

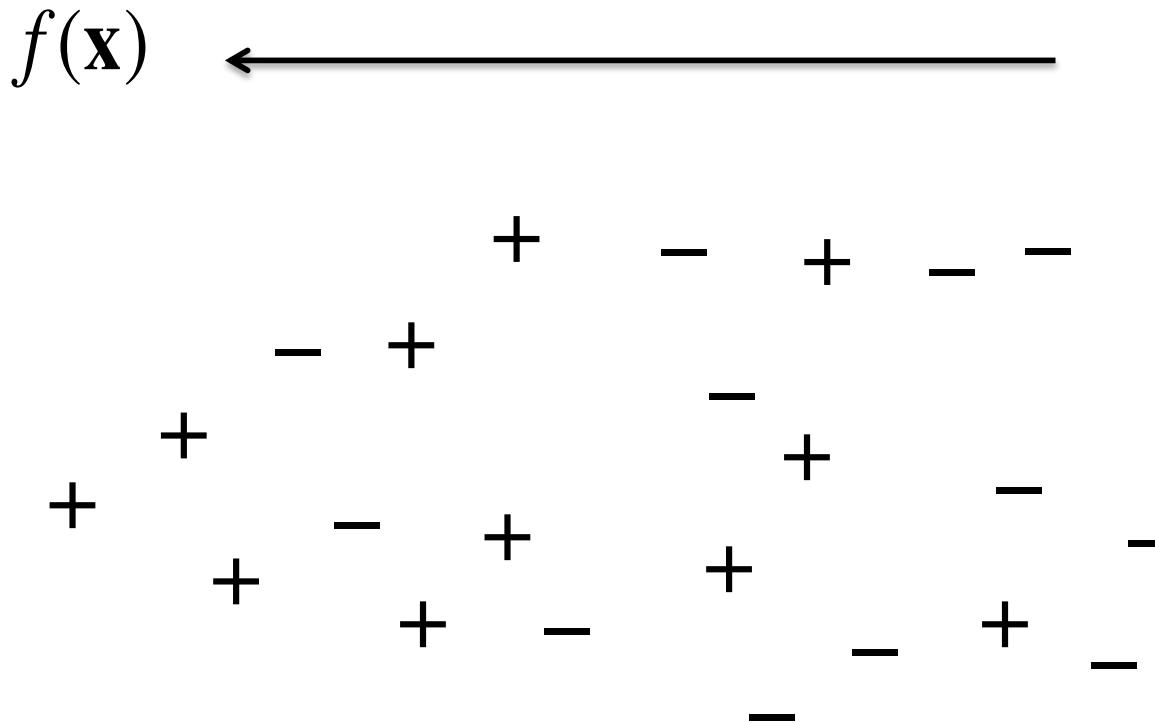
# More on ROC Curves for Individual Models

Cynthia Rudin

Machine Learning Course, Duke

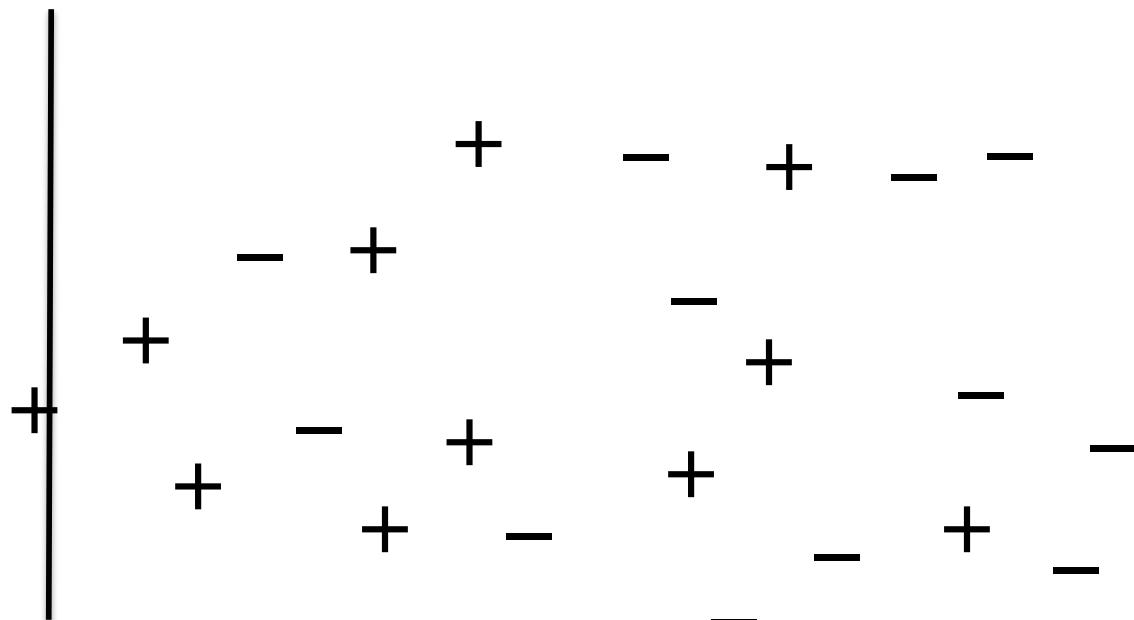
# ROC Curves

- Adjust the decision boundary



# ROC Curves

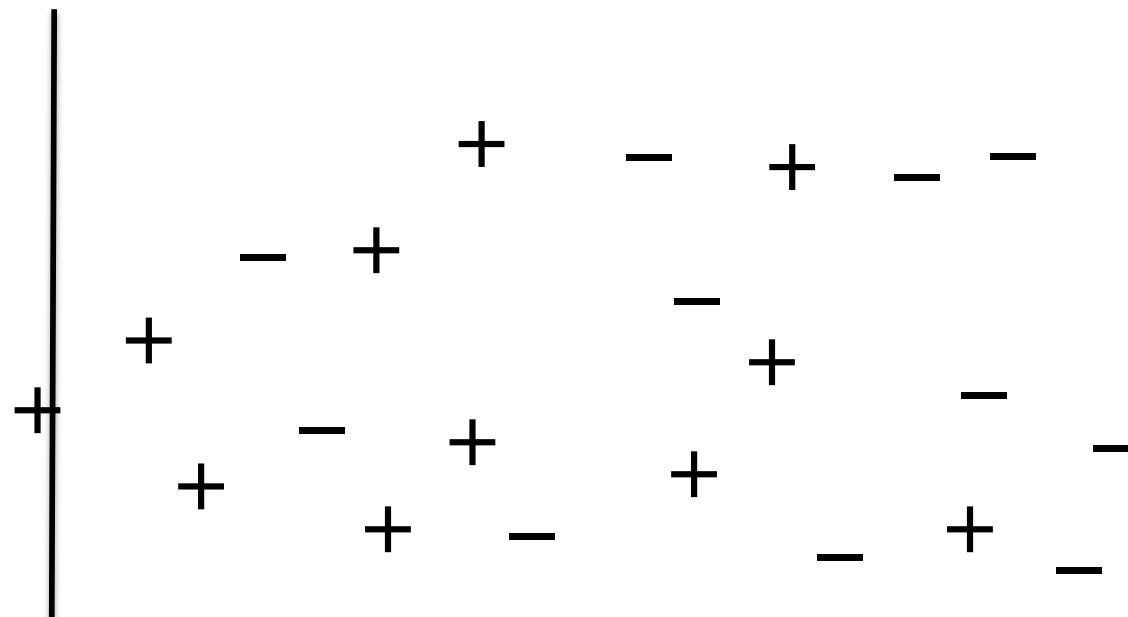
- Adjust the decision boundary



# ROC Curves

- Adjust the decision boundary

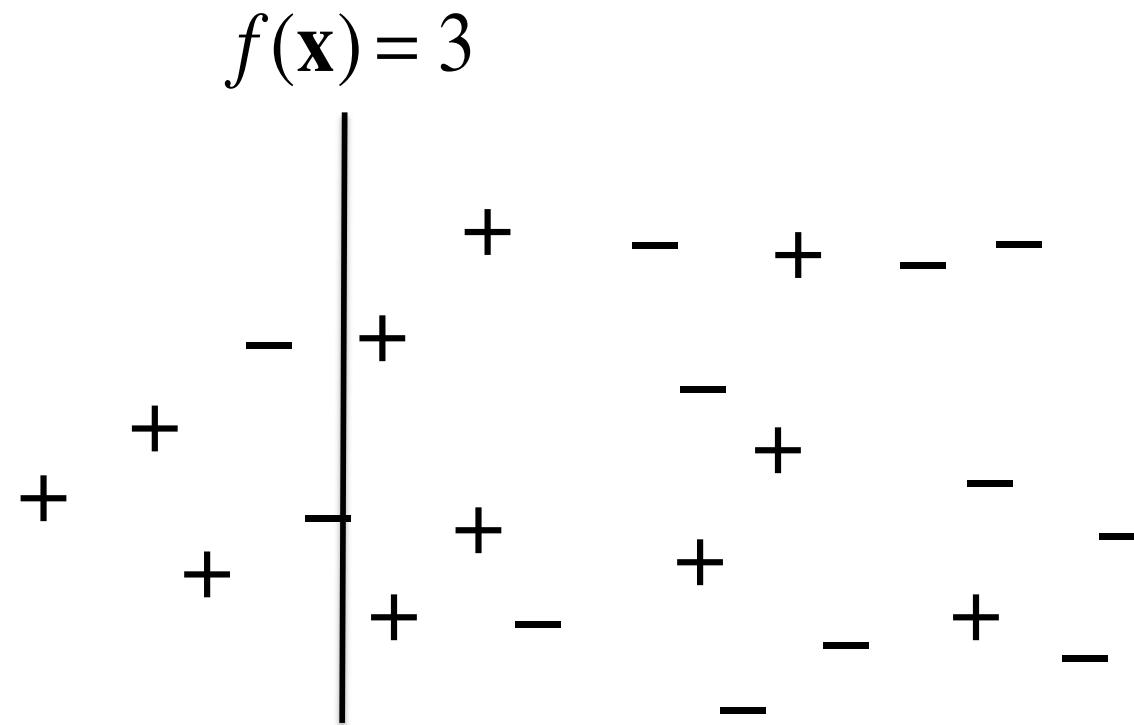
$$f(\mathbf{x}) = 7$$



- TPR = 1/11
- FPR = 0/12

# ROC Curves

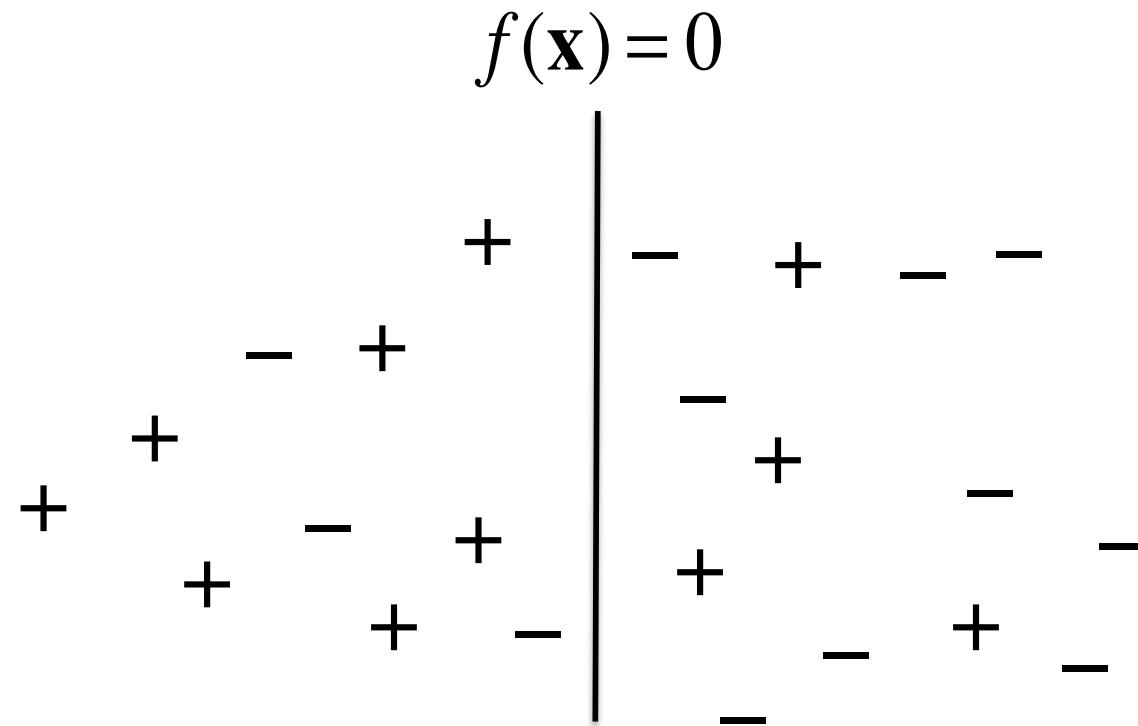
- Adjust the decision boundary



- TPR = 3/11
- FPR = 2/12

# ROC Curves

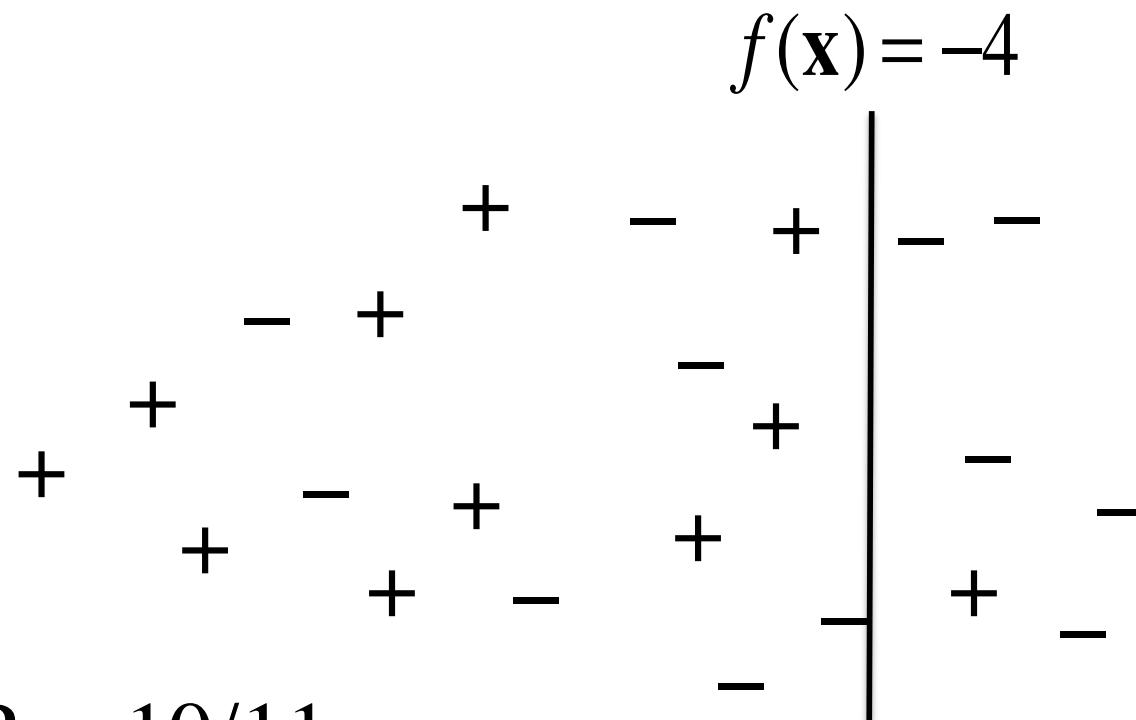
- Adjust the decision boundary



- TPR = 7/11
- FPR = 3/12

# ROC Curves

- Adjust the decision boundary



- TPR = 10/11
- FPR = 7/12

# ROC Curves

- For a particular False Positive Rate (FPR), what is the True Positive Rate (TPR)?



# ROC Curves

- Let's do it without scaling



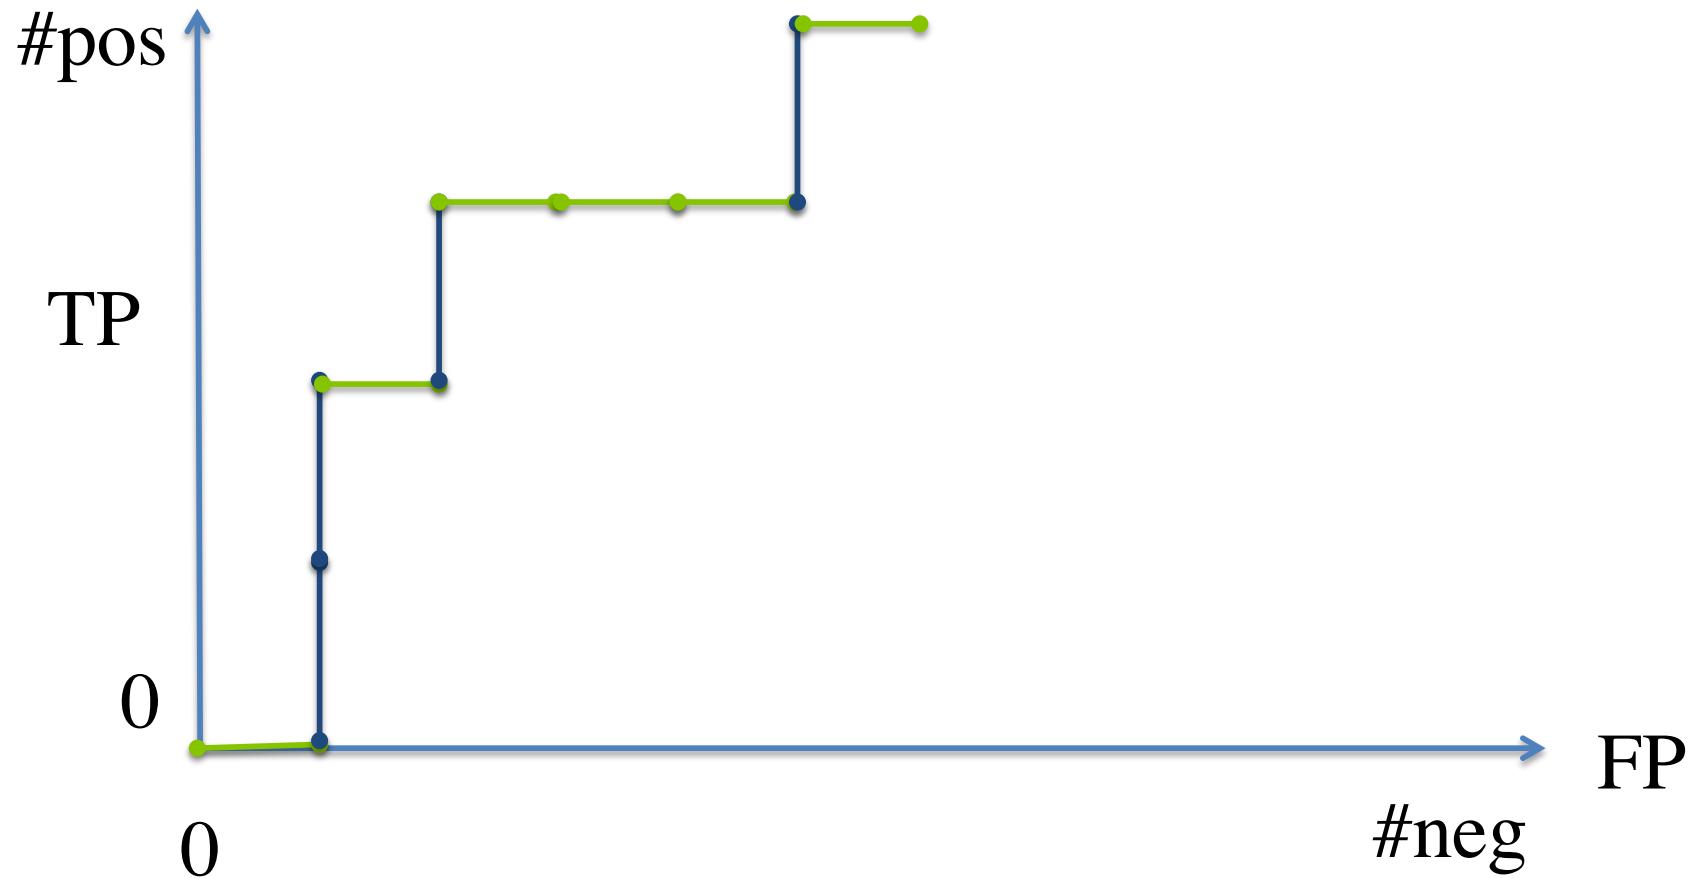
# ROC Curves

- Let's do it without scaling



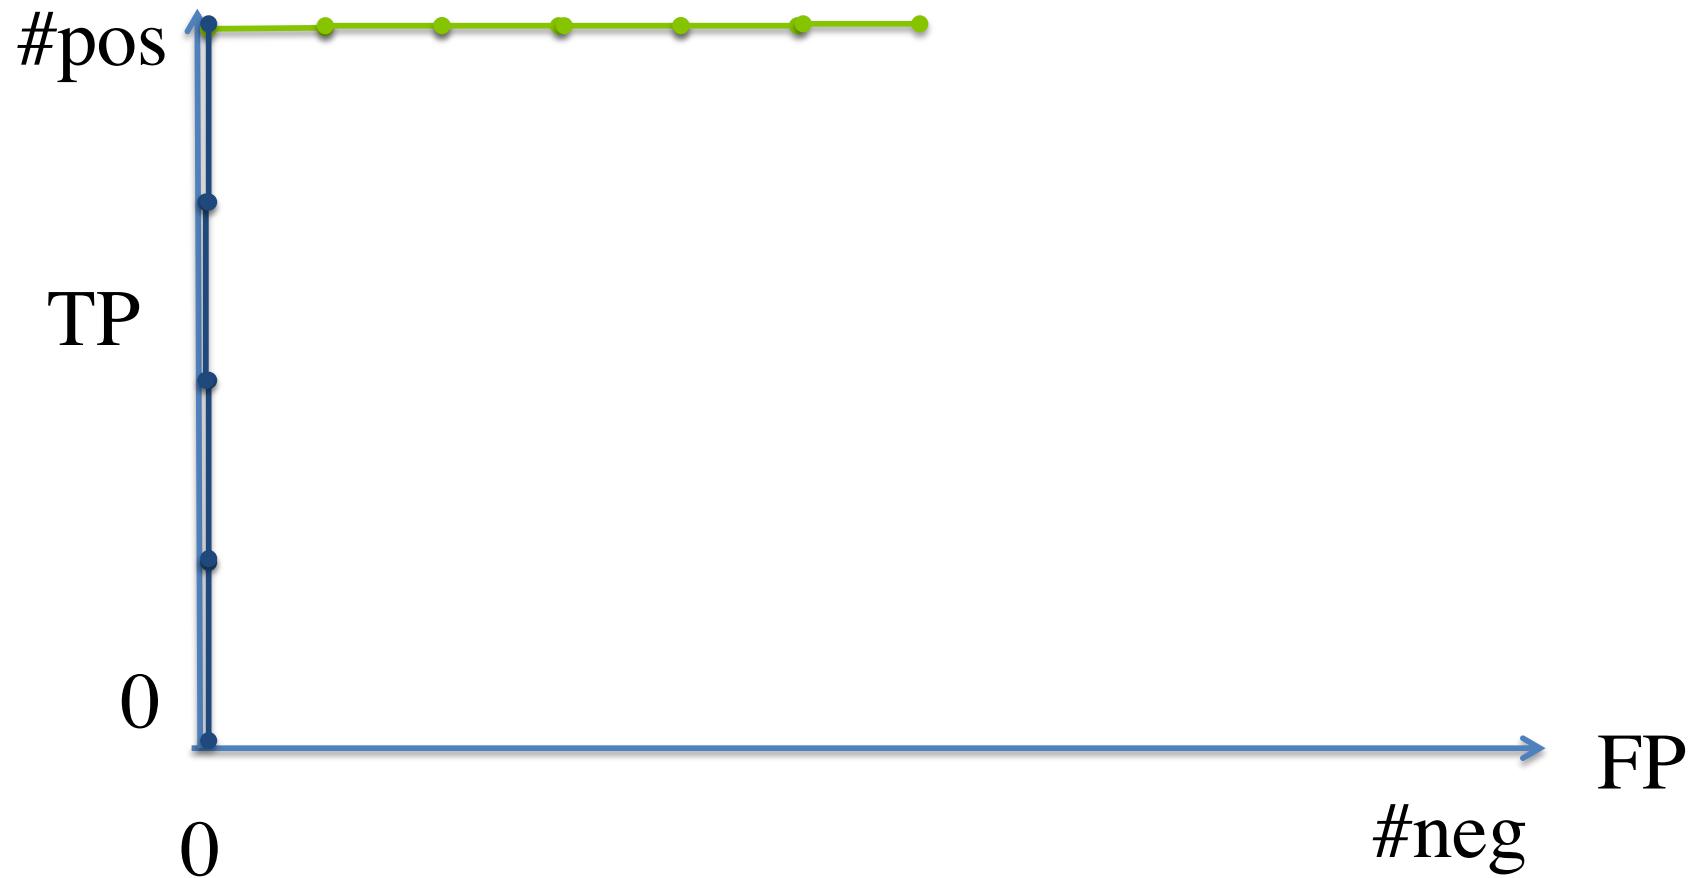
# ROC Curves

- To do this, you need only values of  $f(x)$  and  $y$ .
  - e.g.,  $f(x) = 15, 12, 10, 8, 6, 2, -1, -3, -14, \dots$
  - e.g.,  $y(x) = \textcolor{brown}{-1}, \textcolor{brown}{1}, \textcolor{brown}{1}, \textcolor{brown}{-1}, \textcolor{brown}{1}, \textcolor{brown}{-1}, \textcolor{brown}{-1}, \textcolor{brown}{-1}, \textcolor{brown}{1}, \textcolor{brown}{-1}, \dots$



