

# Machine Learning: Nearest Neighbors

ROB 102: Introduction to AI & Programming

Lecture 11

2021/11/22

# Last time...

**Image classification** is a type of **supervised learning** where we predict the class of an image using labelled data.

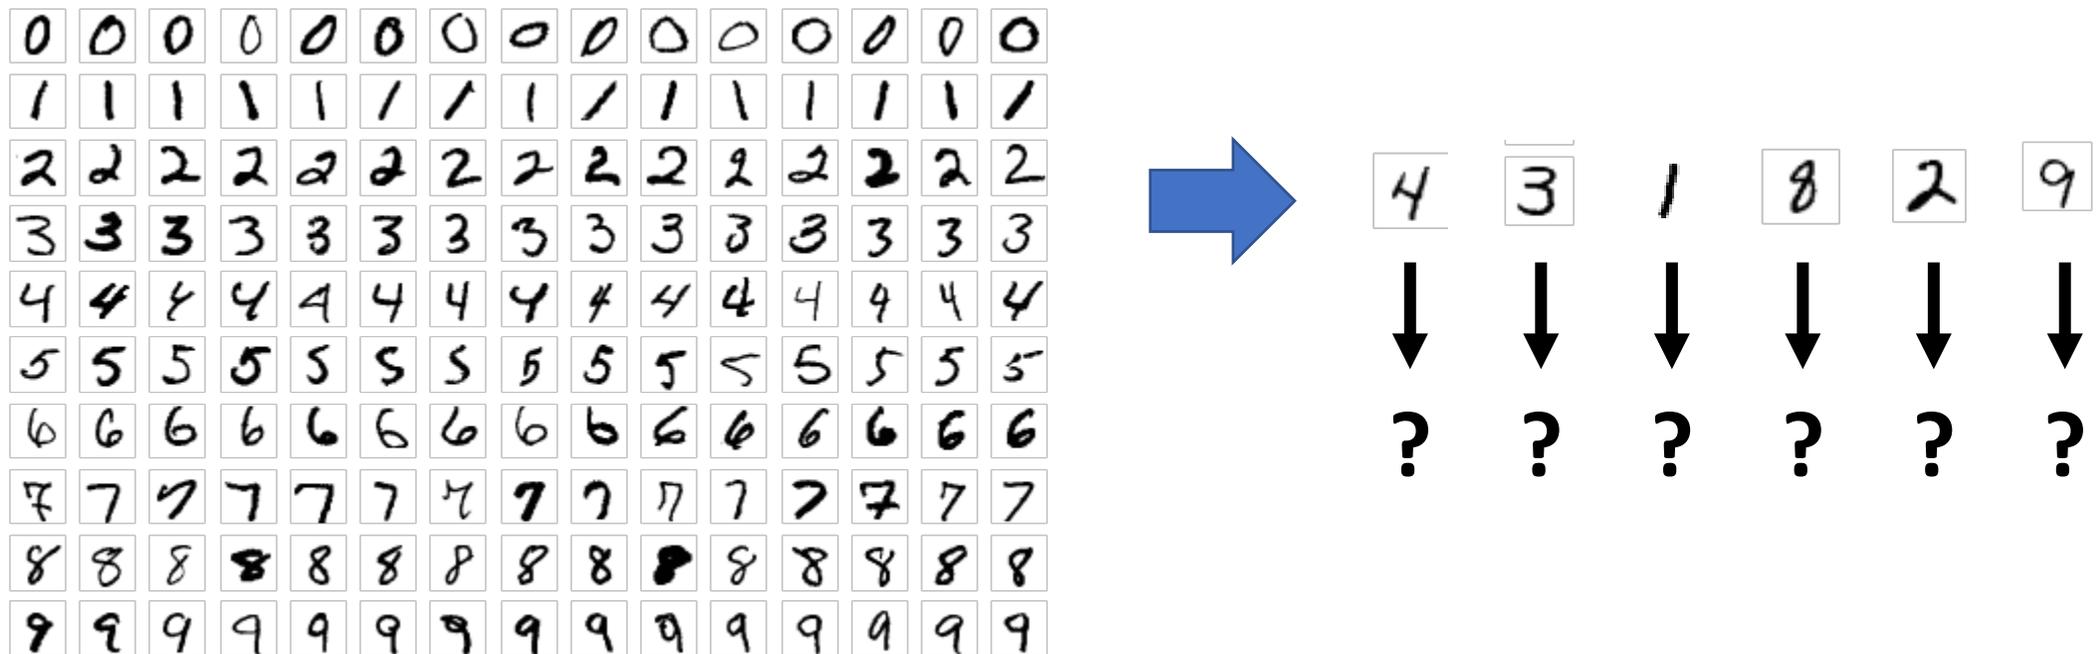


Image: CC4.0 ([link](#))

# Last time...

**Machine Learning Algorithm:**

**Training time:**

Learn a prediction model by optimizing over a labelled dataset.

**Testing time:**

Use the model to perform prediction on new data.

**Data Split:**

**Training set:** Labelled data used for training a machine learning algorithm.

**Test set:** Data used to test the accuracy of the machine learning algorithm.

# Project 4: Machine Learning

Implement three machine learning algorithms to classify images from the MNIST dataset.

1. Nearest neighbors
2. Linear Classifier
3. Neural Network

The assignment instructions are available!

<https://robotics102.github.io/projects/a4>

**Menu**

---

HOME

---

COURSE INFORMATION ▼

---

PROJECTS ▲

- PROJECT 0: INTRO TO C++
- PROJECT 1: WALL FOLLOWING
- PROJECT 2: POTENTIAL FIELD CONTROL
- PROJECT 3: PATH PLANNING
- PROJECT 4: MACHINE LEARNING

---

TUTORIALS ▼

---

**Course Times**

---

Lectures and labs

 MW 10-11:30 AM @ GFL 107



**ROB 102:** Introduction to AI and  
Programming



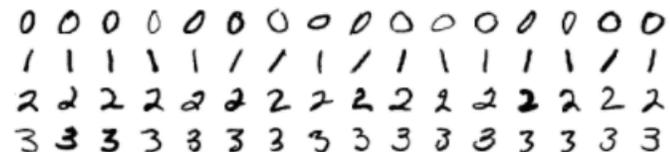
---

## Project 4: Machine Learning for Image Classification

**Due December 10th, 2021 at 11:59 PM.**

In this project, we will use machine learning algorithms to perform **image classification**. Image classification is the task of predicting the class, or label, of an image out of a set of known images. This is a type of *supervised learning*, which refers to algorithms that learn a function from labelled datasets.

We will be writing algorithms to do image classification on the MNIST dataset. MNIST consists of tiny 28×28 pixel images of handwritten digits from 0 to 9. A few example images from each are shown below.



The template code is available! Use the Github Classroom link to join.

The screenshot shows a GitHub repository page for 'robotics102 / machine-learning'. The repository is a private template with 4 unwatched items, 0 stars, and 0 forks. The repository is on the 'main' branch and has 1 branch and 0 tags. The repository description is 'Template code for ROB 102 Project 4: Machine Learning.' The repository contains several files: '.gitignore', 'README.md', 'linear\_classifier.ipynb', 'nearest\_neighbors.ipynb', and 'two\_layer\_neural\_net.ipynb'. The repository is owned by 'janapavlassek' and has 6 commits. The repository is a template and can be used to create a new repository.

robotics102 / machine-learning Private template

Unwatch 4 Star 0 Fork 0

Code Issues Pull requests Actions Projects Security Insights Settings

main 1 branch 0 tags

Go to file Add file Code Use this template

janapavlassek Add the project page link to the readme c0d5e01 3 days ago 6 commits

File	Description	Time
.gitignore	Initialize template files	18 days ago
README.md	Add the project page link to the readme	3 days ago
linear_classifier.ipynb	Add tests and documentation	3 days ago
nearest_neighbors.ipynb	Add tests and documentation	3 days ago
two_layer_neural_net.ipynb	Add tests and documentation	3 days ago

README.md

## ROB 102: Machine Learning

Template code for ROB 102 Project 4: Machine Learning in Julia. See the project description at

About

Template code for ROB 102 Project 4: Machine Learning.

Readme

Releases

No releases published  
[Create a new release](#)

Packages

No packages published  
[Publish your first package](#)

Contributors 2

# Project 4: Machine Learning

Implement three machine learning algorithms to classify images from the MNIST dataset.

1. **Nearest neighbors** (Today!)
2. Linear Classifier
3. Neural Network

# Image Classification on MNIST

Imagine we have 60k labelled images. How can we predict the class of a new image?



# Image Classification on MNIST

Imagine we have 60k labelled images. How can we predict the class of a new image?



**Idea:** This image of a two might be *numerically close* to other images of twos.

# Nearest Neighbors

**Idea:** Given a new image, find the closest image in the training set. Then, assign the same label to the new image.

Test images 



Nearest training images 

# Nearest Neighbors: Project 4.1

In Project 4 (Part 1), you will implement an algorithm to classify images using Nearest Neighbors.

jupyter nearest\_neighbors Last Checkpoint: 5 hours ago (autosaved)

File Edit View Insert Cell Kernel Help

Code

## Part I: Nearest Neighbors

This notebook implements a nearest neighbors classifier.

### Imports

Some imports we'll need.

```
In [2]: 1 using MLDatasets
        2 using Plots
        3 using Images
        4 using MosaicViews
        5 using LinearAlgebra
        6 using Random
        7 using Printf
```

```
num_viz = 10
idx = shuffle(1:N_test)[1:num_viz]
nearest = x_train[indices[idx], :, :]

imgs = [x_test[idx, :, :]; nearest]
# imgs = reshape(imgs', width, height, num_viz * 2)
imgs = MNIST.convert2image(permutdims(imgs, (2, 3, 1)))

for i in 1:num_viz
    img = idx[i]
    @printf("Img %-2d -> Predicted: %-10s True: %s\n", i, y_pred[img], y_test[img])
end

println("\nTop row = test image, bottom row = nearest neighbor.")
mosaicview(imgs, fillvalue=1, nrow=2, npad=3, rowmajor=true)
```

```
Img 1 -> Predicted: 8           True: 3
Img 2 -> Predicted: 8           True: 8
Img 3 -> Predicted: 4           True: 4
Img 4 -> Predicted: 0           True: 0
Img 5 -> Predicted: 8           True: 8
Img 6 -> Predicted: 9           True: 9
Img 7 -> Predicted: 4           True: 4
Img 8 -> Predicted: 2           True: 8
Img 9 -> Predicted: 1           True: 1
Img 10 -> Predicted: 3          True: 3
```

Top row = test image, bottom row = nearest neighbor.

```
Out[36]: 3 8 4 0 8 9 4 8 1 3
          8 8 4 0 8 9 4 2 1 3'
```

# Nearest Neighbors

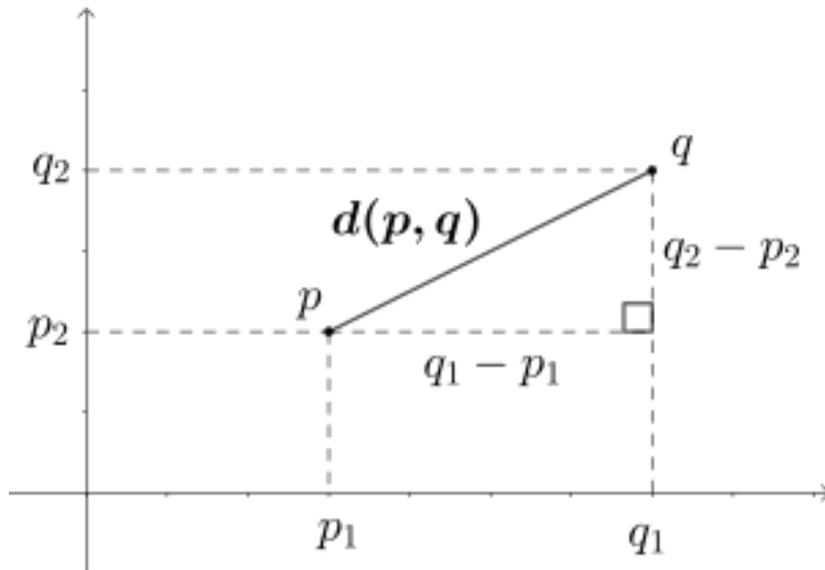
**Idea:** Given a new image, find the closest image in the training set. Then, assign the same label to the new image.

What does “nearest” mean?

distance(  ,  )

# Euclidean Distance

Recall: The Pythagorean Theorem gives us the distance:



In 2D:

$$d(p, q) = \sqrt{(q_1 - p_1)^2 + (q_2 - p_2)^2}$$

distance

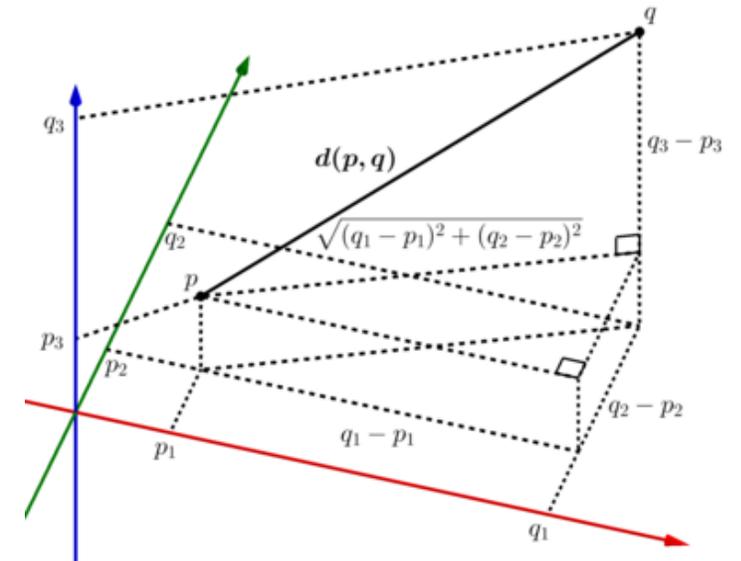
# Euclidean Distance

Recall: The Pythagorean Theorem gives us the distance:

In 3D:

$$d(p, q) = \sqrt{(q_1 - p_1)^2 + (q_2 - p_2)^2 + (q_3 - p_3)^2}$$

↑  
distance



# Euclidean Distance

Recall: The Pythagorean Theorem gives us the distance:

In 3D:

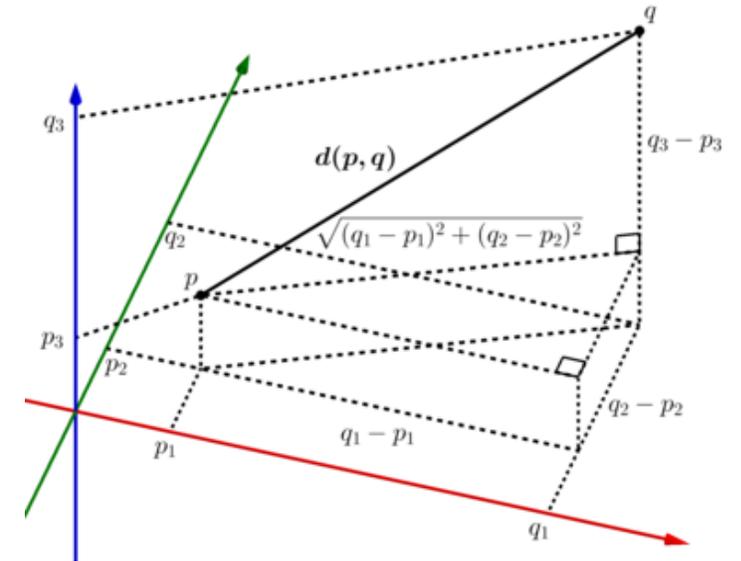
$$d(p, q) = \sqrt{(q_1 - p_1)^2 + (q_2 - p_2)^2 + (q_3 - p_3)^2}$$

