

Cats, Dogs & Machine Learning

By: Activities from [MIT Media Lab](#)

Duration: 2 hours

LEVEL	SUBJECTS	PROVINCES / TERRITORIES	TOOL
Grades 4-6, 7-8	Science & Technology; Social Studies	Across Canada	The Teachable Machine

Overview

Students learn about supervised machine learning through building a cat-dog classifier using Google's Teachable Machine tool, and discuss the reality of algorithmic bias.

Prep Work

- This activity requires WiFi and laptops or chromebooks with webcams
- Introduce the concept of AI: <https://www.canadalearningcode.ca/lessons/learn-like-a-computer/>
- Print out cat-dog images (<http://bit.ly/cat-dog-image-dataset>) and worksheets per pair of learners (<http://bit.ly/algorithmic-bias-activity-sheet>).
- Divide learners into pairs, with one computer per pair.

Materials

- Slides (<http://bit.ly/mit-supervised-learning-slides> and

Key Coding Concepts

- ✓ Algorithms
- ✓ Debugging
- ✓ Machine Learning

Terminology

Algorithms are step-by-step sets of operations to be performed to help solve a problem

Artificial Intelligence is the ability of a machine or computer program to think and make decisions.

<http://bit.ly/mit-algorithmic-bias-slides>)

- Worksheet
- Laptops or chromebooks + chargers
- Dog/cat cards

Lesson

Review

Ask: What do autocorrect, Alexa, and Snapchat filters have in common? (A: They all use AI)

Ask: What is AI? (See Terminology for definition)

These computer programs or machines can be trained to learn from experience, and even learn by themselves in order to make predictions based off of the information (or data) that we give them. This is called “machine learning.”

Ask: Why do you think technology like this is important? *(Potential answers: To get answers to questions more quickly (DNA matching, voice-assistant technology), to make things work more efficiently (predicting traffic, timing of stop lights))*

Today, we are learning about machine learning by training a computer program to classify cats and dogs.

Activity 1: Intro to Supervised Machine Learning

Use the “Intro to Supervised Machine Learning” slides (<http://bit.ly/mit-supervised-learning-slides>) to introduce the activity. Use presenter mode to view the instructor notes as you teach, or print the slides with notes ahead of time.

Slide 10: Demo how to use the Teachable Machine before having learners open their computers, or play this video instead: <https://youtu.be/3BhkeY974Rg>

Machine learning provides the ability for a machine or program to learn from data or experience.

Curricular Connections

Automation, Computer technology, Trial and error, Problem-solving, Bias.

References

The Teachable Machine by Google.

<https://teachablemachine.withgoogle.com/>

Activities and slides from: An Ethics of Artificial Intelligence Curriculum for Middle School Students was created by Blakeley H. Payne with support from the MIT Media Lab Personal Robots Group, directed by Cynthia Breazeal.

<https://docs.google.com/document/d/1e9wx9oBg7CR0s5O7YnYHVmX7H7pnlTfoDxNdrSGkp60/edit>

**Be sure to note the difference between training data and test data in your demo by saying, “And now I want to test how well my algorithm does with data that is similar but slightly different to what it has seen before”*

Give students time to each try the demo, ~5-8 minutes. Walk around and prompt students to tell you about their training and test datasets. What happens when they change their pose? What happens when both partners are in the frame?

Have learners close their computers and discuss the following as a class:

- What happens if you only train one class? (*A: Only one prediction*)
- What happens as you increase your dataset? (*A: It becomes more accurate*)
- What happens when your test dataset is different from your training dataset? (*A: It might not be as accurate*)

Great. Now we are going to build a cat-dog classifier using the teachable machine.

Activity 2: Intro to Algorithmic Bias

Use the “Intro to Algorithmic Bias” slides to facilitate the activity.

(<http://bit.ly/mit-algorithmic-bias-slides>)

With your partner, I want you to build a machine that classifies cats and dogs.

Important: While training your classifier, try to only include ONE photo of each of the training images.

Hand out printed cat and dog images and Algorithmic Bias worksheets to each pair

Ask: You’ll notice that the Teachable Machine is made up of three parts. What are each of these three parts called? (*A: (1) dataset, (2) learning algorithm, and (3) prediction. Learners will describe in their own words. Have them guess before sharing the official terms. Have them write down the answer on their worksheet.*)

Give students time to train and test their classifier.

(Note: If learners get frustrated by accidentally taking more than one photo and having to reset multiple times, they can continue with a few extra photos. Just have them keep this in mind while recording answers in their worksheet.)

Ask: How are your classifiers working?

(Note: Learners should notice that the classifier works better on cats than dogs).

Discuss as a class or in small groups:

- Is this classifier useful if it only works well on just cats?
- Why do you think it works better on cats vs dogs?
- How could we make it better with our training data? (If students have trouble, ask them to notice similarities in the dataset, e.g., *dogs were really fluffy and cat-like/not as diverse*)

When algorithms, specifically artificial intelligence systems, have outcomes that are unfair in a systematic way, we call that algorithmic bias. We would say that our cat-dog classifier shows algorithmic bias and that it is biased towards cats since it works really well for them and biased against dogs since it doesn't work as well for them.

Give students time to re-curate their datasets using the recurring dataset images.

Ask: What did you do to make it work better?

