Lesson 1
How Do Machines Learn?

In this lesson, we will...

- Explore the basics of machine learning and understand how machines can learn from data.
- Introduce three types of machine learning: supervised, unsupervised, and reinforcement learning.
- Discover real-world examples and applications of machine learning.

Warm Up

Pick one app shown below:

1. Describe what types of content the app recommends to users.
2. Explain how you think the app decides what to recommend.

The three types of machine learning we'll be focusing on today:

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Difference between
Supervised , Unsupervised and Reinforcement Learning

Guess The Machine Learning "Type"

I'm going to present a few scenarios. For each one, decide which type of machine learning you ***think*** it may be similar to:

- **Supervised Learning**
- **Unsupervised Learning**
- **Reinforcement Learning**

We'll vote, hear some explanations, and then I'll reveal the answer!

Scenario 1

A friend sends you 128 random songs in a playlist. You decide to organize and group similar songs. At the start, you don't know what categories to use, but you want the songs to be grouped in a way that makes sense.

By the end, you've categorized the songs into 8 different genre categories.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 1

A friend sends you 128 random songs in a playlist. You decide to organize and group similar songs. At the start, you don't know what categories to use, but you want the songs to be grouped in a way that makes sense.

By the end, you've categorized the songs into 8 different genre categories.

Supervised Learning	**Unsupervised Learning**	Reinforcement Learning

Scenario 2

You're trying to teach a friend how to recognize different types of flowers. You show them many pictures of flowers and tell them the name of each one as you show each.

After a while, they start to recognize the flowers and can name them on their own.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 2

You're trying to teach a friend how to recognize different types of flowers. You show them many pictures of flowers and tell them the name of each one as you show each.

After a while, they start to recognize the flowers and can name them on their own.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 3

You're learning to recognize different types of rocks in a science class. Your teacher shows you pictures of rocks and tells you the type of each one (e.g., igneous, sedimentary, metamorphic).

After studying the labeled examples, you can identify the rock types on your own.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 3

You're learning to recognize different types of rocks in a science class. Your teacher shows you pictures of rocks and tells you the type of each one (e.g., igneous, sedimentary, metamorphic).

After studying the labeled examples, you can identify the rock types on your own.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 4

You train a dog to sit, roll over, and stay by rewarding them with a treat when they do it correctly and snapping your fingers loudly when they don't.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 4

You train a dog to sit, roll over, and stay by rewarding them with a treat when they do it correctly and snapping your fingers loudly when they don't.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 5

You are tasked with finding a hidden object in the classroom. After each step, a classmate yells out "hot" or "cold" to guide your search.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 5

You are tasked with finding a hidden object in the classroom. After each step, a classmate yells out "hot" or "cold" to guide your search.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 6

A teacher presents multiple scenarios of machine learning real-world analogies to their students without telling the students which type of machine learning each scenario represents.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Scenario 6

A teacher presents multiple scenarios of machine learning real-world analogies to their students without telling the students which type of machine learning each scenario represents.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

How can we define and describe each?

Supervised Learning	Unsupervised Learning	Reinforcement Learning

Supervised Learning - in machines

Suppose you want to train an AI to distinguish between dogs and cats. The machine uses **labeled** data to learn patterns and features.

