

# CS 260: Foundations of Data Science

Prof. Thao Nguyen

Fall 2025



HAVERFORD  
COLLEGE

# Admin

- Sit somewhere new
- Lab 3 due tonight at midnight
- Lab 4 posted (due next Tuesday)
  - pair-programming required, **different partner**

## CAMPUS READ 2025



*A student-only  
book signing will  
be held in  
Lutnick 200 from  
6:15–6:45!*

A special book talk and community discussion with marine biologist, policy expert, and writer, Dr. Ayana Elizabeth Johnson.

[RSVP here](#)



**Tuesday,  
September 30**



**7:00 p.m.**



**Roberts Hall,  
Marshall  
Auditorium**

# Outline for today

- Recap SGD (stochastic gradient descent)
- Introduction to classification
  - Decision tree models
  - Probabilistic interpretation
- Evaluation Metrics
  - Confusion matrices
  - Precision and recall
  - ROC curves

# Outline for today

- Recap SGD (stochastic gradient descent)
- Introduction to classification
  - Decision tree models
  - Probabilistic interpretation
- Evaluation Metrics
  - Confusion matrices
  - Precision and recall
  - ROC curves

# Stochastic Gradient Descent for Linear Regression

Key Idea: take the derivative of **one datapoint** at a time and use that to update  $w$

```
set  $w = 0$  vector
```

```
while cost  $J(w)$  still changing (or max iter reached):
```

```
    shuffle data points
```

```
    for  $i = 1 \dots n$ :
```

```
         $w \leftarrow w - \text{alpha}(\text{derivative of } J(w) \text{ wrt } x_i)$ 
```

```
    store  $J(w)$ 
```

# Mini-quiz for linear regression

For each of the following terms/descriptions, write out the corresponding equation:

- 1) Linear regression **model**
- 2) Linear regression **cost function**
- 3) **Gradient** of cost function wrt one datapoint
- 4) **Gradient descent** weight vector update

# Mini-quiz for linear regression

①  $\hat{y} = \vec{w} \cdot \vec{x}$

②  $J(\vec{w}) = \frac{1}{2} \sum_{i=1}^n (\underbrace{y_i}_{\text{truth}} - \underbrace{\vec{w} \cdot \vec{x}_i}_{\text{pred}})^2$

③  $\nabla_{\vec{x}_i} J(\vec{w}) = (\vec{w} \cdot \vec{x}_i - y_i) \vec{x}_i$   
"derivative"  $\nwarrow$   $\searrow$

④  $\vec{w} \leftarrow \vec{w} - \alpha (\nabla_{\vec{x}_i} J(\vec{w}))$



# Outline for today

- Recap SGD (stochastic gradient descent)
- Introduction to classification
  - Decision tree models
  - Probabilistic interpretation
- Evaluation Metrics
  - Confusion matrices
  - Precision and recall
  - ROC curves

# Binary classification examples

- Transactions that indicate credit card fraud
- Accounts that are bots
- Detecting which scans show tumors
- Prenatal test for Down's Syndrome
- Finding genes under natural selection
- Regions of the environment that contains the object the robot is searching for


In all these examples, we are trying to find unusual items ("needle in a haystack") -- we call these *positives*

# Introduction to Classification

## Introduction to Classification

Single feature models  $\vec{y}$

$X_{\text{fever}}$




training data

+ disease  
- no disease

$$P_{\text{pos}} = \frac{6}{8}$$

$P_{\text{pos}} = \text{prob of positive classification}$

model: decision tree with 1 feature ("stumps")



# Introduction to Classification

new idea: use probabilities  
to classify test example

$$\vec{X}_{\text{test}} = [\dots \overset{\text{fever}}{N} \dots]^T$$

$$\text{threshold} = 0.5$$

$$\hat{Y}_{\text{test}} = -$$

(no disease)

$$\text{threshold} = 0.25$$

$$\hat{Y}_{\text{test}} = +$$


disease

$$p_{\text{pos}} \geq \text{threshold} \rightarrow \text{classify } +$$

# Handout 7



# Handout 7



# Outline for today

- Recap SGD (stochastic gradient descent)
- Introduction to classification
  - Decision tree models
  - Probabilistic interpretation
- **Evaluation Metrics**
  - Confusion matrices
  - Precision and recall
  - ROC curves

# Goals of Evaluation

- Think about what metrics are important for the problem at hand
- Compare different methods or models on the same problem
- Common set of tools that other researchers/users can understand



# Training and Testing

(high-level idea)

- **Separate** data into “**train**” and “**test**”
  - $n$  = num training examples
  - $m$  = num testing examples
- **Fit** (create) the model using **training data**
  - e.g., sea\_ice\_1979-2012.csv
- **Evaluate** the model using **testing data**
  - e.g., sea\_ice\_2013-2020.csv

# Confusion Matrices

accuracy: (classification)

$$acc = \frac{\# \text{ correct}}{m} = \frac{1}{m} \sum_{i=1}^m \mathbb{1}(y_i = \hat{y}_i)$$

Note: all the same model, different thresholds!