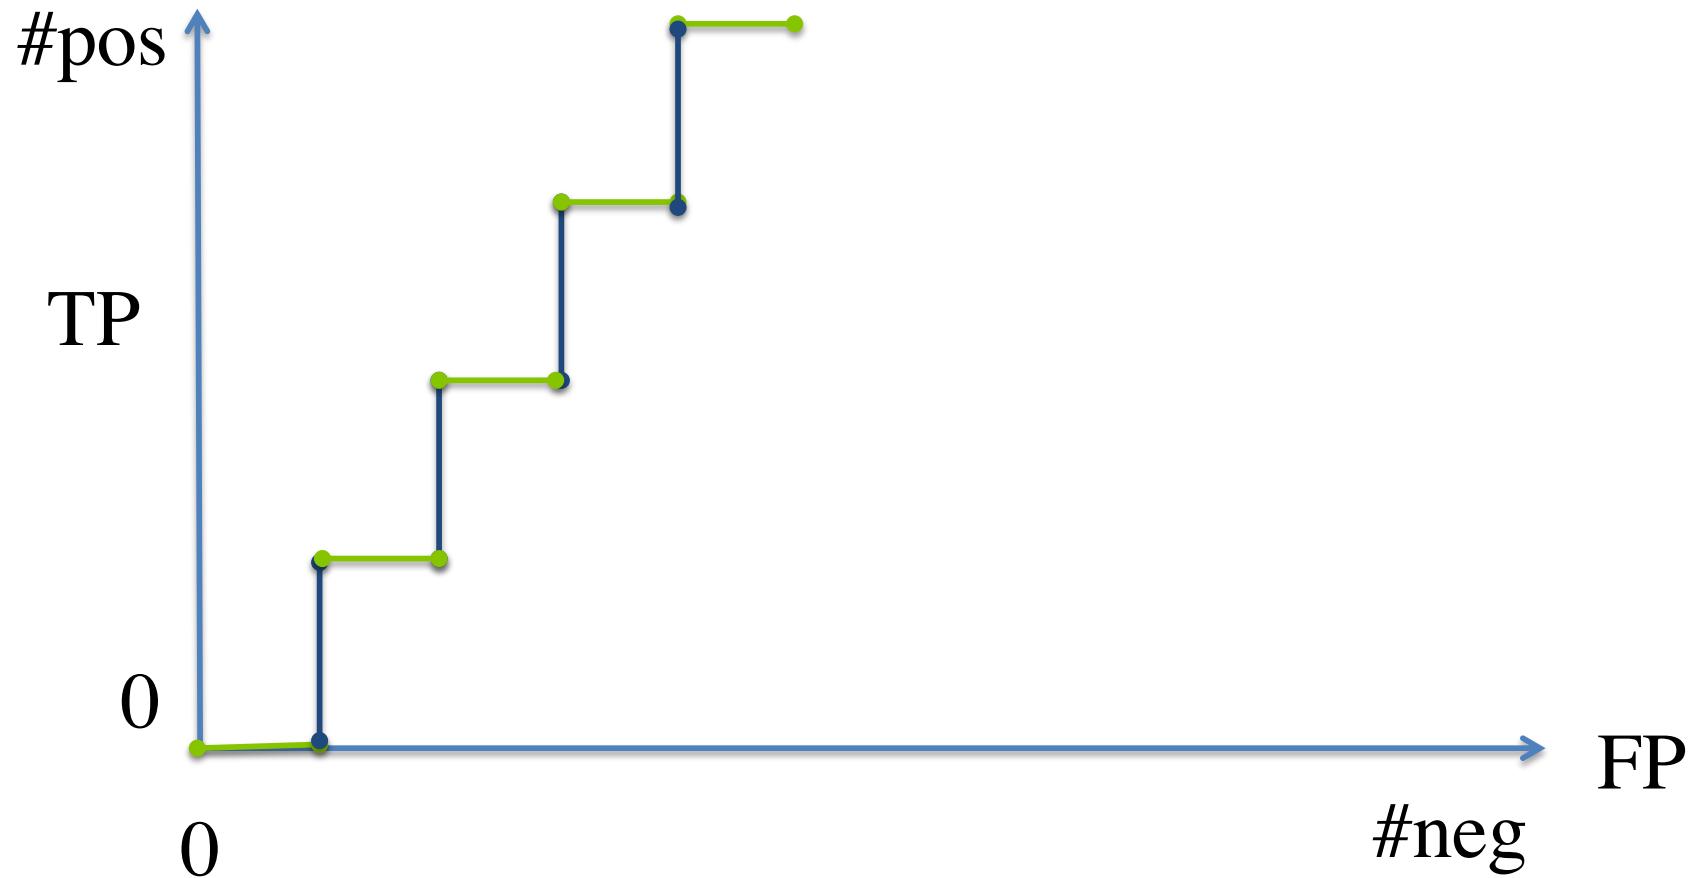
# ROC Curves

- If  $f(x)$  is really good:
  - e.g.,  $f(x) = 15, 12, 10, 8, 6, 2, -1, -3, -14, \dots$
  - e.g.,  $y(x) = 1, 1, 1, 1, -1, -1, -1, -1, -1, \dots$



# ROC Curves

- If  $f(x)$  is really bad:
  - e.g.,  $f(x) = 15, 12, 10, 8, 6, 2, -1, -3, -14, \dots$
  - e.g.,  $y(x) = -1, 1, -1, 1, -1, 1, -1, 1, -1, \dots$



# ROC Curves

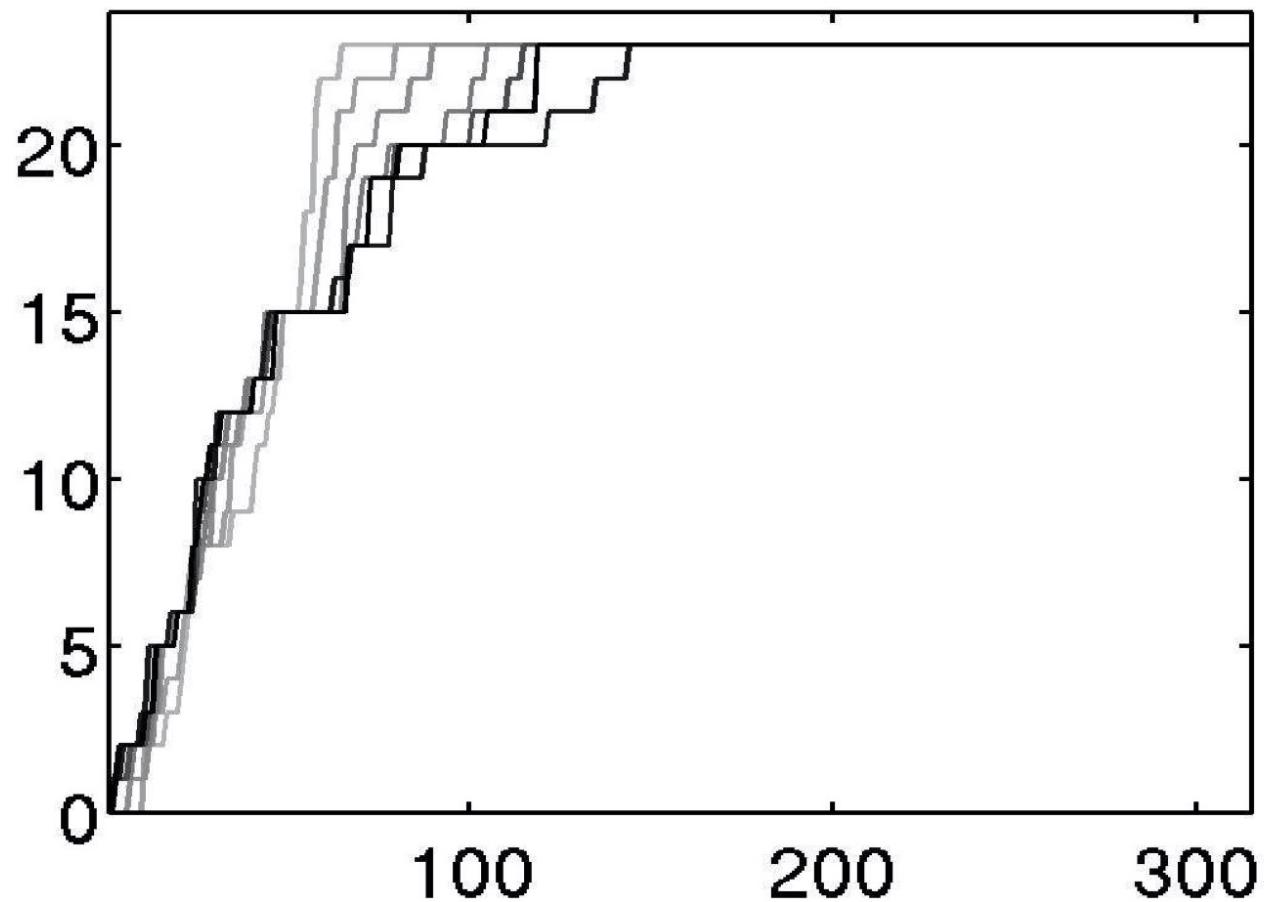
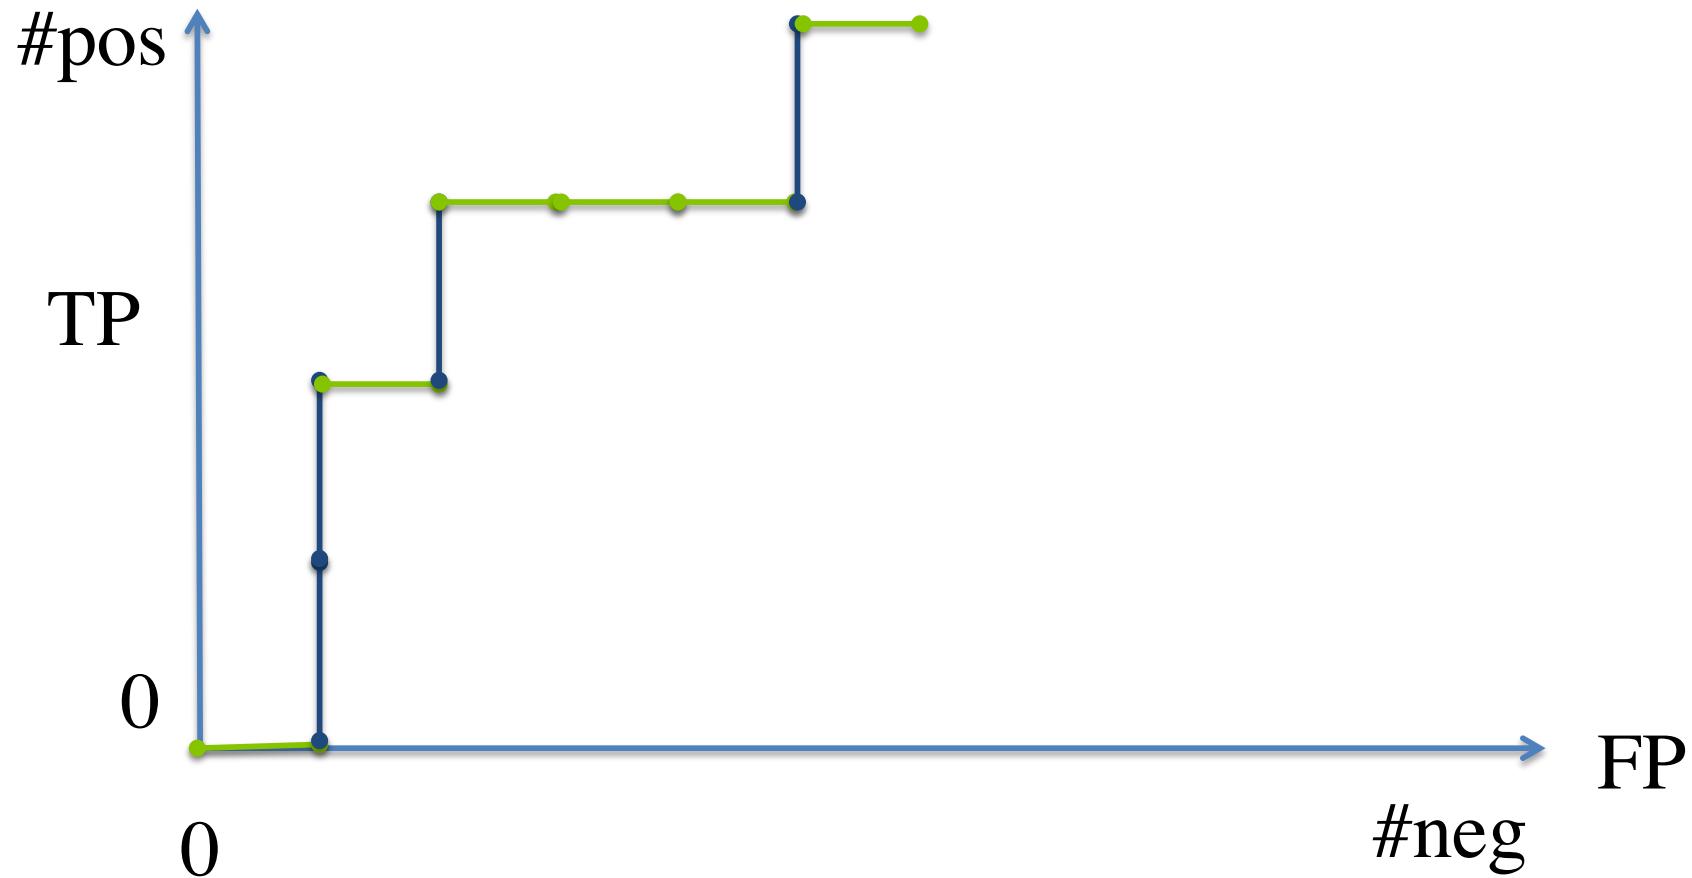