If students say they used less training data, prompt them to think about if it is better to have more vs less data?

So we've seen firsthand how algorithmic bias can occur in our supervised machine learning systems. Now I want us to take some time to watch a video about how it can happen in the real world.

Play Gender Shades facial detection video from slides.

Closing

Discuss the following questions with the class:

What problem did Joy identify in the video?

- *[wait for students to say facial recognition system did not recognize darker, female faces as well as lighter, male faces]*
- *If students need prompting, ask if the technology Joy talked about worked the same for all people.*

Why is this a problem?

- *[wait for students to mention unequal user experiences between groups - "technology doesn't work unless it works for everyone"]*
- *If students need prompting, ask them if they would all like to be able to use Snapchat filters. Is it wrong if they did not all have access to that technology?*

How does Joy suggest we can fix this problem?

- *[wait for students to mention better dataset curation]*
- *If students need prompting, ask them how they improved their classification algorithms during the Teachable Machines activity.*

Assessment

Learning Outcomes:

1. Understand the basic mechanics of artificial intelligence systems:

Know that artificial intelligence is a specific type of algorithm and has three specific parts: dataset, learning algorithm, and prediction.

- a. Understand the problem of classification in the supervised machine learning context.
- b. Understand how the quantity of training data affects the accuracy and robustness of a supervised machine learning model.

2. Understand that all technical systems are socio-technical systems.

Understand that socio-technical systems are not neutral sources of information and serve political agendas:

Know the term “algorithmic bias” in the classification context.

- a. Understand the effect training data has on the accuracy of a machine learning system.
- b. Recognize that humans have agency in curating training datasets.
- c. Understand how the composition of training data affects the outcome of a supervised machine learning system.

Extensions

Spend more time on discussions. Optional: if time permits and the classroom environment is safe. Ask: How might you find images to better curate your dataset?

- [Prompt students to think about where photos exist - social media, mug shots, ID cards. Which of these would be okay to use as sources of information? Which might lead to more bias? Which might breach privacy?]

Visit <https://machinelearningforkids.co.uk> for additional Machine Learning activities.

Visit <https://aischool.microsoft.com/machine-learning/learning-paths/ml-crash-course> for a deeper explanation of machine learning.

Image Datasets

Below are a set of images provided to students in order for them to build their cat-dog classifiers. Three different datasets are included:

Dataset	Description
Initial Training Dataset	<p>These are the images students should use to “teach” their machine learning model which image is a cat and which image is a dog.</p> <p>Note that there are many more cats and that the cats are more diverse in appearance than the dogs. This means that the classifier will more accurately classify cats than dogs.</p>
Test Dataset	<p>These are the images that students should use to test their classifier after training. Students should show these images to their model and record if their classifier predicts if the image is of a dog or a cat.</p> <p>Note: Students <u>should not</u> use these images to teach their classifier. If an image is used to train a classifier, the machine will have already recorded the corresponding label for the particular image. Showing this image to the machine during the testing phase will not measure how well the model generalizes.</p>
Recurating dataset	<p>This is a large assortment of images students can use to make their training dataset of cats and dogs larger and more diverse.</p>

The test dataset should be used twice, once for testing students’ initial classifier and again for testing their recurated dataset.

The images are arranged below in a way that should make printing them off and cutting relatively easy (with two images per page, organized by dataset). Printing images allows students to hold them up in front of a webcam to train.



Initial Training Dataset

Dogs









Cats















©Warren Photographic



©Warren Photographic





Test Dataset







©Warren Photographic





Recurating Dataset

Dogs









Cats









Name: _____ Date: _____


Introduction to Algorithmic Bias Activity Sheet

Description

In this exercise, you will learn about the three components of an artificial intelligence (AI) system. You will then learn about the role of training data in an AI system.

Instructions

1. Go to: <https://teachablemachine.withgoogle.com/>
2. Train Green using the dog images from your initial training dataset (try to only use one photo per dog)
3. Train Purple using the cat images from your initial training dataset (again, only one photo per cat)



Question 1:

(a) For the dog training dataset, record the following:

How many images are included?

How are the images similar?

How are the images different?

(b) For the cat training dataset, record the following:






How many images are included?



How are the images similar?

How are the images different?

Question 2: Train your classifier on the two training datasets.

Once your classifier is finished, test your dataset with cards given to you containing the following image. Fill in the table on the next page about your testing dataset:

Image	Classification	Confidence Score	Correct?
			
			
			
			
			

Question 3: Which class did your classifier work better on? (Circle one)

Cats

Dogs

Why do you think that is?

Question 4: With your group, use the photos on the tables to re-curate your training dataset. Record the following:

A. For the dog training dataset, record the following:

a. How many images are included?

b. How are the images similar?

c. How are the images different?



B. For the cat training dataset, record the following:




a. How many images are included?





b. How are the images similar?

c. How are the images different?

Question 5: Train your new classifier on your two new training datasets.

Once your classifier is finished, test your dataset with cards given to you containing the following image. Fill in the table on the next page about your testing dataset:

Image	Classification	Confidence Score	Correct?
			
			
			

Question 6: Did your new algorithm work... (circle one)

Better for dogs

The same for both
cats and dogs

Better for cats

Question 7: Explain your answer to Question 5:

