



USC

# CSCI 567: Machine Learning

## Discussion – Evaluation Metrics

**Wang (Bill) Zhu**

**Nov 10**

# Overview

- Thresholding
- Confusion matrix
- Accuracy
- Precision and Recall
- ROC and AUC
- Calibration

# Thresholding

- Binary classification:  $y = f(x), y \in \{0, 1\}$
- A logistic regression model outputs a probability in  $(0, 1)$
- Choose a threshold to convert it to a binary value
- 0.5 is not always the best
- *Why? Depends on the evaluation metrics.*

# Confusion Matrix – Tumor Prediction

- Use 2x2 confusion matrix to separate out different kinds of errors
- Class-imbalanced setup: 9% of examined tumors are malignant, 91% benign

<b>True Positives (TP)</b>  Reality: Malignant ML predicted: Malignant	<b>False Positives (FP)</b>  Reality: Benign ML predicted: Malignant Type-1 Error
<b>False Negatives (FN)</b>  Reality: Malignant ML predicted: Benign Type-2 Error	<b>True Negatives (TN)</b>  Reality: Benign ML predicted: Benign

# Evaluation Metrics: Accuracy – Can Be Misleading

- Accuracy is the fraction of predictions our model got right

<b>True Positives (TP)</b>  Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b>  Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b>  Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b>  Reality: Benign ML predicted: Benign Number of TN results: 90

# Evaluation Metrics: Accuracy – Can Be Misleading

- Accuracy is the fraction of predictions our model got right

- $$\text{Accuracy} = \frac{TP+TN}{TP+FP+FN+TN}$$

<b>True Positives (TP)</b>  Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b>  Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b>  Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b>  Reality: Benign ML predicted: Benign Number of TN results: 90

# Evaluation Metrics: Accuracy – Can Be Misleading

- Accuracy is the fraction of predictions our model got right
- $$\text{Accuracy} = \frac{TP+TN}{TP+FP+FN+TN}$$
- How about a model that predicts negative all the time?

<b>True Positives (TP)</b> Reality: Malignant ML predicted: Malignant Number of TP results: <u>1</u>	<b>False Positives (FP)</b> Reality: Benign ML predicted: Malignant Number of FP results: <u>1</u>
<b>False Negatives (FN)</b> Reality: Malignant ML predicted: Benign Number of FN results: <u>8</u>	<b>True Negatives (TN)</b> Reality: Benign ML predicted: Benign Number of TN results: <u>90</u>

## Exercise (2 mins)

**In which of the following scenarios would suggest that the ML model is doing a good job?**

- A. A deadly, but curable, medical condition afflicts .01% of the population. An ML model uses symptoms as features and predicts this affliction with an accuracy of 99.99%.
- B. An expensive robotic chicken crosses a very busy road a thousand times per day. An ML model evaluates traffic patterns and predicts when this chicken can safely cross the street with an accuracy of 99.99%.
- C. In the game of roulette, a ball is dropped on a spinning wheel and eventually lands in one of 38 slots. Using visual features (the spin of the ball, the position of the wheel when the ball was dropped, the height of the ball over the wheel), an ML model can predict the slot that the ball will land in with an accuracy of 50%.



# Evaluation Metrics: Precision and Recall

- What proportion of positive identifications was actually correct?

- Precision =  $\frac{TP}{TP+FP}$

<b>True Positives (TP)</b> Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b> Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b> Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b> Reality: Benign ML predicted: Benign Number of TN results: 90

# Evaluation Metrics: Precision and Recall

- What proportion of positive identifications was actually correct?

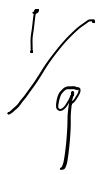
- Precision =  $\frac{TP}{TP+FP}$

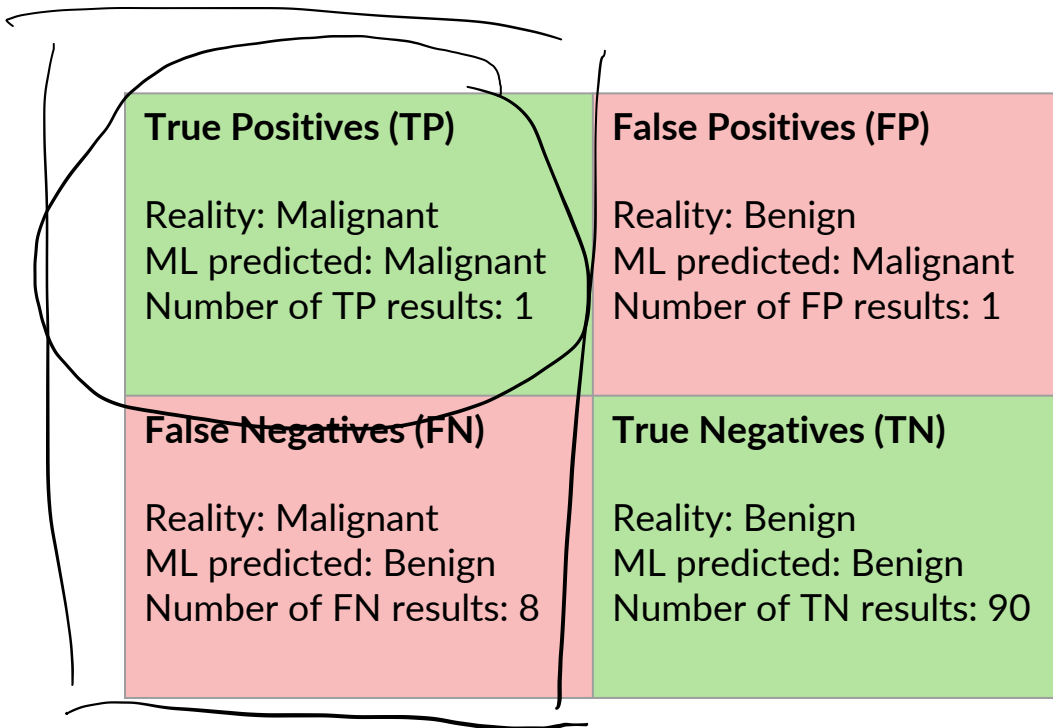
- 0.5

<b>True Positives (TP)</b>  Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b>  Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b>  Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b>  Reality: Benign ML predicted: Benign Number of TN results: 90

# Evaluation Metrics: Precision and Recall

- What proportion of actual positives was identified correctly?

- Recall =  $\frac{TP}{TP+FN}$  



<b>True Positives (TP)</b> Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b> Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b> Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b> Reality: Benign ML predicted: Benign Number of TN results: 90

# Evaluation Metrics: Precision and Recall

- What proportion of actual positives was identified correctly?

- $\text{Recall} = \frac{TP}{TP+FN}$

- 0.11

<b>True Positives (TP)</b>  Reality: Malignant ML predicted: Malignant Number of TP results: 1	<b>False Positives (FP)</b>  Reality: Benign ML predicted: Malignant Number of FP results: 1
<b>False Negatives (FN)</b>  Reality: Malignant ML predicted: Benign Number of FN results: 8	<b>True Negatives (TN)</b>  Reality: Benign ML predicted: Benign Number of TN results: 90

# Precision and Recall: A Tug of War

Hard (but possible)


- Cannot improve both at the same time by changing threshold

- $$\text{Precision} = \frac{TP}{TP+FP}, \text{ Recall} = \frac{TP}{TP+FN}$$

$\downarrow$   $\downarrow$   
 $\overline{TP}$   $\overline{FP}$

- $$F1 = \frac{1}{\text{precision}^{-1} + \text{recall}^{-1}}$$

Actual Neg



Actual pos

## Exercise (2 min)

**Consider a classification model that separates email into two categories: "spam" or "not spam." If you raise the classification threshold, what will happen to precision?**


- A. Probably increase.
- B. Probably decrease.
- C. Definitely increase.
- D. Definitely decrease.

