Lecture 28

Prediction IV

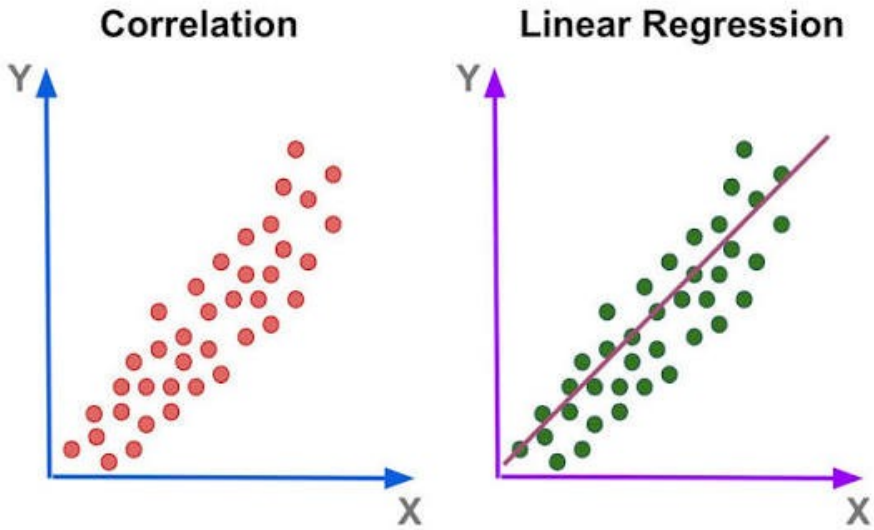

Announcements

- Course project directed questions are out!

- Office hours every Tuesday (noon to 4pm) and Friday (11am to 1pm)
  - Link to book appointments is on the course website

Today

Bias-Variance Tradeoff

Evaluating a classifier

Accessing Model Performance

There are several metrics for accessing model performance

- We'll talk about two: squared error (regression) and accuracy (classification)

Goal: want to know how good predictions are when we predict on new data

Splitting Data

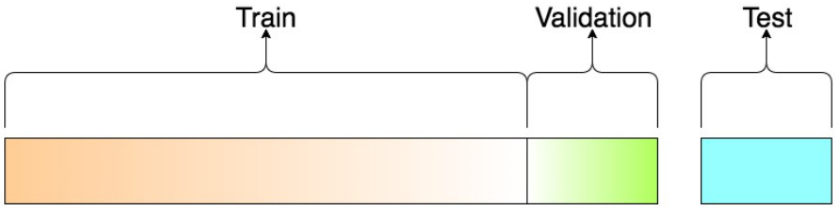

We typically split out dataset into train, validate, and test

Training Data

- Data we use to train a model (e.g. solving for parameters of linear regression)

Validation Data

- Data we use for model selection (when we're choosing between multiple models)

Test Data

- Data used to measure performance on unseen data (how well the model generalizes)
- The model never sees this data during training

What is good?

Want: a model that gives high test data performance (low test error)

Why not settle for a model that has high training data performance?

- i.e. low error on the training data
- There's no guarantee that the method with the highest training data performance will have the highest test data performance
- In many cases, it will not (*overfitting*)