Image taken from P-Norm Push, JMLR, 2009

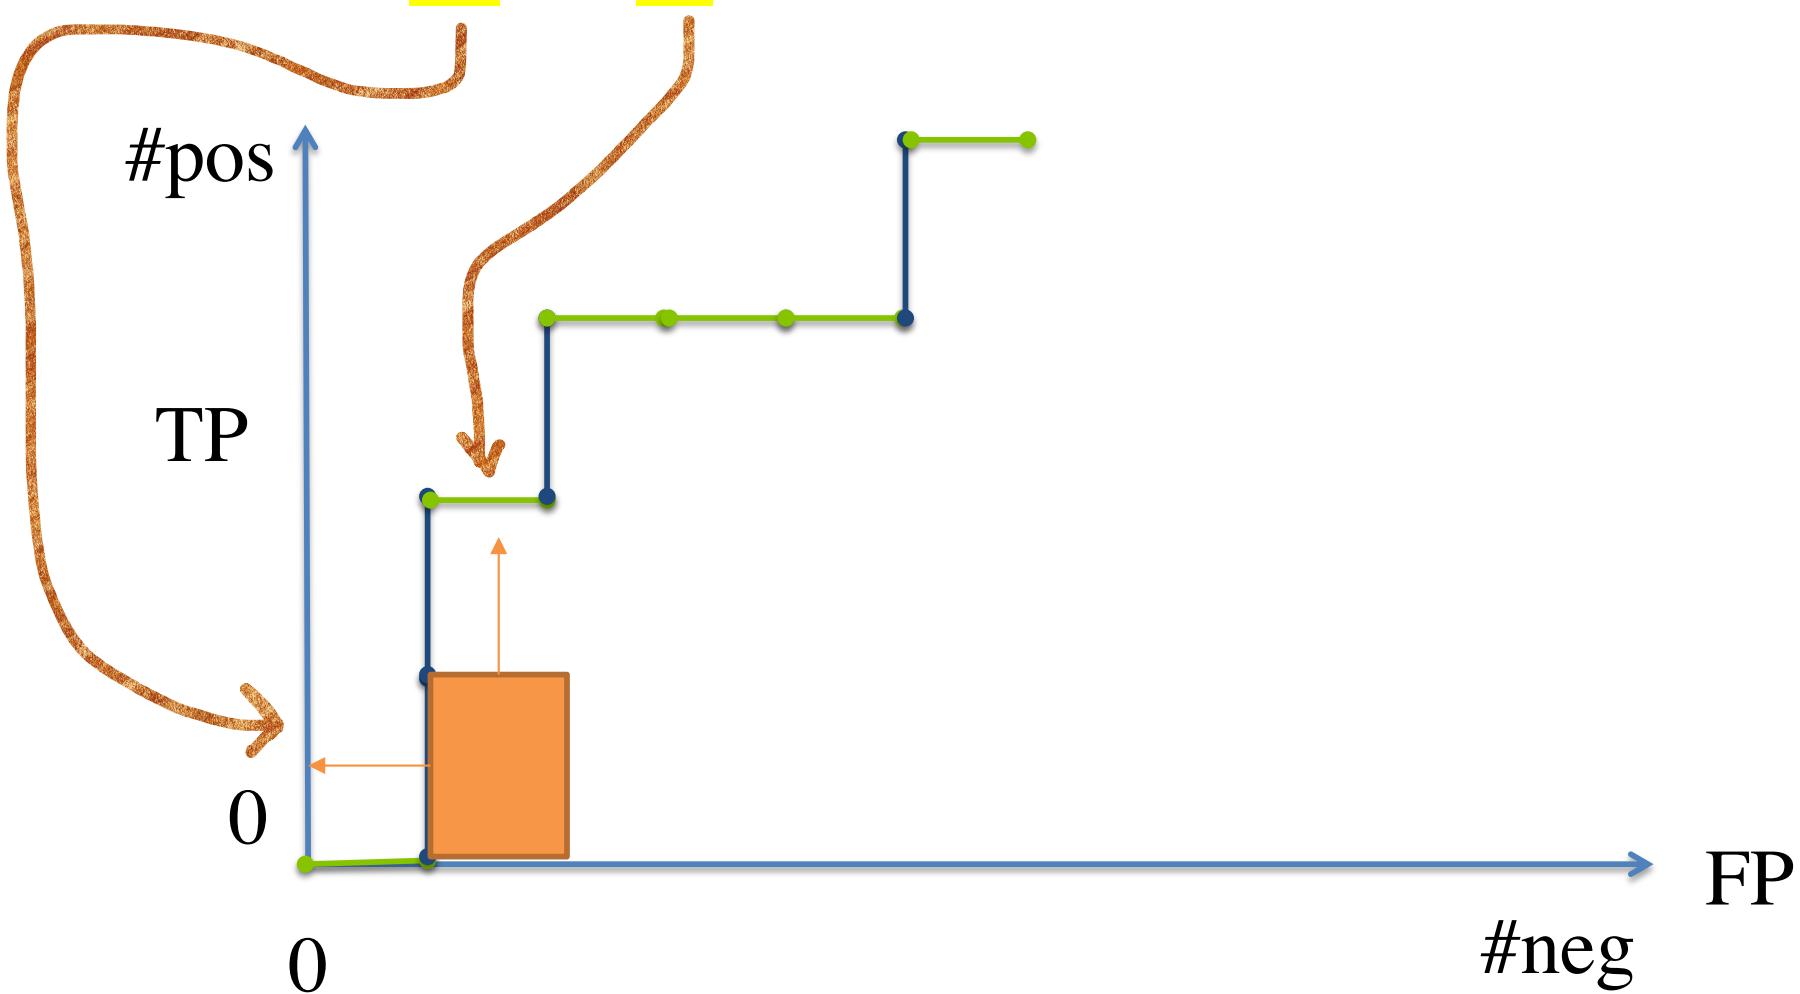
# The AUC is a rank statistic / U-statistic

- $f(x) = 15, 12, 10, 8, 6, 2, -1, -3, -14, \dots$
- $y(x) = -1, 1, 1, -1, 1, -1, -1, -1, 1, -1, \dots$



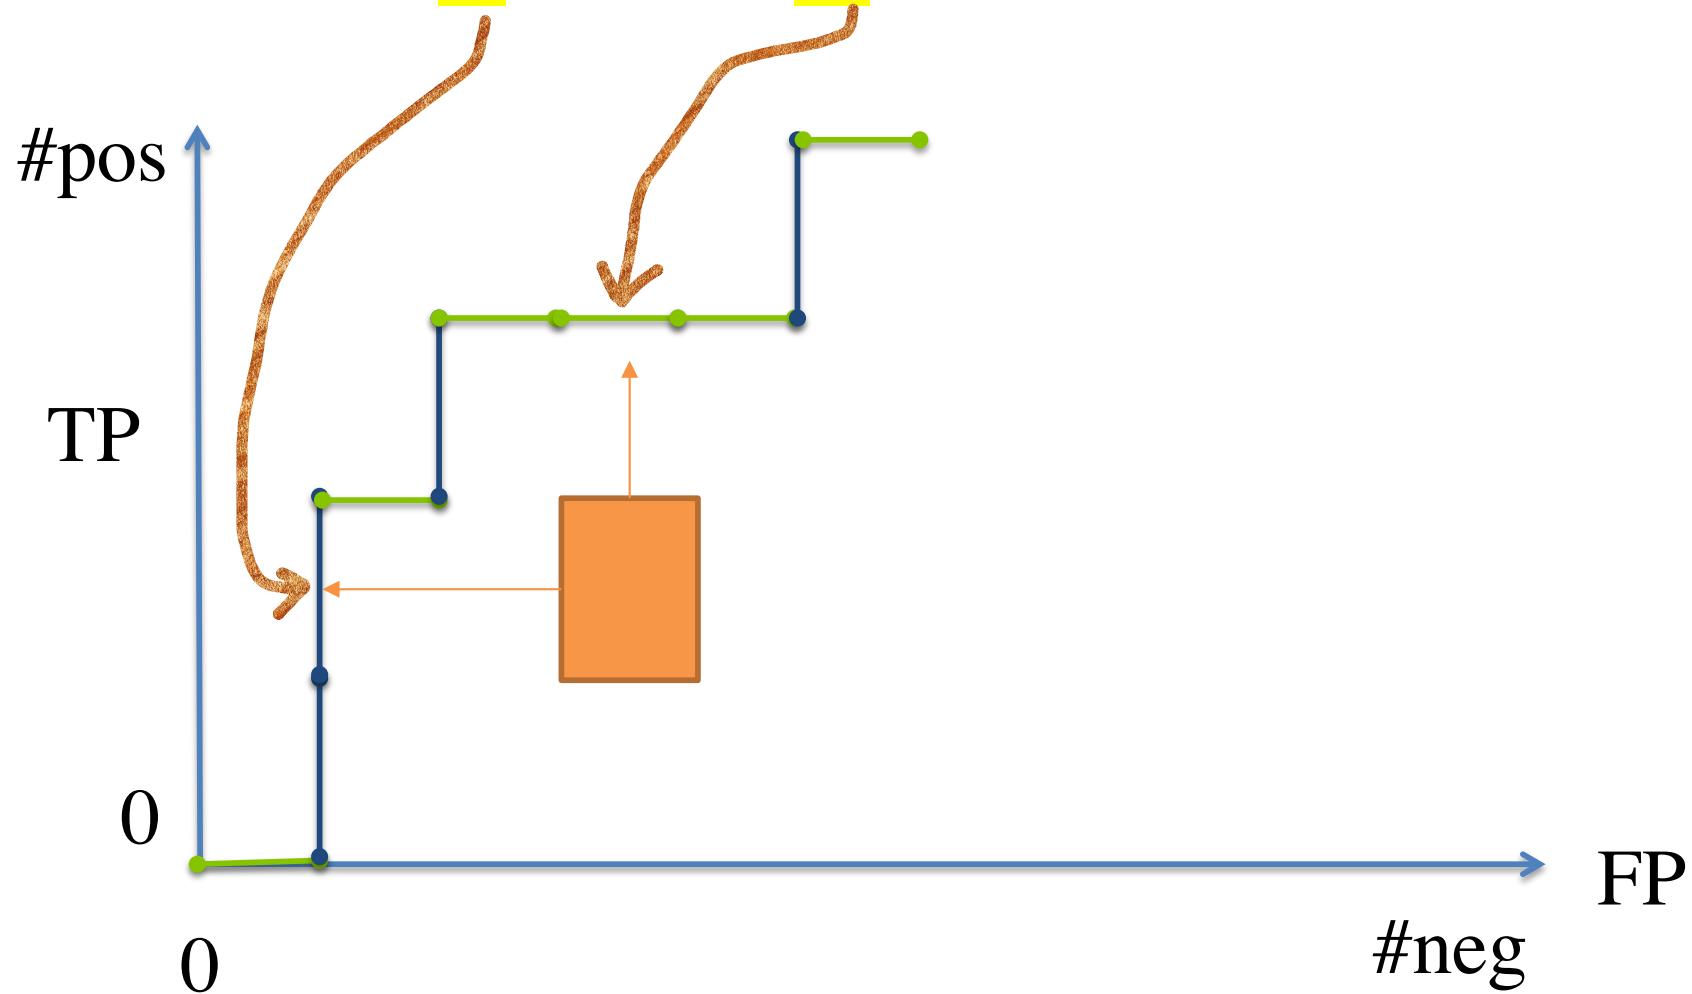
# The AUC is a rank statistic / U-statistic

- $f(x) = 15, 12, 10, 8, 6, 2, -1, -3, -14, \dots$
- $y(x) = -1, 1, 1, -1, 1, -1, -1, -1, 1, -1, \dots$



