMISTIC MODEL (happy face icon)
Finds all positive sentences, but includes many false positives

79 ©2018 Emily Fox STAT/CSE 416: Intro to Machine Learning

Tradeoff precision and recall

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Can we tradeoff precision & recall?

Low precision,
high recall

High precision,
low recall

Optimistic Model
Predict almost
everything as positive



Pessimistic Model
Predict positive only
when **very** sure

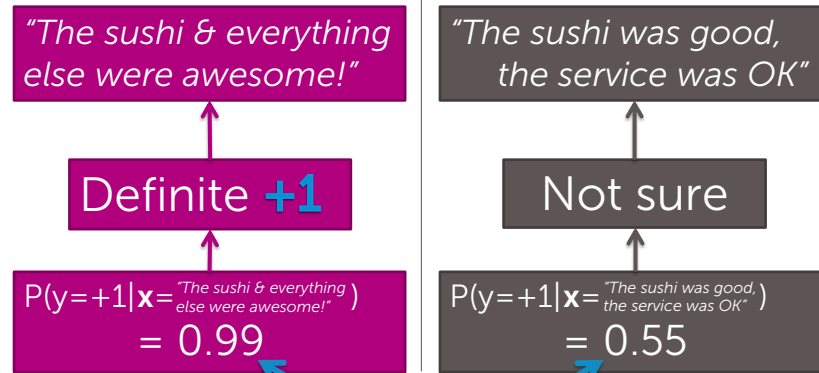


81

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

How confident is your prediction?



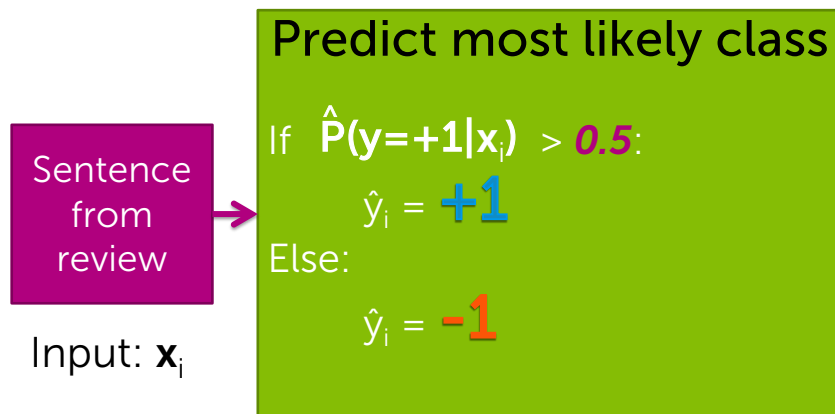
Can be used to tradeoff precision and recall

82

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Basic classifier

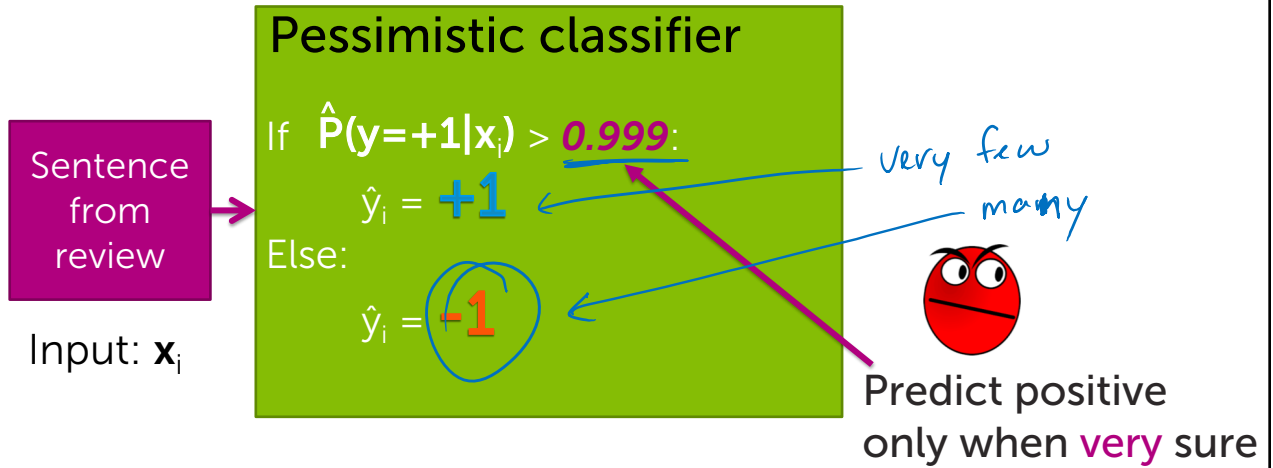


83

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Pessimistic: High precision, low recall