**Consider two models—A and B—that each evaluate the same dataset. Which one of the following statements is true?**

- A. If model A has better recall than model B, then model A is better.
- B. If model A has better precision and better recall than model B, then model A is probably better.
- C. If Model A has better precision than model B, then model A is better.

# Recap: HW1 Q2.3

- What if FP has a cost of 1 but FN has a cost of 3? Decision boundary still 0.5? (Hint: insert the costs into the loss function)



**2.3 (2pts)** When  $F(x; A, k, b)$  is a CDF, we can interpret  $F(x; A, k, b)$  as the probability of  $x$  belonging to the class 1 (in a binary classification problem). Suppose we know  $F(x; A, k, b)$  and want to predict the label of a datapoint  $x$ . We need to decide on a threshold value for  $F(x; A, k, b)$ , above which we will predict the label 1 and below which we will predict  $-1$ . Show that setting the threshold to be  $F(x; A, k, b) \geq 1/2$  minimizes the classification error.

**Ans.** Expected error for a given threshold ( $T$ ) for a given  $x$  is

$$\begin{aligned} \text{Expected loss given } x &= p(y = 0|x)I[F(x) \geq T] + p(y = 1|x)I[F(x) < T] \\ &= \underbrace{[1 - F(x)]I[F(x) \geq T]}_{\text{FP}} + \underbrace{[F(x)]I[F(x) < T]}_{\text{FN}} \end{aligned}$$

The above is minimized at  $T = 0.5$ .

$$\Rightarrow \text{FP} + \text{FN} =$$

## Recap: HW1 Q2.3

- What if FP has a cost of 1 but FN has a cost of 3? Decision boundary still 0.5? (Hint: insert the costs into the loss function)
- $1 - T = 3T, T = 0.25$

$$FP = 3 \quad FN = 1 \quad \Rightarrow 75$$

**2.3 (2pts)** When  $F(x; A, k, b)$  is a CDF, we can interpret  $F(x; A, k, b)$  as the probability of  $x$  belonging to the class 1 (in a binary classification problem). Suppose we know  $F(x; A, k, b)$  and want to predict the label of a datapoint  $x$ . We need to decide on a threshold value for  $F(x; A, k, b)$ , above which we will predict the label 1 and below which we will predict  $-1$ . Show that setting the threshold to be  $F(x; A, k, b) \geq 1/2$  minimizes the classification error.

**Ans.** Expected error for a given threshold ( $T$ ) for a given  $x$  is


$$\begin{aligned} \text{Expected loss given } x &= p(y = 0|x)I[F(x) \geq T] + p(y = 1|x)I[F(x) < T] \\ &= [1 - F(x)]I[F(x) \geq T] + [F(x)]I[F(x) < T] \end{aligned}$$

The above is minimized at  $T = 0.5$ .



# Application: Contrast set

Q: Is there *at least* 1 image with exactly 2 dark bottles on a counter.




Expected A: True

Vision-  
Language  
Model

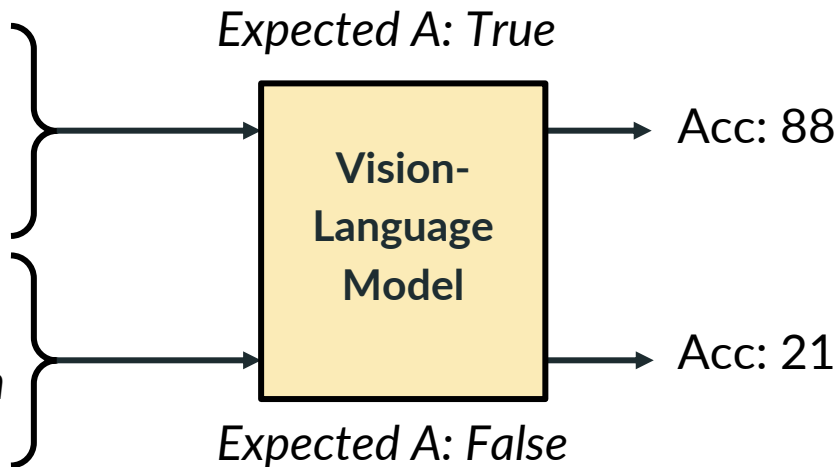
Acc: 88

# Application: Contrast set

Q: Is there **at least** 1 image with exactly 2 dark bottles on a counter.