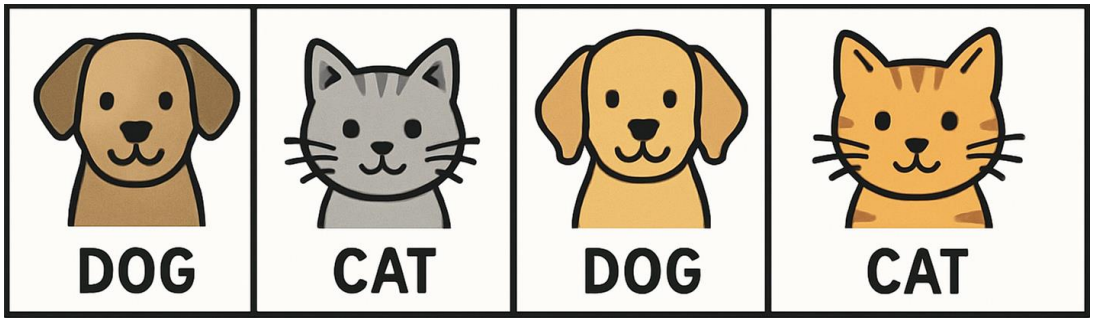

labeled images of dogs and cats

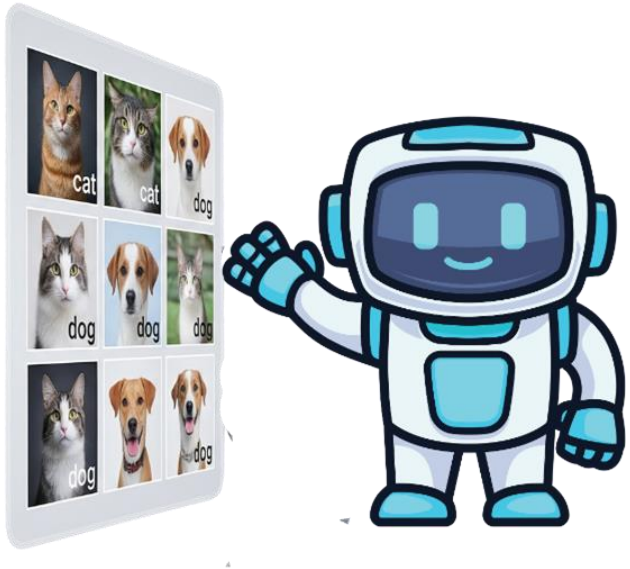

Supervised Learning - in machines

Once the training is complete, it can look at a new picture it's never seen before and **predict** “dog” or “cat.”

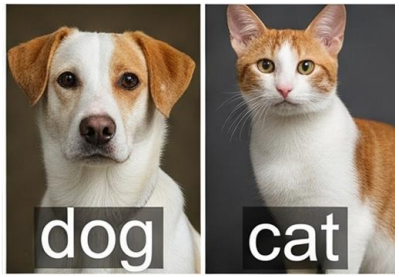

This process is called **supervised learning** because the AI is being "supervised" by the labeled data.

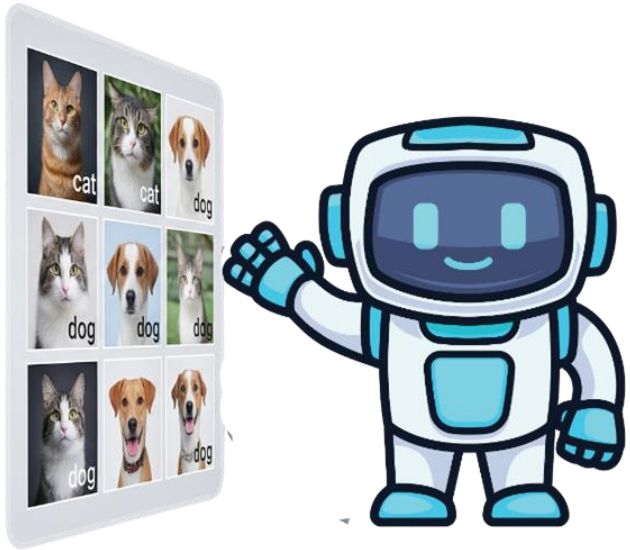

Unsupervised Learning - in machines

Music apps use **unsupervised learning** to create playlists you will like by analyzing your listening history and grouping you with users who likely have similar tastes.

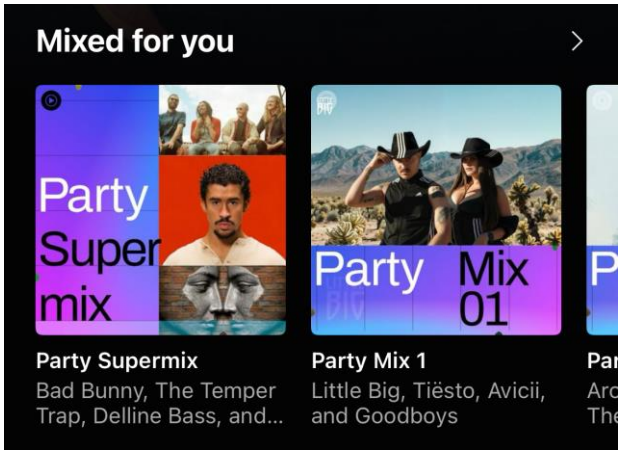

The AI can discover patterns in the data, such as genres, moods, or artists, to generate personalized playlists.

Reinforcement Learning - in machines

Reinforcement learning enables machines to learn tasks by trial and error, guided by a reward system. This approach is revolutionizing fields like gaming (chess, poker, and video games) and autonomous, self-driving vehicles.

Reinforcement Learning - in machines

By exploring millions of possibilities, learning from mistakes, and adapting to feedback, the AI can discover innovative strategies, often surpassing human expertise.

