INTRODUCTION TO DATA SCIENCE

This lecture is
based on course by E. Fox and C. Guestrin, Univ of Washington

21/11/2017

WFAiS UJ, Informatyka Stosowana
II stopień studiów

Evaluating Classifiers

Using reviews to promote my restaurant

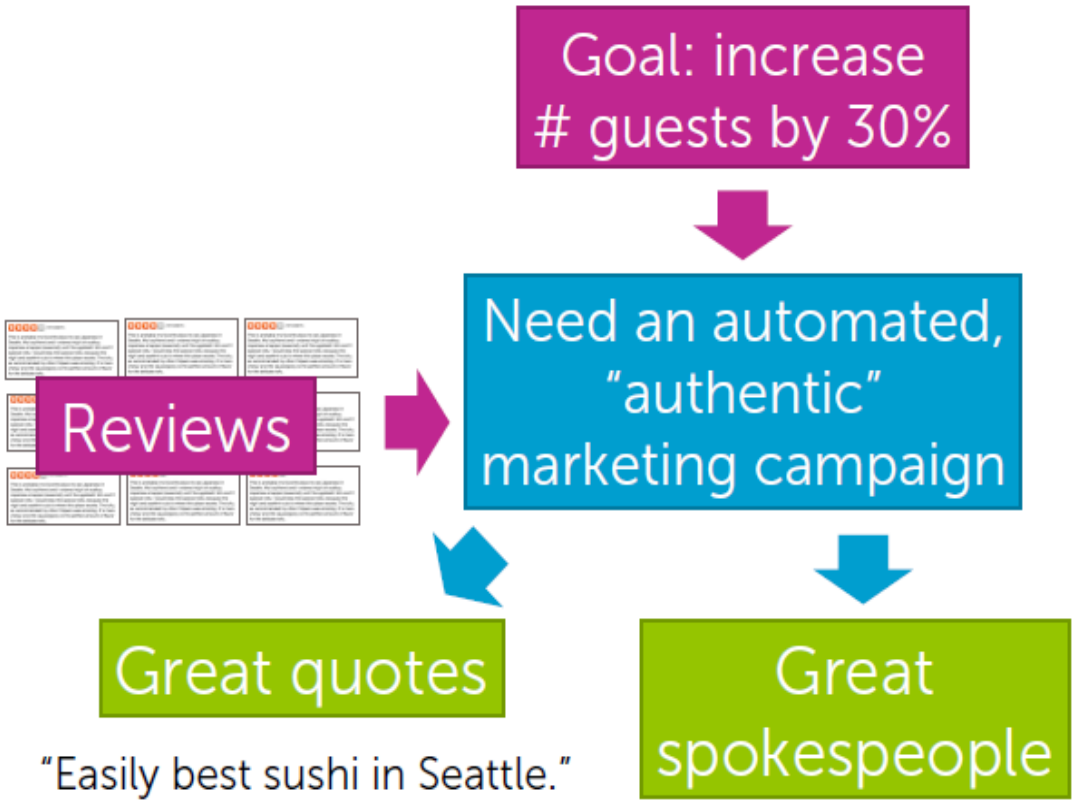

Classifiers

How do I find sentences with positive sentiment?

All reviews
for my restaurant

Classifiers

Use the sentiment classifier model!

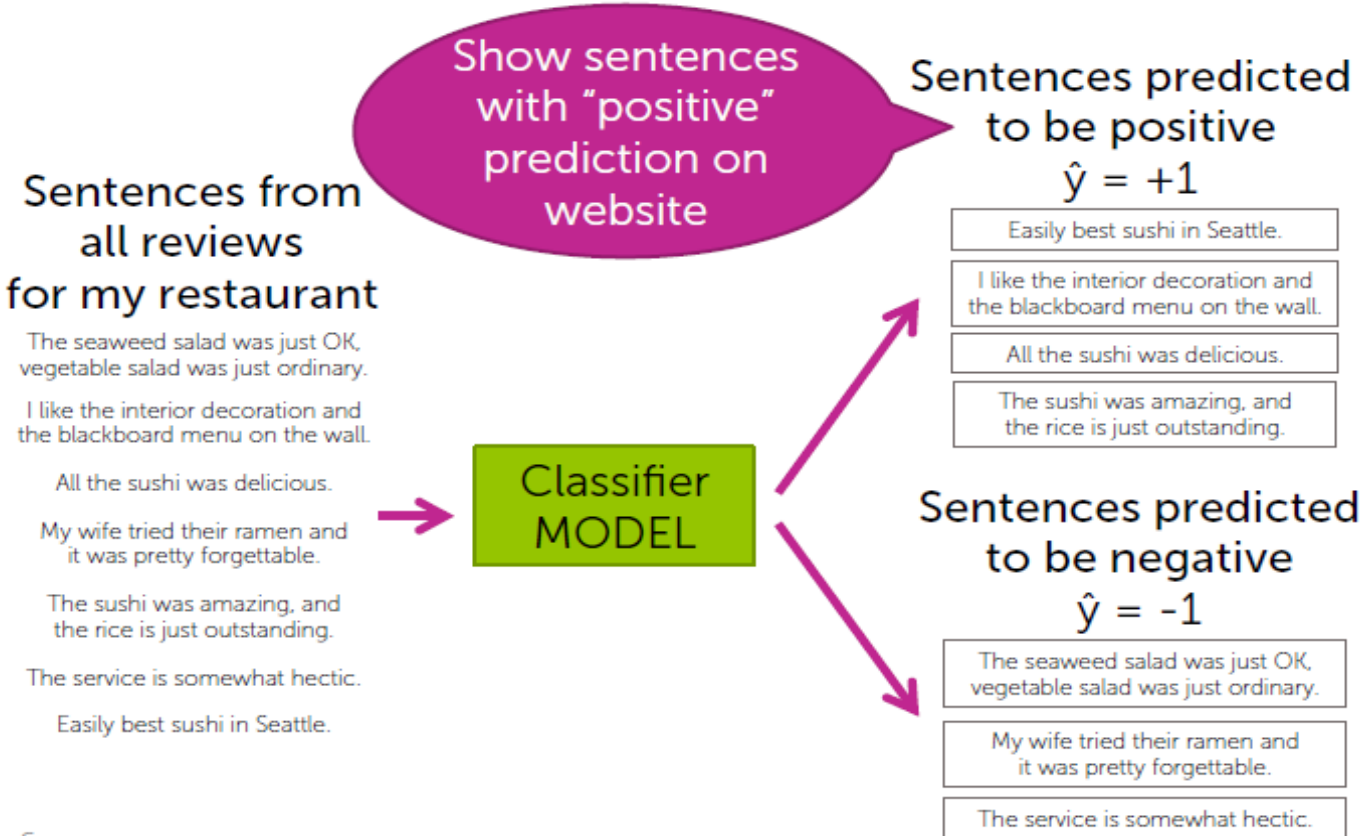

What does it mean for the classifiers to be good?

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...

What does it mean for the classifiers to be good?

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

Precision & Recall

Automated marketing campaign cares about something else...