Confusion matrix

test data.

80

negatives, 20 positives

m = 100

(A)

|   |             |    |
|---|-------------|----|
|   | <u>pred</u> |    |
|   | -           | +  |
| - | 65          | 15 |
| + | 7           | 13 |

→ 80  
→ 20

(B)

|   |             |    |
|---|-------------|----|
|   | <u>pred</u> |    |
|   | -           | +  |
| - | 50          | 30 |
| + | 1           | 19 |

→ 80  
→ 20

(C)

|   |             |   |
|---|-------------|---|
|   | <u>pred</u> |   |
|   | -           | + |
| - | 76          | 4 |
| + | 11          | 9 |

$$acc = \frac{65 + 13}{100} = \frac{78}{100}$$

acc = .69  
low thresh

high thresh

# Confusion Matrices

|            |          | Predicted class     |                     |
|------------|----------|---------------------|---------------------|
|            |          | Negative            | Positive            |
| True class | Negative | True negative (TN)  | False positive (FP) |
|            | Positive | False negative (FN) | True positive (TP)  |


# Confusion Matrices

|            |          | Predicted class                |  |                                    |
|------------|----------|--------------------------------|--|------------------------------------|
|            |          | Negative                       | Positive                                 |                                    |
| True class | Negative | True negative (TN)             | False positive (FP)<br>“false alarm”     | N (total number of true negatives) |
|            | Positive | False negative (FN)<br>“miss”  | True positive (TP)                       | P (total number of true positives) |
|            |          | N* (what we said was negative) | P* (what we said was positive “flagged”) |                                    |

# Confusion Matrices

|            |          | Predicted class                 |  |   |
|------------|----------|---------------------------------|--|---|
|            |          | Negative                        | Positive                               |   |
| True class | Negative | True negative (TN)<br>✓         | False positive (FP)<br>“false alarm” ✗ | N |
|            | Positive | False negative (FN)<br>“miss” ✗ | True positive (TP)<br>✓                | P |
|            |          | N*                              | p*                                     |   |