In N-D:

$$d(p, q) = \sqrt{\sum_{i=1}^N (q_i - p_i)^2}$$

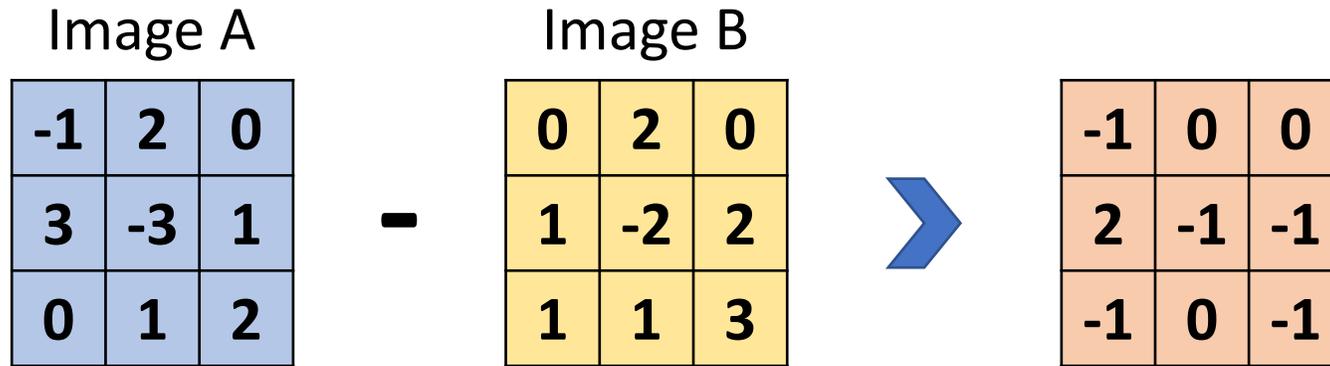
distance  $\uparrow$

$$= \sqrt{(q_1 - p_1)^2 + (q_2 - p_2)^2 + \dots + (q_N - p_N)^2}$$



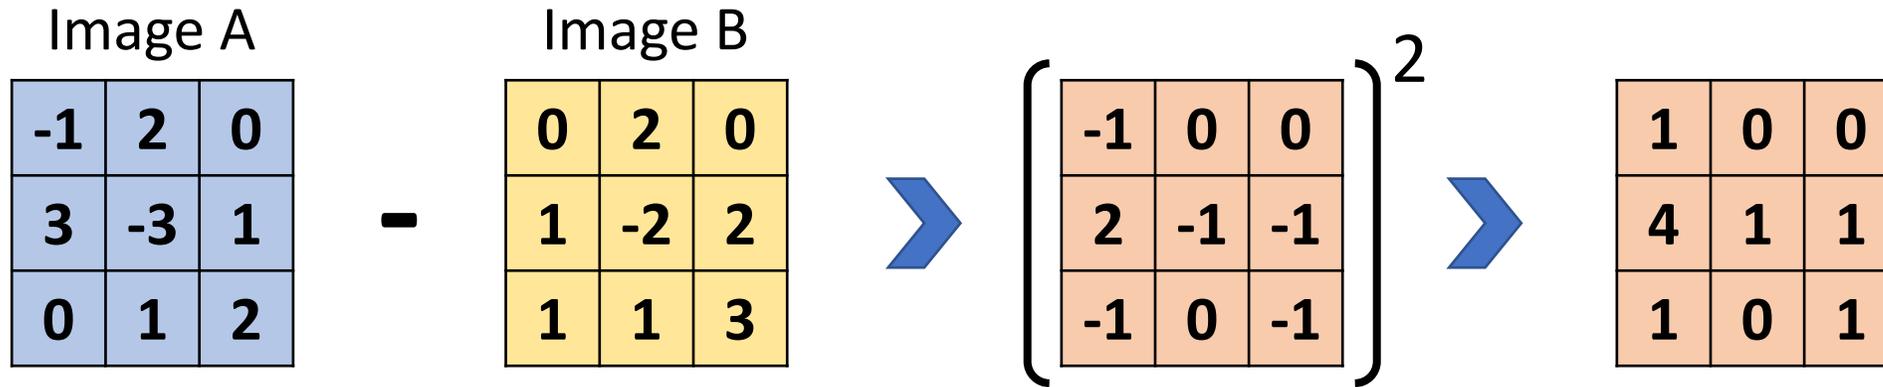
# Euclidean Distance: Example

$$d(p, q) = \sqrt{\sum_{i=1}^N (q_i - p_i)^2}$$



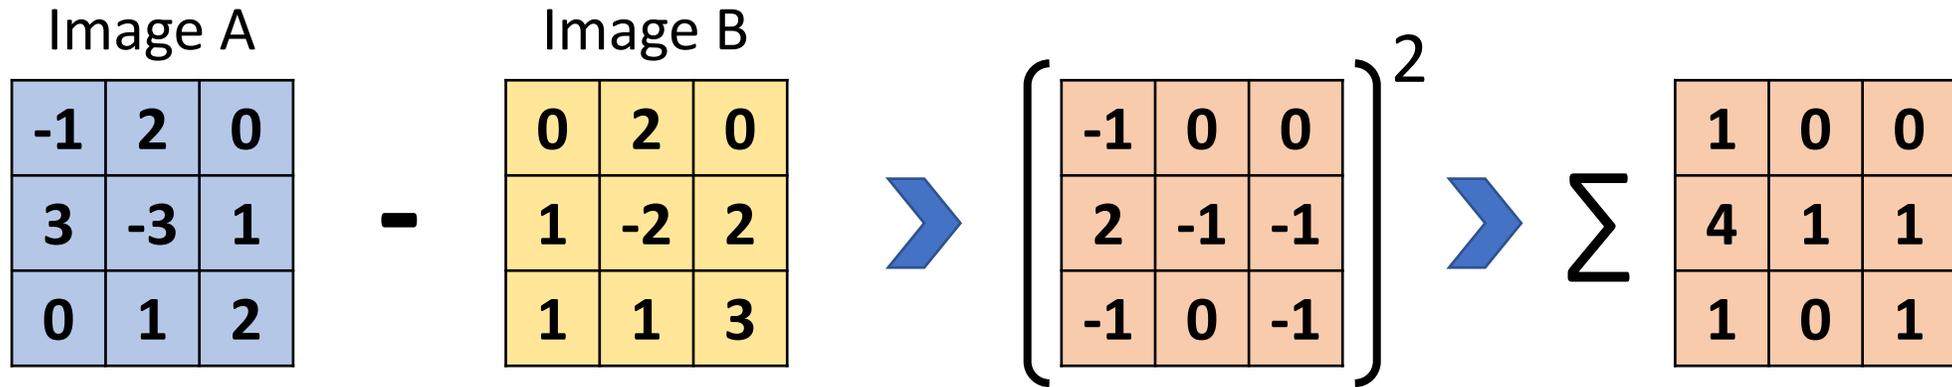
# Euclidean Distance: Example

$$d(p, q) = \sqrt{\sum_{i=1}^N (q_i - p_i)^2}$$



# Euclidean Distance: Example

$$d(p, q) = \sqrt{\sum_{i=1}^N (q_i - p_i)^2}$$



$$1 + 0 + 0 + 4 + 1 + 1 + 1 + 0 + 1 = 9$$

# Euclidean Distance: Example

$$d(p, q) = \sqrt{\sum_{i=1}^N (q_i - p_i)^2}$$

Image A

-1	2	0
3	-3	1
0	1	2

-

Image B

0	2	0
1	-2	2
1	1	3



-1	0	0
2	-1	-1
-1	0	-1

<sup>2</sup>

Σ

1	0	0
4	1	1
1	0	1



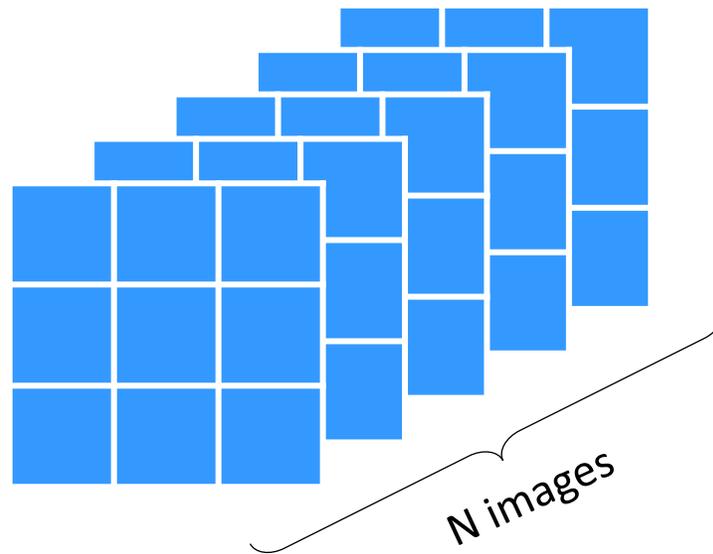
$$1 + 0 + 0 + 4 + 1 + 1 + 1 + 0 + 1 = 9$$



$$\text{distance}(A, B) = \sqrt{9} = 3$$

# Images as Matrices

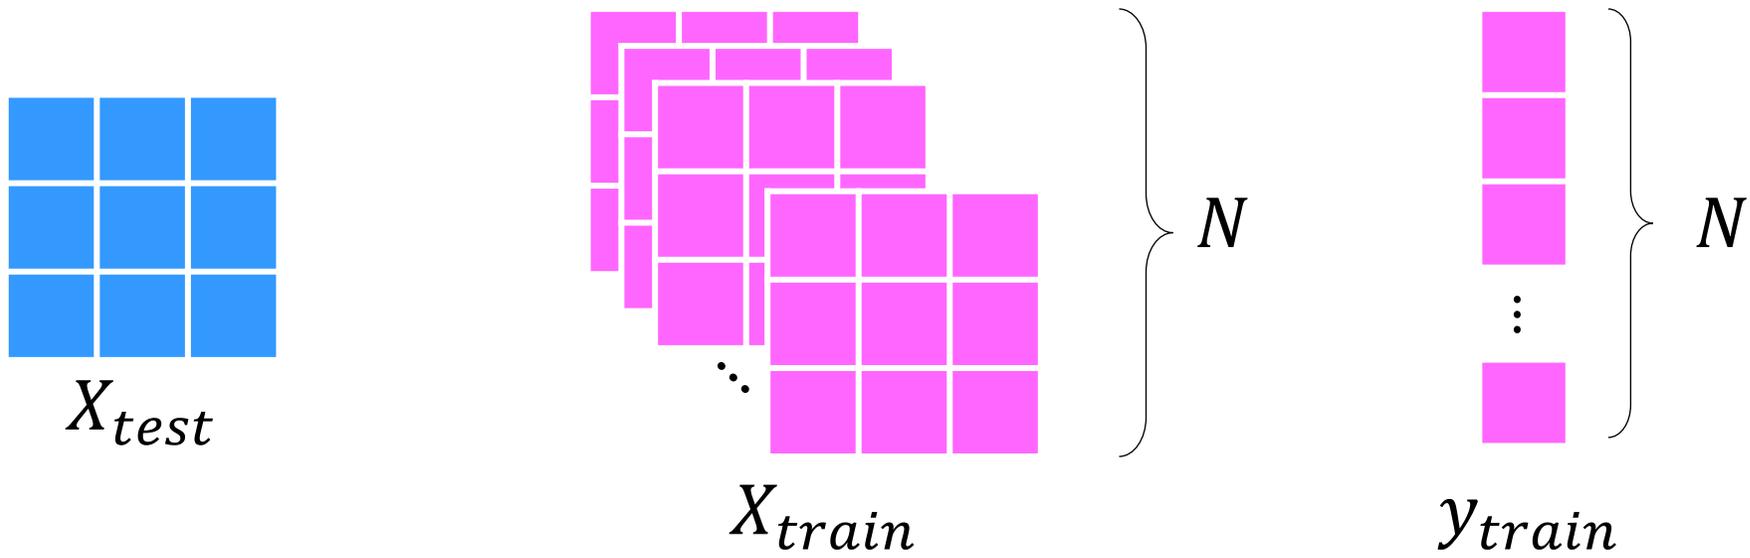
If we have many images, we will stack them up in a vector of matrices, or a 3D matrix of size  $N \times W \times H$ .



In Project 4, the data will be stored in matrices like these. This is a convenient representation, but we also like big matrices because computers are very good at dealing with them.

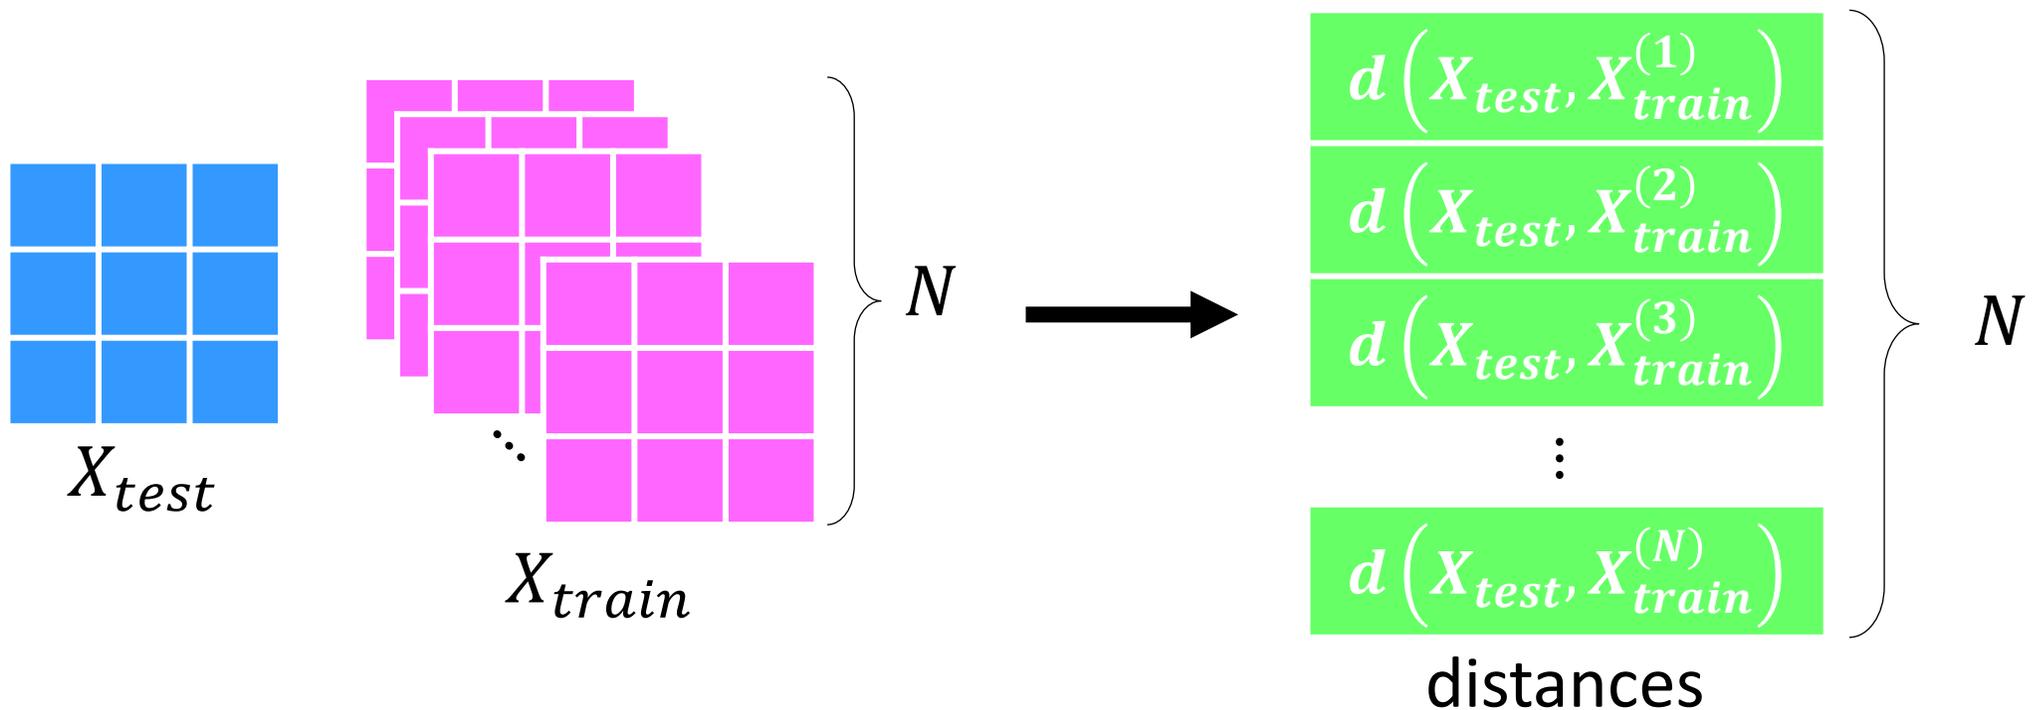
# Nearest Neighbors

Back to the nearest neighbors algorithm. Say we have a matrix of  $N$  training images and a test image we want to classify.



# Nearest Neighbors

Back to the nearest neighbors algorithm. Say we have a matrix of  $N$  training images and a test image we want to classify.