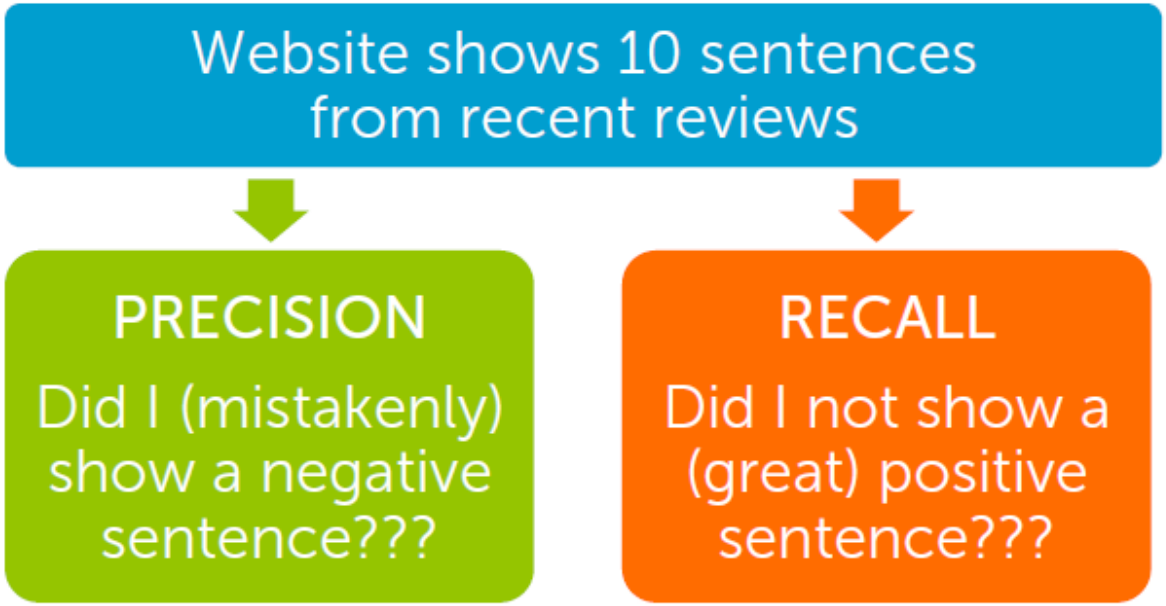

Accuracy doesn't capture these issues well...

Precision & Recall

Precision:
Fraction of positive predictions that are actually positive

Recall:
Fraction of positive data predicted to be positive

Precision

What fraction of the positive predictions are correct?