Model Flexibility

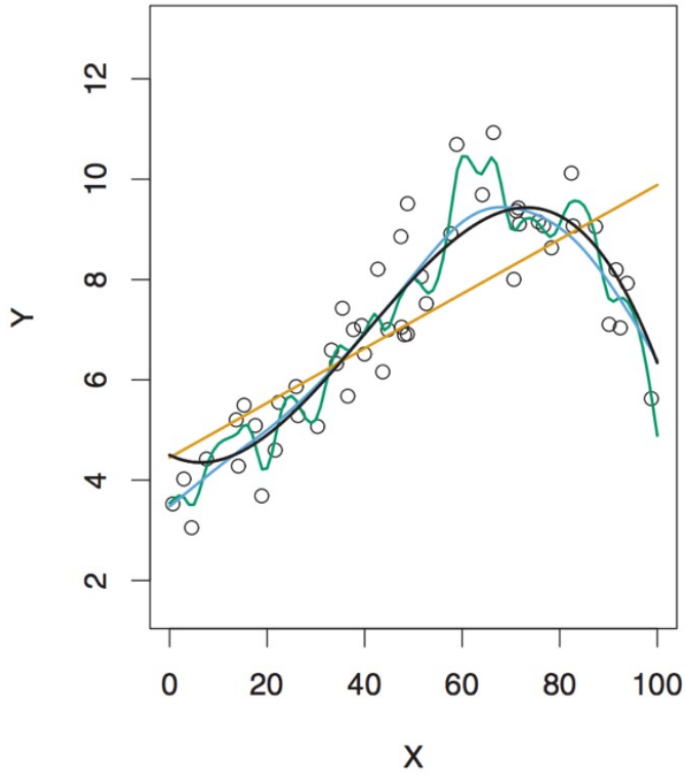

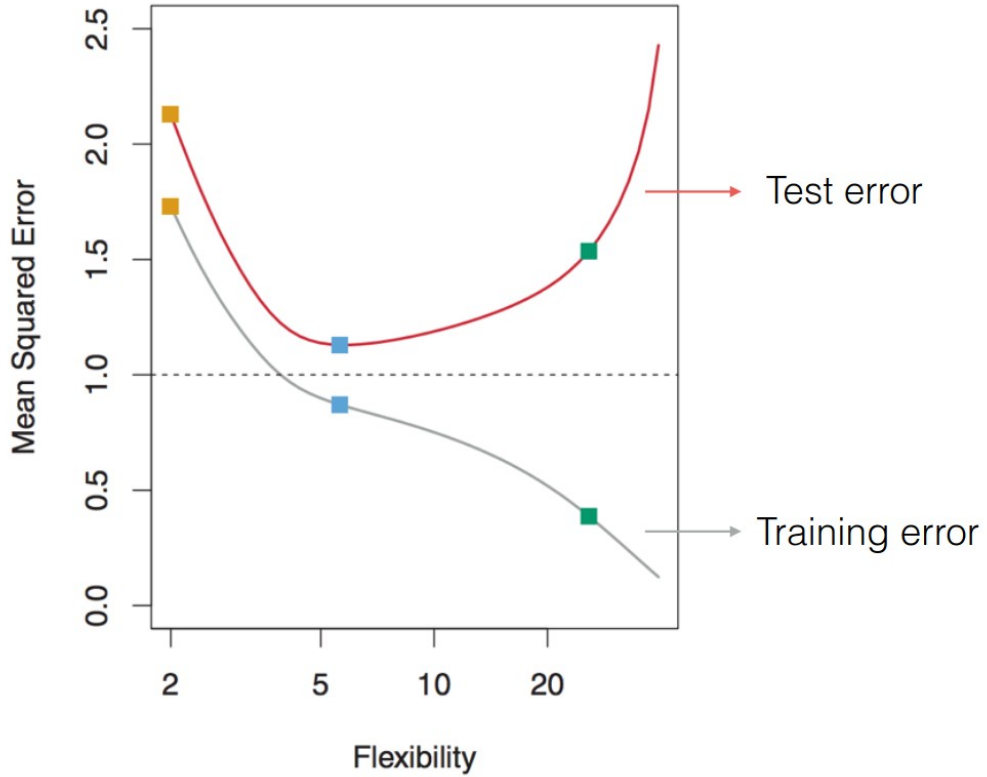

Model Flexibility

As model flexibility increases, training error decreases

- Flexibility allows a model to better *fit* the training data
- Some of the variability in the training data may be noise, however

The test error may or may not decrease

If the test error is much larger than the training error, then the model is **overfitting**