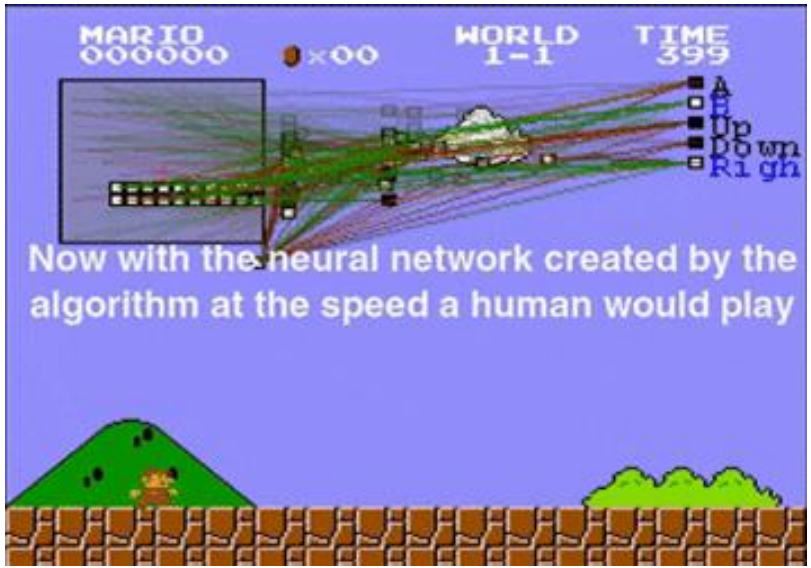

Recap

Supervised Learning	Unsupervised Learning	Reinforcement Learning
The AI system is shown lots of labeled examples and learns to recognize patterns associated with each label.		

Recap

Supervised Learning	Unsupervised Learning	Reinforcement Learning
The AI system is shown lots of labeled examples and learns to recognize patterns associated with each label.	The AI system looks for patterns in the data and groups or organizes it on its own, without being given labels or categories in advance.	

Recap

Supervised Learning	Unsupervised Learning	Reinforcement Learning
The AI system is shown lots of labeled examples and learns to recognize patterns associated with each label.	The AI system looks for patterns in the data and groups or organizes it on its own, without being given labels or categories in advance.	The AI system is given a goal and learns by taking actions and receiving rewards or penalties based on the outcome. Over time, it gets better at achieving the goal.

Supervised vs Unsupervised vs Reinforcement

Pick 3 of the examples below. What type of machine learning is being used to help the AI learn or make decisions: **supervised**, **unsupervised**, or **reinforcement**? Explain your thinking.

An AI robot that learns to stack blocks by trying different actions and getting rewarded when it succeeds.	A photo app that recognizes all of the pictures containing each of your friends' faces in your photo library.	An AI model trained on thousands of labeled X-rays to detect signs of disease.
A recommendation engine that groups movies into similar themes based on viewing patterns.	An AI robot dog that learns to walk by getting rewarded when it doesn't fall after taking a step.	AI that groups online shoppers into categories based on their browsing history and purchase habits.
An AI that learns to play Super Mario World and beat levels as fast as possible.	An AI that organizes a library of songs into different playlists based on similarities in sound, without being told the genres.	An AI that can detect and filter spam emails.

In the next lesson, we'll look inside the "brain" of an AI to explore how machines learn from experience - using something called a **neural network**.

Lesson 2
Inside A Neural Network

In this lesson, we will...

- Learn how neural networks process information through a simulated activity.
- Understand the general structure of a neural network.
- Explore how neural networks can be used to represent complex data and make predictions.

Warm Up

In the last lesson, we saw that machines can learn in different ways - from finding patterns in labeled or unlabeled data, or by being rewarded. Imagine an AI looking at millions of images, or playing a video game and earning points.

What do you think is happening inside the AI that helps it learn or get better?

Let's do an activity that shows how a machine might take in information, process it, and learn - just like AI!

Warm Up Activity

I have a **secret image** with a short, descriptive caption.

We're going to play a game to see if **<u>one person</u>** can guess the caption describing the image based on limited clues.

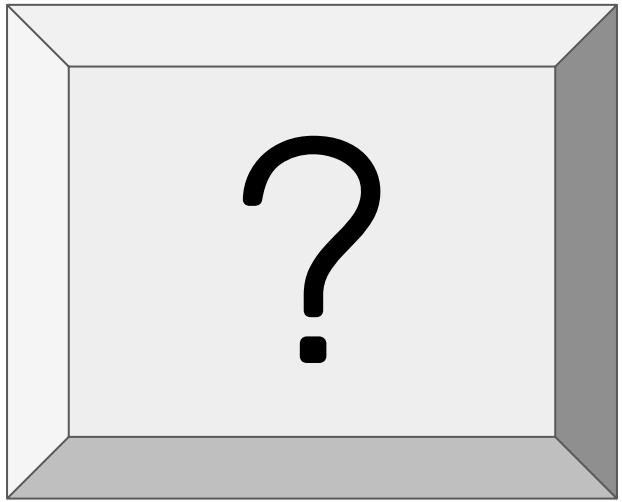

Warm Up Activity

For example, if this is my image:

The caption might be: **a fish riding a bike down a hill**

Warm Up Activity

You will be assigned to 1 of 3 groups:

1. **input group (approximately ⅔ of the class)**

   will each receive 4 blank index cards

1. **hidden group (approximately ⅓ of the class)**

   will each receive 2 blank index cards

1. **output group (just 1 person)**

   will receive 1 blank index card

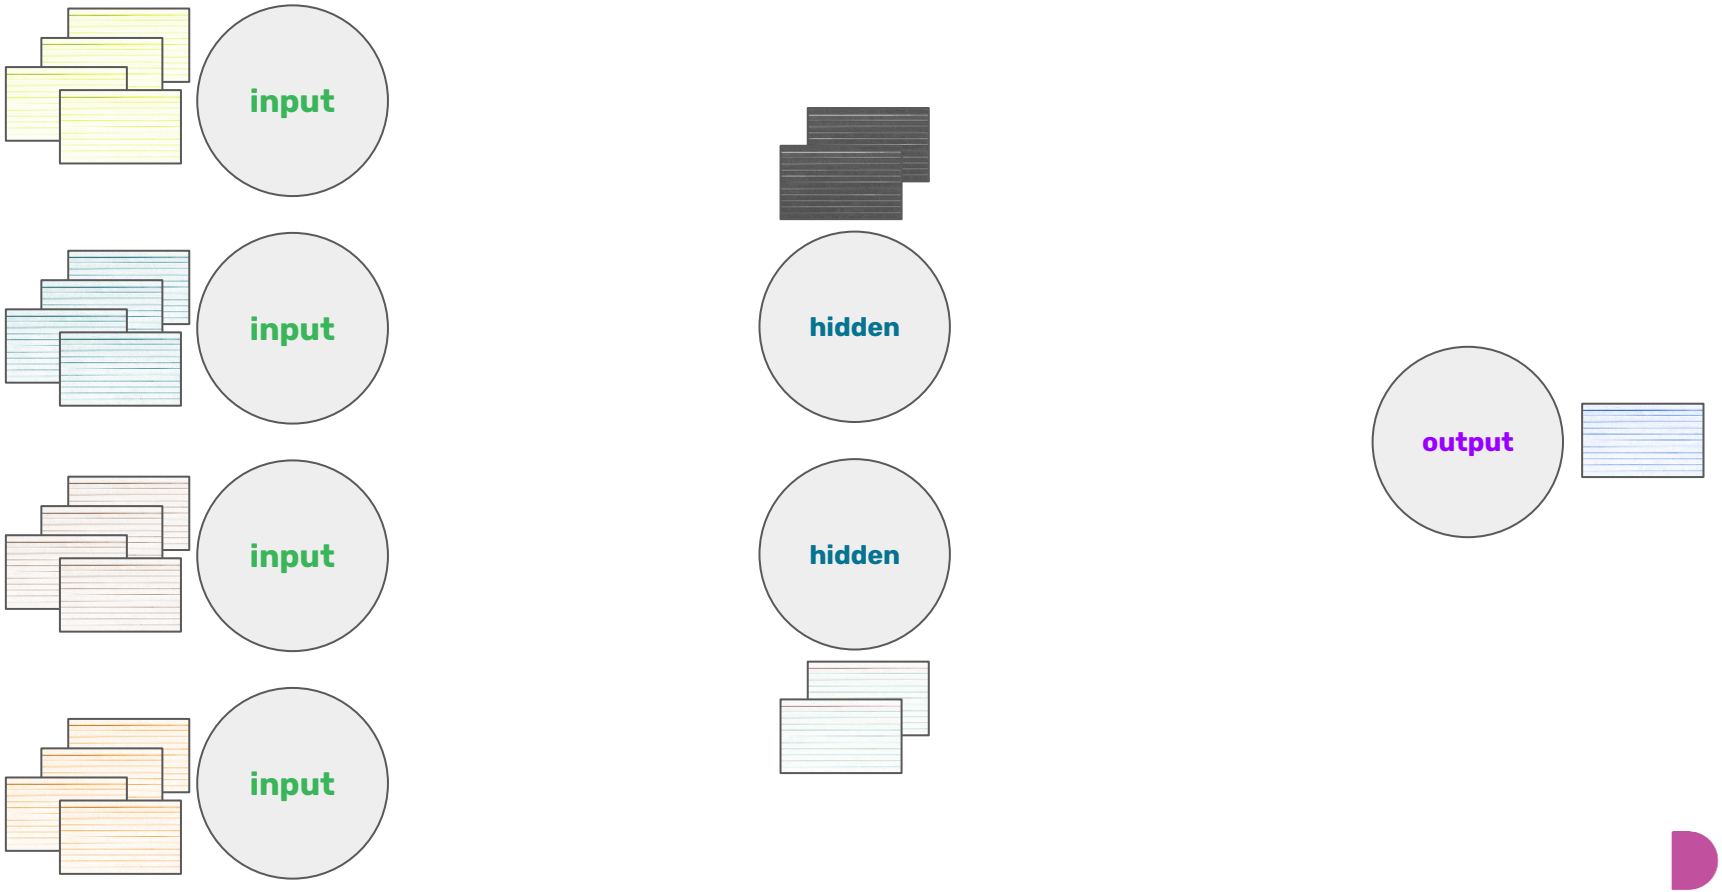
input
input
input
input
hidden
hidden
output