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

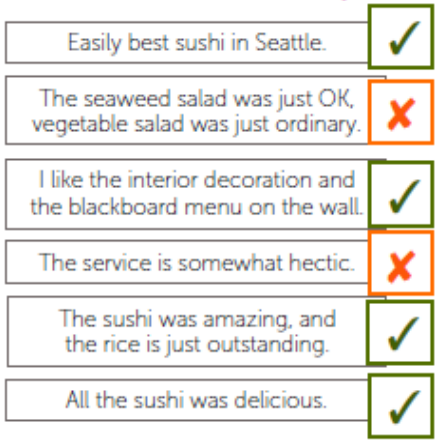

Only 4 out of 6 sentences predicted to be positive are actually positive

Precision

Precision: Fraction of positive predictions that are actually positive

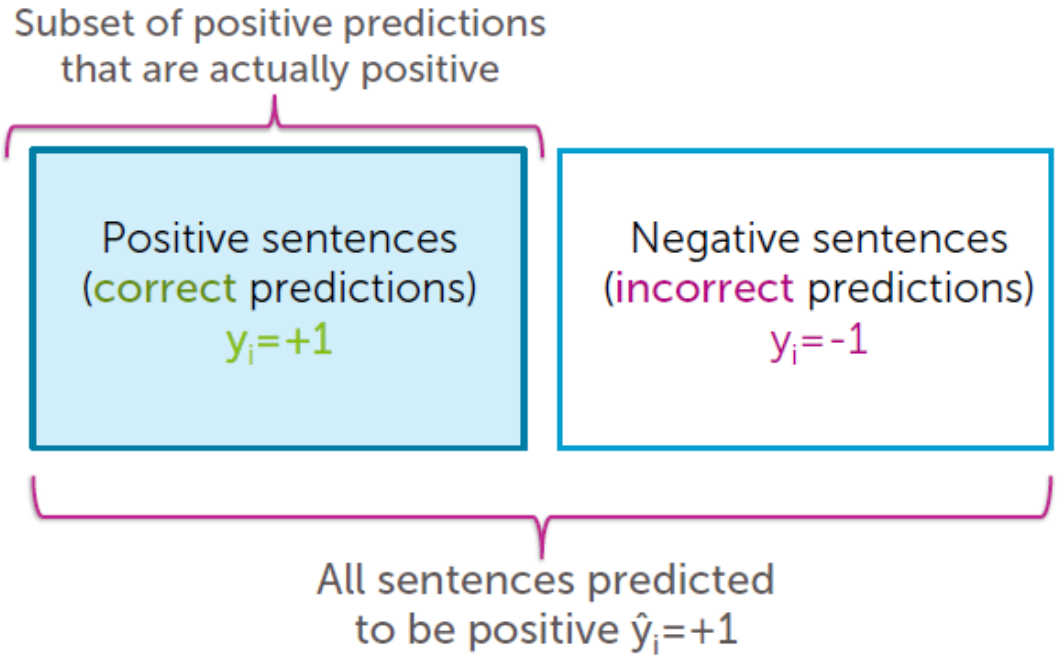

Precision

Types of error: *Review*

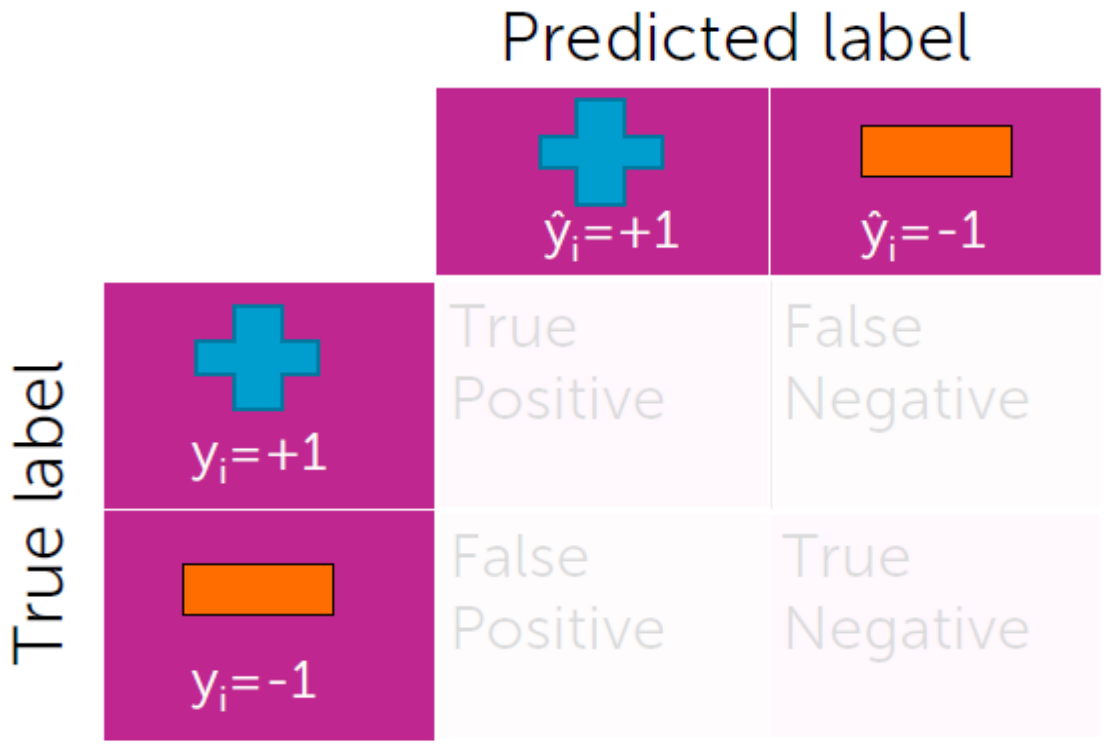

Precision

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

Precision

Why precision is important

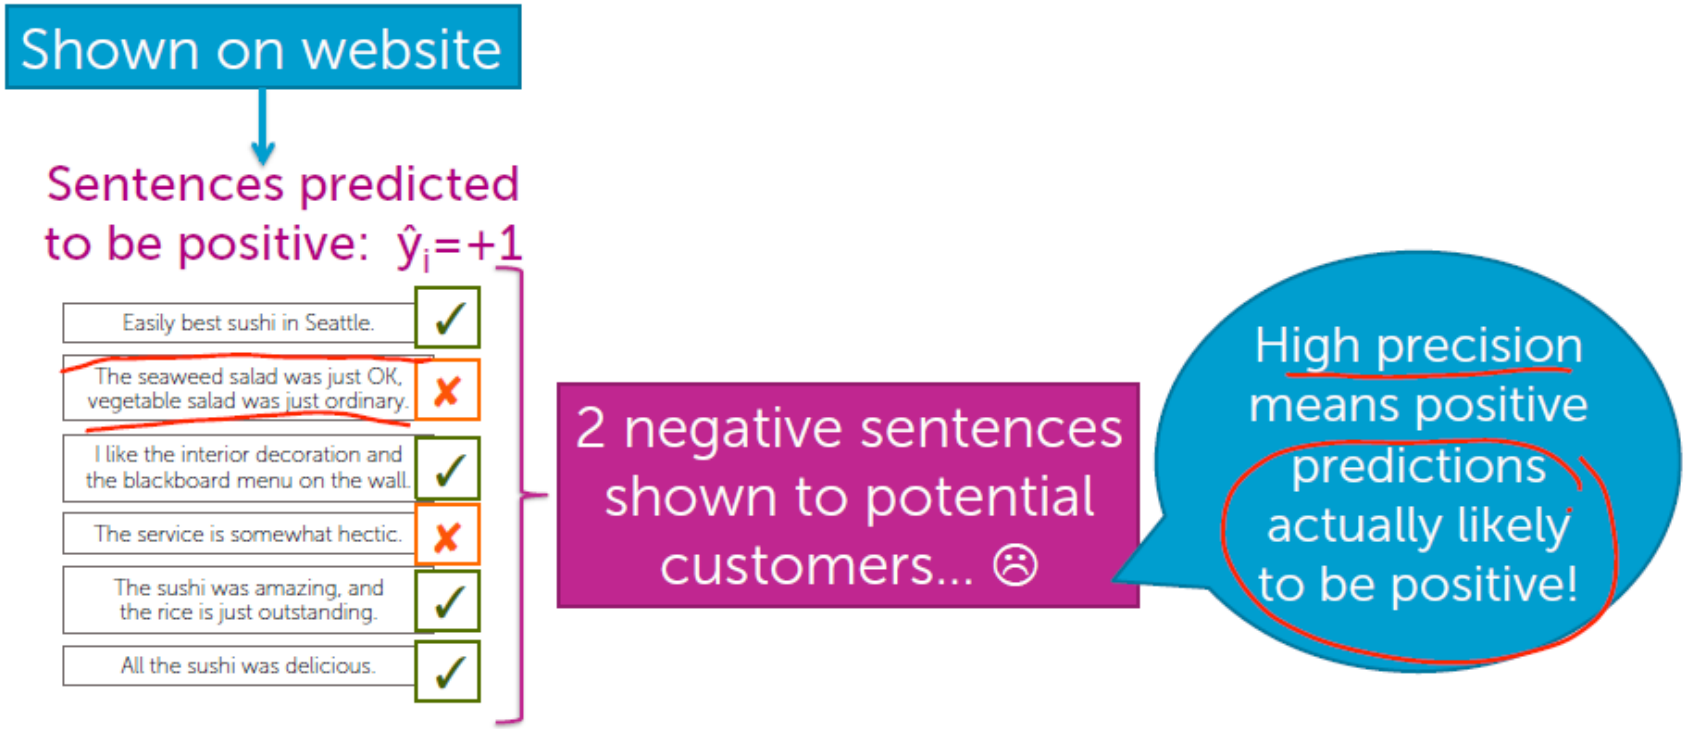

Recall

Did I find all the positive sentences?

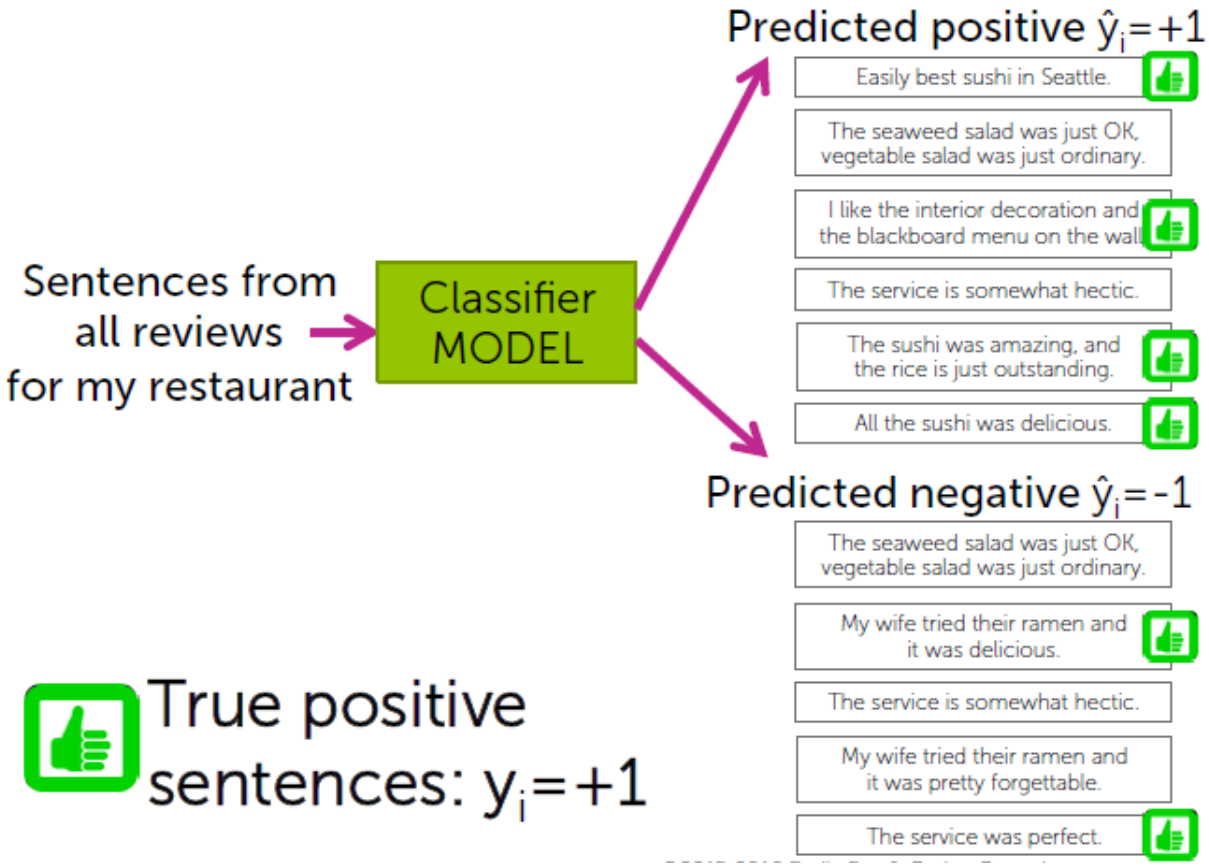

Recall

What fraction of positive sentences were missed out?