Model Flexibility

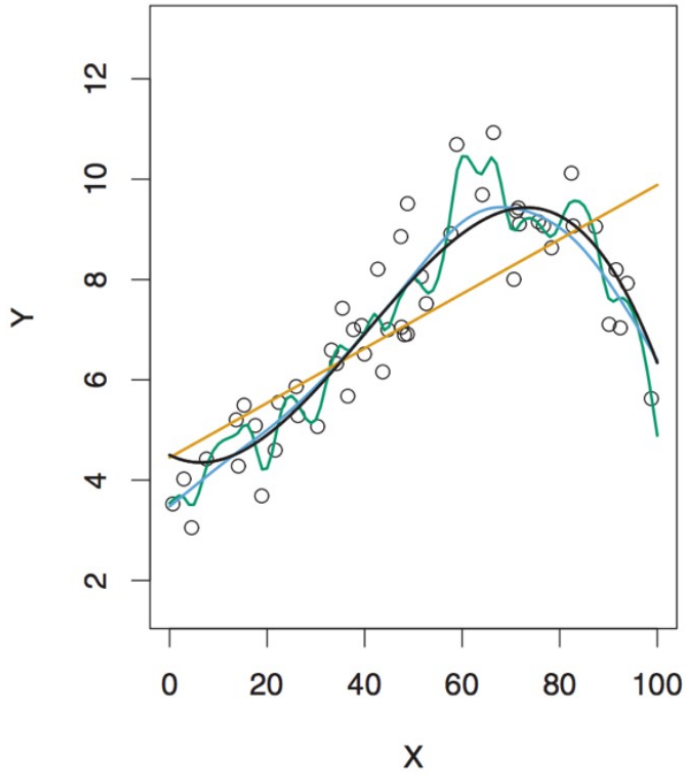

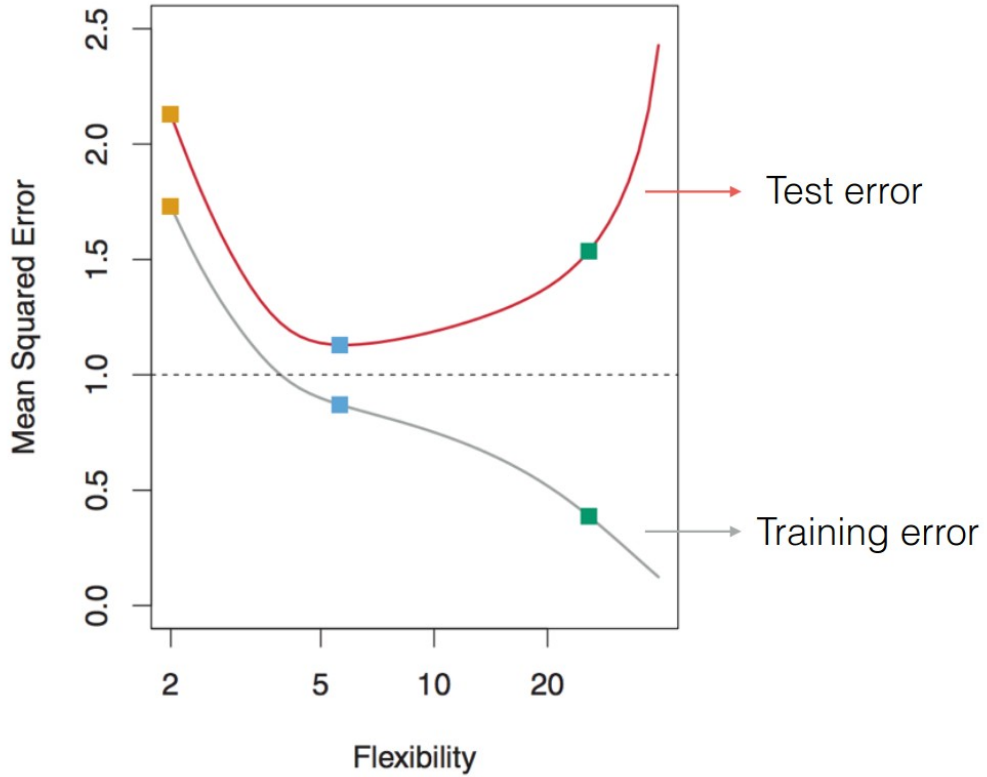

Bias and Variance

Expected test data error = $\texttt{Variance} + \texttt{Bias}^2 + \sigma^2$

- On average, how does our model perform on new data
- **`Bias`**: error caused by simplifying assumptions built into the model
- **`Variance`**: how much of the learned function will change if trained on different training data

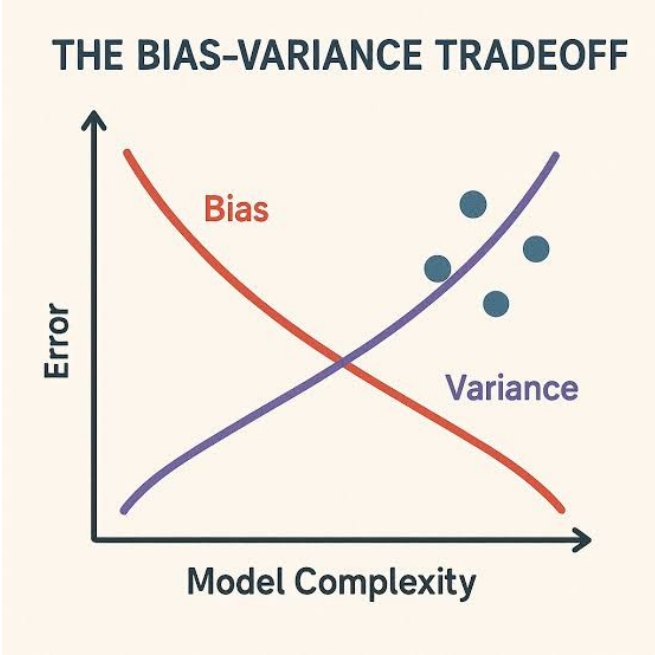

Bias and Variance

Generally, more flexible methods have more variance and less bias (and vice versa)

The **best method** for a prediction task will balance the two types of error to achieve the **lowest test error**

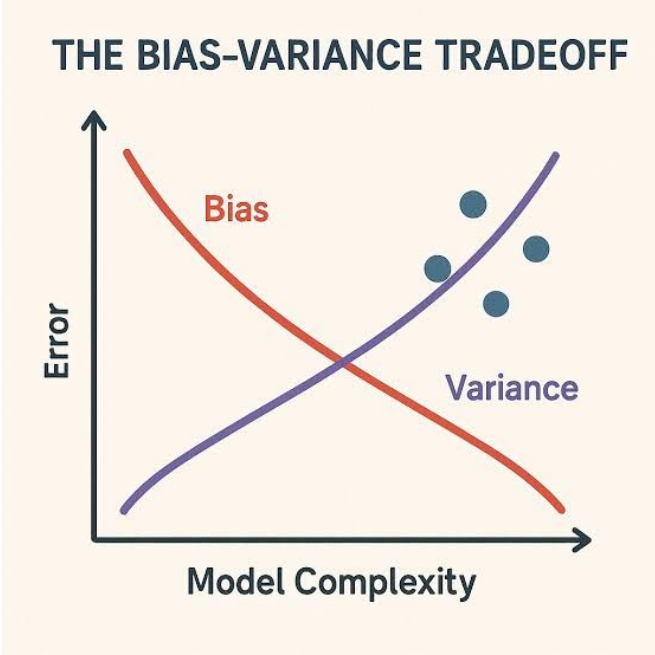

Bias and Variance

Generally, more flexible methods have more variance and less bias (and vice versa)

The **best method** for a prediction task will balance the two types of error to achieve the **lowest test error**

In statistics, we usually *only care* about unbiased models

- Statistics typically tries to learn about the world
- In ML, we're just trying to predict well