# Nearest Neighbors

The smallest distance to the test image is given by:

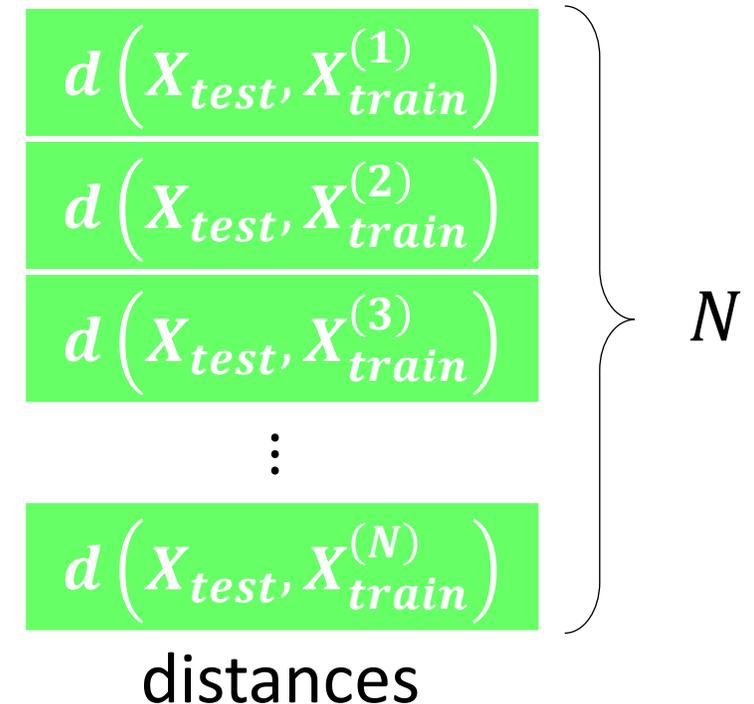
$$\text{minimum}(\text{distances})$$

Let's say the test image is closest to train image  $X_{train}^{(7)}$ . Then:

$$\text{argmin}(\text{distances}) = 7$$

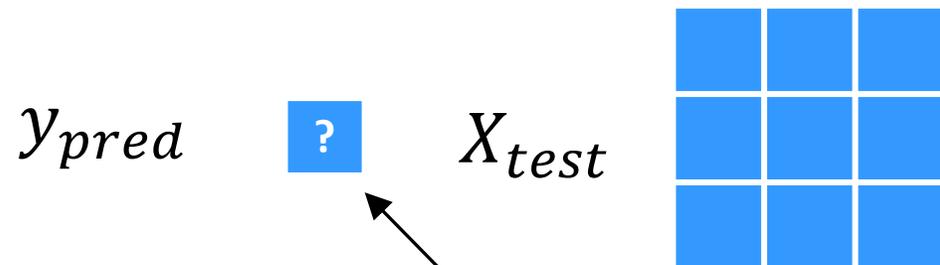
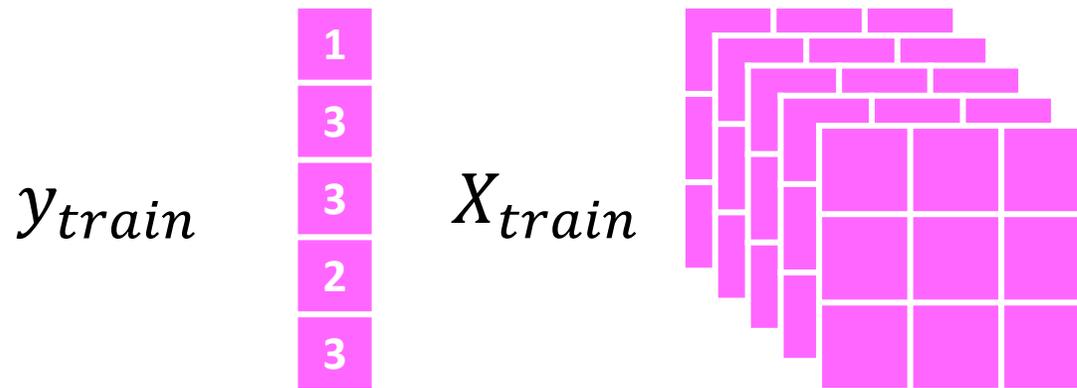
The function `argmin` gives the argument that minimizes `distances`. So, we can predict:

$$y_{pred} = y_{train}[7]$$



# Nearest Neighbors

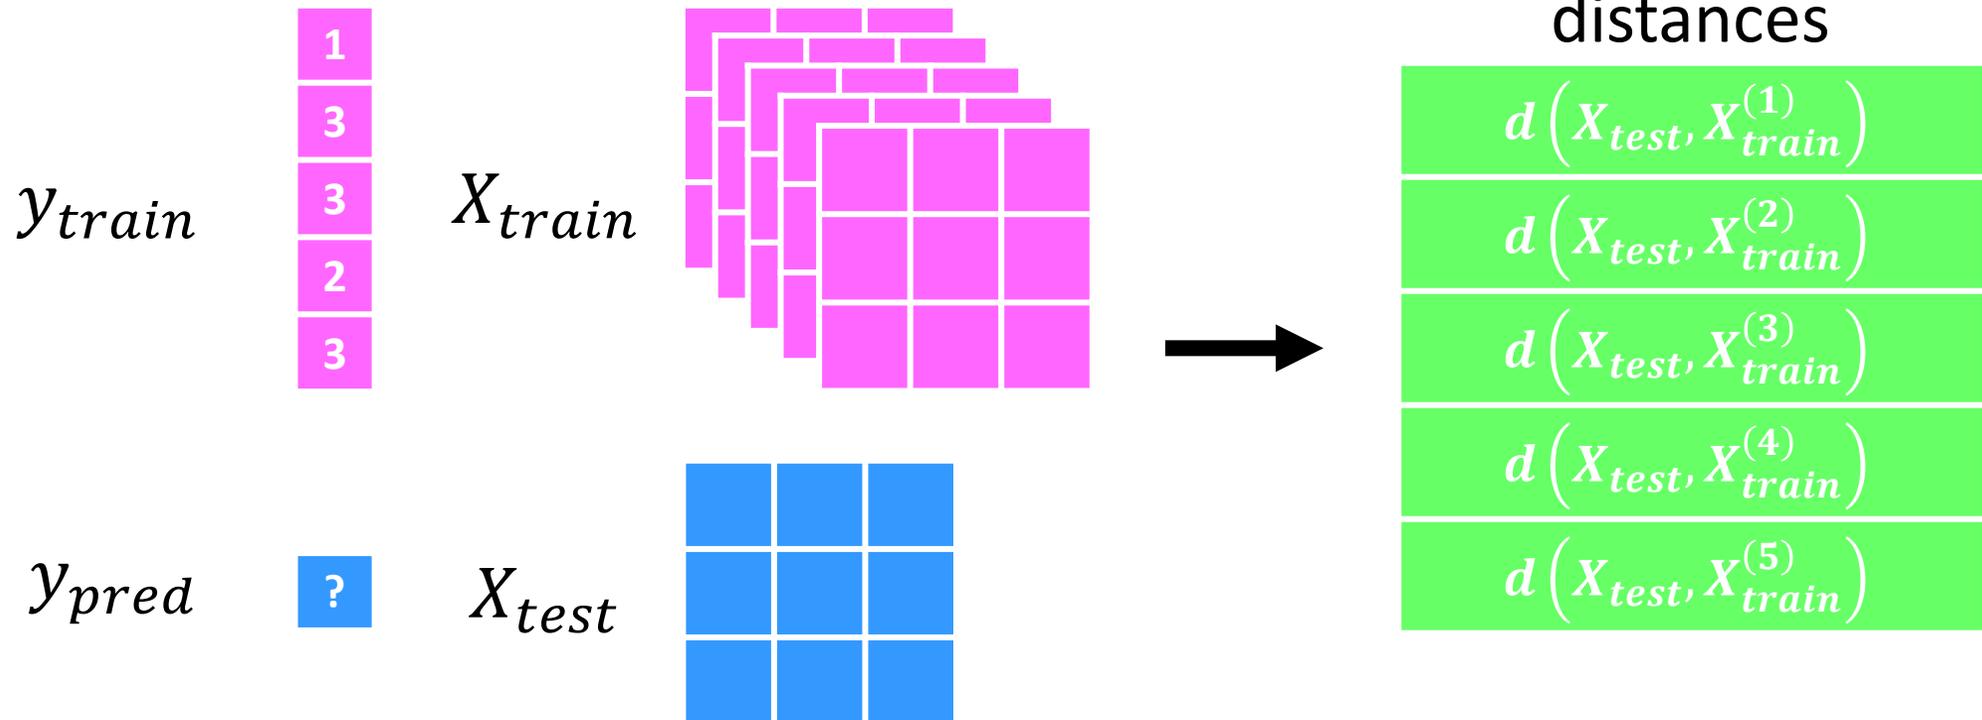
A small example:



**Goal:** Predict test image label,  $y_{pred}$

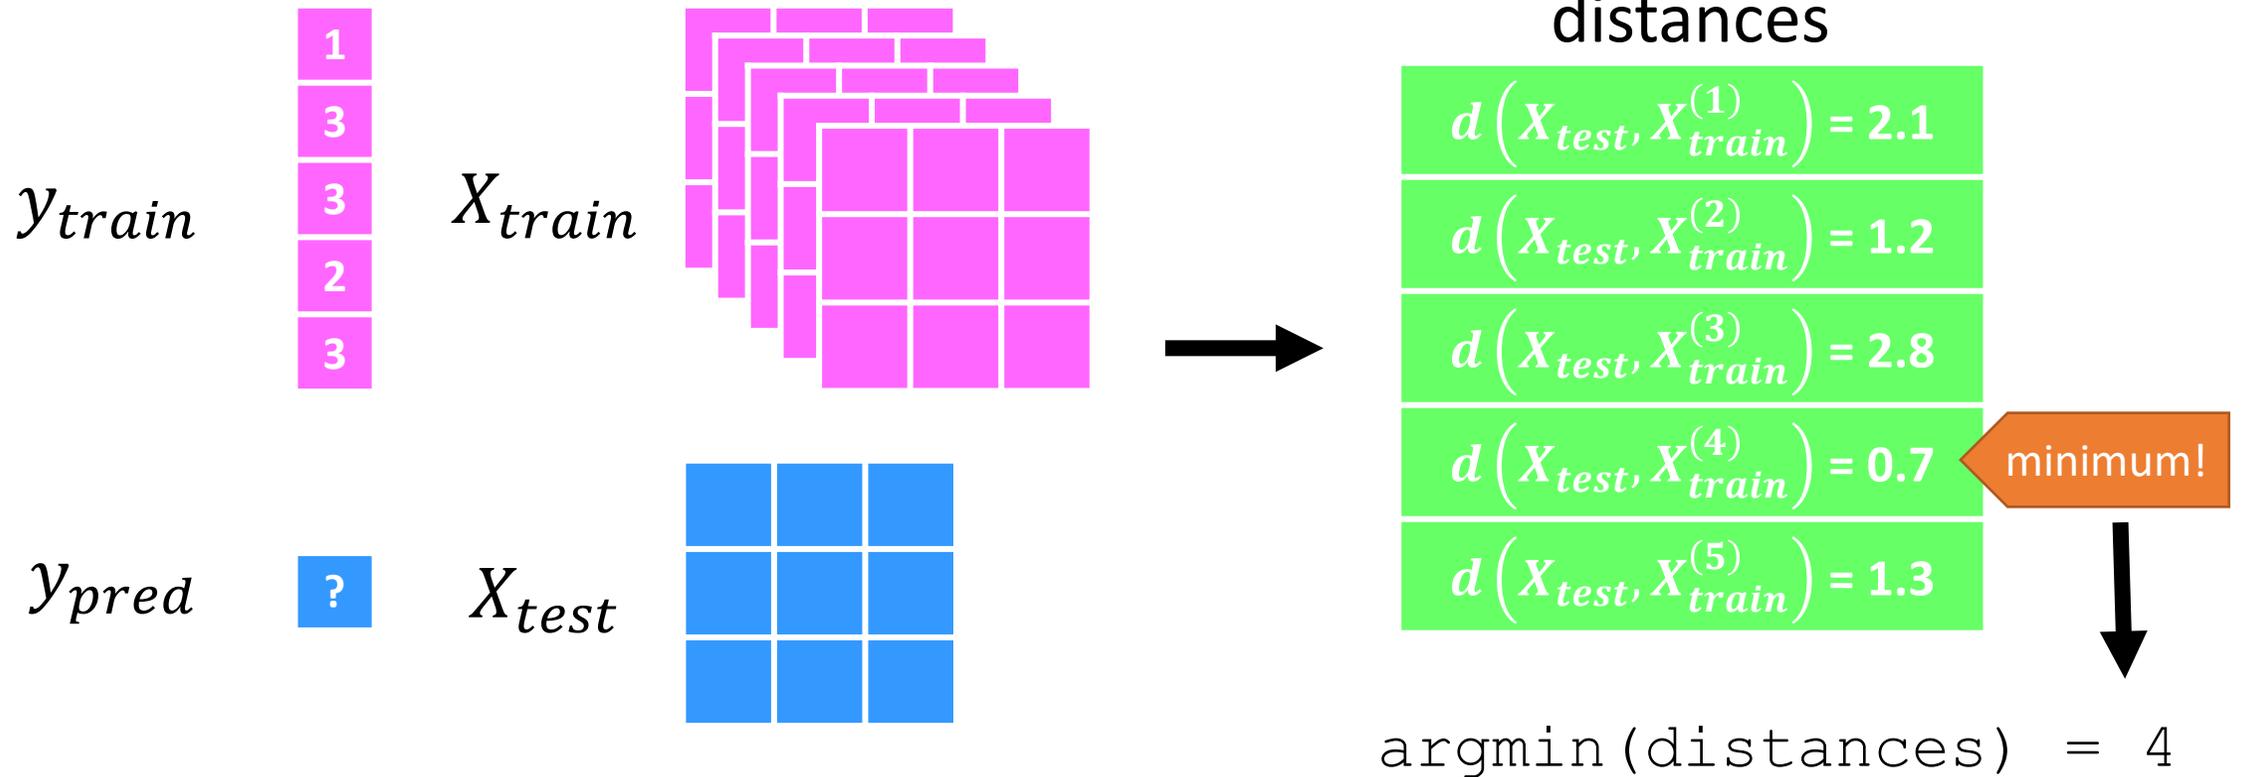
# Nearest Neighbors

A small example:



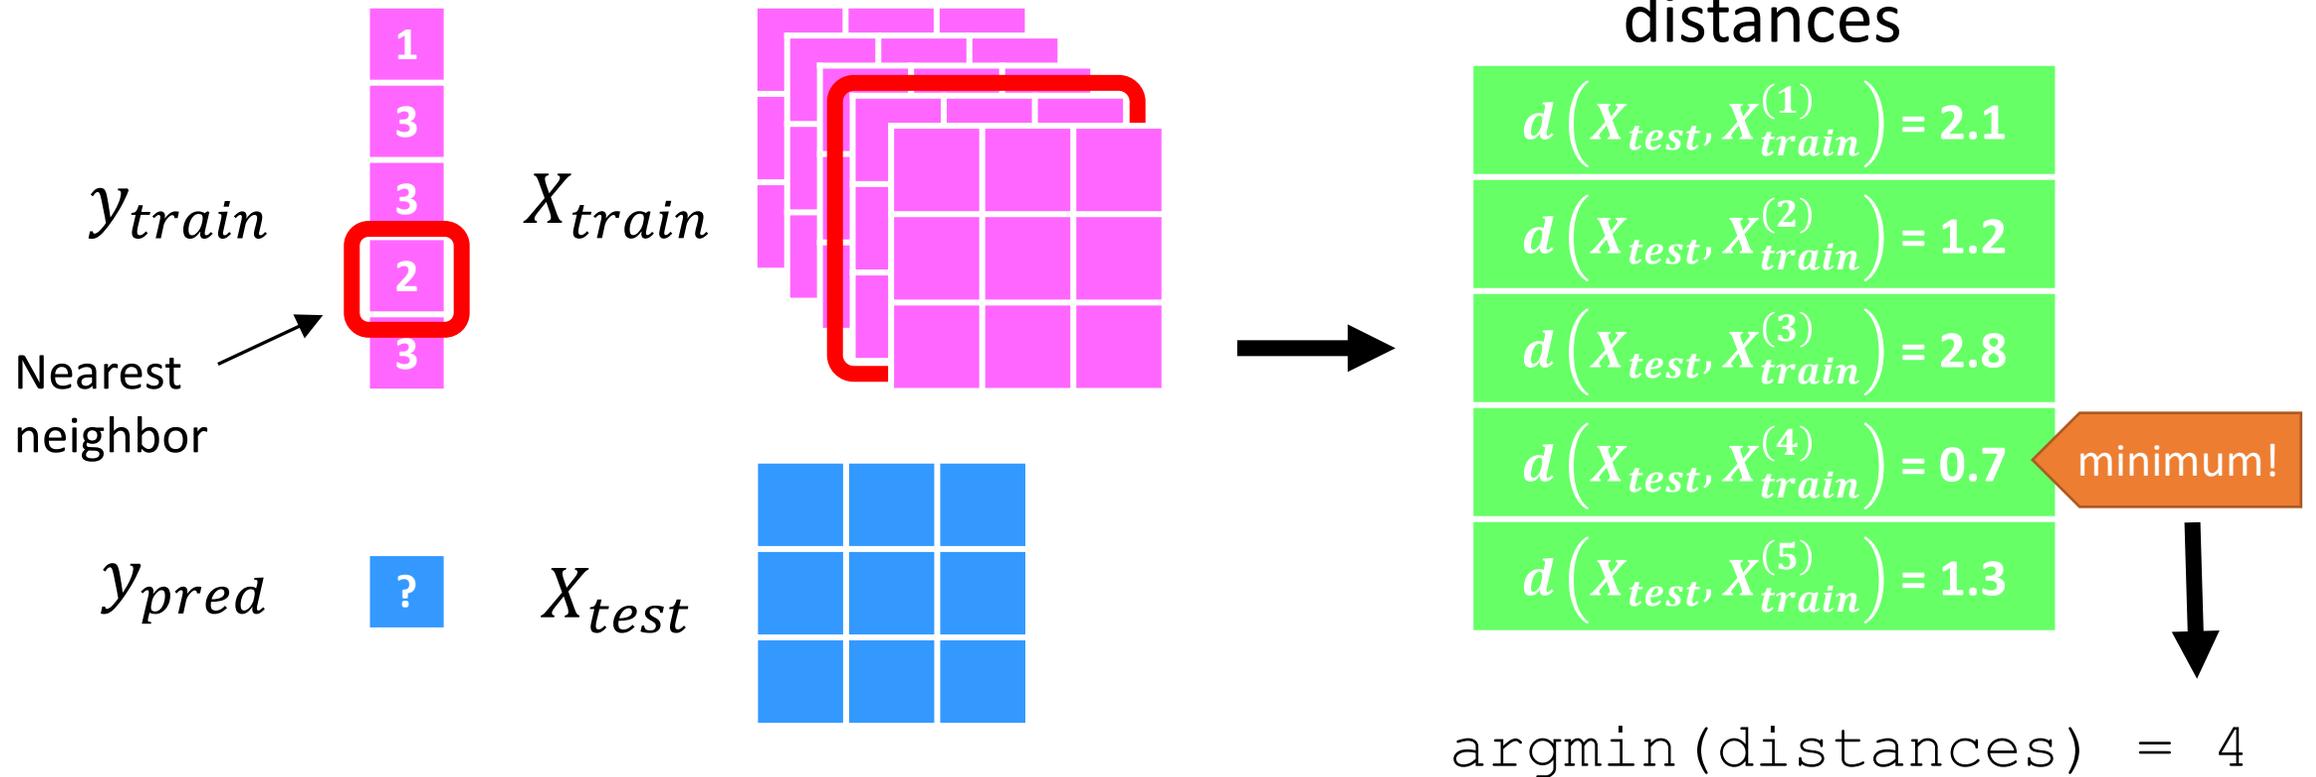
# Nearest Neighbors

A small example:



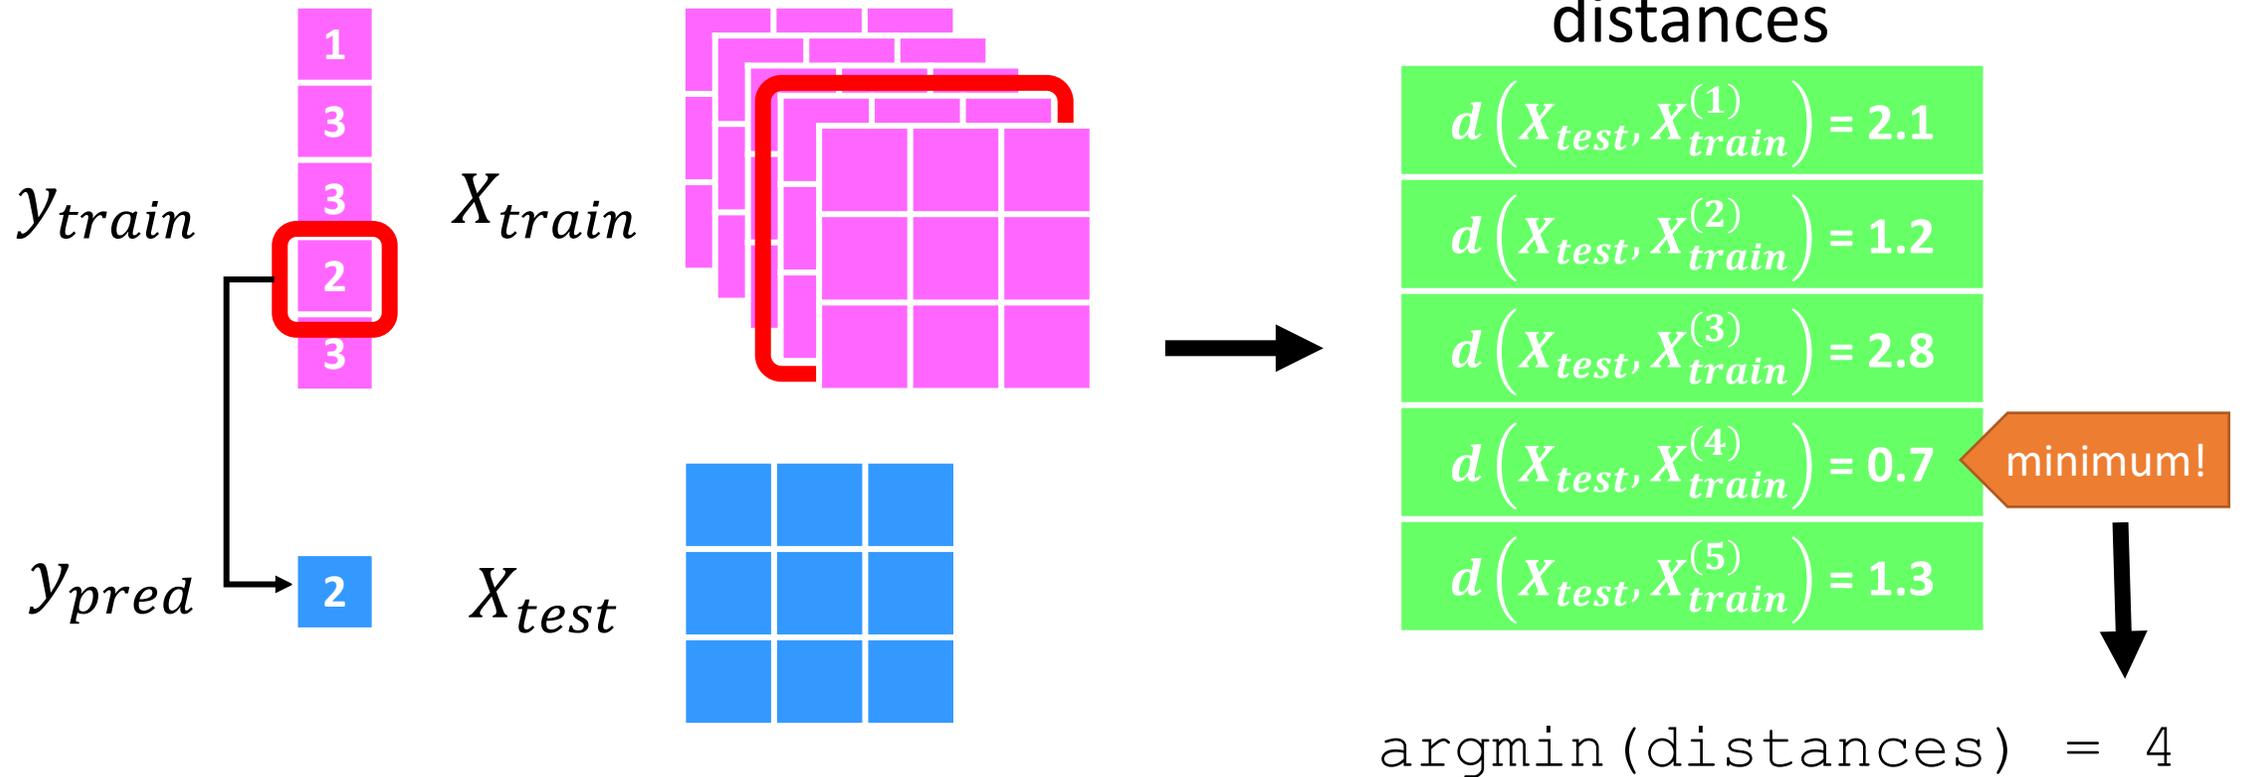
# Nearest Neighbors

A small example:



# Nearest Neighbors

A small example:



# Nearest Neighbors: Algorithm

## Training time:

Save the data,  $(X_{train}, y_{train})$ .

**Testing time:** Given  $N_{test}$  test images and  $N_{train}$  training images:

For each test image

```
for i = 1:Ntest do:
    distances = [0, ..., 0] (vector of length Ntrain)
    for j = 1:Ntrain do:
        distances[j] = distance(Xtest[i], Xtrain[j])
    nearest_idx = argmin(distances)
    y_pred[i] = y_train[nearest_idx]
```

# Nearest Neighbors: Algorithm

## Training time:

Save the data,  $(X_{train}, y_{train})$ .

**Testing time:** Given  $N_{test}$  test images and  $N_{train}$  training images:

```
for i = 1:Ntest do:
```

```
    distances = [0, ..., 0] (vector of length  $N_{train}$ ) ← Initialize distances  
                                     to zero
```

```
    for j = 1:Ntrain do:
```

```
        distances[j] = distance(Xtest[i], Xtrain[j])
```

```
    nearest_idx = argmin(distances)
```

```
    y_pred[i] = y_train[nearest_idx]
```

# Nearest Neighbors: Algorithm

## Training time:

Save the data,  $(X_{train}, y_{train})$ .

**Testing time:** Given  $N_{test}$  test images and  $N_{train}$  training images:

```
for i = 1:Ntest do:
    distances = [0, ..., 0] (vector of length Ntrain)
    {
        for j = 1:Ntrain do:
            distances[j] = distance(Xtest[i], Xtrain[j])
        nearest_idx = argmin(distances)
        y_pred[i] = y_train[nearest_idx]
    }
```

Calculate distance between current test image and each train image

# Nearest Neighbors: Algorithm

## Training time:

Save the data,  $(X_{train}, y_{train})$ .

**Testing time:** Given  $N_{test}$  test images and  $N_{train}$  training images:

```
for i = 1:Ntest do:  
    distances = [0, ..., 0] (vector of length  $N_{train}$ )  
    for j = 1:Ntrain do:  
        distances[j] = distance(Xtest[i], Xtrain[j])  
    nearest_idx = argmin(distances) ← Find the index of the  
    ypred[i] = ytrain[nearest_idx] nearest neighbor
```

# Nearest Neighbors: Algorithm

## Training time:

Save the data,  $(X_{train}, y_{train})$ .

**Testing time:** Given  $N_{test}$  test images and  $N_{train}$  training images:

```
for i = 1:Ntest do:
```

```
    distances = [0, ..., 0] (vector of length  $N_{train}$ )
```

```
    for j = 1:Ntrain do:
```

```
        distances[j] = distance(Xtest[i], Xtrain[j])
```

```
    nearest_idx = argmin(distances)
```

```
    y_pred[i] = y_train[nearest_idx] ← Assign the nearest neighbor's label to the current test image
```

# Exercise!

0	1	0
2	-1	0
3	-1	1

Image A  
label: 3

1	-4	2
0	1	0
-3	1	0

Image B  
label: 1

-1	0	2
-2	3	0
3	-1	1

Image C  
label: 2

Training data

1	0	4
-2	1	0
3	-1	3

Test Image

**Exercise:** Classify this image using nearest neighbors

# Exercise: Solution

0	1	0
2	-1	0
3	-1	1

1	-4	2
0	1	0
-3	1	0

-1	0	2
-2	3	0
3	-1	1

Image A  
label: 3

Image B  
label: 1

Image C  
label: 2

Training data

1	0	4
-2	1	0
3	-1	3

Test Image

$$\text{distance}(\text{test}, A) = \sqrt{42} = 6.48$$

$$\text{distance}(\text{test}, B) = \sqrt{73} = 8.54$$

$$\text{distance}(\text{test}, C) = \sqrt{16} = 4$$

# Exercise: Solution

0	1	0
2	-1	0
3	-1	1

Image A  
label: 3

1	-4	2
0	1	0
-3	1	0

Image B  
label: 1

-1	0	2
-2	3	0
3	-1	1

Image C  
label: 2

Training data

1	0	4
-2	1	0
3	-1	3

Test Image

$$\text{distance}(\text{test}, \text{A}) = \sqrt{42} = 6.48$$

$$\text{distance}(\text{test}, \text{B}) = \sqrt{73} = 8.54$$

$$\text{distance}(\text{test}, \text{C}) = \sqrt{16} = 4$$

**Nearest Neighbor:** Image C

**Predicted Label = 2**

# How do we evaluate our model?

**Accuracy:** Percentage of correct classifications made by the model.

$$\text{accuracy} = \frac{\# \text{ correct predictions}}{N_{test}}$$

Total number of  
data points tested

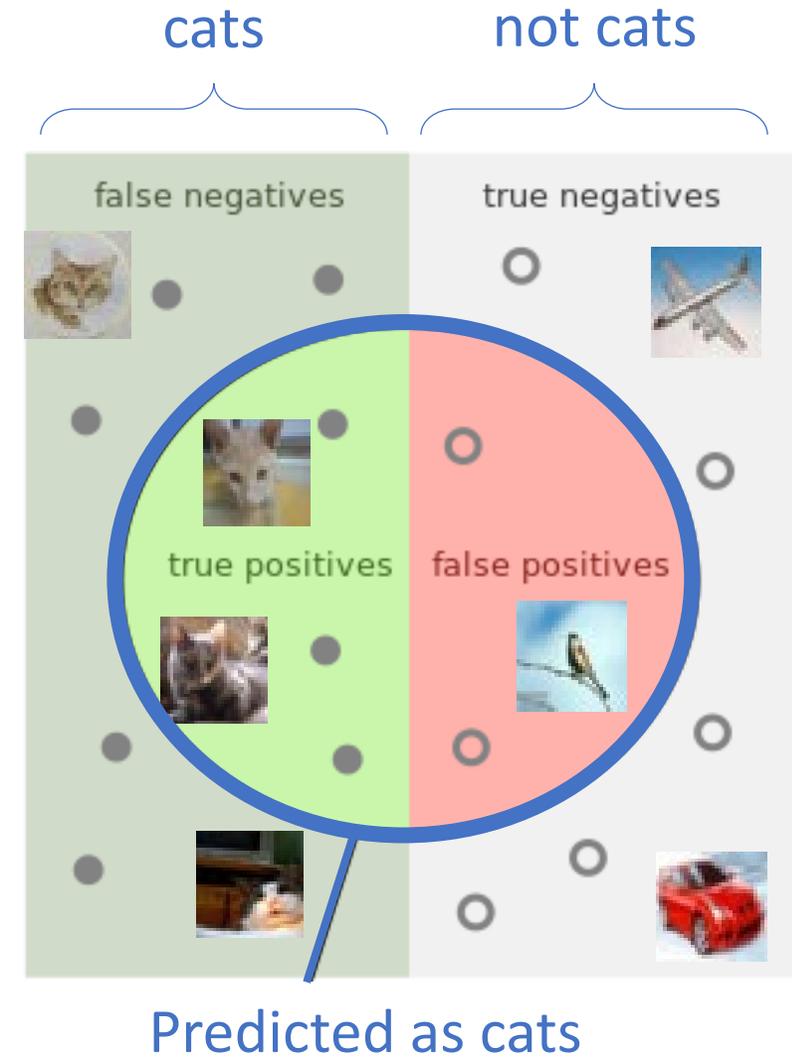


Quick, easy to interpret measure of how good the prediction is. Doesn't show why / how we're failing.

# Evaluation: Types of Error

Say we have a binary classification problem, where a data point can be classified as one of two options.

Ex: Cat detector (1 = cat, 0 = not cat)



# Evaluation: Types of Error

Say we have a binary classification problem, where a data point can be classified as one of two options.