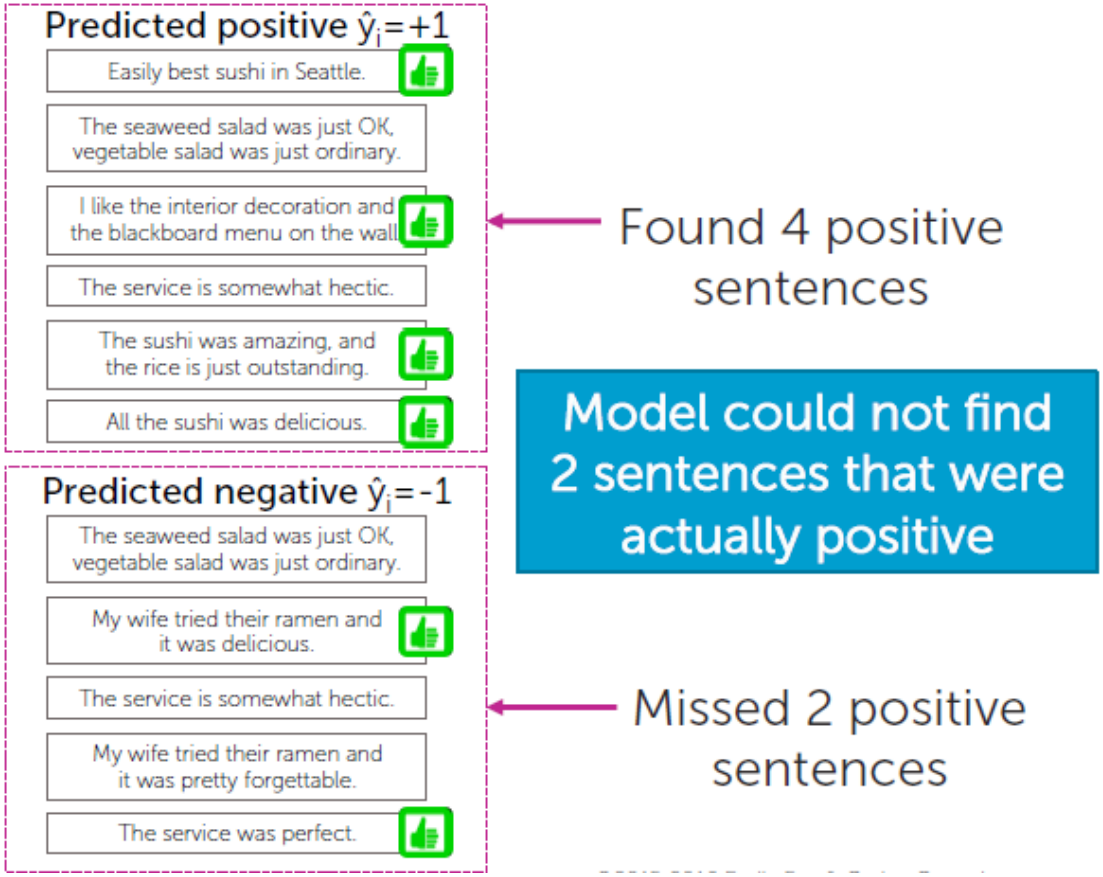

Recall

Recall: Fraction of positive data predicted to be positive

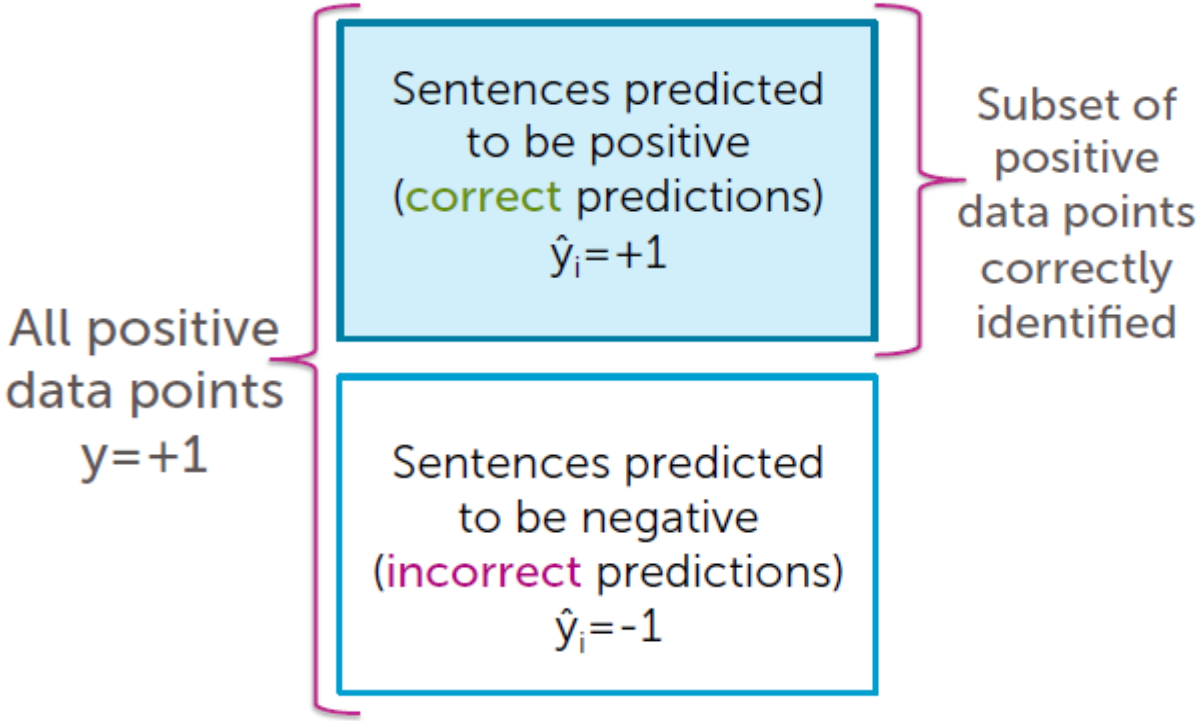

Recall

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

Recall

Why is recall important?