First Steps - Input Group

1. I'm going to show an image **ONLY** to the students in the **input** group.

2. Students in the **input** group will write down **4 words**, one on each blank index card, describing what they see, **WITHOUT** communicating with anyone else.

3. Each student in the **input** group will then hand **one card** to each person in the **hidden** group, until they run out of cards.

Next Steps - Hidden Group

4. Students in the **hidden** group will look at all of the index cards they receive from the **input** group. The **hidden** group **does not** get to see the image.

5. They will write **2 words**, one on each blank index card they have, **WITHOUT** communicating with anyone else.

6. Each student in the **hidden** group will then hand their 2 index cards directly to the **output** person.

Final Steps - Output Group (1 Person)

7. The **output** person will look at all of the index cards they received from the **hidden** group.

8. Then, they will write a caption on their index card, predicting what is the original image that I showed **only** to students in the **input** group.

9. We'll then compare the **output** person's caption to the *actual* image caption.

EVERYONE MUST REMAIN SILENT THROUGHOUT THE ENTIRE PROCESS!

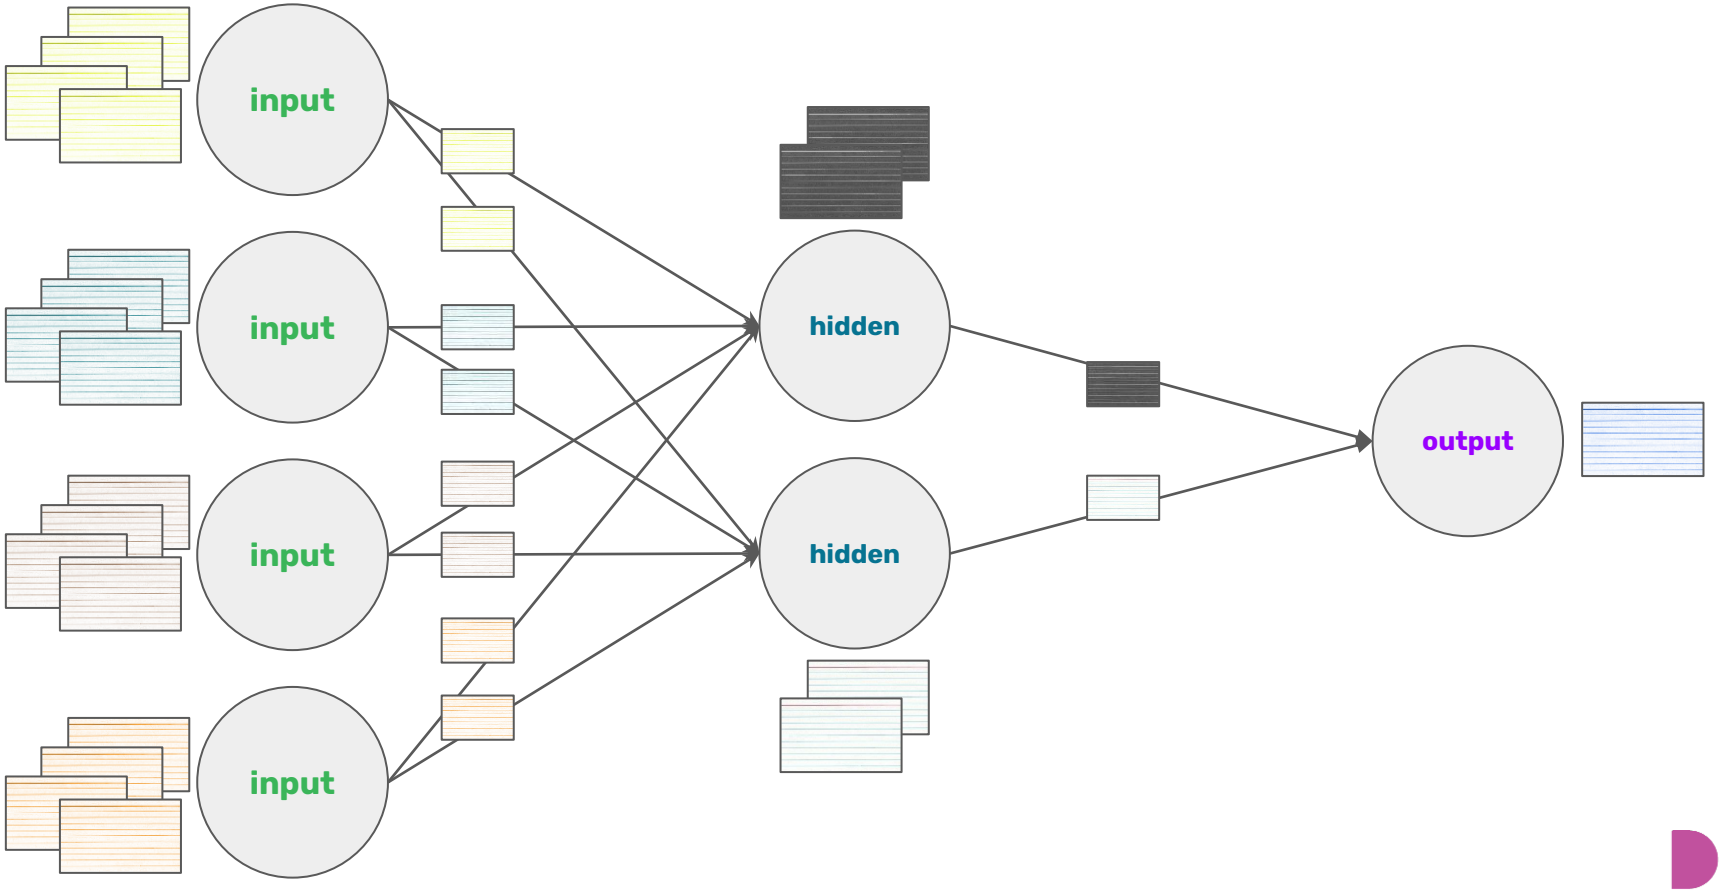
input
input
input
input
hidden
hidden
output