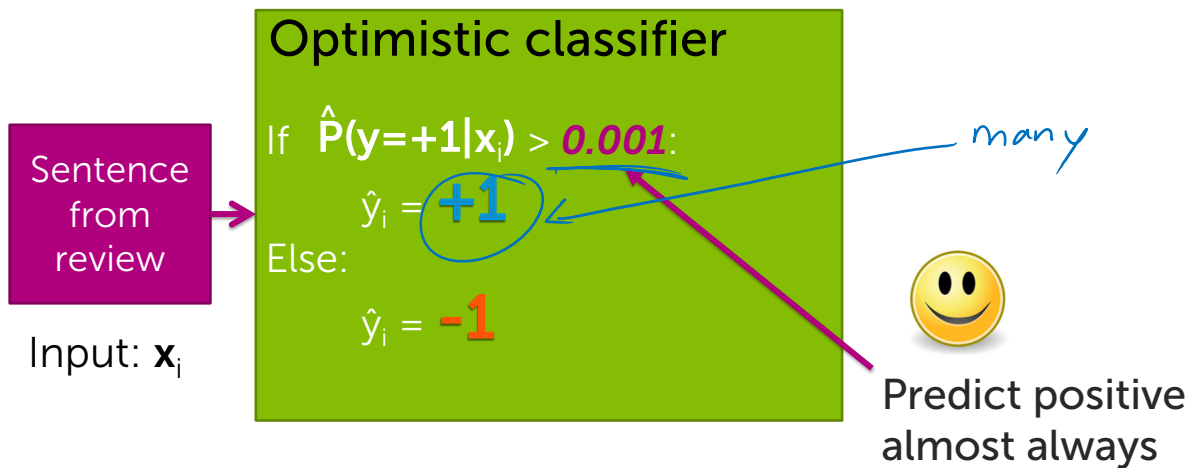


84

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Optimistic: Low precision, high recall




85

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Prediction probability threshold

Probability t above which model predicts true

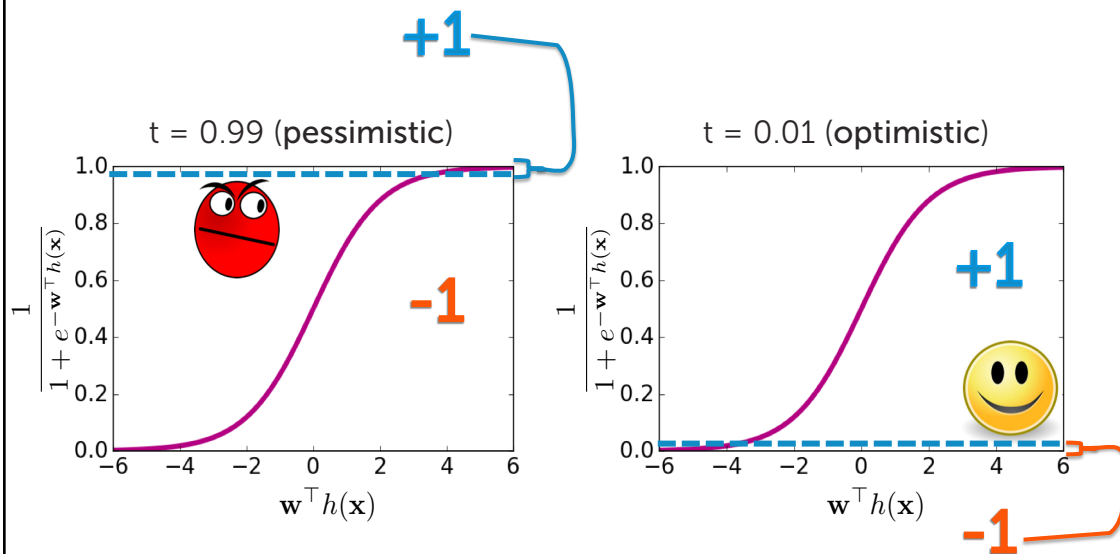


86

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Example threshold values

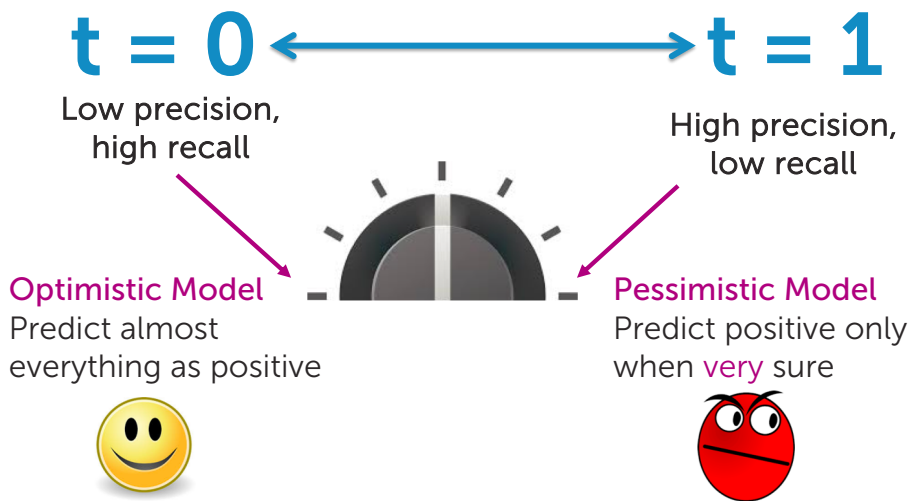


87

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Tradeoff precision & recall with threshold