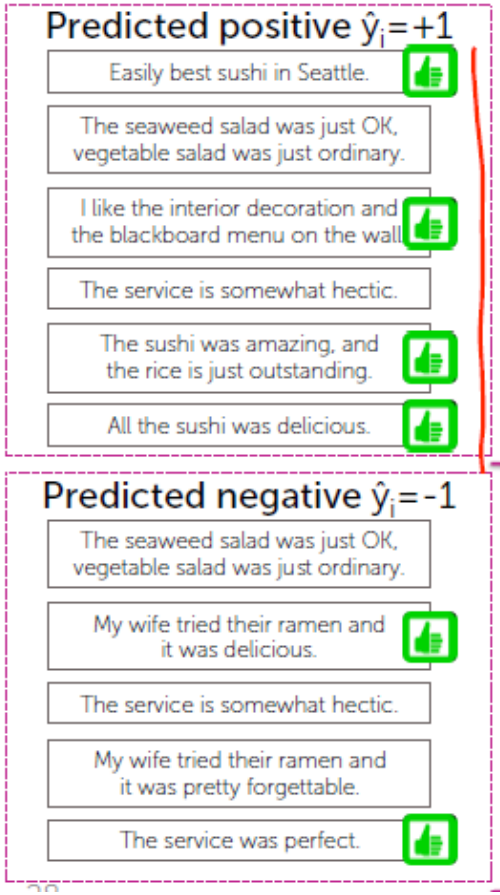

Want to show positive sentences on website

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

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

Precision-recall extremes

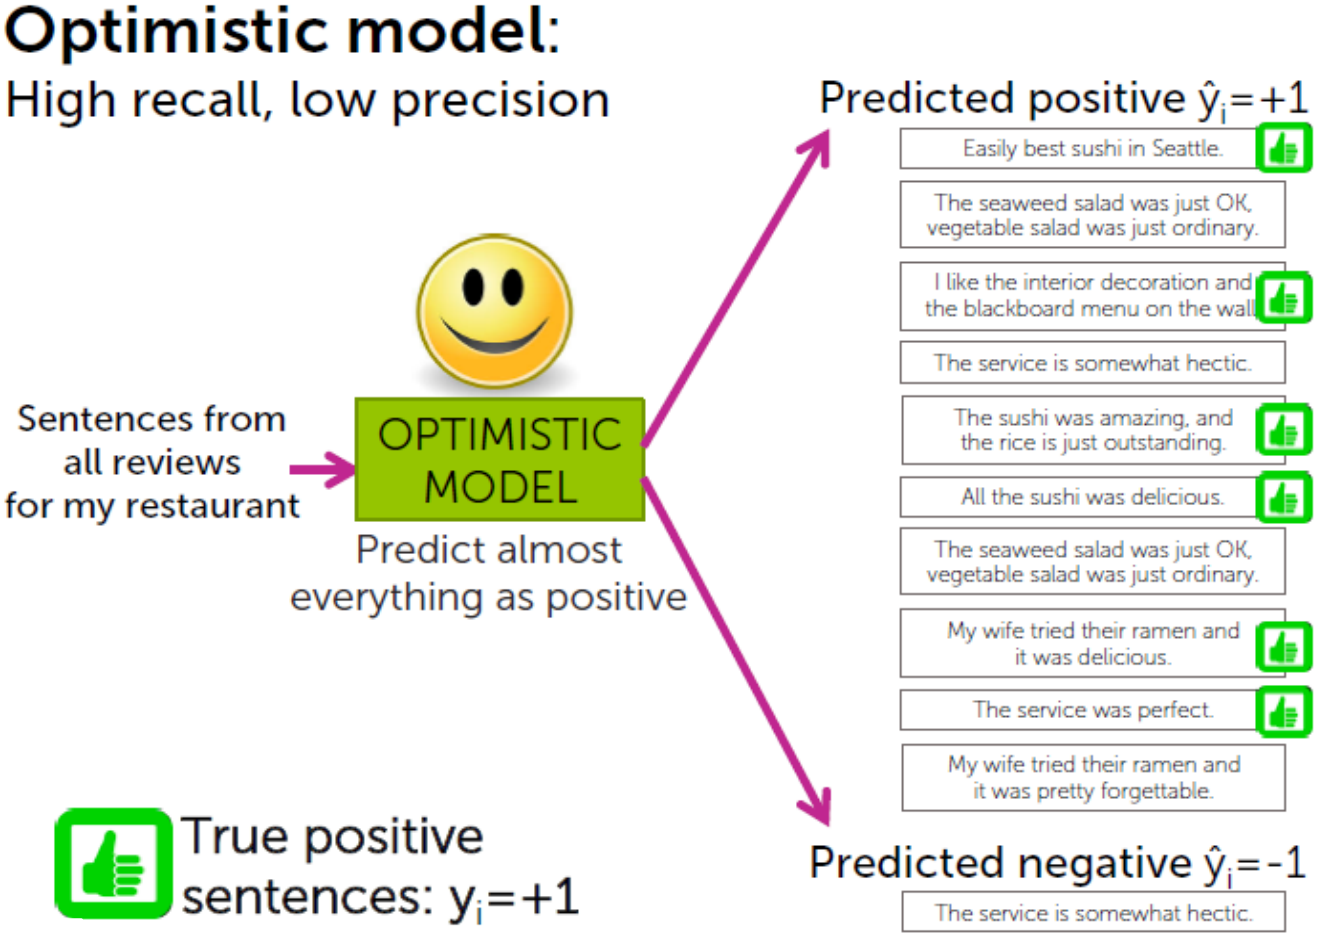

Precision-recall extremes

Pessimistic model:

High precision, low recall

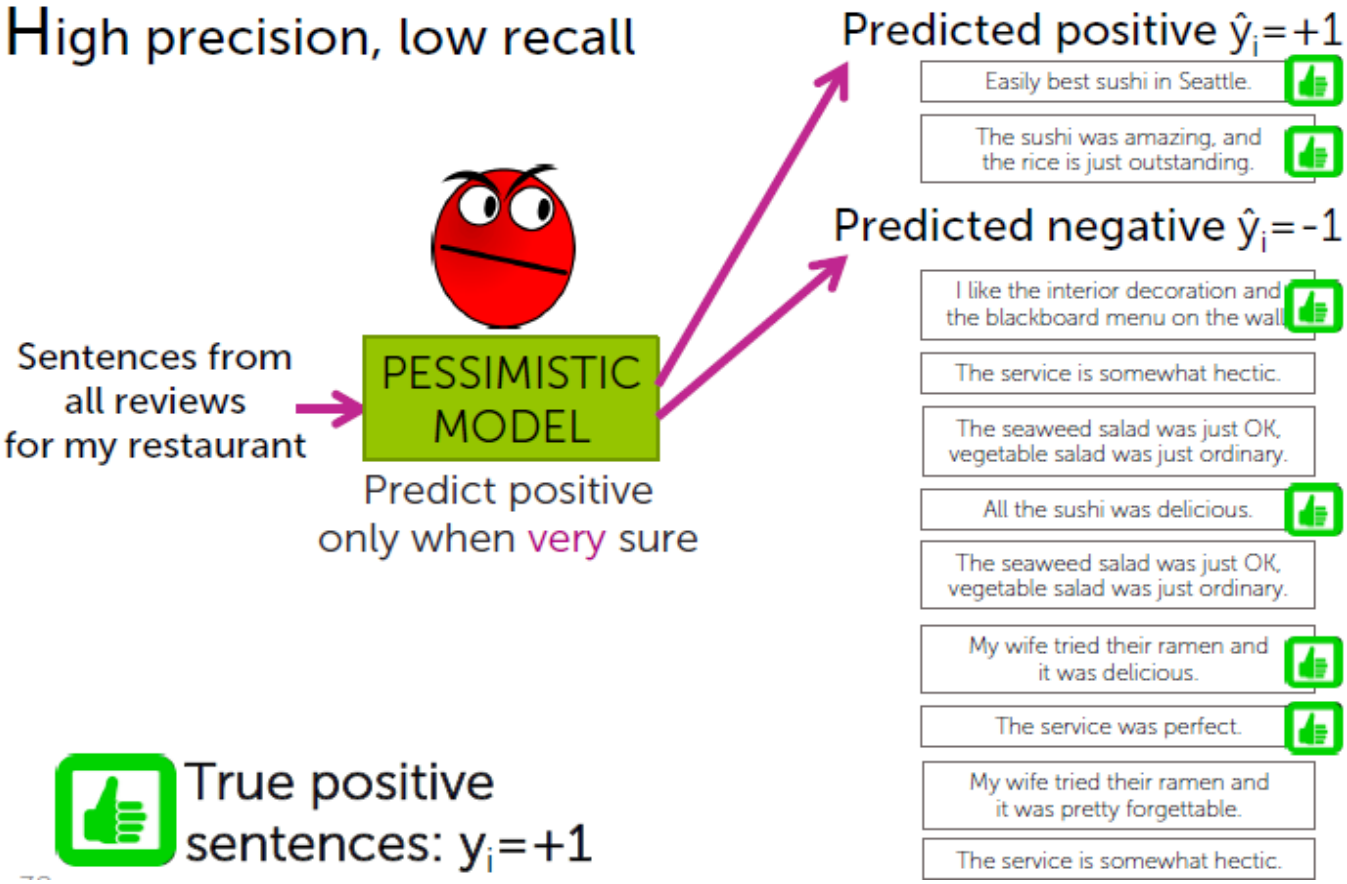

Balancing precision & recall

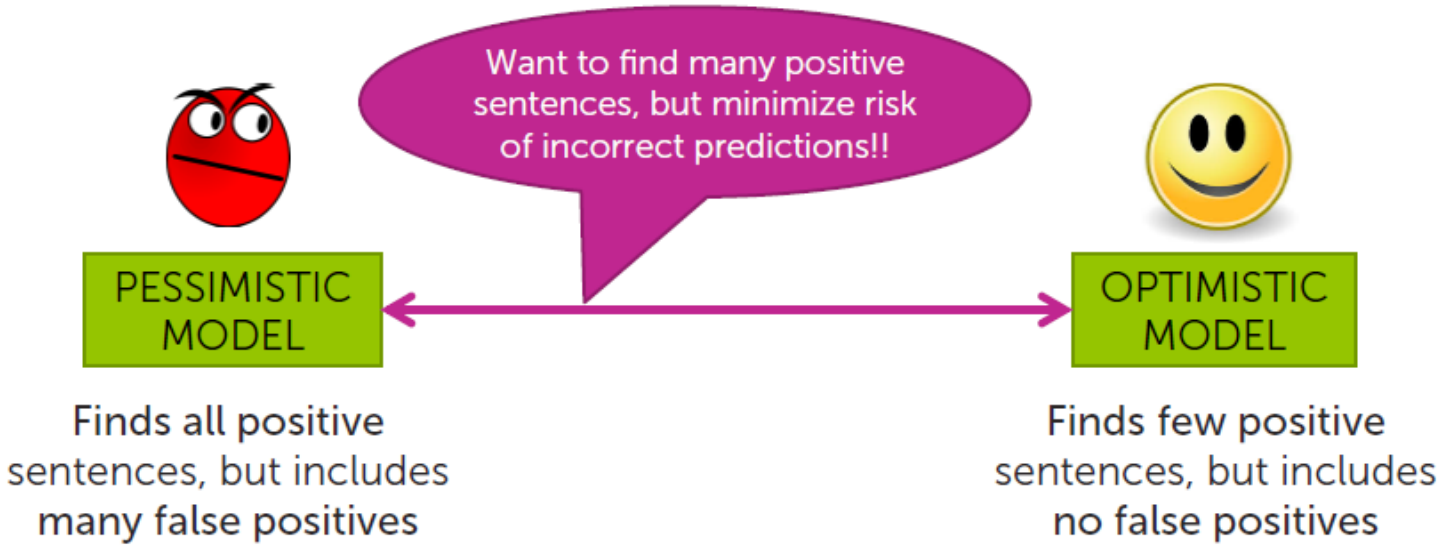

Can we tradeoff precision & recall?

Low precision, high recall

High precision, low recall

Optimistic Model
Predict almost everything as positive

Pessimistic Model
Predict positive only when very sure

How confident is your prediction?

