part 2 | image classification

AI designers create image classification models (also known as image classifiers) to sort images. In this lesson, students will learn about the three steps of creating image classification models. They are:

1. collect and label the data (images),
2. train the model,
3. test the model.

For the main activity, students will act out the process of creating an image classification model with provided images of cats and dogs. They will:

1. sort the images into *cat* and *dog* categories,
2. make rules for each category (e.g. what do all the pictures have in common?),
3. use the rules to classify new images.

After, students will repeat these three steps when they use Teachable Machine to create a classifier for sorting images of doctors and nurses. Students round out the lesson talking about **data bias**, asking, *how can we reduce data bias in our models?*

Learning targets

Students will be able to:

- List the three steps of creating an image classification model. (Associated Standards)
- Define **data bias** and articulate some causes and remedies of data bias in image classification. (Associated Standards)

Students will...

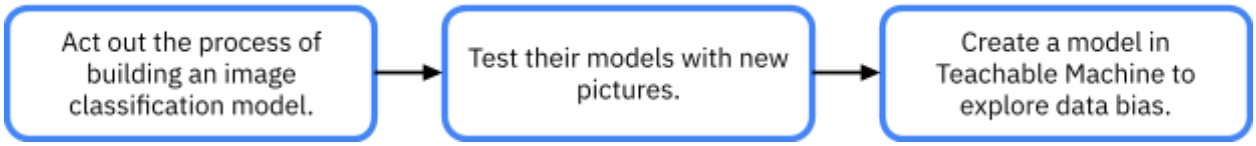

Key vocab

- Class: a labeled group of images that share a common feature.
- Model: A finished classifier, trained on thousands of inputs.

- Data bias: when an image classification model makes incorrect predictions because it doesn't have a diverse enough group of input images.

Materials

- Scissors and glue sticks

Prep

- Prepare the cat/dog datasets. Either print the printable datasets or distribute the digital version to students.
    - ↳ Both files contain the same six datasets. Each pair of students should receive one dataset.
- Print the dog and cat test photos, one per table group.
- Ensure students have access to Teachable Machine.

---

warm-up | whole class | 3-5 min

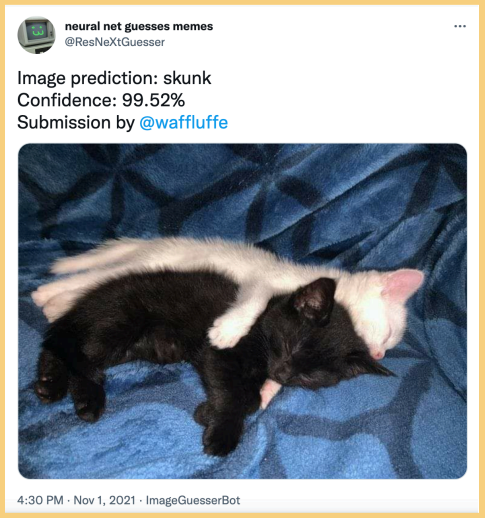

Show students the picture of two kittens in the classroom slides. Ask:

- What animal do you see in the picture?
- How confident are you, on a scale of 1-100?
- What do you think a computer would see? Why?

Once students have responded, reveal that an AI model actually classified the picture as a skunk!

Prompt students to think-pair-share:

- Why do you think the computer was tricked?

intro | whole-class | 5-7 min

Preview the day's learning targets with students.

Remind students that image classification is an AI process that tells pictures apart. To train an image classification model, a designer follows three steps:

1. Collect and sort the data

    - Humans collect thousands of pictures and sort them into groups called **classes** (for example, cats vs. dogs). Some of these photos are saved for the testing in step 3.
2. Train the model
    - The neural network trains the **model** by taking the pictures it was given and finding patterns in pictures of the same group.
3. Test the model
    - Both the computer and the human designer test the model by giving it new pictures to classify.

Today, you'll be acting out the steps of creating an image classification model!

vocab spotlight

A **class** is a labeled group of images that share a common feature.

A **model** is a finished classifier, trained on thousands of inputs.

check for understanding

Let's consider an example. A designer is making an image classification model for apples vs. oranges.

Which step of the process does each following action fall under? Allow students to discuss in groups or as a class.

- Action: The designer holds up a fruit to the camera and sees if the model correctly classifies it.
    - ↳ Step 3 - Test the model
- The designer gathers images of apples and oranges and labels each one "apple" or "orange."
    - ↳ Step 1 - Collect and label the data
- The model studies the images and finds patterns for each group.
    - ↳ Step 2 - Train the model

act it out | small groups | 18-20 min

Explain: In this activity, we are going to act out the process of making an image classification model that can tell the difference between a cat picture and a dog picture.