Ready to see the actual image and caption?

a Chihuahua playing basketball against 4 black cats

How did we do?

How close was the **output** person's caption to:

a Chihuahua playing basketball against 4 black cats?

Reflection and Analysis

1. Let's ask the **output** person to choose which index cards would have been most helpful to predict the correct caption.
2. What can the **hidden** and **input** groups learn from the chosen index cards?
3. What can the **hidden** and **input** groups do differently on their cards to improve the **output** person's prediction next time?

Let's Try Again

This time, each person in the **input group** will only see part of the full image, not the whole thing.

This is more like what happens in a neural network: each **input** **node** only processes a tiny piece of the full input, and the network has to work together to make sense of it!

Some saw this:

Some saw this:

Some saw this:

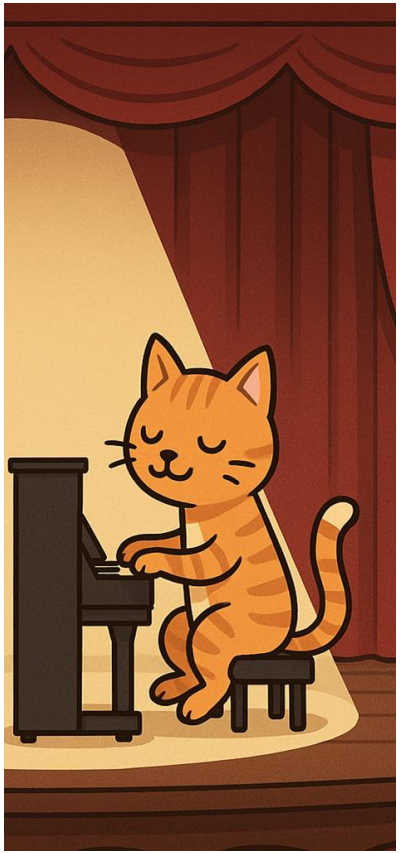

Some saw this:

Ready to see the actual image and caption?

A squirrel juggles while a cat plays the piano under a spotlight on stage.