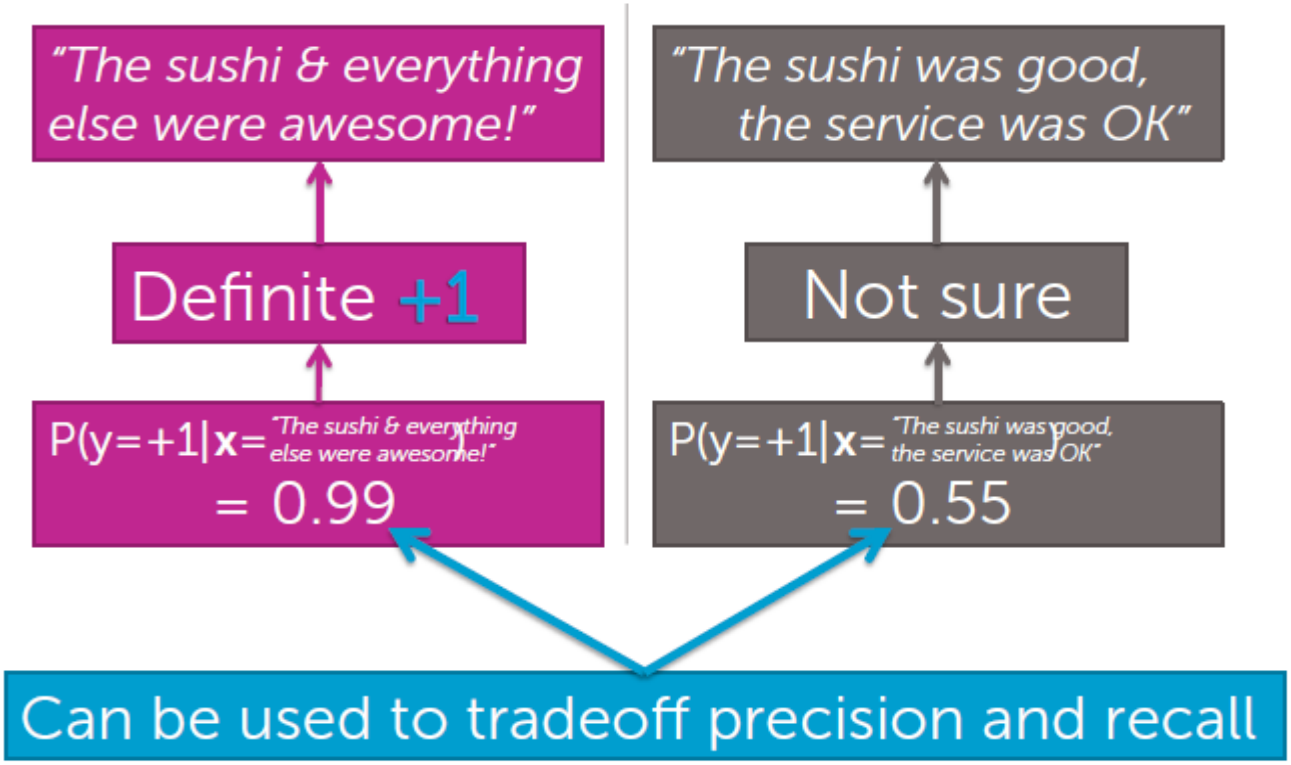

Tradeoff precision & recall with threshold

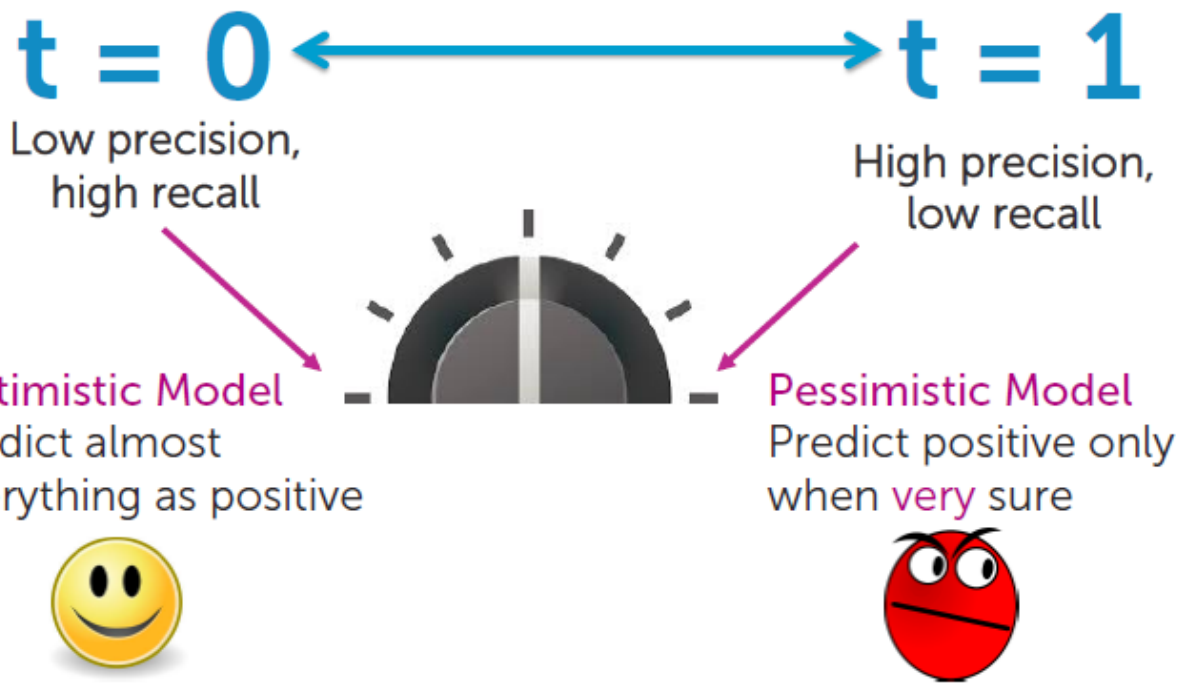

The precision-recall curve

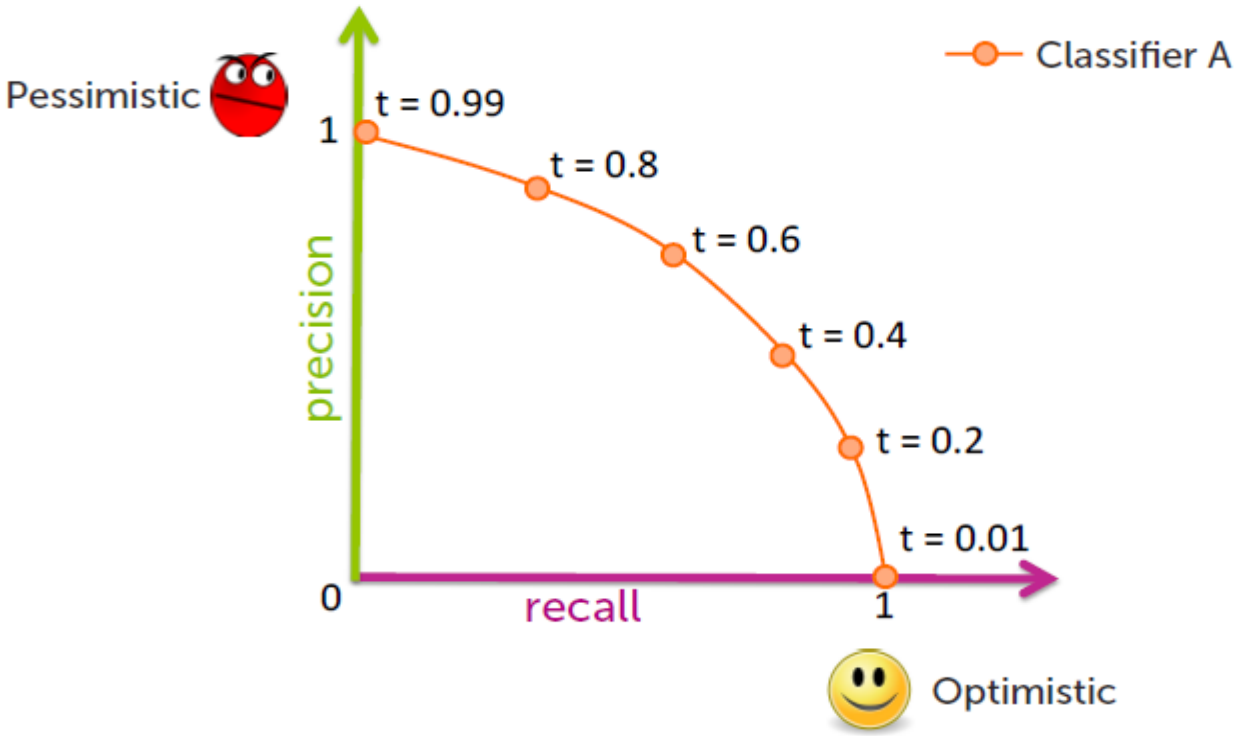

What does the perfect algorithm look like?