# Confusion Matrices

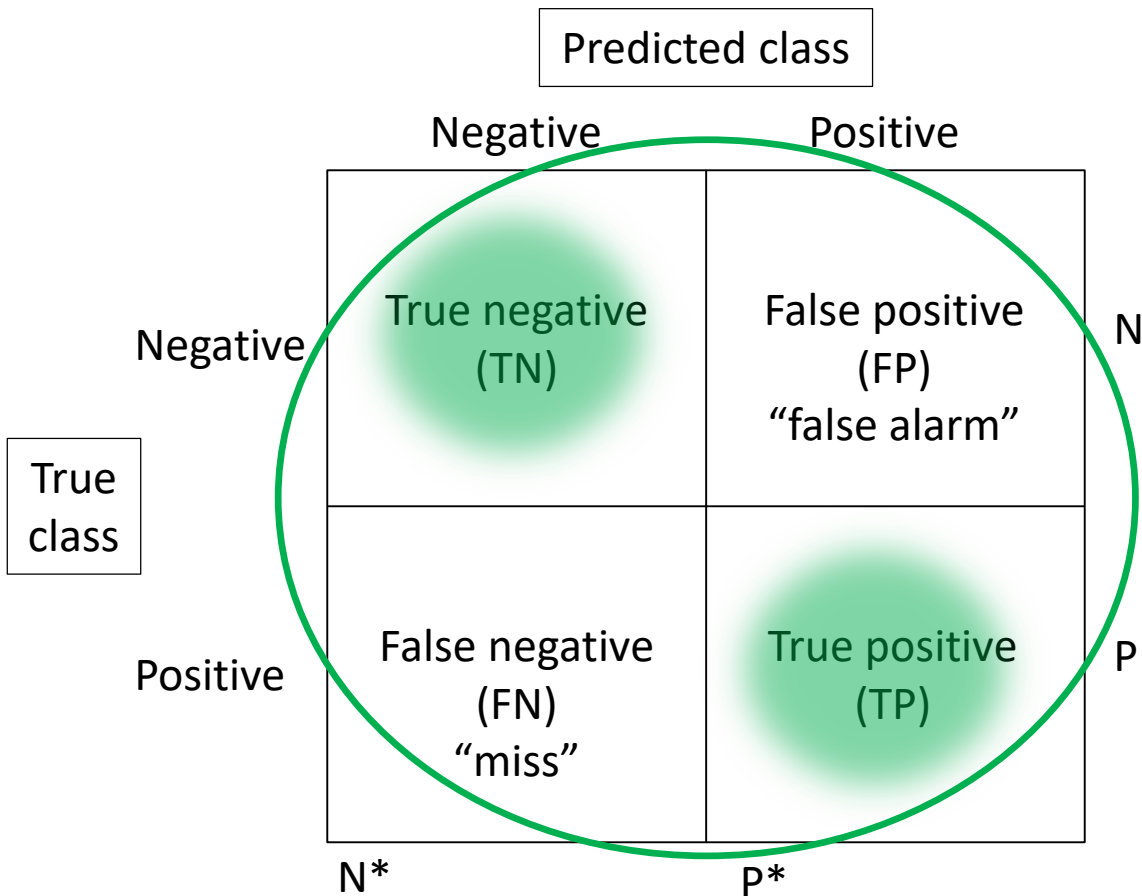


Error:

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

$$= (FN+FP)/(N+P)$$

# Confusion Matrices



Accuracy = 1-Error:

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

$$= (TN+TP)/(N+P)$$

# Confusion Matrices

|            |          | Predicted class               |                                      |   |
|------------|----------|-------------------------------|--------------------------------------|---|
|            |          | Negative                      | Positive                             |   |
| True class | Negative | True negative (TN)            | False positive (FP)<br>"false alarm" | N |
|            | Positive | False negative (FN)<br>"miss" | True positive (TP)                   | P |
|            |          | N*                            | p*                                   |   |

Precision:

$$TP/(FP+TP) = TP/P^*$$



# Confusion Matrices

|            |          | Predicted class               |                                      |   |
|------------|----------|-------------------------------|--------------------------------------|---|
|            |          | Negative                      | Positive                             |   |
| True class | Negative | True negative (TN)            | False positive (FP)<br>"false alarm" | N |
|            | Positive | False negative (FN)<br>"miss" | True positive (TP)                   | P |
|            |          | N*                            | p*                                   |   |

Recall  
(True Positive Rate):

$$TP/(FN+TP) = TP/P$$

# Confusion Matrices

|            |          | Predicted class               |                                      |   |
|------------|----------|-------------------------------|--------------------------------------|---|
|            |          | Negative                      | Positive                             |   |
| True class | Negative | True negative (TN)            | False positive (FP)<br>"false alarm" | N |
|            | Positive | False negative (FN)<br>"miss" | True positive (TP)                   | P |
|            |          | N*                            | p*                                   |   |

False Positive Rate:

$$FP/(TN+FP) = FP/N$$

# Precision and Recall


- Precision: of all the “flagged” examples, which ones are actually relevant (i.e., positive)?

(Purity)

- Recall: of all the relevant results, which ones did I actually return?

(Completeness)


# Precision and Recall



$P=6$  (number of images that are actually Sara)

- Precision?
- Recall?


# Precision and Recall



$P=6$  (number of images that are actually Sara)

- Precision =  $TP/(FP+TP) = 3/5$
- Recall?