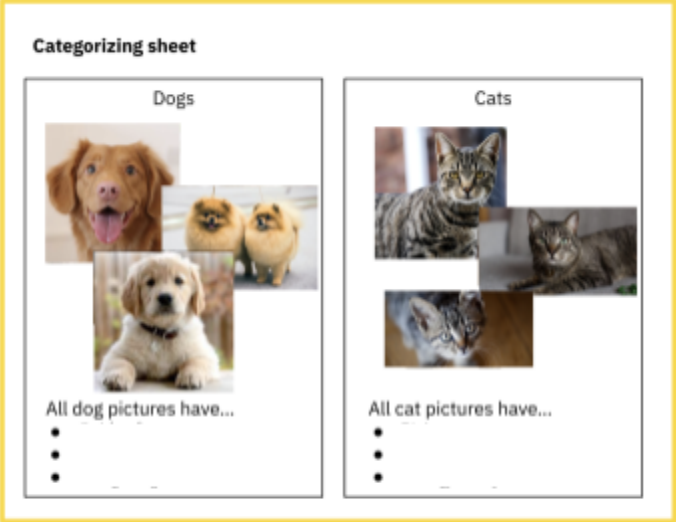

Example student work after step 1

Step 1: Collect and label the data.

- Explain: let's begin! The data has already been collected for you, so all you need to do is sort (label) it.
- Distribute one categorizing sheet and one dataset from the cat/dog datasets to each pair of students and prompt students to cut out the images and glue them in the correct category.

Step 2: Train the model.

- Prompt students to write down common patterns within each group.
- Have students consider:
  - ↳ What items are common in each class?
  - ↳ What colors show up often?
  - ↳ What shapes are always present?
  - ↳ What backgrounds are common within classes?
  - ↳ What other patterns do you see in each group?
- Based on the patterns students notice, have them write rules about each class (e.g., "all the cat pictures have flowers"). By creating rules, students have now trained their models, just like a neural network creates a set of mathematical rules when training an image classification model.

facilitator note

Encourage students to stay true to their role as "AI" and *only* make rules based on the pictures they see! A real image classification model only has the images it's given, unlike humans with our many years of experience seeing cats and dogs.

Step 3: Test the model.

- Show students new images on the classroom slides. Prompt groups to shout out the class that represents their classification, based only on the rules they generated. For each example, students should select the category where the most rules are met.

Reflect

- Prompt students to reflect on their classifications:
  - ↳ What test pictures worked well for your model?
  - ↳ What pictures did your model classify incorrectly?

responsible design | pairs | 13-15 min

Prompt students to open Teachable Machine and recreate their dog and cat classes without adding or changing any images. Once their models are trained on Teachable Machine, prompt students to test their models with test photos. Ask students:

- How did your model you built classify this image?
- How confident is the model?
- Why was your prediction right or wrong?
- How could you improve your dataset to be more effective in this case?

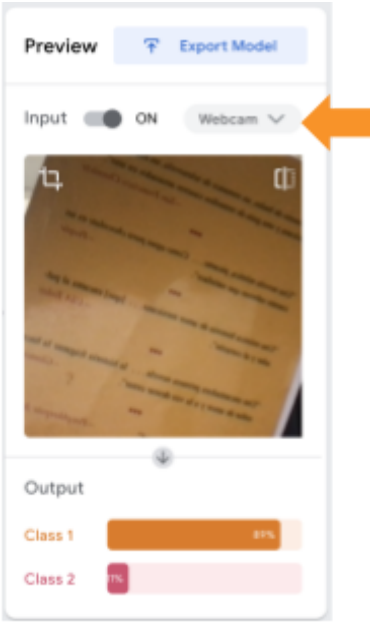

Change the input type to allow photos to be uploaded.

Explain that many real-life computer vision systems (like ones that identify people in crowds) perform worse for people with deeper skin complexions, for women, and for gender-nonconforming people. This is **data bias**.

vocab spotlight

Data bias happens when an image classification model makes incorrect predictions because it doesn't have a diverse enough group of input images.

Ask:

- How is data bias created?

- Whose job is it to fix data bias?
- How can we find out if we have data bias in our own projects?
- How can we fix data bias in our projects?

If time allows, prompt students to share their datasets and retrain their models to be less prone to data bias.

> **facilitator note**
>
> Access the data bias teacher support for more information about the topic, prompts for further discussion, and links to external resources.

exit ticket | independent | 3-5 min

Prompt students to complete the exit ticket.

> **facilitator note**
>
> The exit ticket question about bias could also work well as a class discussion. Consider revisiting it in Part 3.