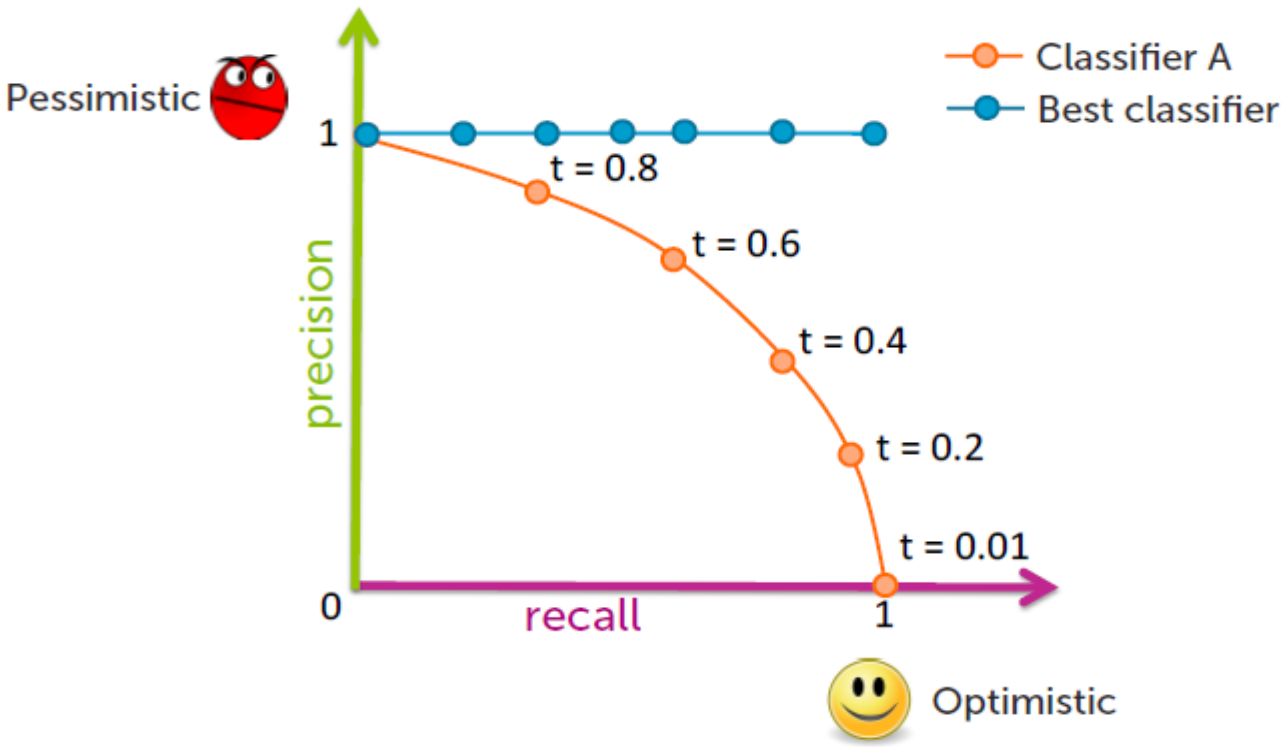

Which classifier is better? A or C?

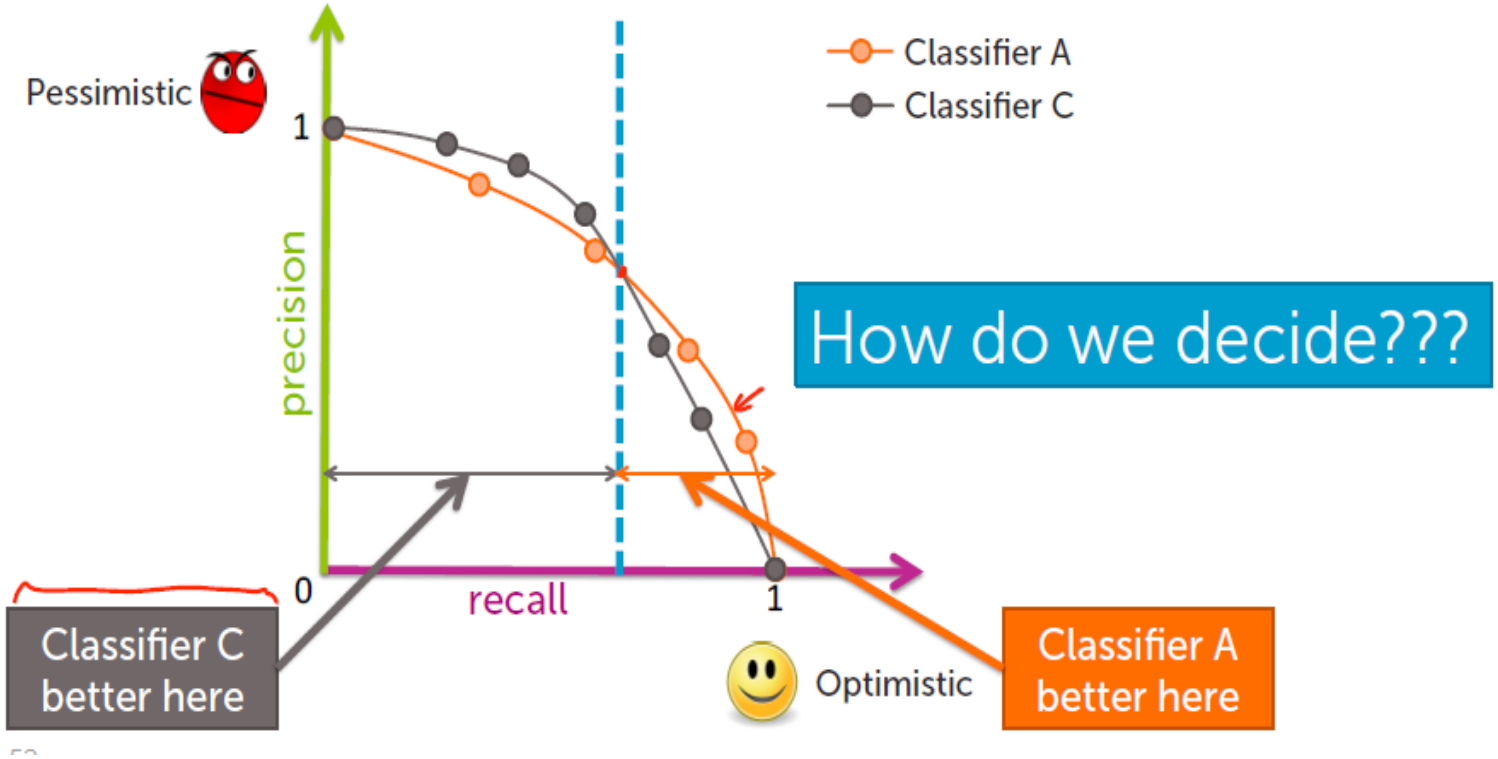

Compare algorithms

- Often, reduce precision-recall to single number to compare algorithms
  - F1 measure, area-under-the-curve (AUC),...

Precision at k

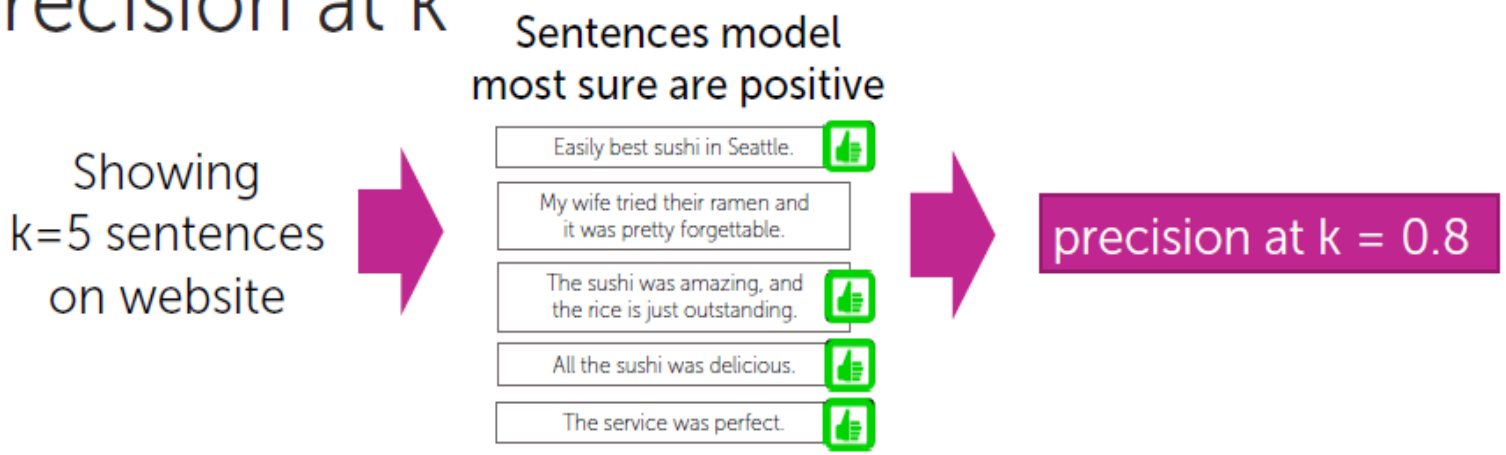

What you can do now

- 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