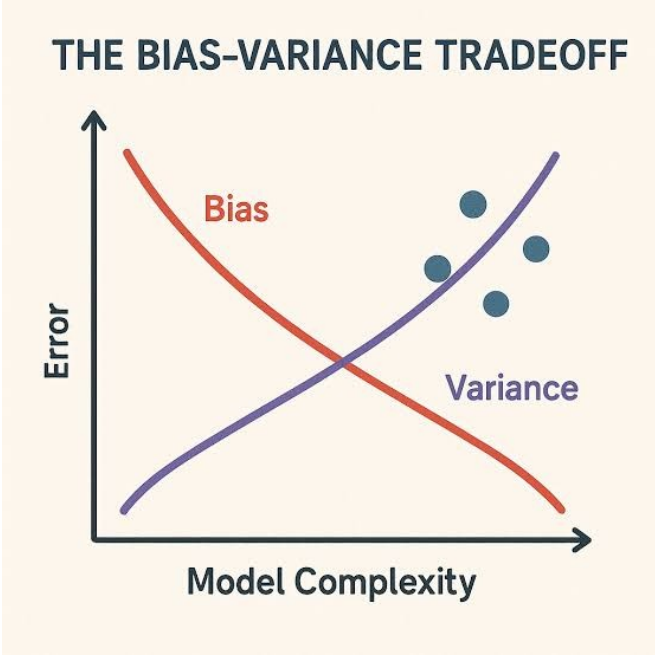

Wrapping up linear regression

Nice properties

- Simple model
- Solves for coefficients that minimize MSE on training data
- Interpretable coefficients

Weak properties

- May be too simple to make accurate predictions
- Sensitive to outliers in the training data due to MSE

Classifiers

Target is categorical

- e.g. predicting whether or not any of the items ordered will arrive damaged
- Target: `damaged` or `not damaged`

How should we measure error?

Classifiers

Target is categorical

- e.g. predicting whether or not any of the items ordered will arrive damaged
- Target: `damaged` or `not damaged`

How should we measure error?

- Accuracy?

$$Accuracy = \frac{examples\ correct}{total\ examples}$$

Classifiers

Target is categorical

- e.g. predicting whether or not any of the items ordered will arrive damaged
- Target: `damaged` or `not damaged`

How should we measure error?

$$Accuracy = \frac{examples\ correct}{total\ examples}$$

- Accuracy?
- High accuracy is good, but we may care more about being correct when the items are actually damaged than when they are not