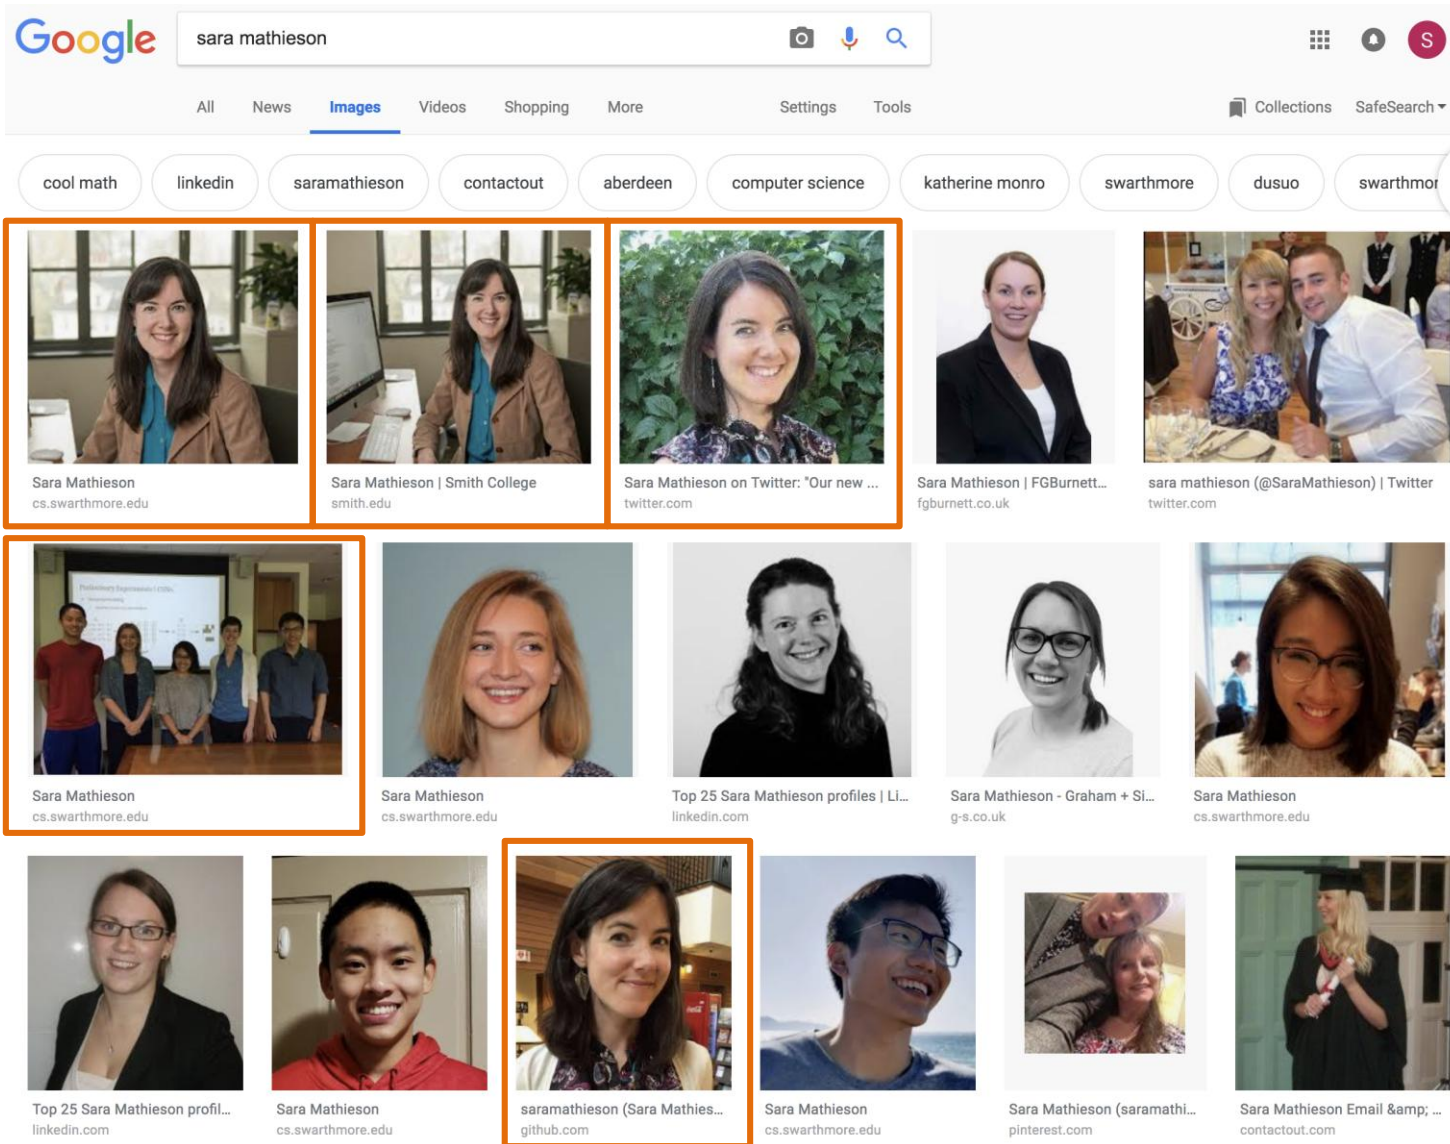
# Precision and Recall



$P=6$  (number of images that are actually Sara)

- Precision =  $TP/(FP+TP) = 3/5$
- Recall =  $TP/(FN+TP) = 3/6$

# Precision and Recall



$P=6$  (number of images that are actually Sara)

- Precision =  $5/16$
- Recall =  $5/6$



# Outline for today

- Recap SGD (stochastic gradient descent)
- Introduction to classification
  - Decision tree models
  - Probabilistic interpretation
- Evaluation Metrics
  - Confusion matrices
  - Precision and recall
  - ROC curves *Next time!*