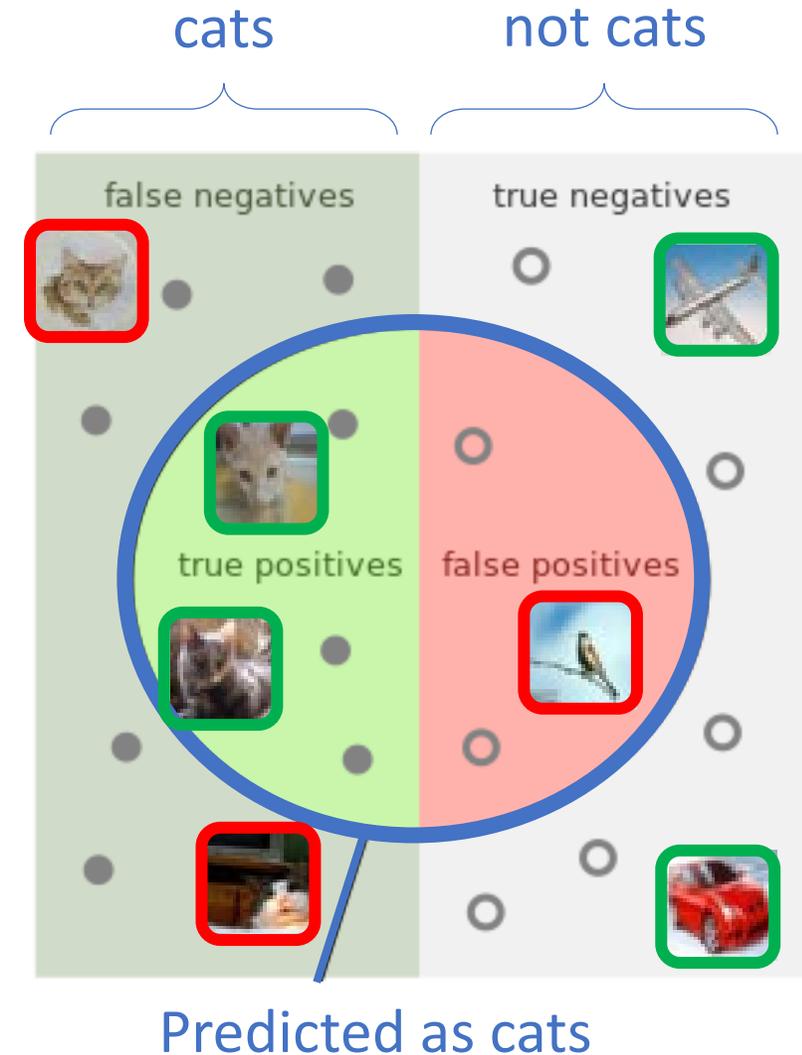
Ex: Cat detector (1 = cat, 0 = not cat)

**True positive:** Cat correctly classified as cat.

**False positive:** Non-cat incorrectly classified as cat.

**True negative:** Non-cat correctly classified as not cat.

**False negative:** Cat incorrectly classified as not cat.



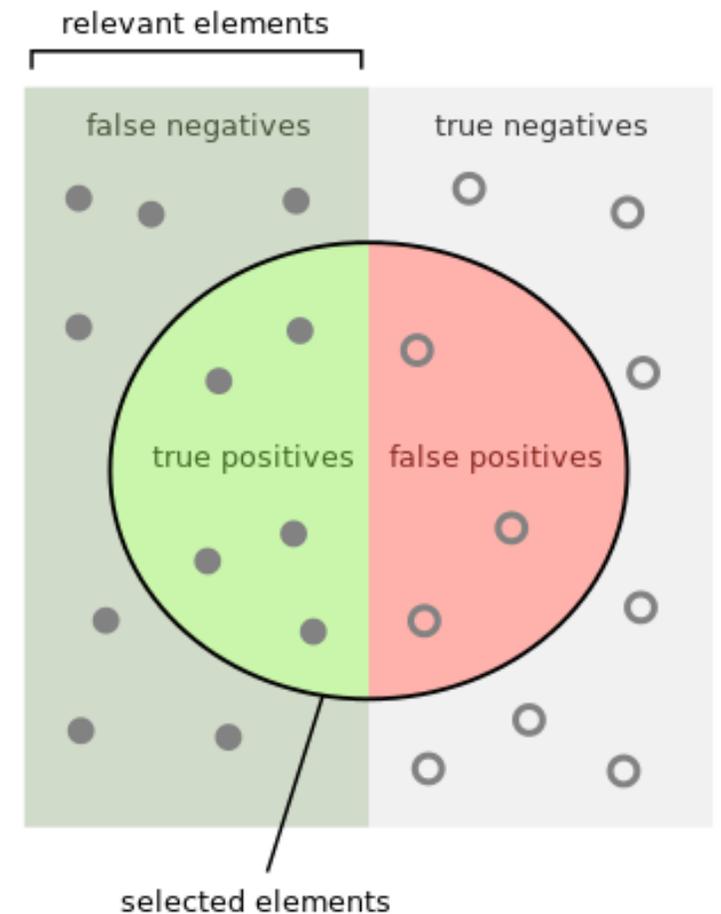
# Evaluation: Precision & Recall

**Precision:** How valid the results are.

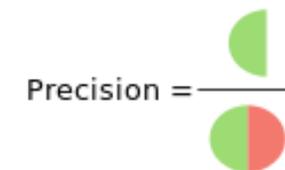
$$\text{precision} = \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false positives}}$$

**Recall:** How complete the results are.

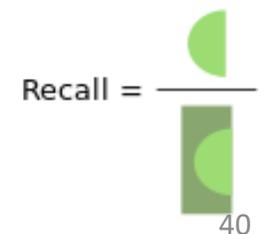
$$\text{recall} = \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false negatives}}$$



How many selected items are relevant?



How many relevant items are selected?



# Evaluation: Precision & Recall

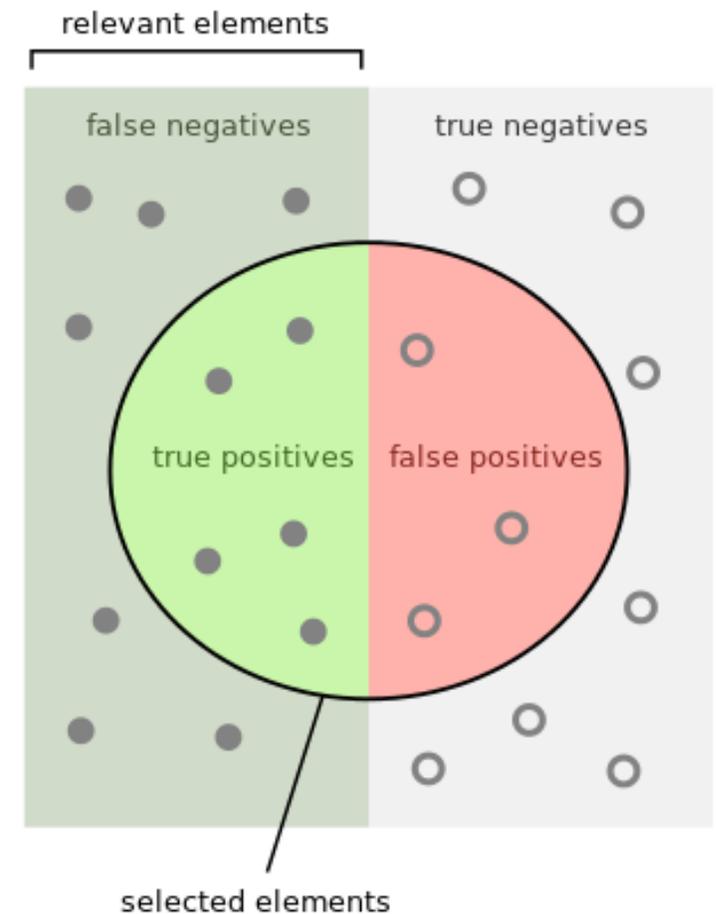
**Precision:** Helpful when it's important to have few false positives.

- Ex: A search engine should not show any irrelevant results, but it's okay to miss some relevant ones.

**Recall:** Helpful when it's important to have few false negatives.

- Ex: If a cancer detection algorithm gives a false negative, that's VERY bad! If there are some false positives, that's not so bad.

The metric chosen depends on the application!



How many selected items are relevant?

$$\text{Precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}$$

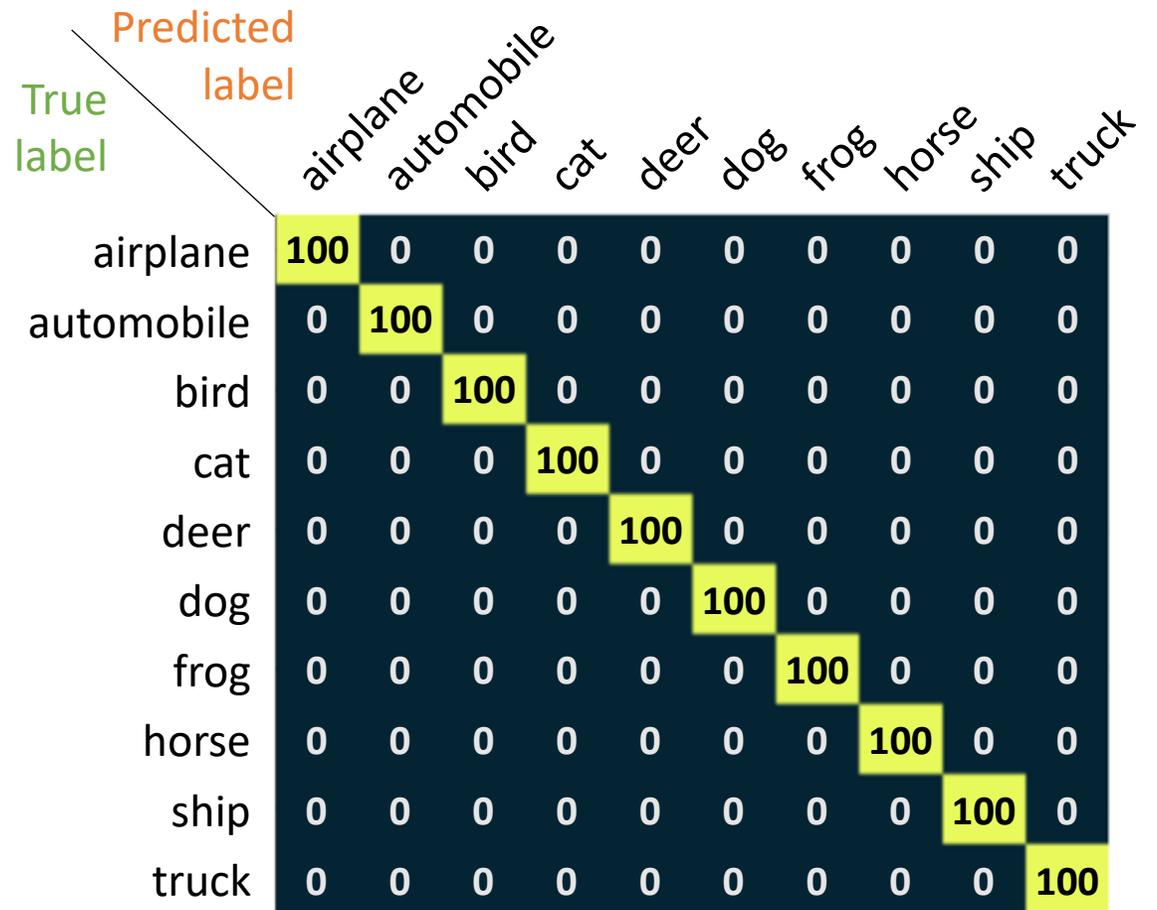
How many relevant items are selected?

$$\text{Recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}$$

# Evaluation: Confusion Matrix

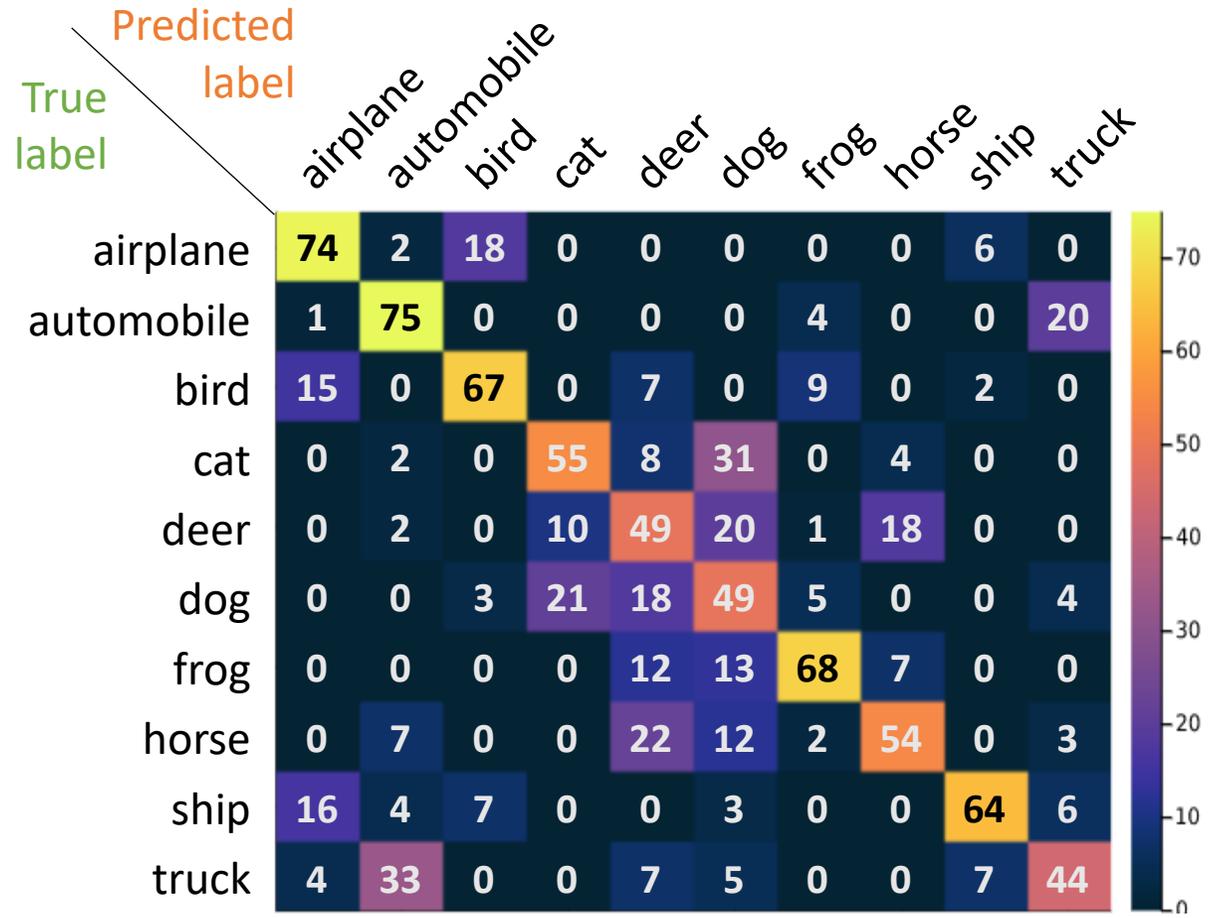
A heatmap of classification labels vs true labels.

For a perfect classifier, the confusion matrix looks like this.



# Evaluation: Confusion Matrix

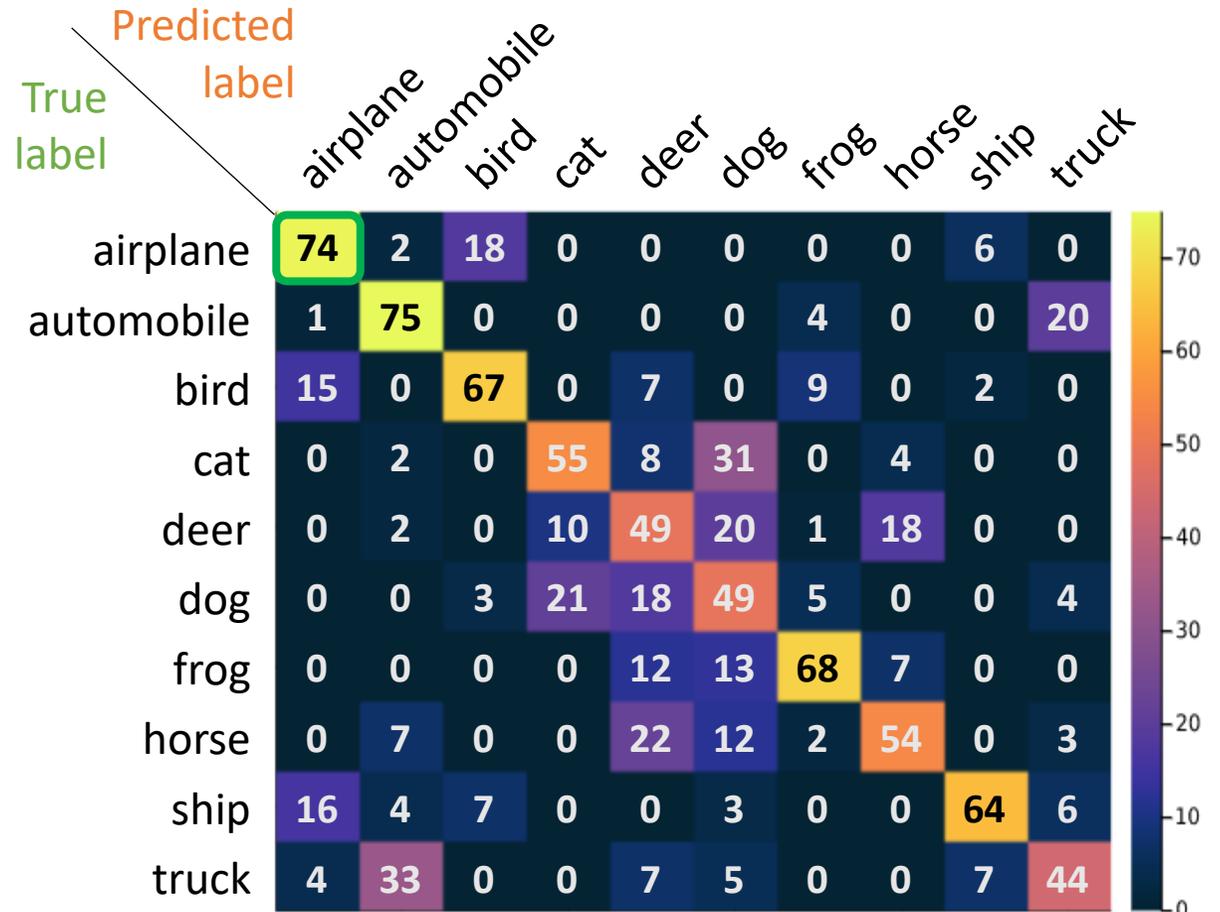
A heatmap of classification labels vs true labels.



# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

# airplanes classified as airplanes

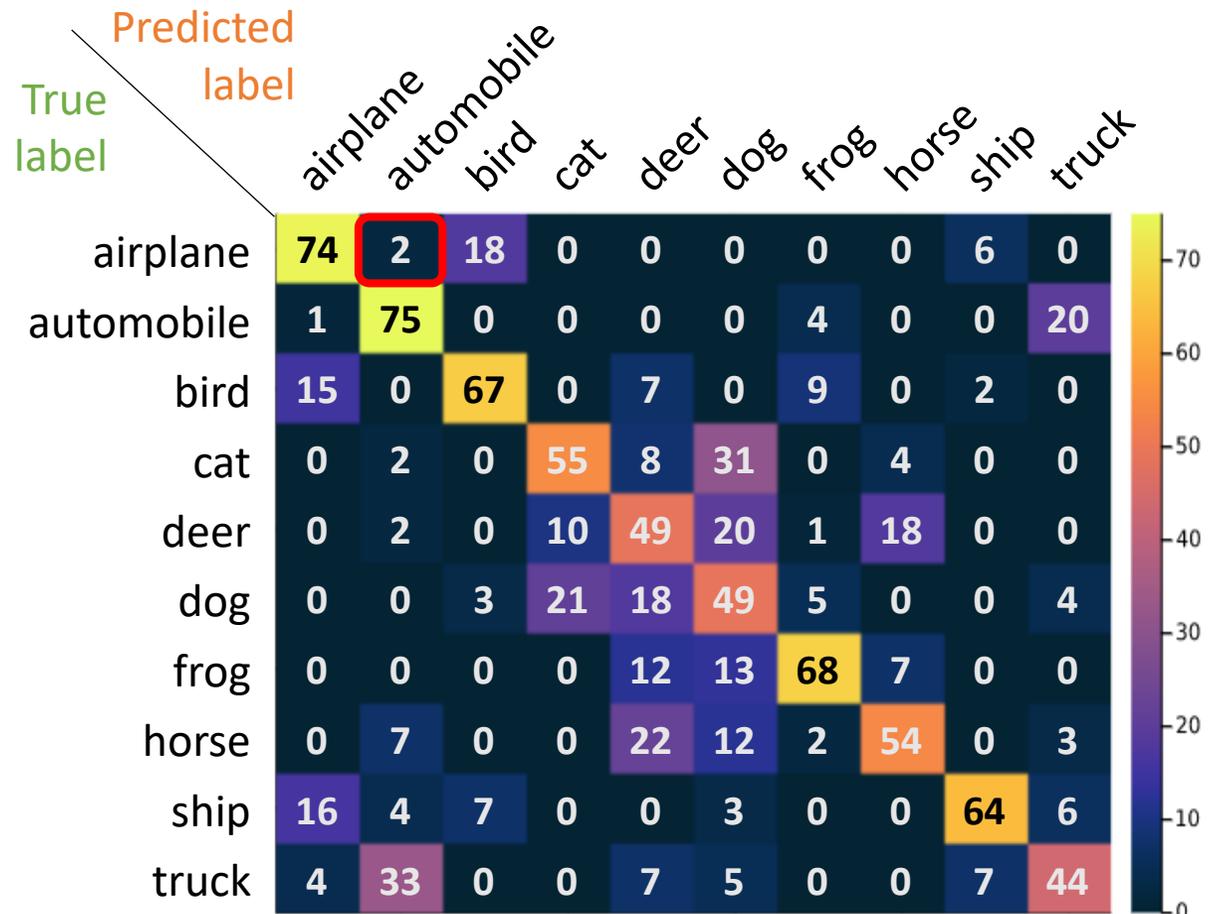


# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

# airplanes classified as airplanes

# airplanes classified as automobiles



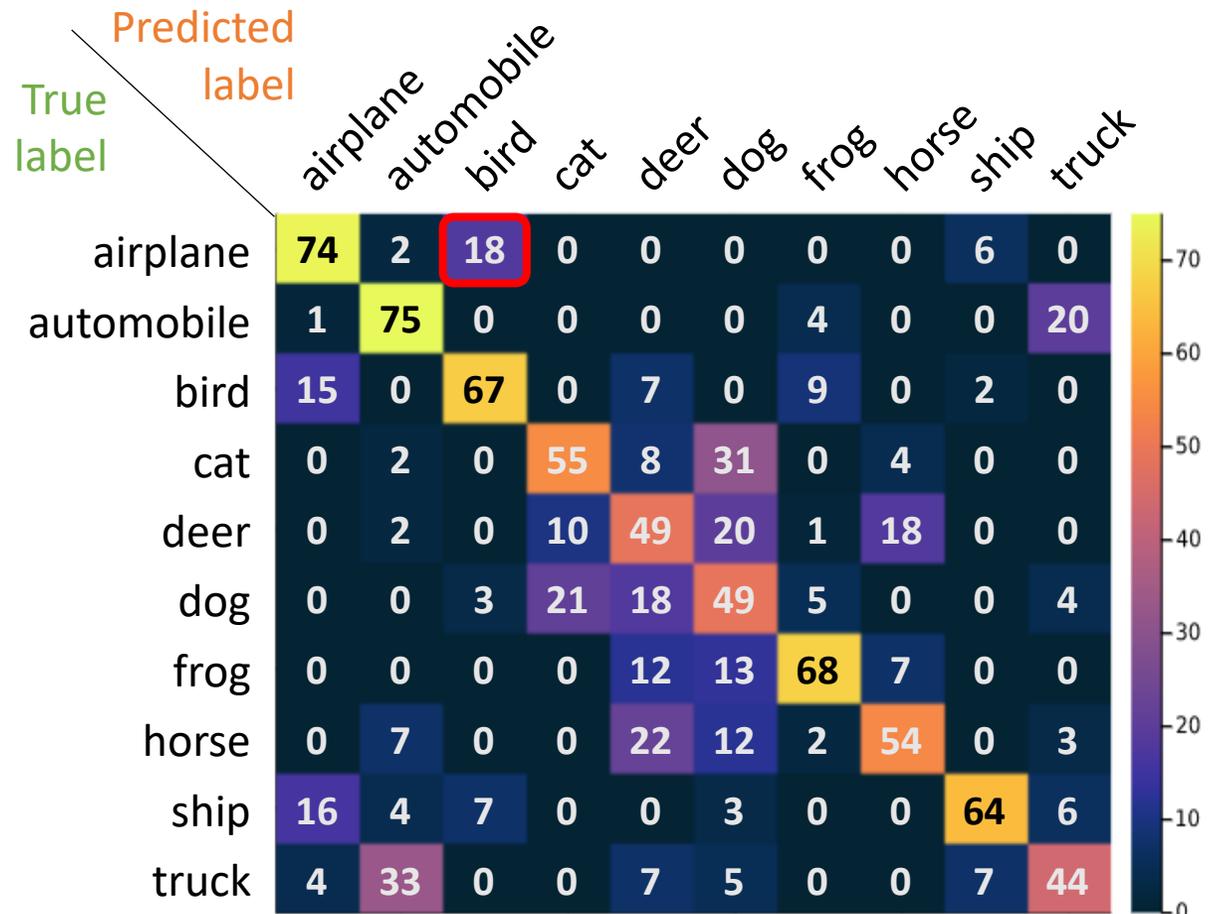
# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

# airplanes classified as airplanes

# airplanes classified as automobiles

# airplanes classified as birds



# Evaluation: Confusion Matrix

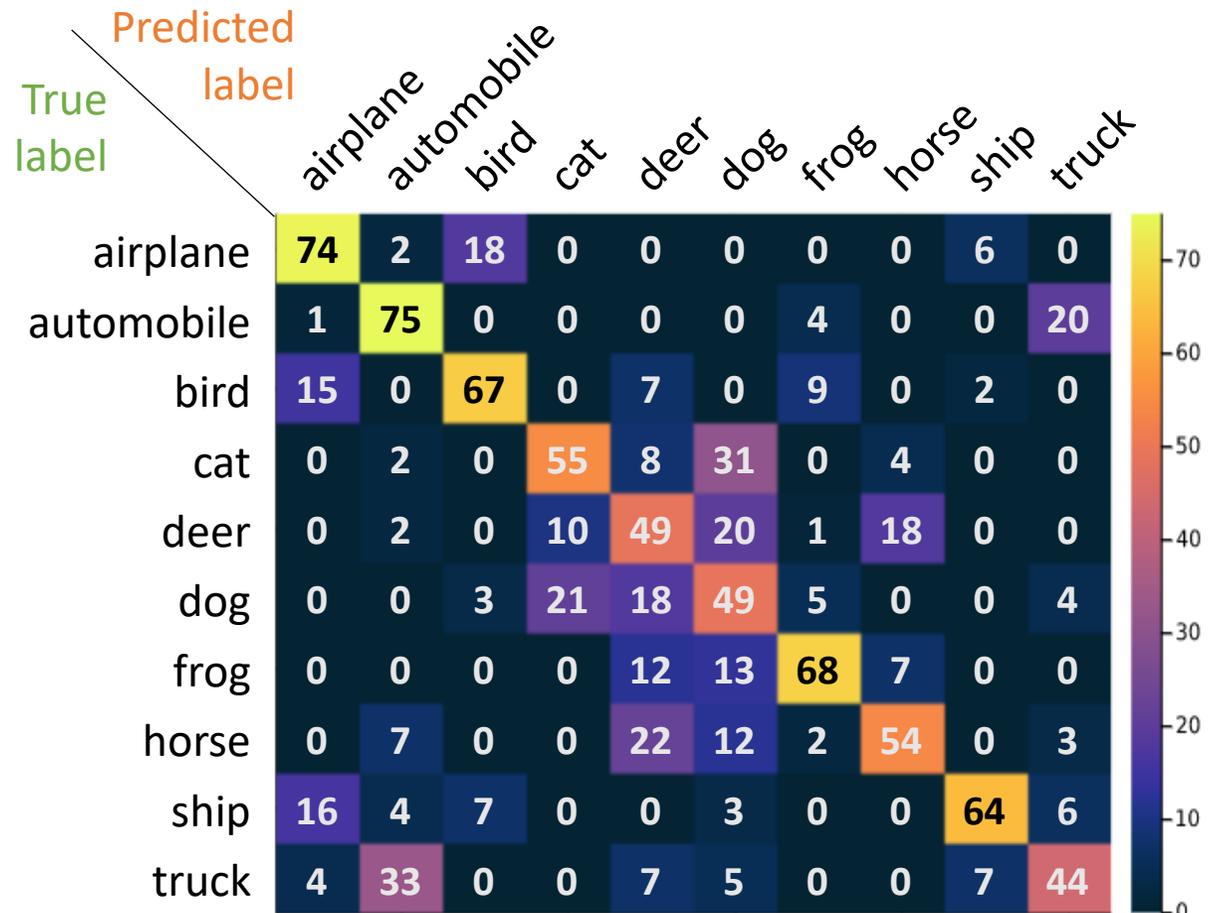
A heatmap of classification labels vs true labels.

# airplanes classified as airplanes

# airplanes classified as automobiles

# airplanes classified as birds

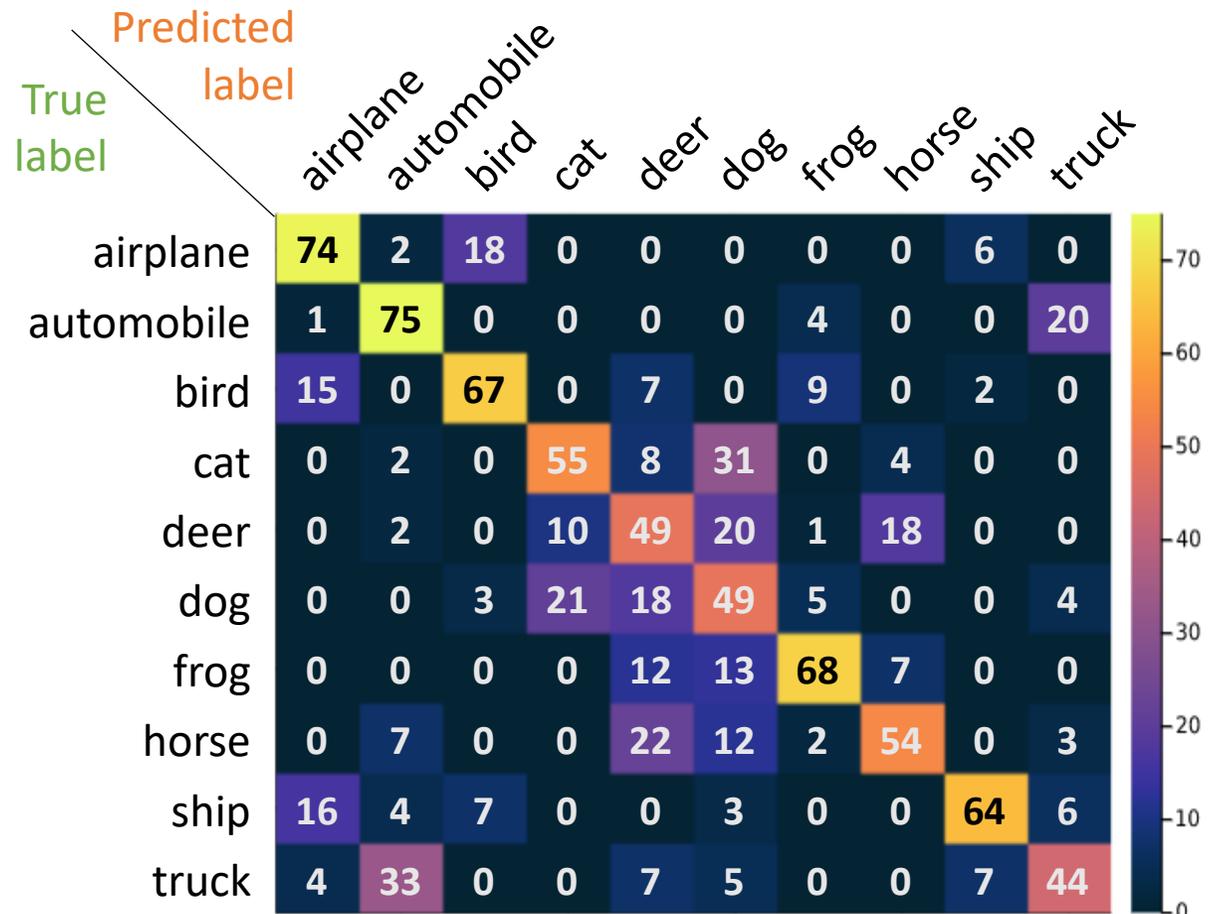
etc...



# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

The confusion matrix gives us more insight about where the algorithm is failing.