Contrast Q: Is there **less than** 1 image with exactly 2 dark bottles on a counter.



# Application: Contrast set

Q: Is there **at least** 1 image with exactly 2 dark bottles on a counter.



Contrast Q: Is there **less than** 1 image with exactly 2 dark bottles on a counter.

Expected A: True

Vision-  
Language  
Model

Acc: 88


Expected A: False

Acc: 21

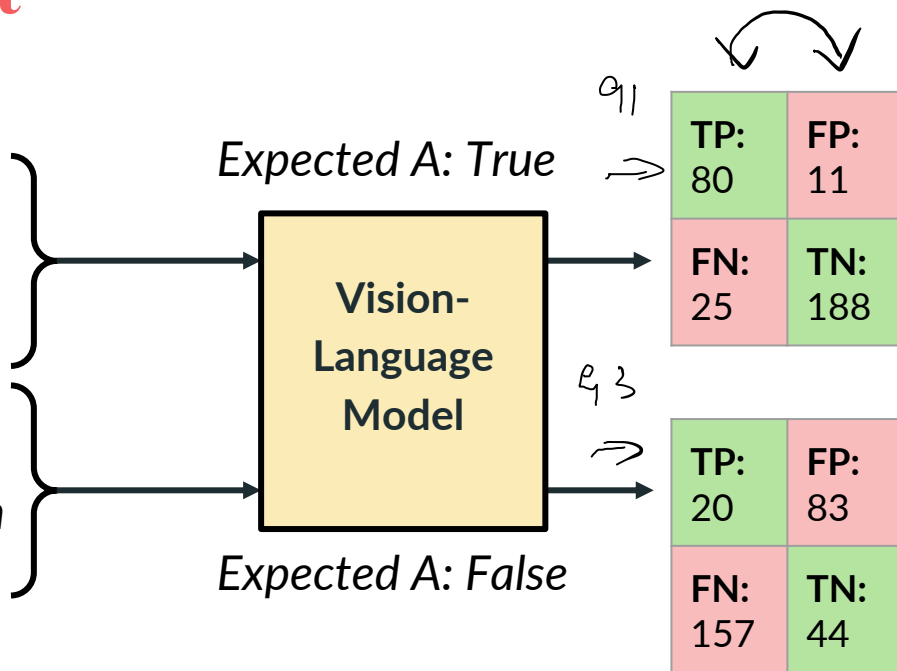
What does this tell us? Contrast Qs are hard? They have low correlation/grounding on images? The VL model is bad?

# Application: Contrast set

Q: Is there **at least** 1 image with exactly 2 dark bottles on a counter.




Contrast Q: Is there **less than** 1 image with exactly 2 dark bottles on a counter.