Confusion Matrix

Idea: count data points by model behavior

True Positive (TP): predicted true correctly

True Negative (TN): predicted false correctly

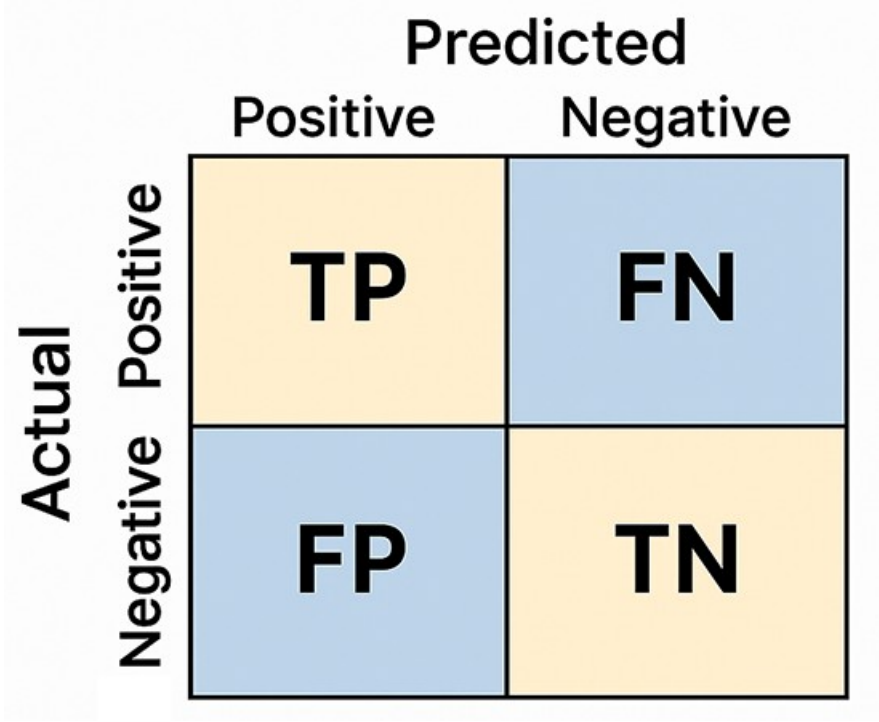

Confusion Matrix

Idea: count data points by model behavior

True Positive (TP): predicted true correctly

True Negative (TN): predicted false correctly

False Positive (FP): predicted true incorrectly

False Negative (FN): predicted false incorrectly

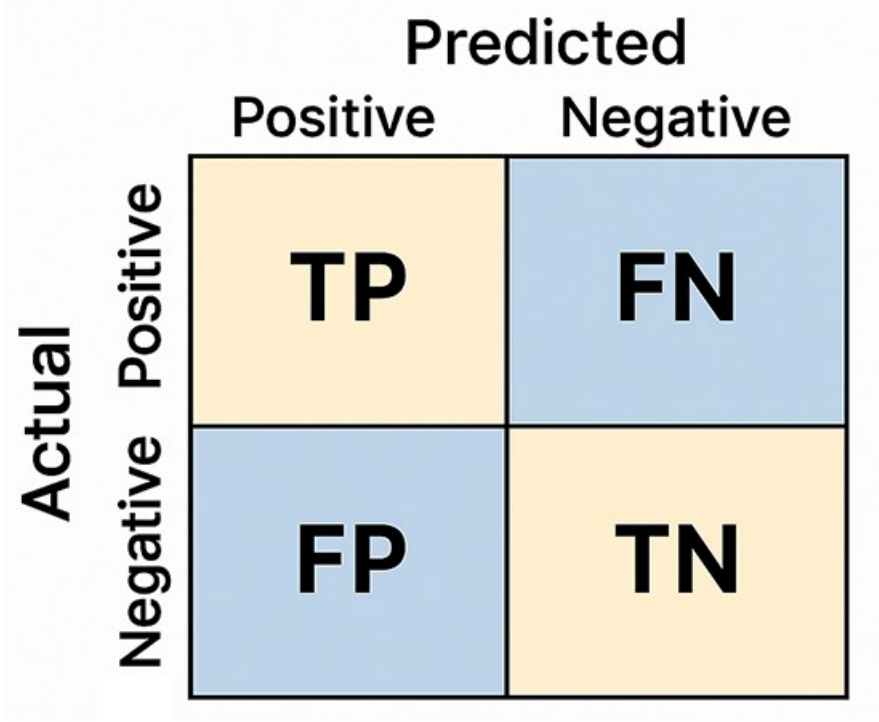

Confusion Matrix

Idea: count data points by model behavior

True Positive (TP): predicted true correctly

True Negative (TN): predicted false correctly

False Positive (FP): predicted true incorrectly

False Negative (FN): predicted false incorrectly

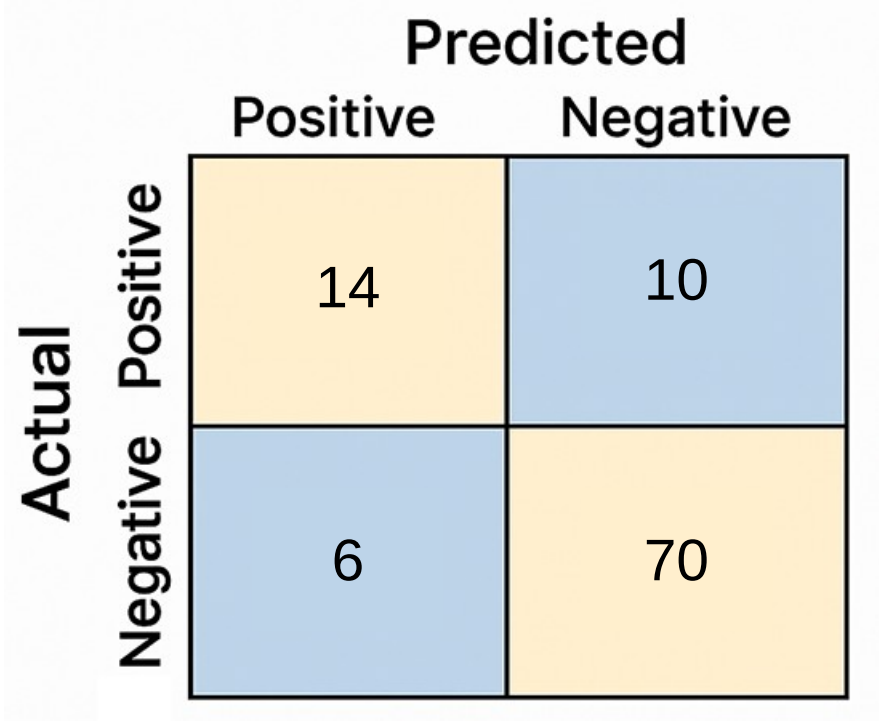

Confusion Matrix

Idea: count data points by model behavior

True Positive (TP): predicted true correctly

True Negative (TN): predicted false correctly

False Positive (FP): predicted true incorrectly

False Negative (FN): predicted false incorrectly

Accuracy =

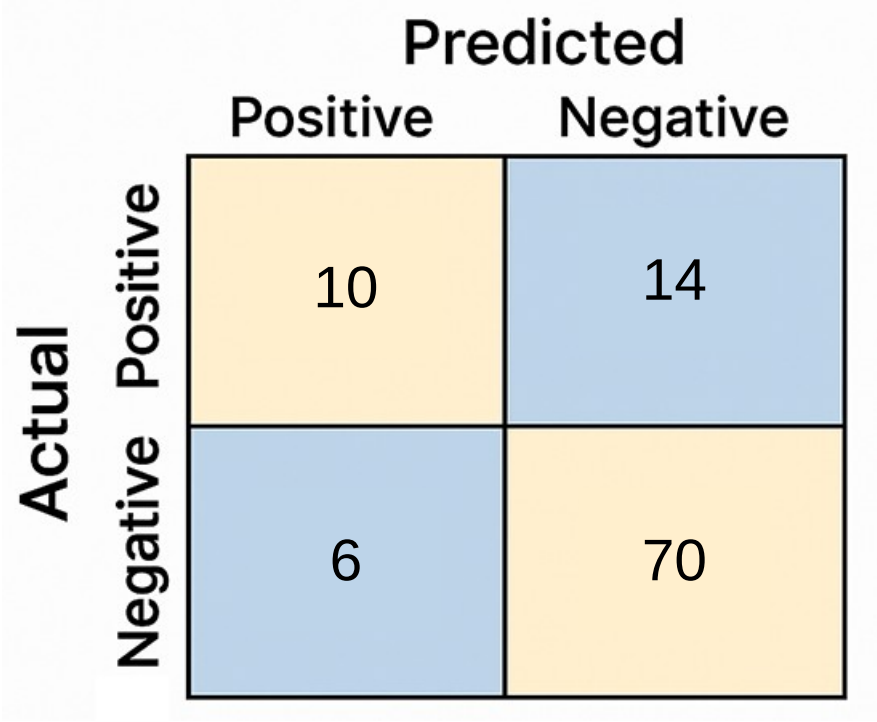

Confusion Matrix

Idea: count data points by model behavior

True Positive (TP): predicted true correctly

True Negative (TN): predicted false correctly

False Positive (FP): predicted true incorrectly

False Negative (FN): predicted false incorrectly

Accuracy = (TP + TN) / Total

= (10 + 70) / 100 = 80%

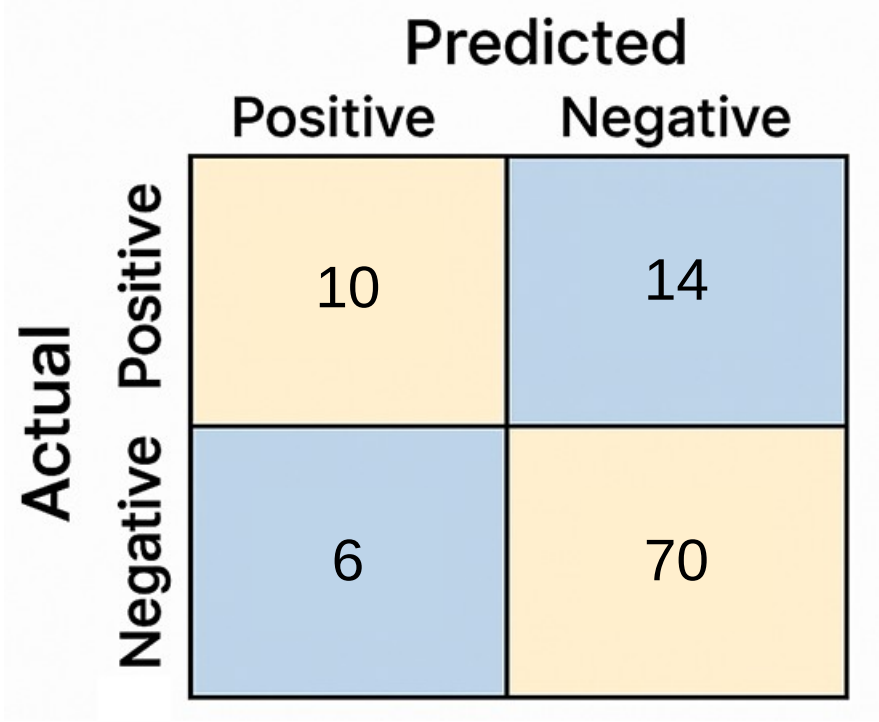

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

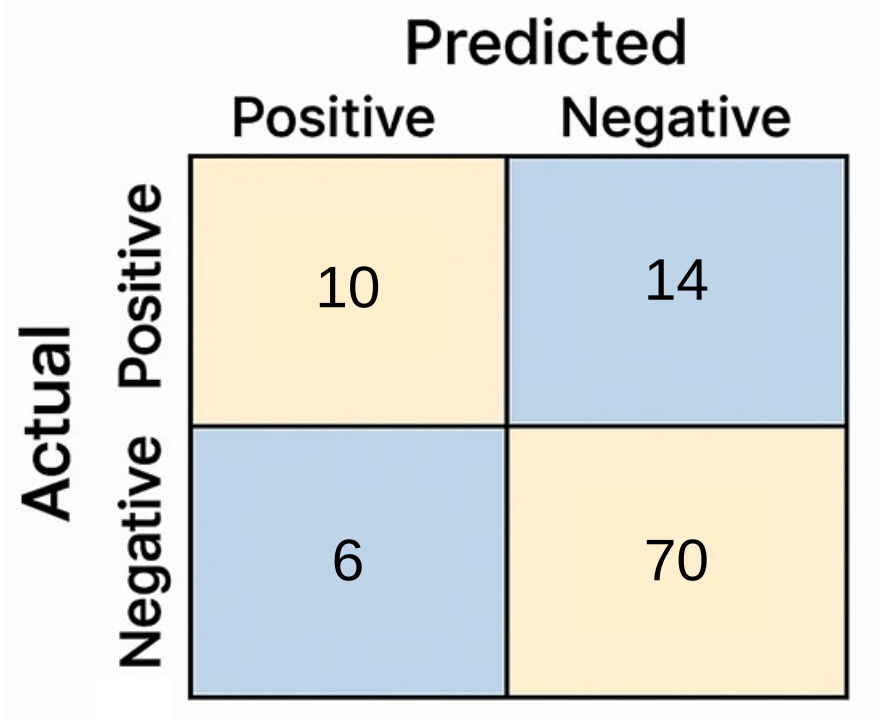

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

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

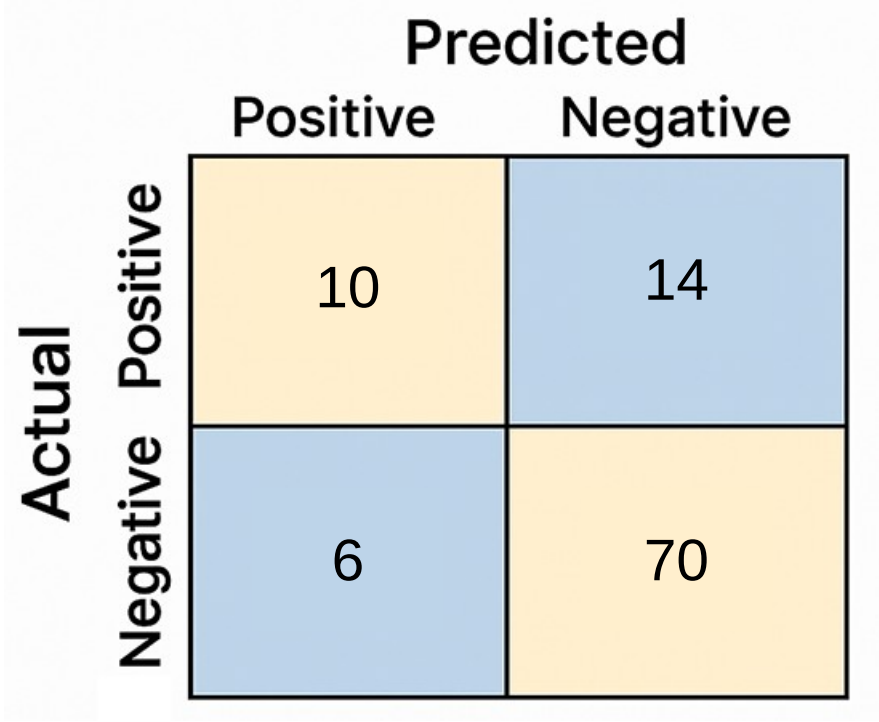

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

$$Precision = \frac{TP}{TP + FP} = 62.5\%$$

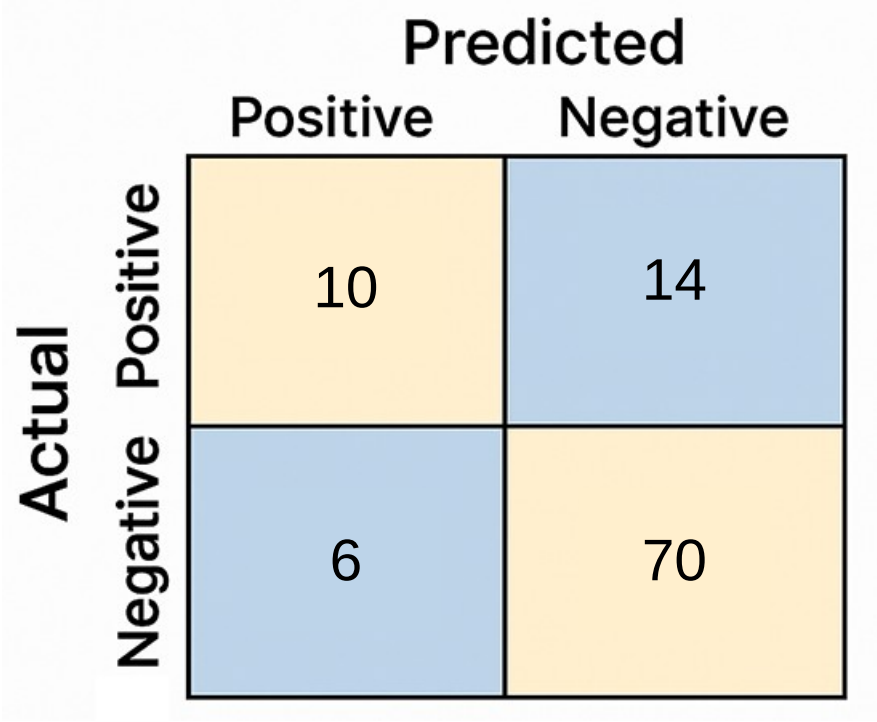

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

$$Precision = \frac{TP}{TP + FP} = 62.5\,\%$$

Recall: out of all the actual true items in the dataset, how many did the model find?