88

©2018 Emily Fox


STAT/CSE 416: Intro to Machine Learning

Precision-recall curve

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

The precision-recall curve




90

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What does the perfect algorithm look like?




91

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Which classifier is better? A or B?

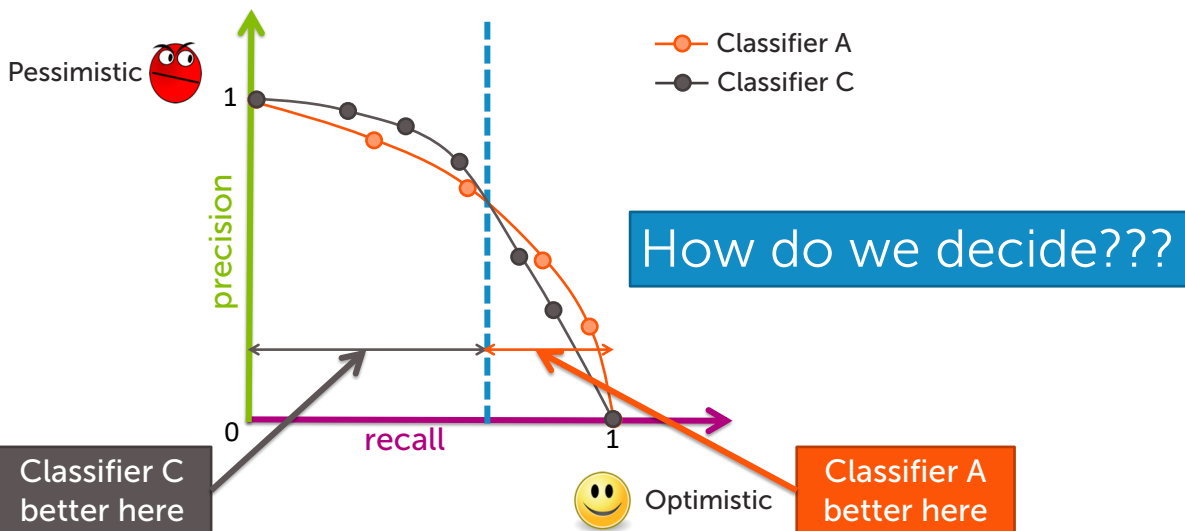


92

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Which classifier is better? A or C?



93

©2018 Emily Fox

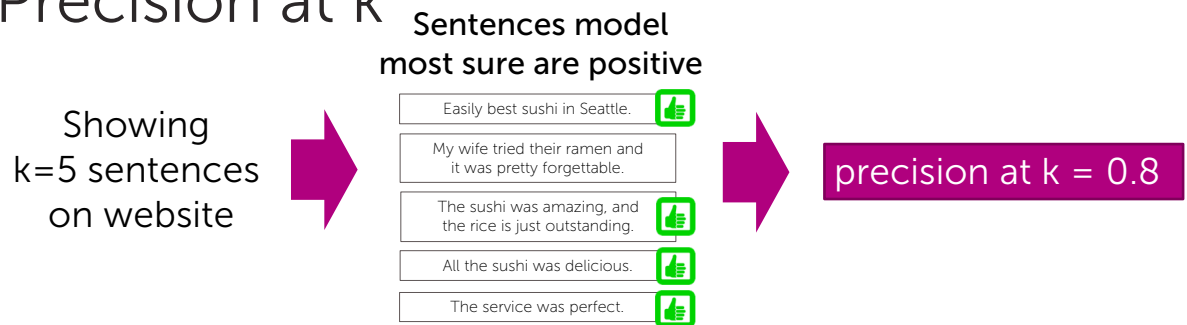
STAT/CSE 416: Intro to Machine Learning

Compare algorithms

Often, reduce precision-recall to single number to compare algorithms

- F1 measure, area-under-the-curve (AUC),...

Precision at k



94

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

Summary of precision-recall

©2018 Emily Fox

STAT/CSE 416: Intro to Machine Learning

What you can do now...

- Classification accuracy/error are not always right metrics
- **Precision** captures fraction of positive predictions that are correct
- **Recall** captures fraction of positive data correctly identified by the model
- Trade-off **precision** & **recall** by setting probability thresholds
- Plot **precision-recall** curves.
- Compare models by computing precision at **k**