# The AUC is a rank statistic / U-statistic

- $f(x) = 15, 12, \boxed{10}, 8, 6, 2, \boxed{-1}, -3, -14, \dots$
- $y(x) = -1, 1, \boxed{1}, -1, 1, -1, \boxed{-1}, -1, 1, -1, \dots$



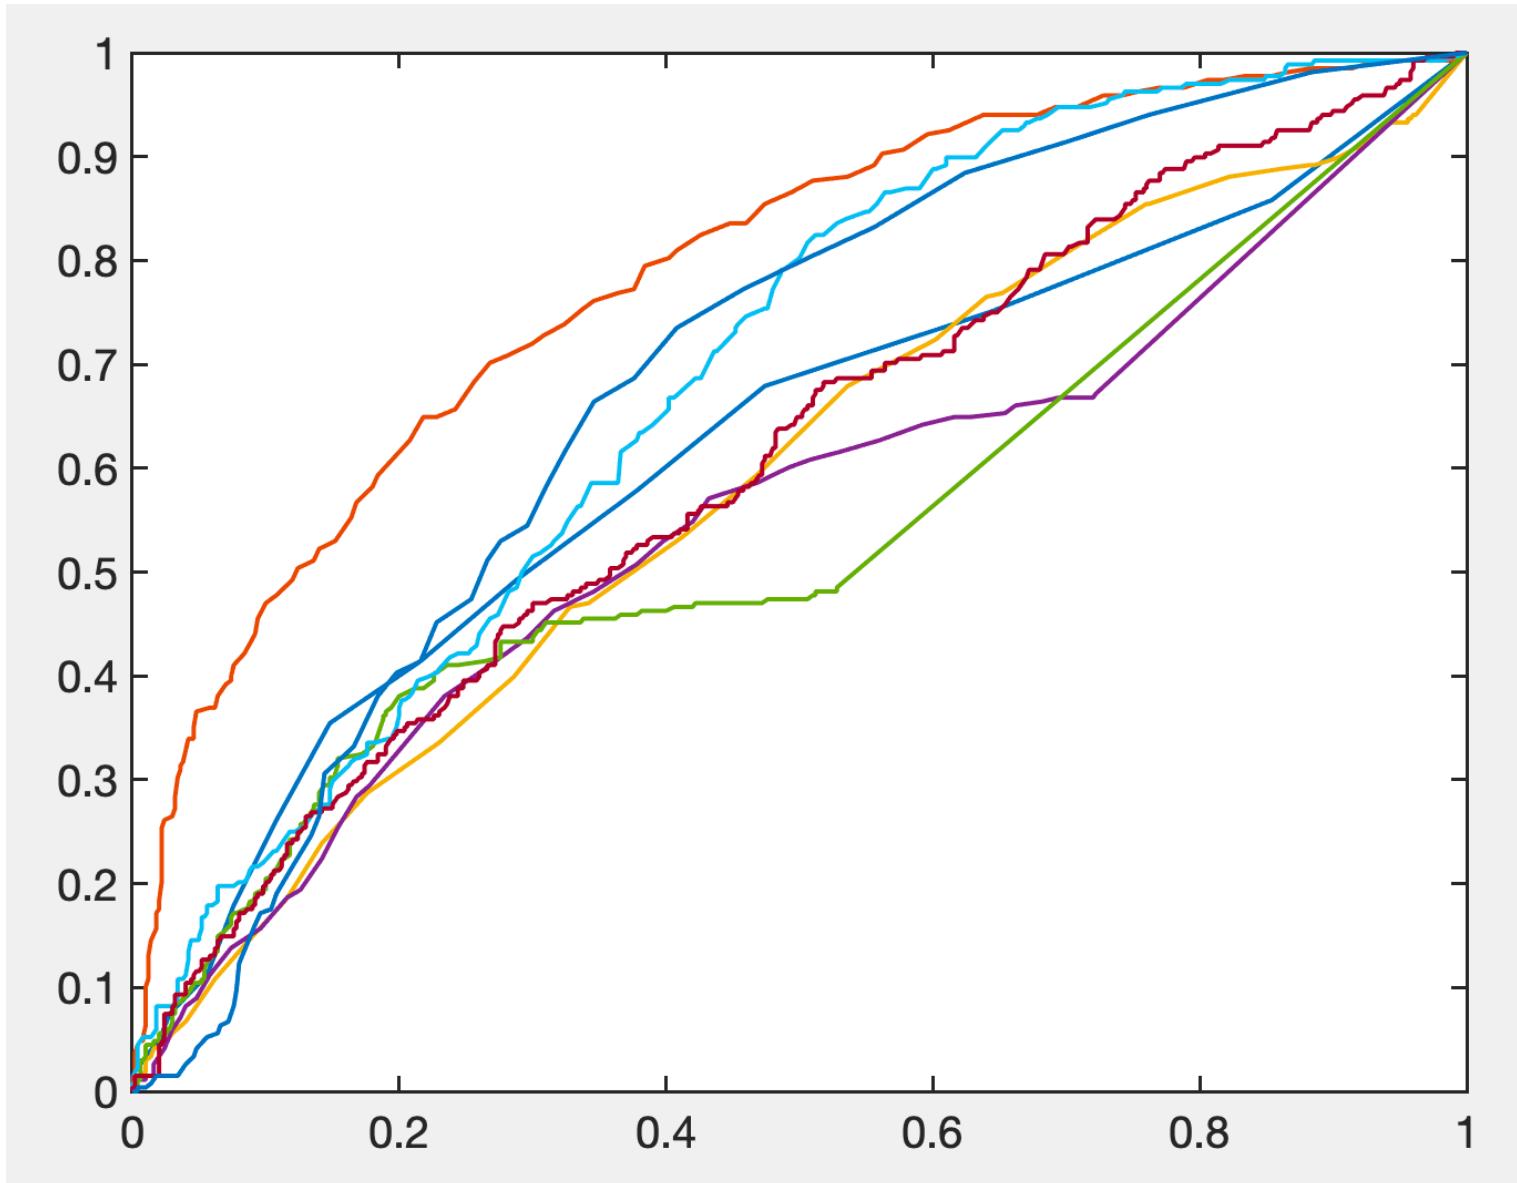
# The AUC is a rank statistic / U-statistic

- $\text{AUC} = \frac{1}{n_+ n_-} \sum_{i=1}^{n_+} \sum_{k=1}^{n_-} 1_{f(x_i) > f(x_k)}$   
(here,  $n_+$  is the number of positives,  $n_-$  is number of negatives)
  - = Fraction of correctly ranked positive-negative pairs.
  - =  $1 - (\text{fraction of misranked positive-negative pairs})$

Thus, AUC optimization is equivalent to supervised bipartite ranking.

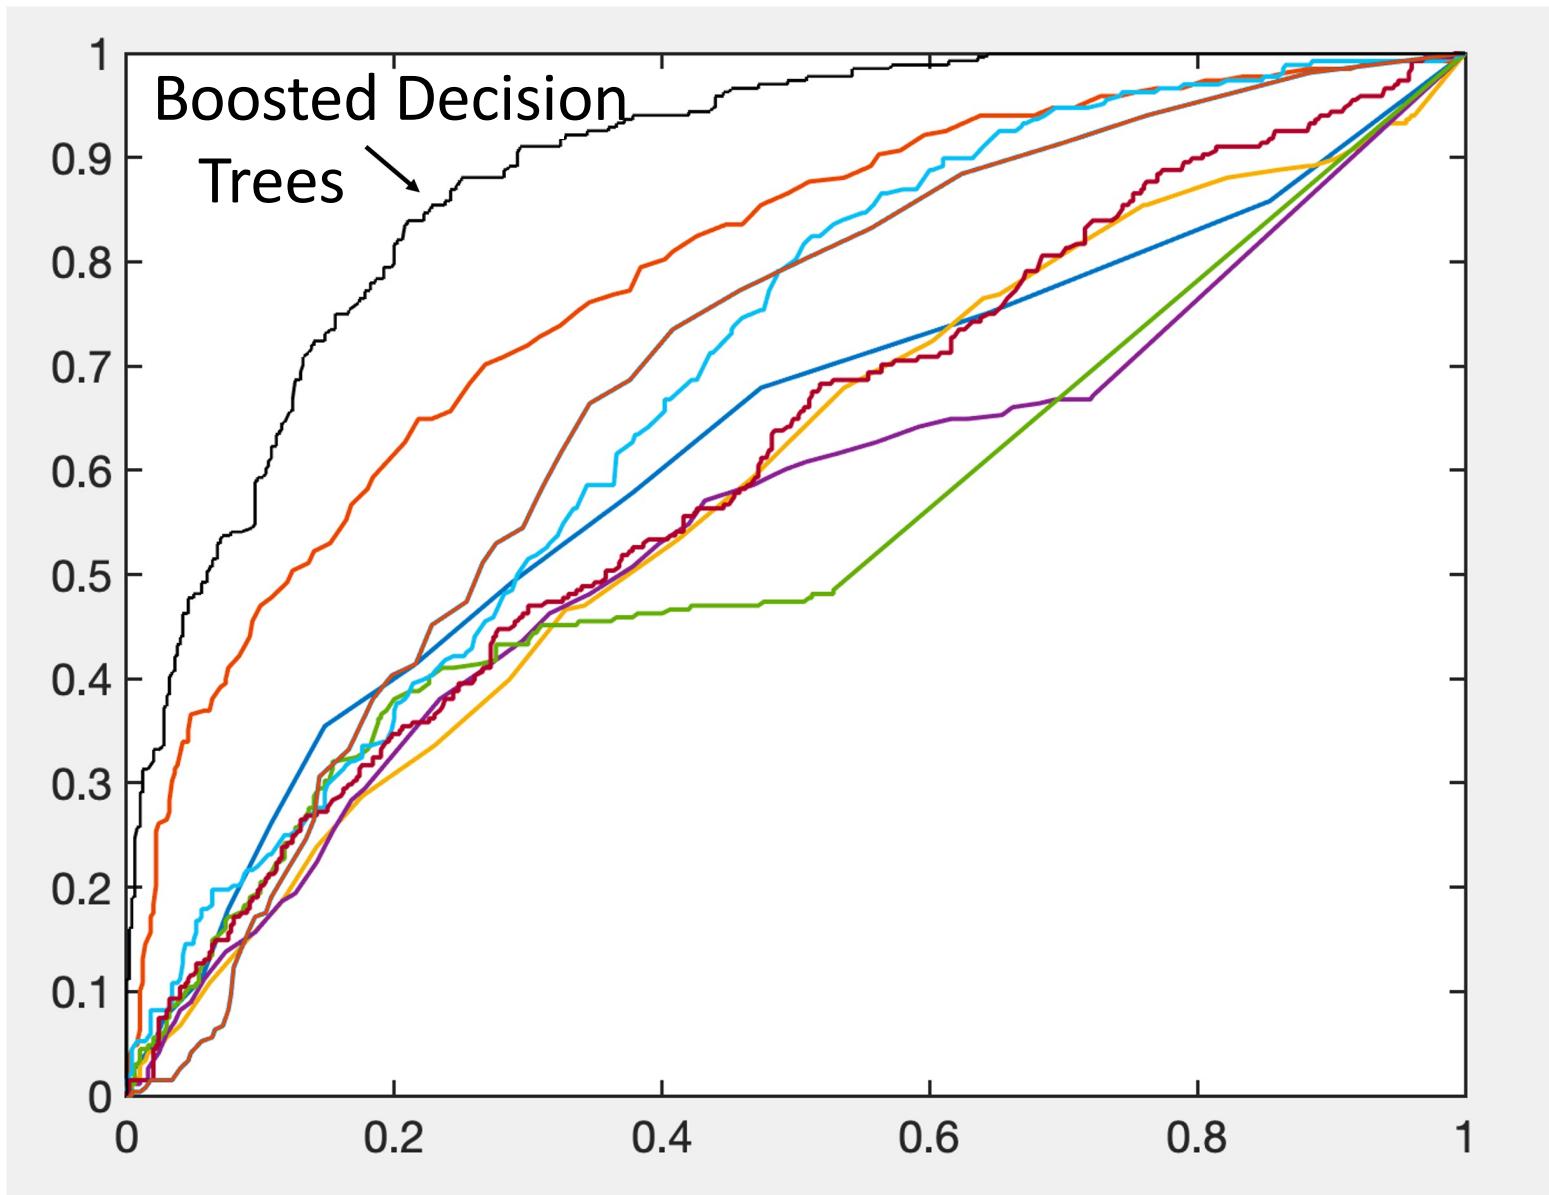
# ROC Curves for Data Exploration

Pima Indians Diabetes Data  
8 features



# ROC Curves for Data Exploration

Pima Indians Diabetes Data  
8 features



# Imbalanced Data

□

Cynthia Rudin

Machine Learning Course, Duke

# Evaluation (from earlier lecture)

Many ways to evaluate a classifier:

- Confusion matrix (TP, TN, FP, FN)
- Accuracy / misclassification error
- Precision, Recall, F1-score
- ROC curves, AUC/AUROC