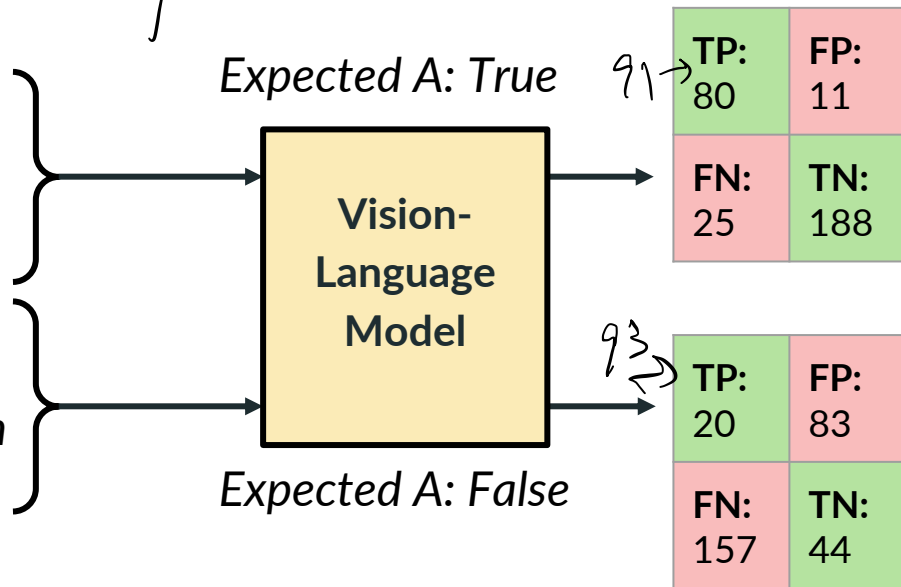



# Application: Contrast set

Q: Is there **at least** 1 image with exactly 2 dark bottles on a counter.



Contrast Q: Is there **less than** 1 image with exactly 2 dark bottles on a counter.



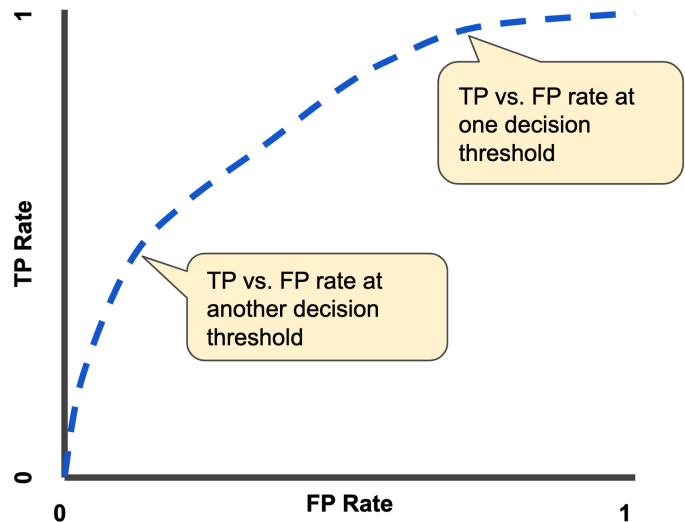
What does this tell us? (Probably) the model is over-stable on its prediction.

# A ROC Curve

- Each point is the TP and FP rate at one decision threshold



- $$\text{TPR (Recall)} = \frac{TP}{TP+FN}$$

- $$\text{FPR} = \frac{FP}{FP+TN}$$



# Evaluation Metrics: AUC (AUROC)


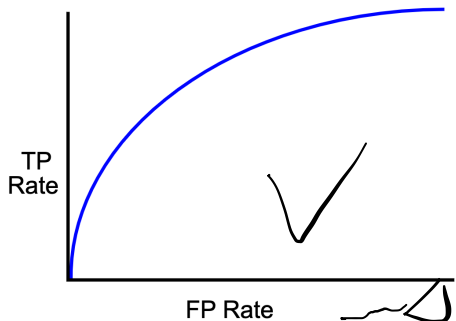

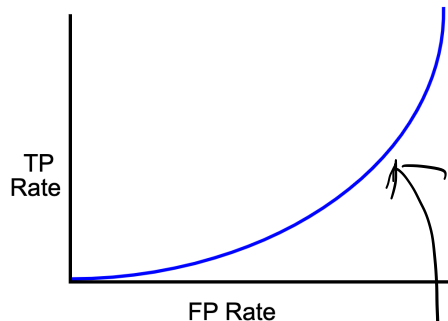

- AUC: “Area under the ROC Curve”
- The probability that the model ranks a random positive example more highly than a random negative example
- Independent of the threshold



- Actual Negative
- Actual Positive

## Exercise (2 mins)

Which of the following ROC curves produce AUC values greater than 0.5?



✓  
5 big



# Calibration

- Prediction bias = average of prediction - average of labels *≠ bias in  $(W, b)$*
- Zero bias alone does not mean everything is perfect
- It's a great sanity check: incomplete features? noisy data? buggy pipeline?
- Don't fix bias with a calibration layer, fix it in the model