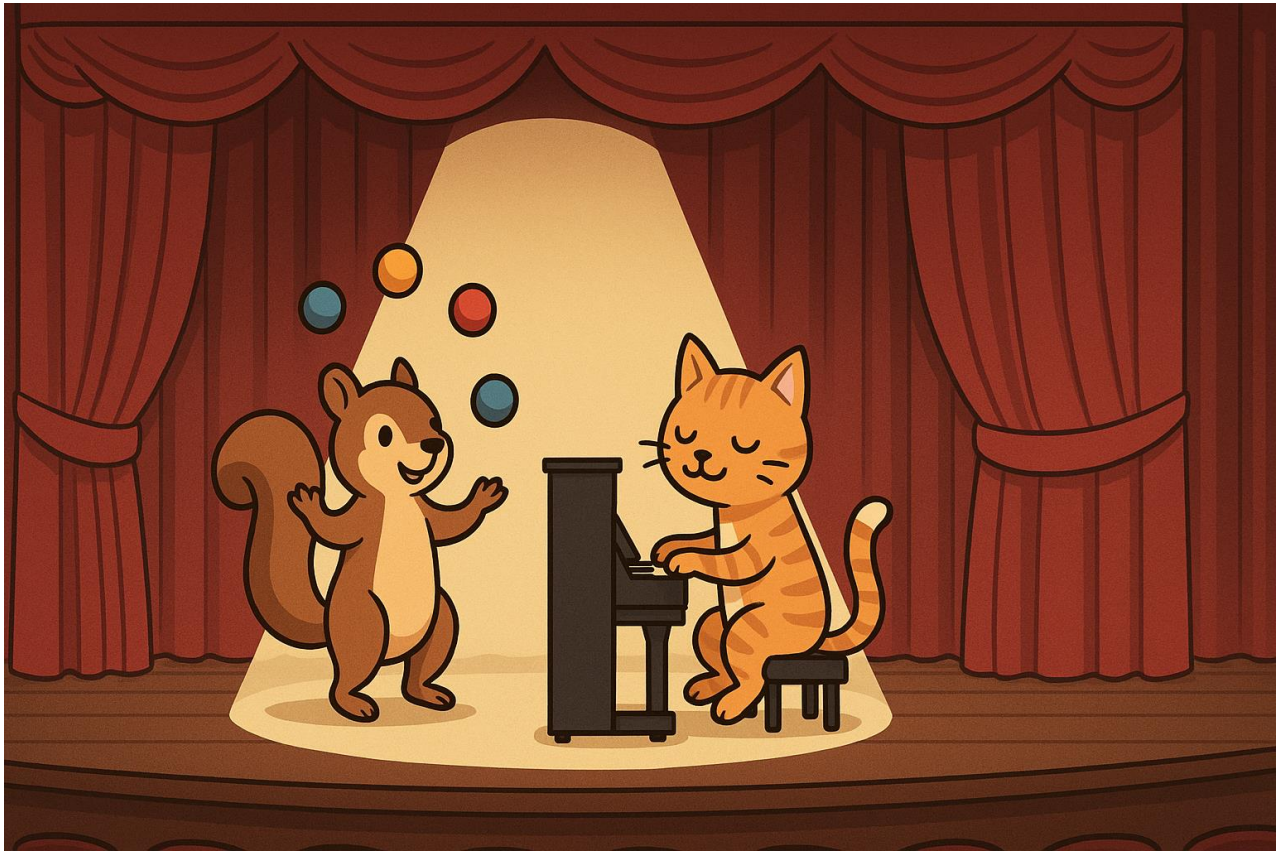

Reflection

1. How close was the **output group**'s guess to the actual caption?
2. Which details may have been *added* or *missed* by the **input** and **hidden** groups? (e.g., did anyone include the spotlight?)
3. What could each group do differently if we were to do a 3rd round?

We've just simulated what happens in a **neural network**!

What is a neural network?

A **neural network** is a type of computer system inspired by how the human brain works. It helps a computer learn patterns and make decisions.