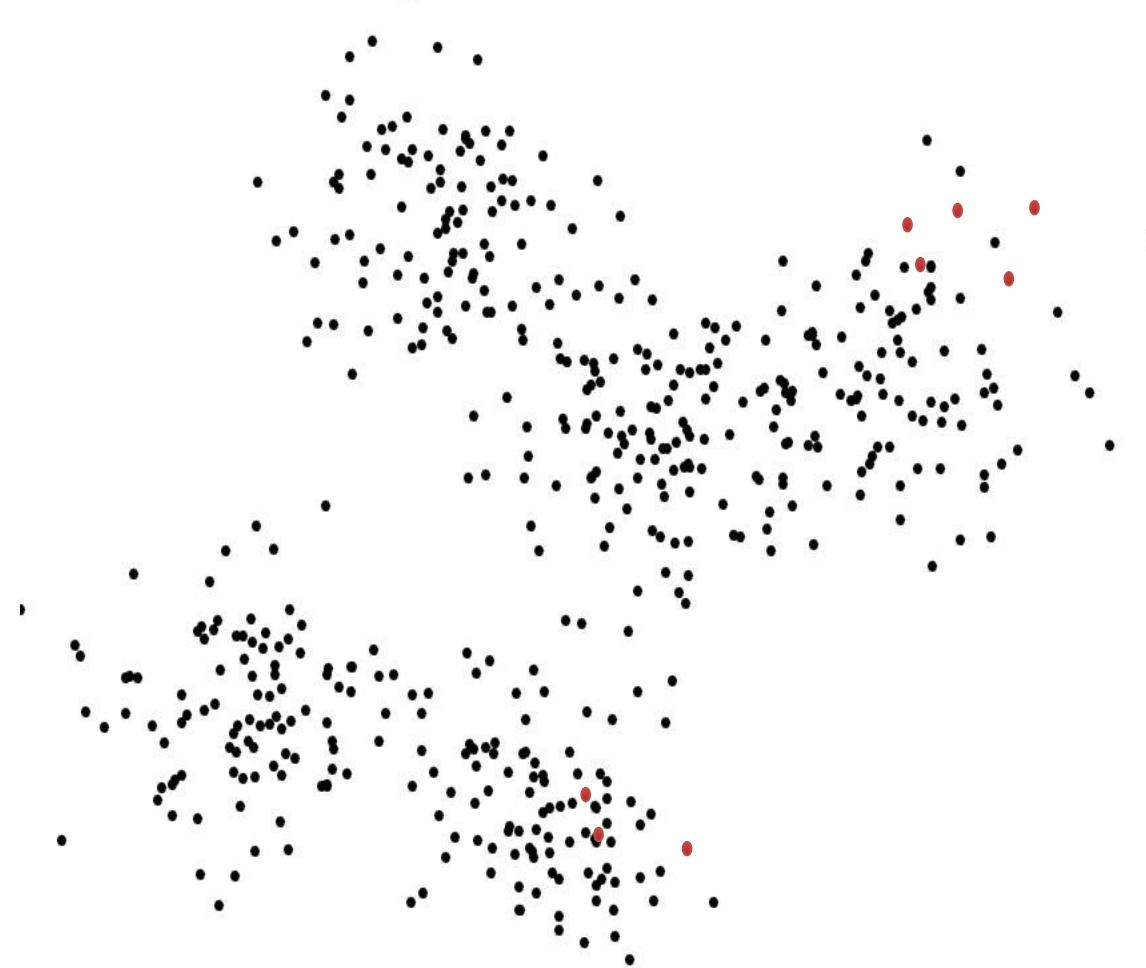
When dataset is balanced, accuracy might be ok...  
but what about when an FP is different than an FN?

# Imbalanced Data

- One of the most annoying and difficult problems in ML.

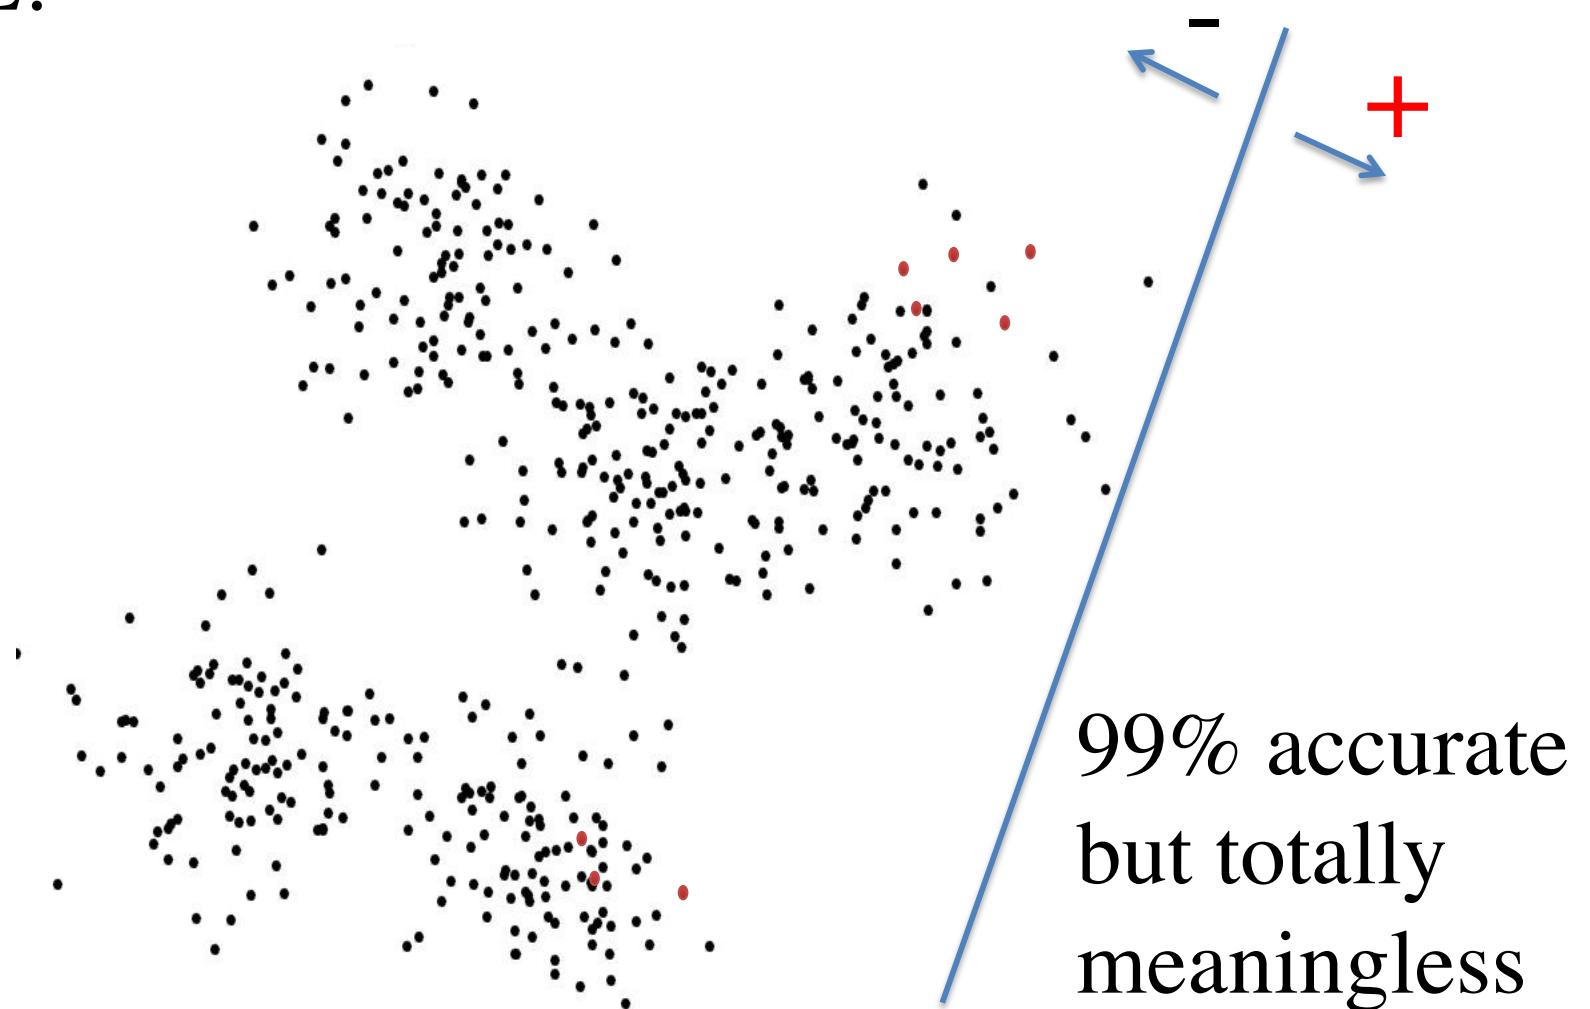
# Imbalanced Data

- One of the most annoying and difficult problems in ML.



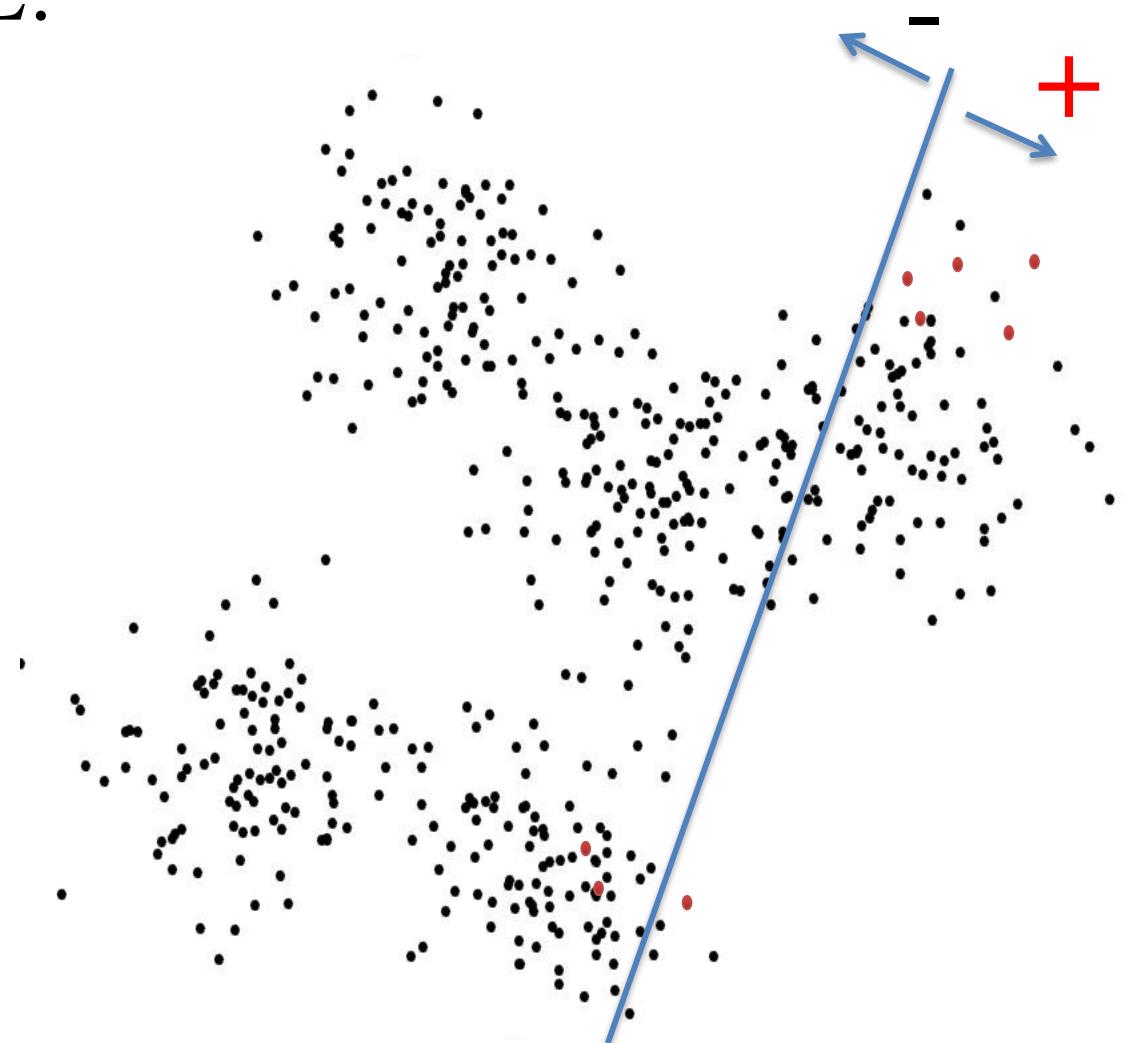
# Imbalanced Data

- One of the most annoying and difficult problems in ML.



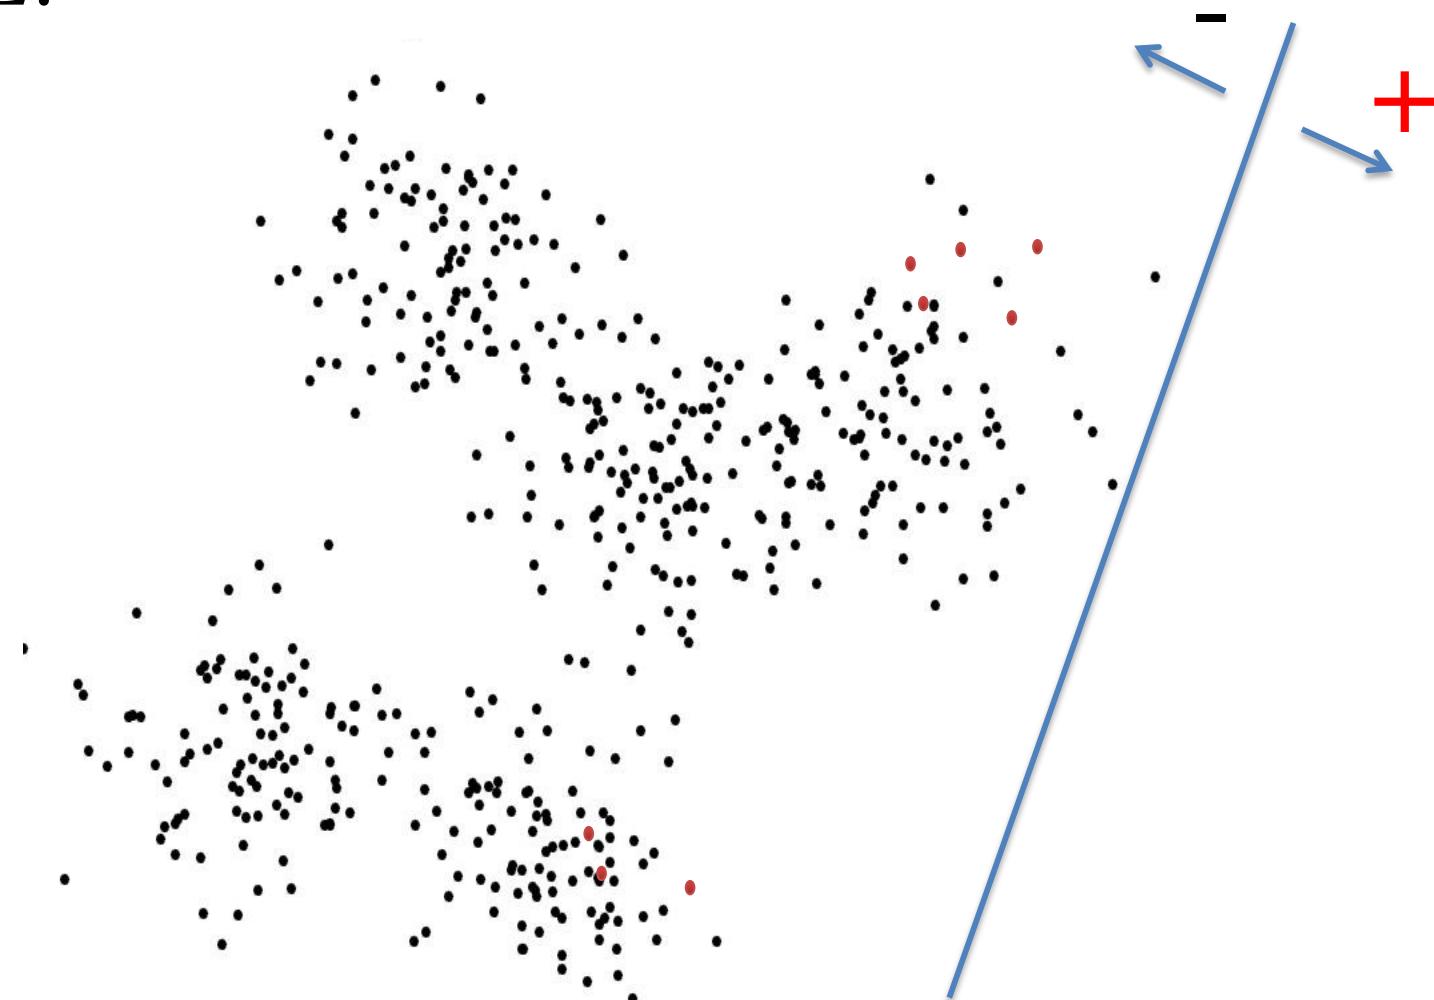
# Imbalanced Data

- One of the most annoying and difficult problems in ML.



# Imbalanced Data

- One of the most annoying and difficult problems in ML.



Starting with the standard approach:

$$\frac{1}{n} \sum_{i=1}^n \ell(y_i f(x_i)) + \text{Regularization}(f)$$

A misclassified positive is worth the same as a misclassified negative.

$$\frac{1}{n} \sum_{i=1}^n \ell(y_i f(x_i)) + \text{Regularization}(f)$$

$$\frac{1}{n} \left( \sum_{\substack{i \text{ positives} \\ i \text{ where } y_i=1}}^n \ell(y_i f(x_i)) + \sum_{\substack{k \text{ negatives} \\ k \text{ where } y_k=-1}}^n \ell(y_k f(x_k)) \right) + \text{Regularization}(f)$$

Let's weigh losses for positives and negatives differently.

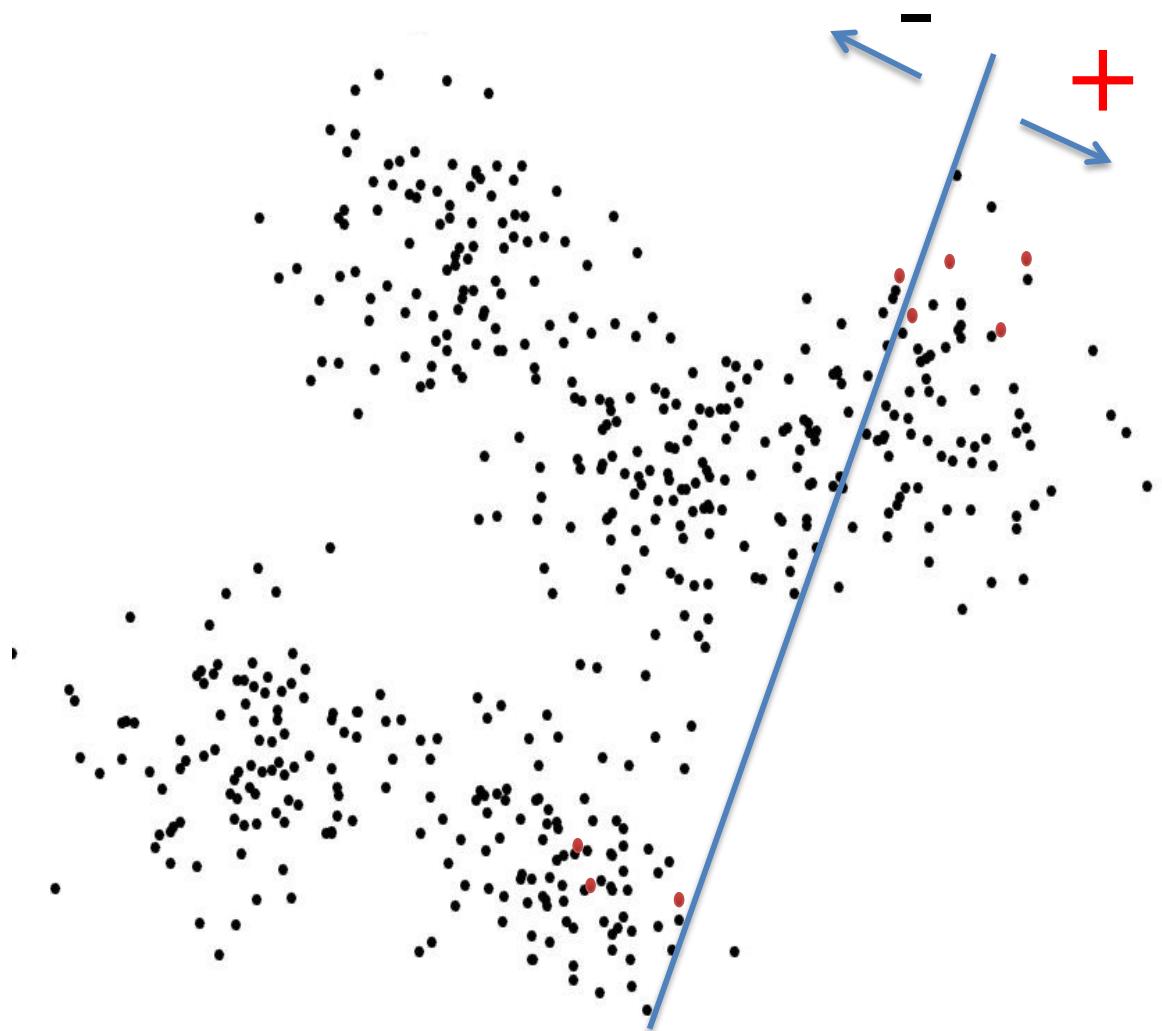
Each positive is worth C times a negative

$$\frac{1}{n} \left( C \sum_{\substack{i \text{ positives} \\ i \text{ where } y_i=1}}^n \ell(y_i f(x_i)) + \sum_{\substack{k \text{ negatives} \\ k \text{ where } y_k=-1}}^n \ell(y_k f(x_k)) \right)$$

+ Regularization( $f$ )

Let's weigh losses for positives and negatives differently.

- Perhaps the algorithm will choose this model now.



Each positive is worth C times a negative

$$\frac{1}{n} \left( C \sum_{\substack{i \text{ positives} \\ i \text{ where } y_i=1}}^n \ell(y_i f(x_i)) + \sum_{\substack{k \text{ negatives} \\ k \text{ where } y_k=-1}}^n \ell(y_k f(x_k)) \right)$$

+ Regularization( $f$ )

# Imbalanced Data

- Don't report plain accuracy.
- Adjust imbalance parameter C to obtain your ideal balance between TP/FP.
- Look at the confusion matrix to assess FP's and FN's separately.



# ROC Curves for Algorithms

□

Cynthia Rudin

Machine Learning Course, Duke

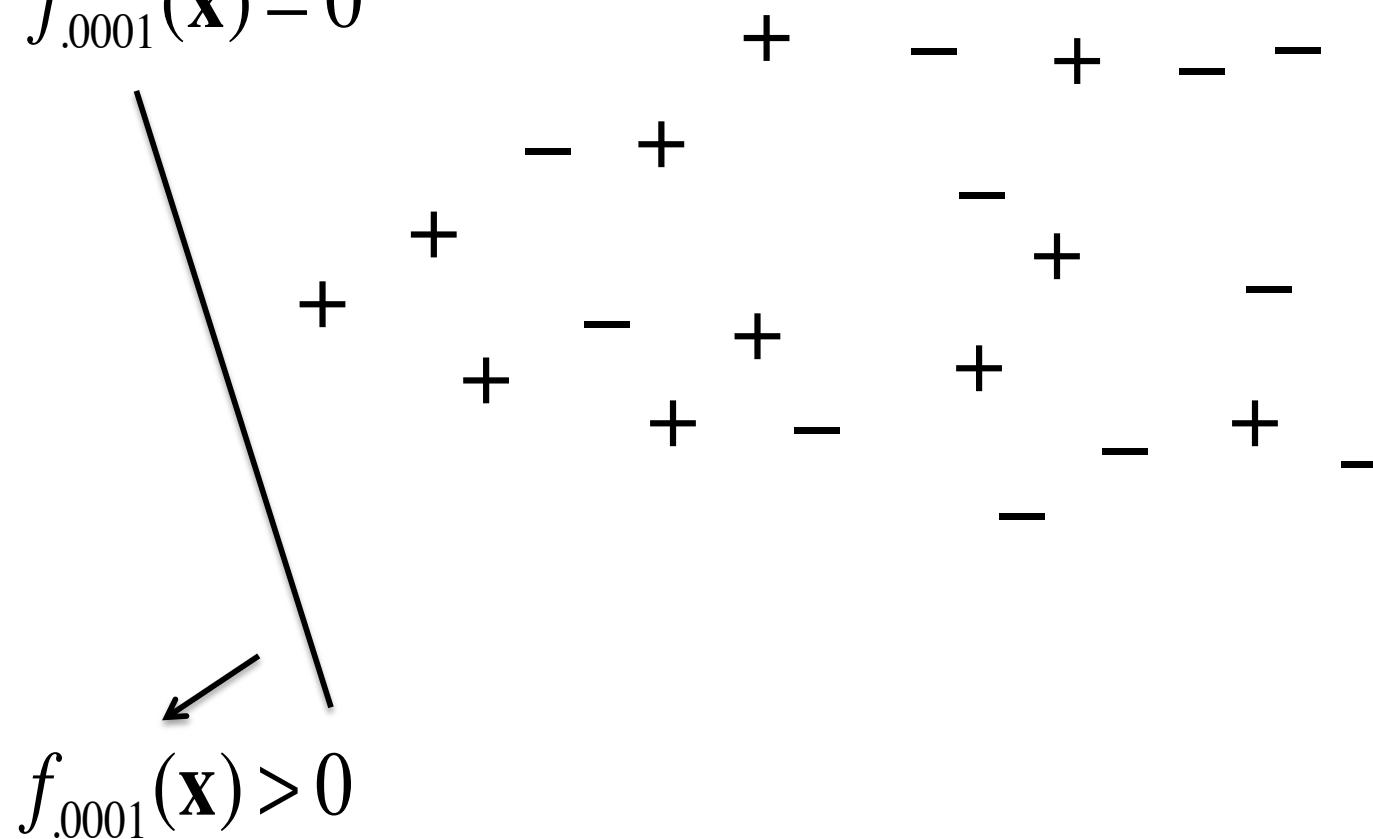
- ROC Curves can be produced in 2 ways:
  - Using a single real-valued classifier. In that case the ROC curve evaluates the **classifier**. (In earlier lectures)
  - Using a single algorithm and sweeping the imbalance parameter across the full range. In that case, the ROC curve evaluates the **algorithm**. (This lecture)