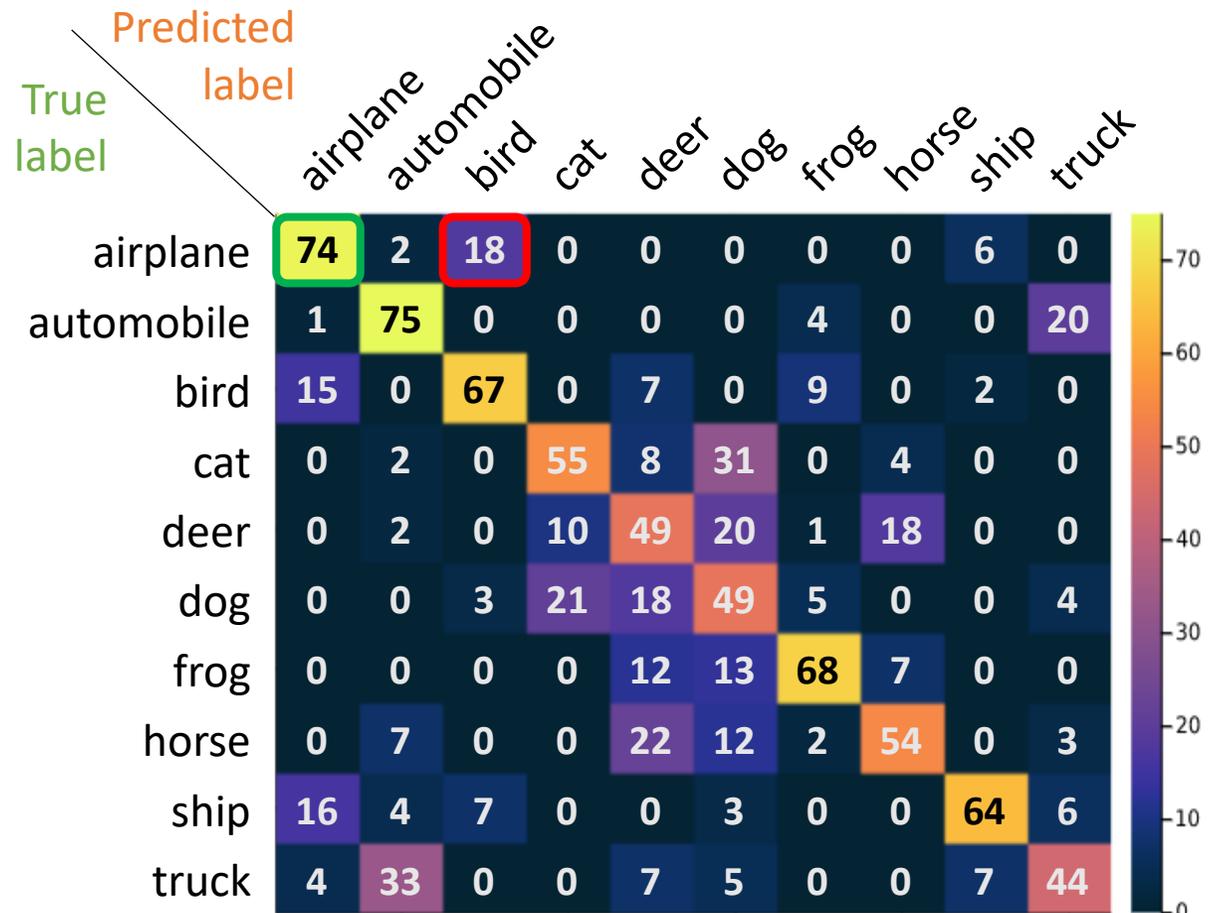


# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

The confusion matrix gives us more insight about where the algorithm is failing.

**Airplanes** are being misclassified as **birds**.

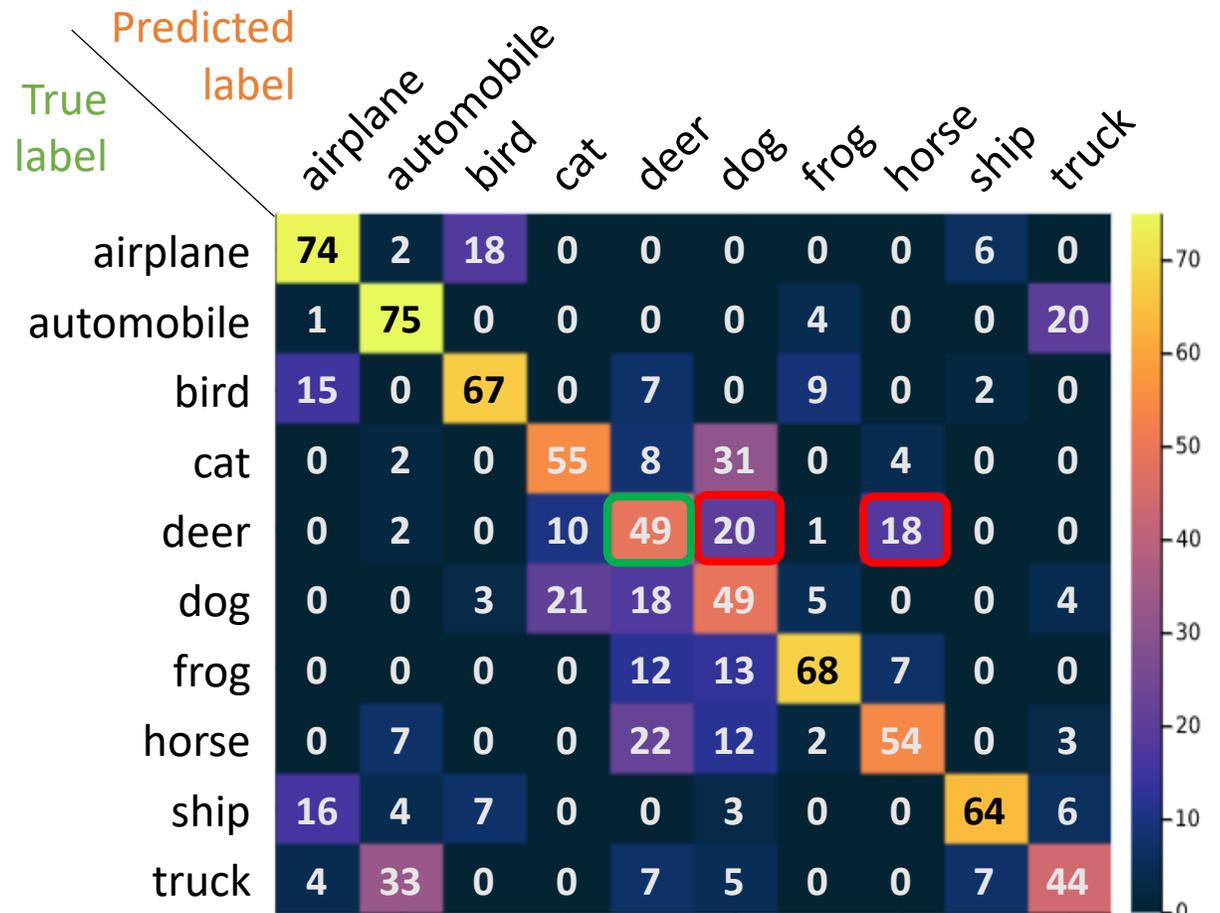


# Evaluation: Confusion Matrix

A heatmap of classification labels vs true labels.

The confusion matrix gives us more insight about where the algorithm is failing.

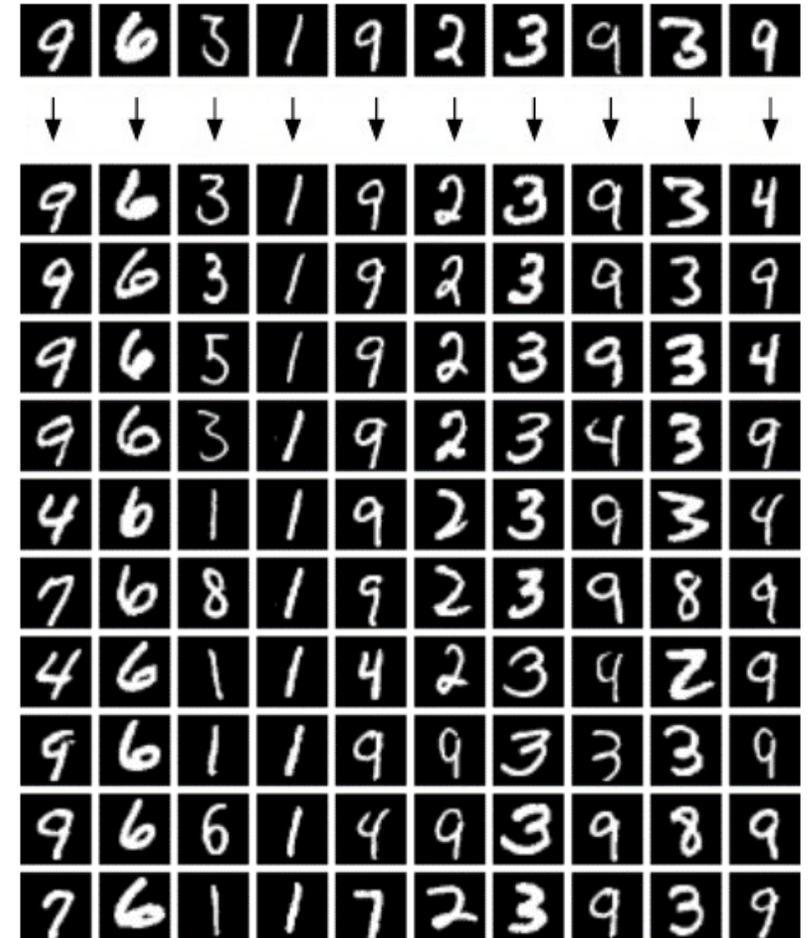
**Deer** are hard to classify. They are being labelled as **dogs** and **horses**.



# k-Nearest Neighbors

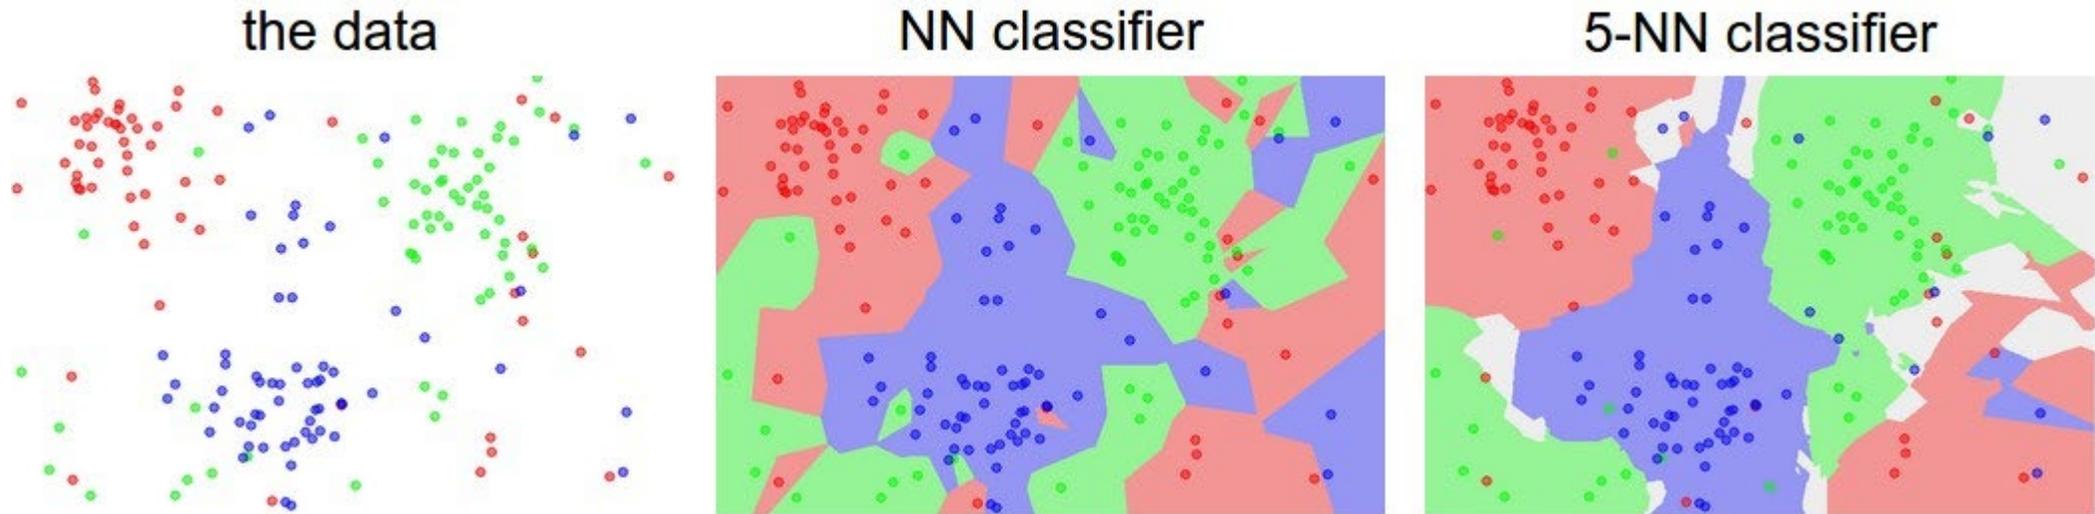
To make our algorithm more robust, we can let the  $k$  nearest neighbors vote on the label for the test image. This is the “ $k$ ” in  $k$ -nearest neighbors (kNN).

Up until now, we were describing 1-NN.



# k-Nearest Neighbors

The decision boundaries for our data change. We are overfitting less.



Experiment with kNN here: <http://vision.stanford.edu/teaching/cs231n-demos/knn/>

# Hyperparameters

How do we pick  $k$ ?

$k$  is an example of a **hyperparameter**: a parameter we choose, which isn't learned.

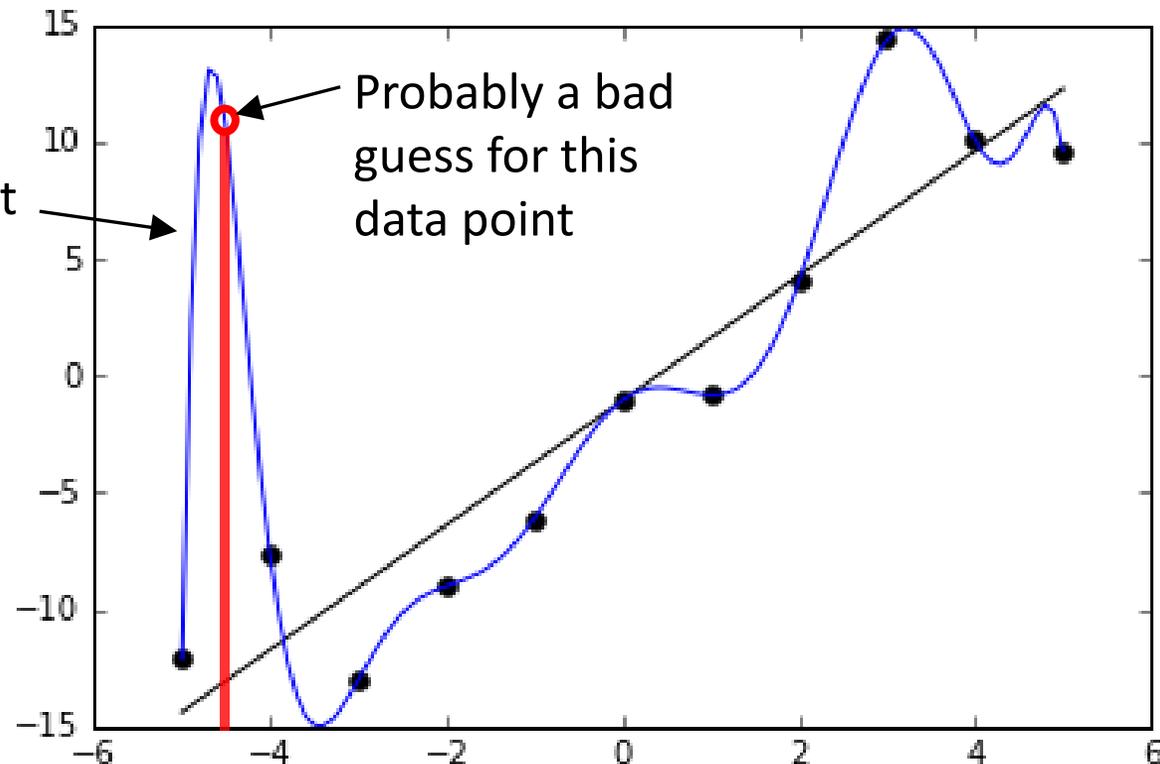
Generally, we need to tune these parameters by trying different values and selecting the best performing ones.

# Overfitting

**Overfitting** happens when we fit a model that corresponds too closely to our data. Overfitting impacts performance on new data.