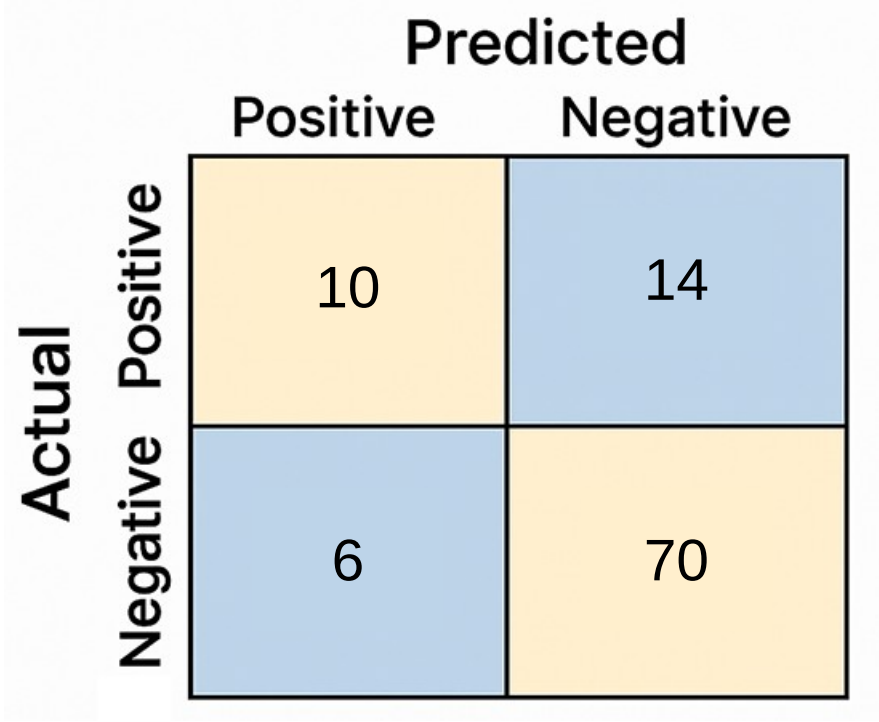

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

$$Precision = \frac{TP}{TP + FP} = 62.5\%$$

Recall: out of all the actual true items in the dataset, how many did the model find?

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

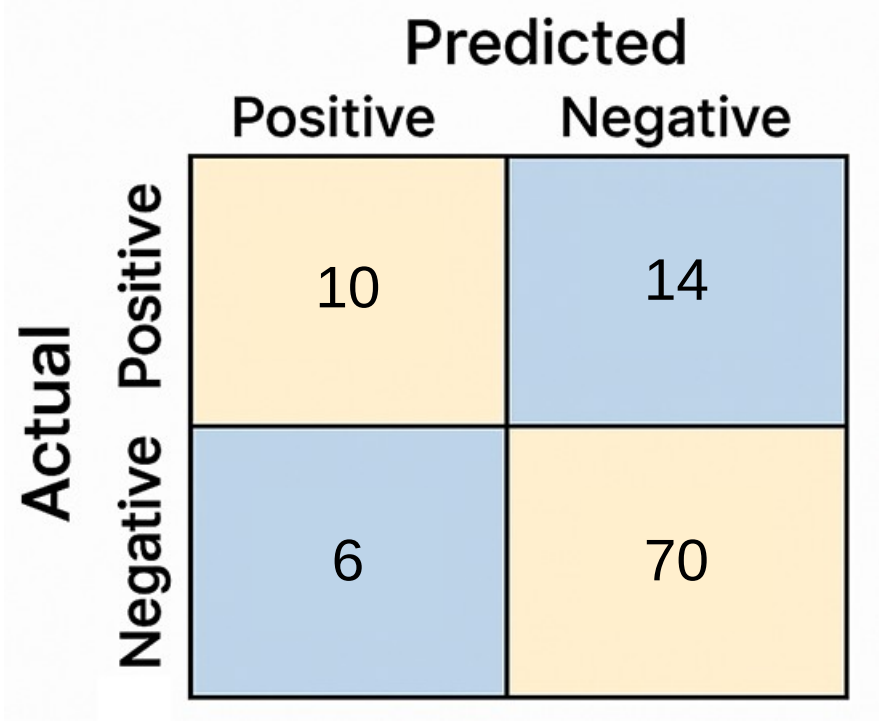

Confusion Matrix

Precision: out of all items the model predicted true, how many were actually true?

$$Precision = \frac{TP}{TP + FP} = 62.5\,\%$$

Recall: out of all the actual true items in the dataset, how many did the model find?

$$Recall = \frac{TP}{TP + FN} = 41.7\,\%$$

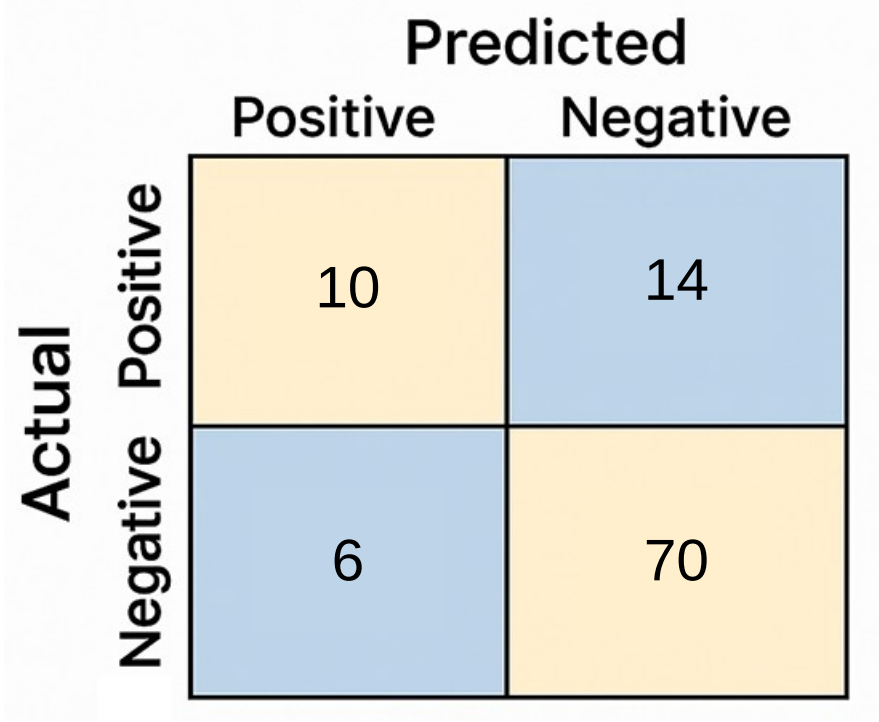

Next Time

Introduce data privacy notions