# ROC Curves for algorithms

- Run the algorithm, sweeping across C values.

$$C = .0001$$

$$f_{.0001}(\mathbf{x}) = 0$$

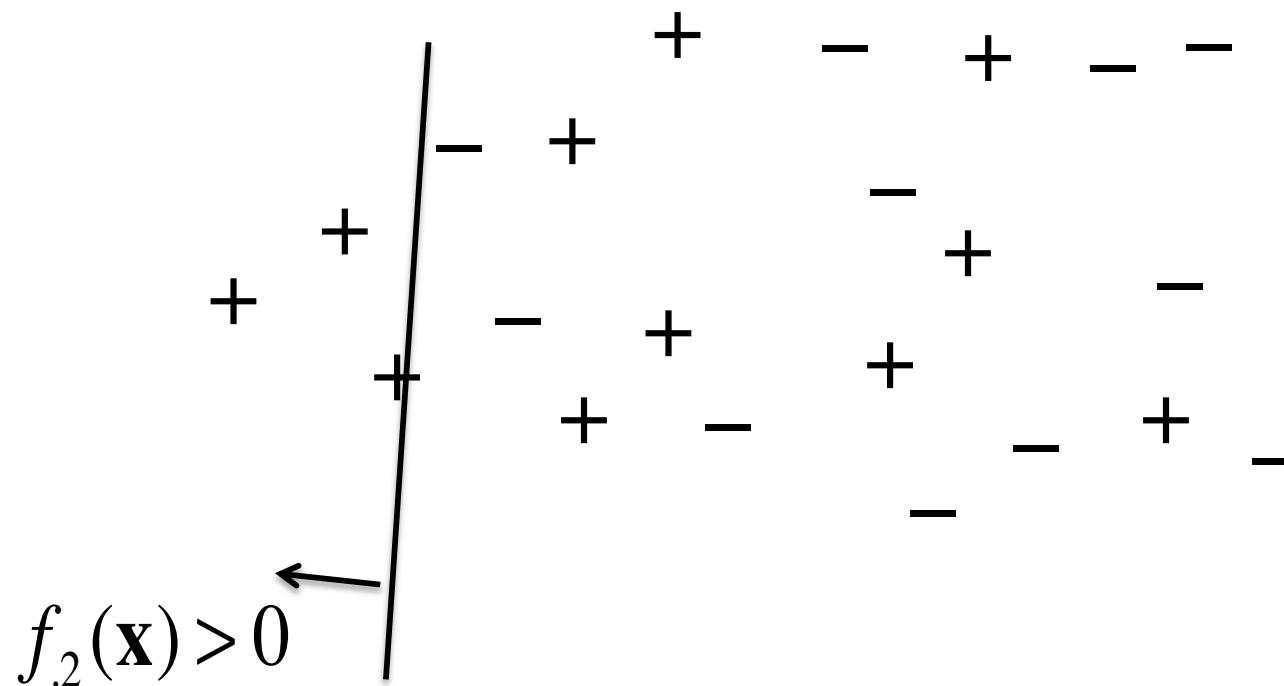


# ROC Curves for algorithms

- Run the algorithm, sweeping across C values.

$$C = .2$$

$$f_2(\mathbf{x}) = 0$$

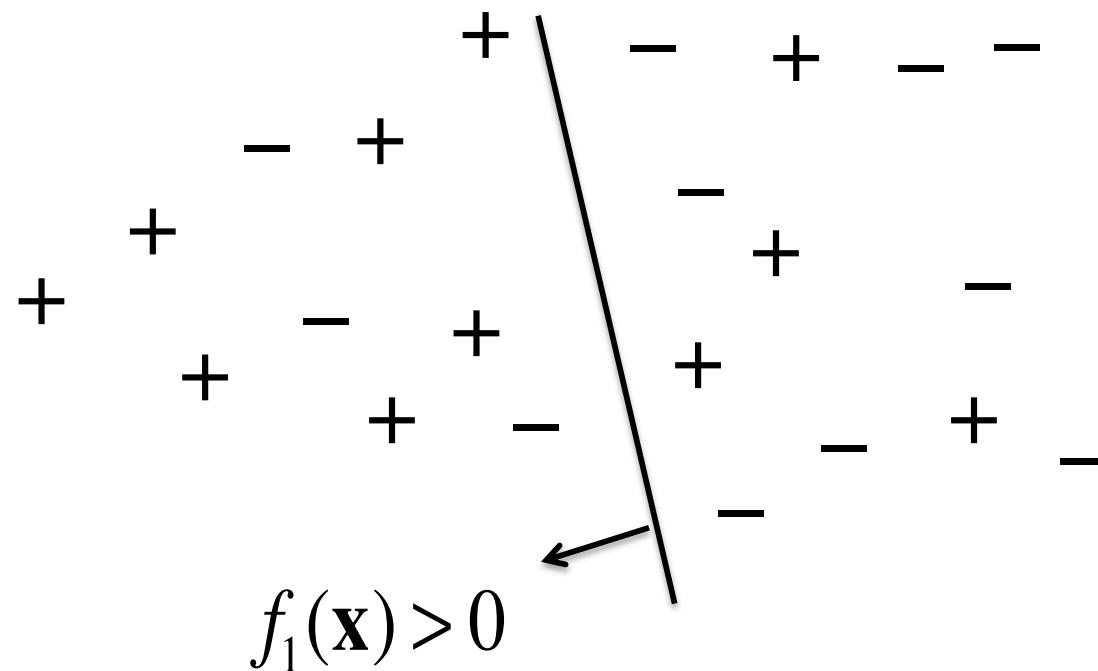


# ROC Curves for algorithms

- Run the algorithm, sweeping across C values.

$$C = 1$$

$$f_1(\mathbf{x}) = 0$$

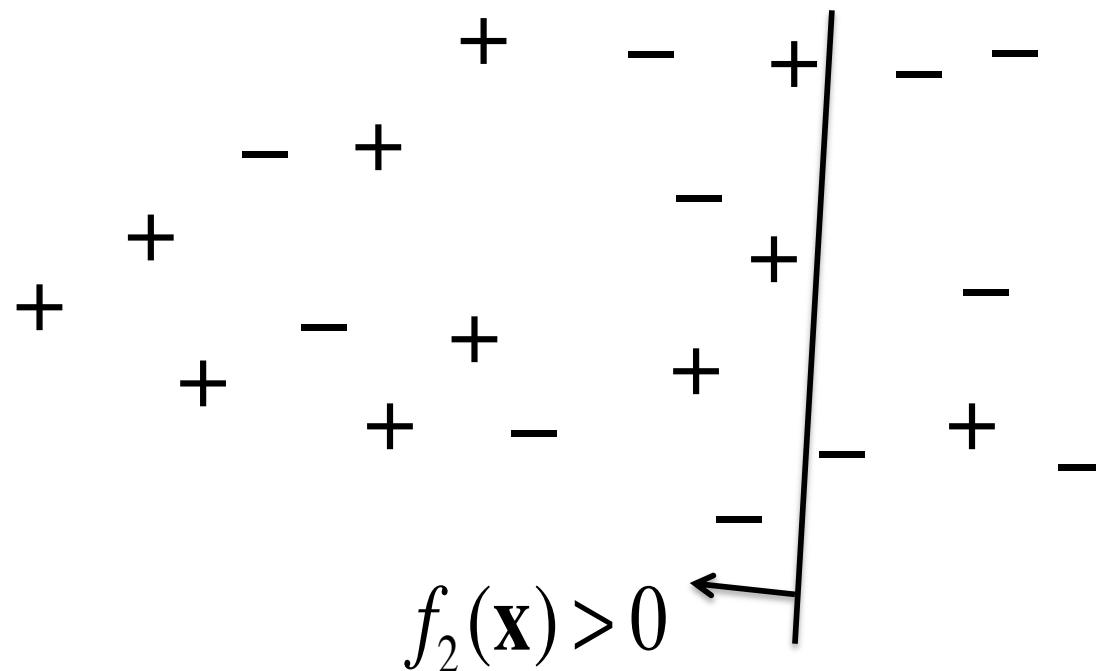


# ROC Curves for algorithms

- Run the algorithm, sweeping across C values.

$$C = 2$$

$$f_2(\mathbf{x}) = 0$$

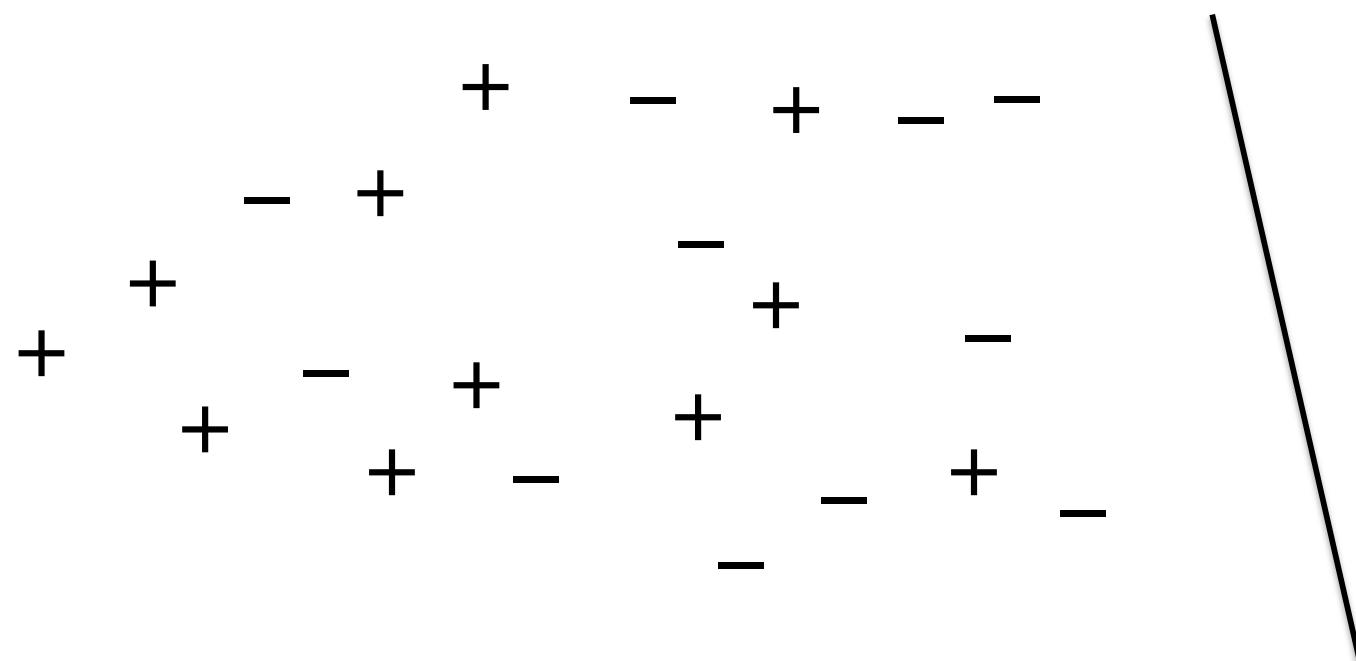


# ROC Curves for algorithms

- Run the algorithm, sweeping across C values.

$$C = 1000$$

$$f_{1000}(\mathbf{x}) = 0$$



# ROC Curves for algorithms

- For a particular False Positive Rate (FPR), what is the True Positive Rate (TPR)?



- ROC Curves can be produced in 2 ways:
  - From a single real-valued classifier. In that case the ROC curve evaluates the **classifier**.
  - From an algorithm, sweeping the imbalance parameter across the full range. In that case, the ROC curve evaluates the **algorithm**.

# ROC Curves