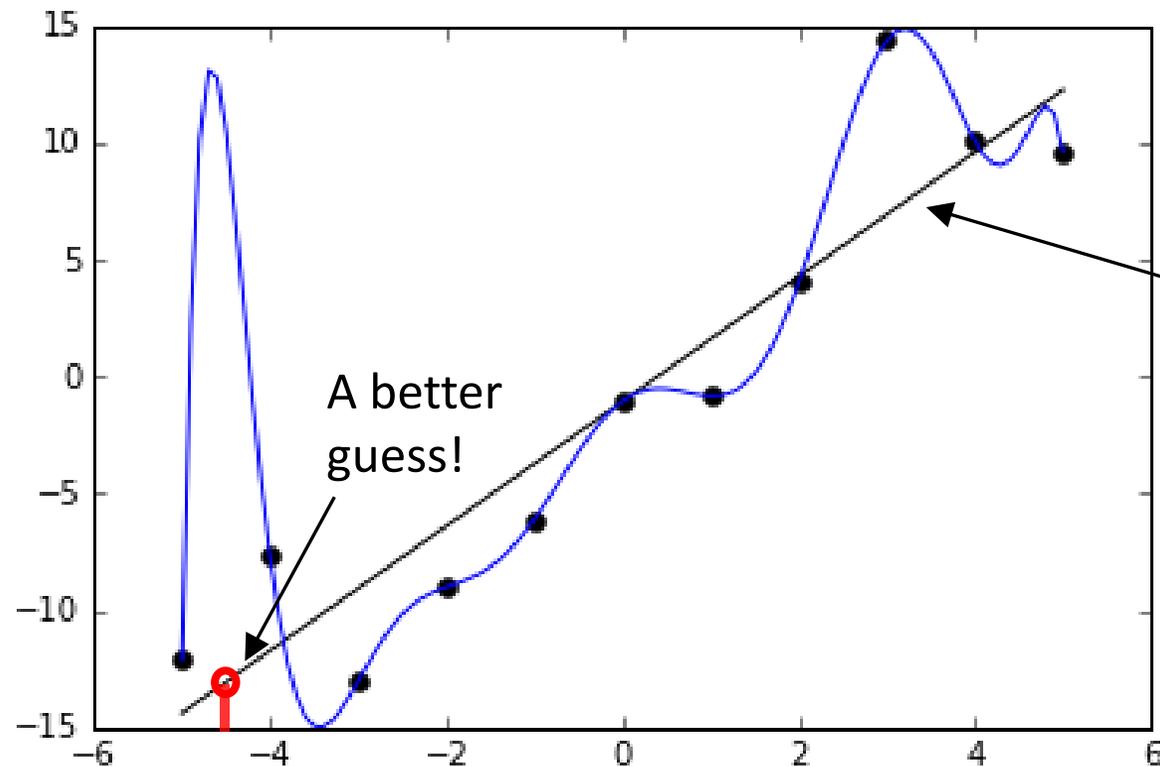
The blue line is a perfect fit for the given data.

But it's not a very good choice of model.



# Overfitting

**Overfitting** happens when we fit a model that corresponds too closely to our data. Overfitting impacts performance on new data.



A better  
guess!

This linear fit is likely a  
better choice.

Avoids overfitting.

# Overfitting

**Overfitting** happens when we fit a model that corresponds too closely to our data. Overfitting impacts performance on new data.

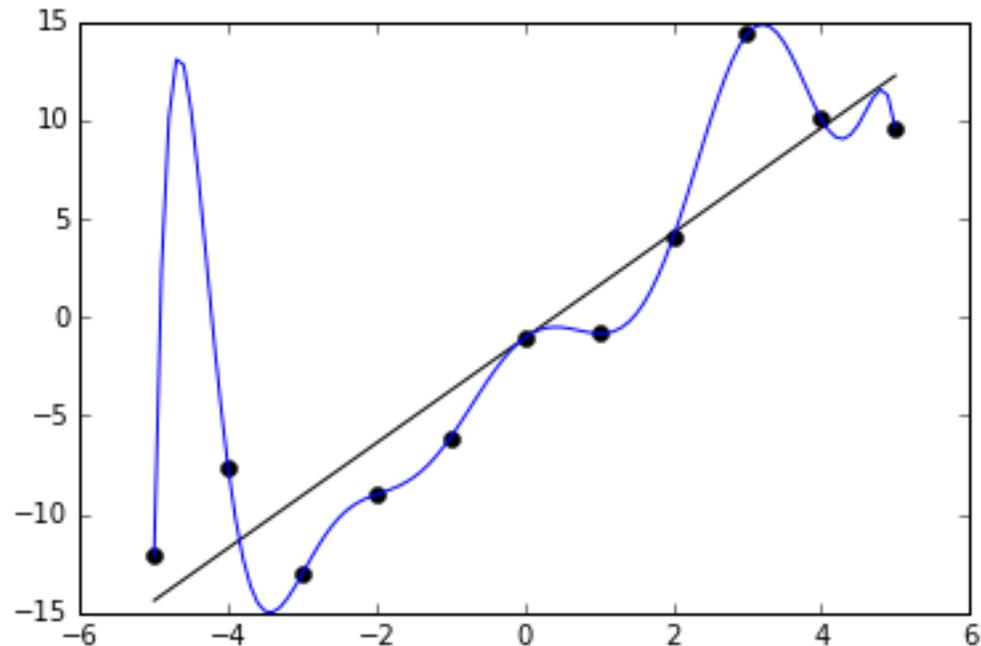


Image: CC 4.0 ([link](#))

A classification example

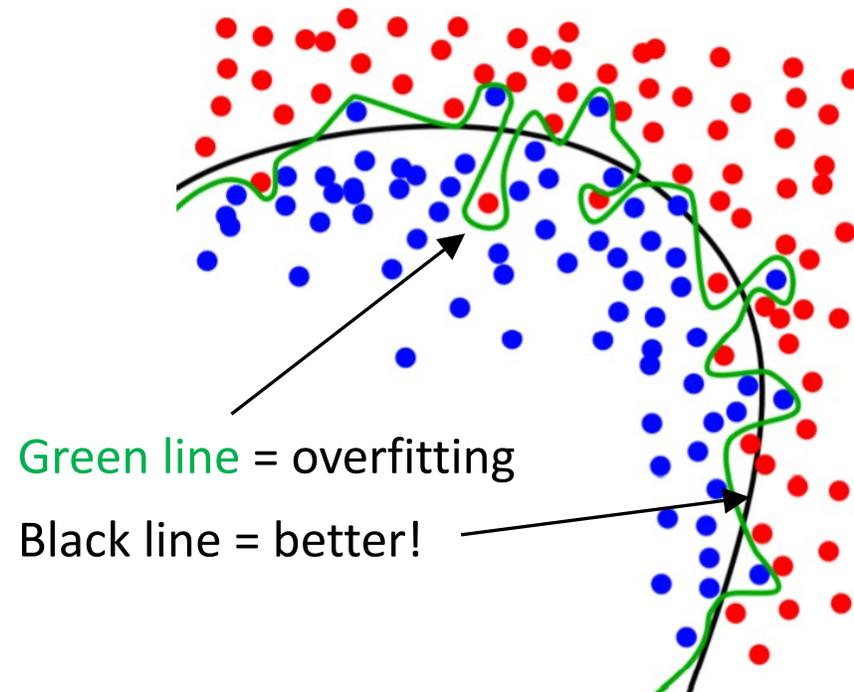


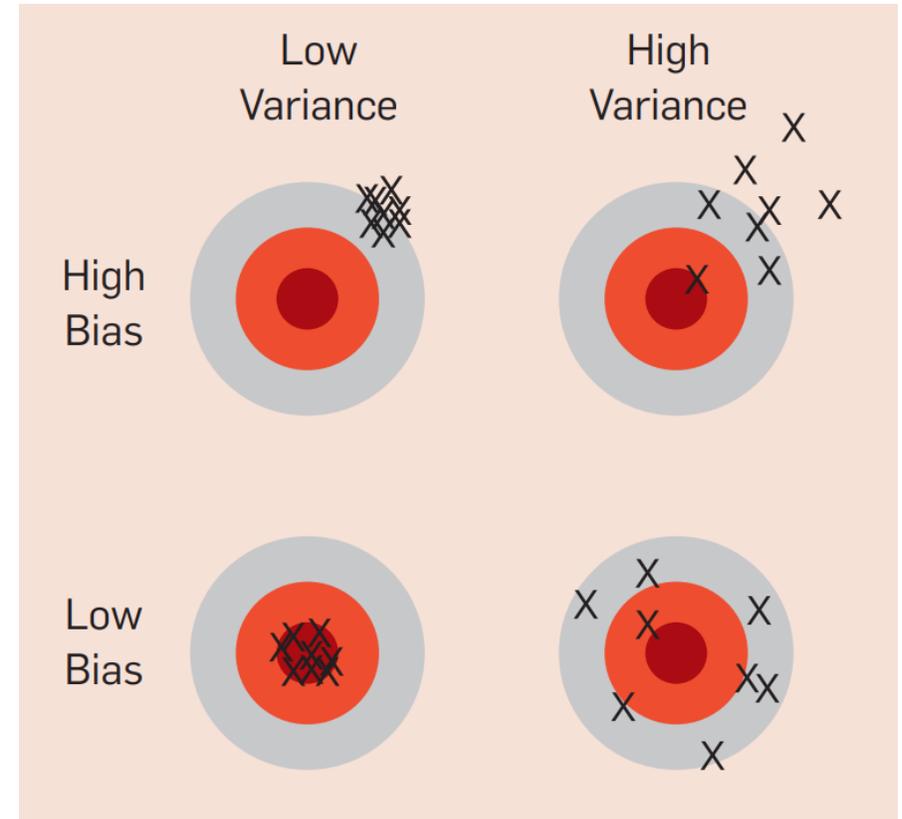
Image: CC 4.0 ([link](#))

# Bias-Variance Tradeoff

**Bias** is error due to deviation from the true value (underfitting).

**Variance** is error due to sensitivity to variations in the data (overfitting).

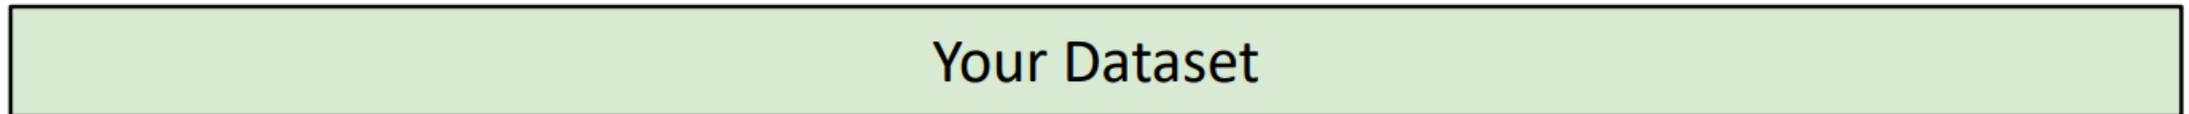
When choosing hyperparameters, we need to tradeoff between both.



# Setting Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

Bad idea 😞



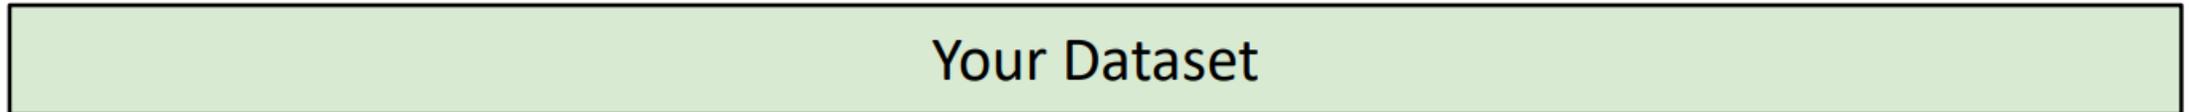
We are minimizing the **training error**. This is basically just “memorizing” the training data (overfitting!).

The training error should be low. Ex: For nearest neighbor,  $k=1$  will give zero training error. **Training error should only be used as a sanity check.**

# Setting Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

Bad idea 😞



**Idea #2:** Split data into **train** and **test**, choose hyperparameters that work best on test data

Better idea 😊



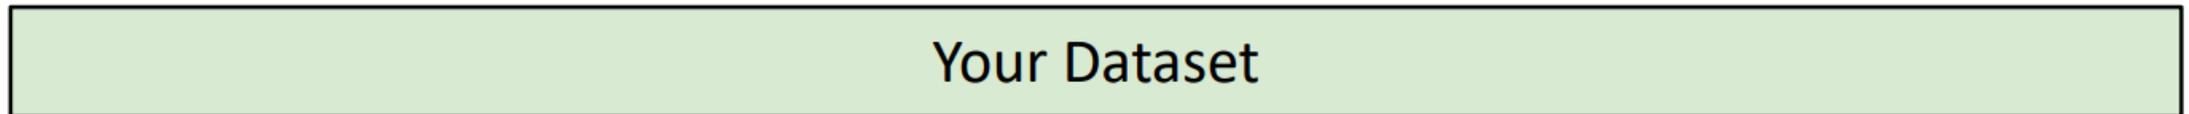
We are minimizing the **testing error**.

This is better, but we still don't know how we'll do on new data.

# Setting Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

Bad idea 😞



**Idea #2:** Split data into **train** and **test**, choose hyperparameters that work best on test data

Better idea 😊

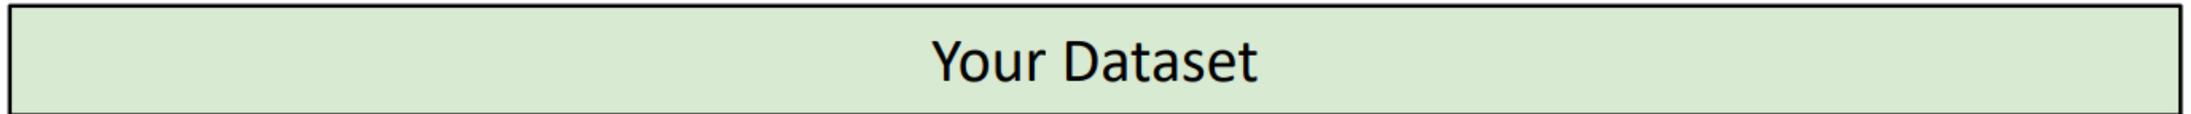


**Idea #3:** Split data into **train**, **val**, and **test**; choose hyperparameters on val and evaluate on test

Good idea 😊

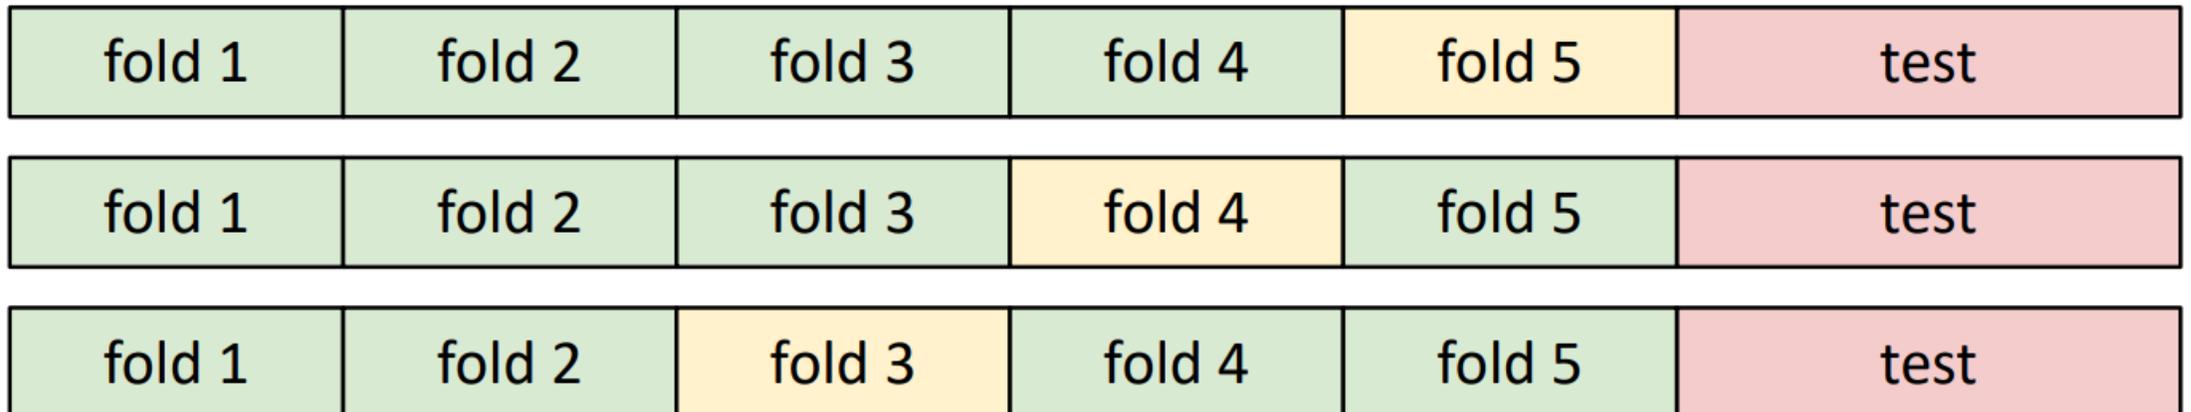


# Setting Hyperparameters



**Idea #4: Cross-Validation:** Split data into **folds**, try each fold as validation and average the results

Best idea! 😊



Useful for small datasets, but (unfortunately) not used too frequently in deep learning

# Project 4: Machine Learning

Implement three machine learning algorithms to classify images from the MNIST dataset.

## 1. **Nearest neighbors** (Today!)

- ✓ How to find the distance between images
- ✓ The nearest neighbors algorithm
- ✓ Evaluating classification algorithms
- ✓ Setting hyperparameters

## 2. Linear Classifier ← Next time!

## 3. Neural Network