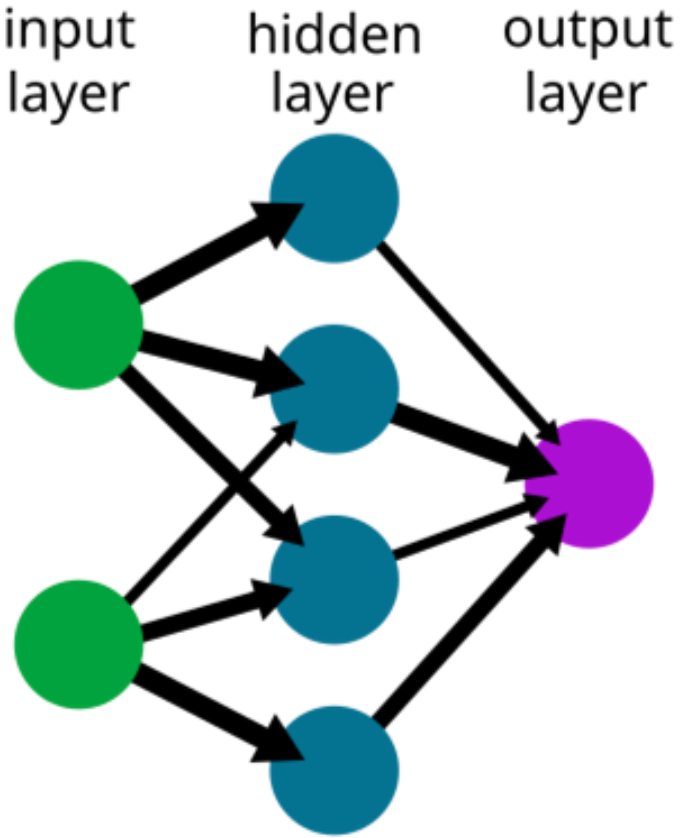

What is a neural network?

It is made up of **layers** of connected **nodes** (also called **neurons**). Each node performs a small task to help with the overall learning process.

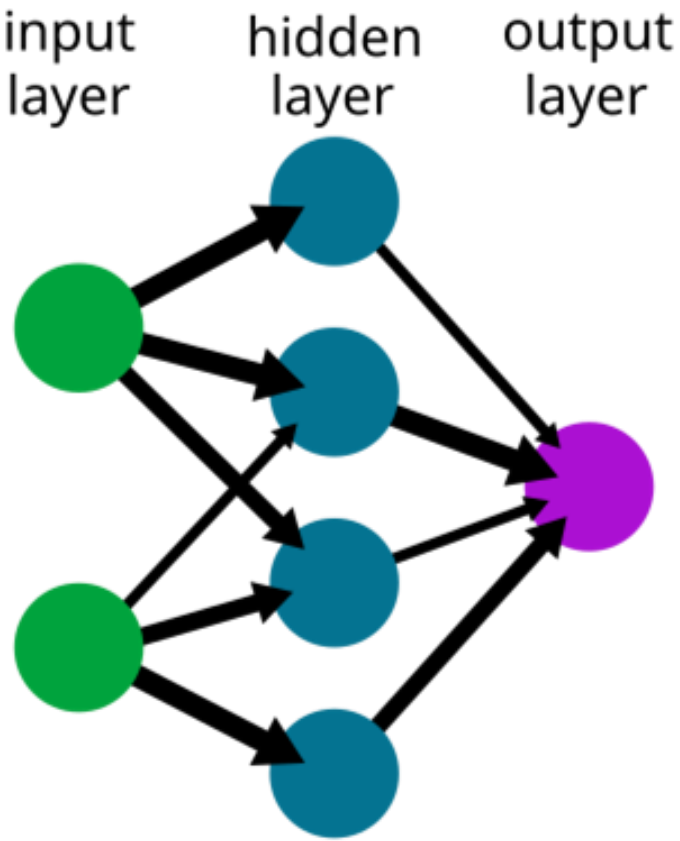

What is a neural network?

The **input layer** is the part of the neural network that receives the raw data, like pixels from an image or words from a sentence, and passes that information to the next layer for processing.

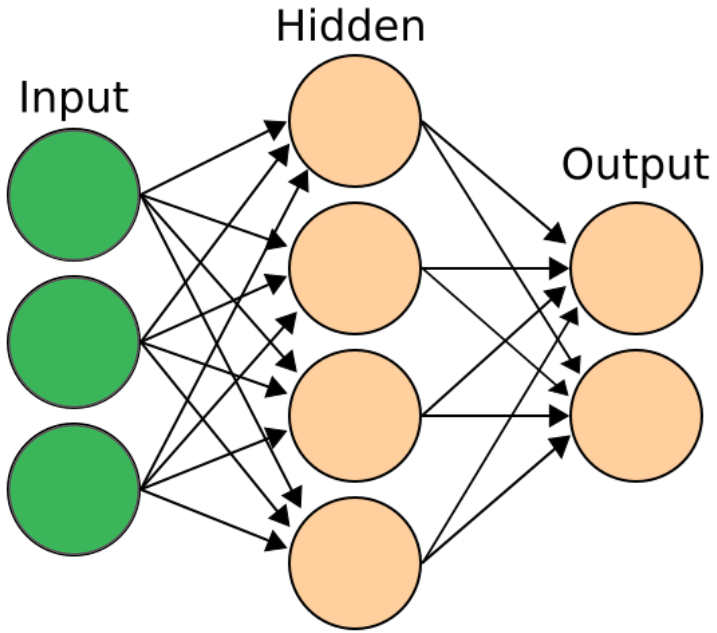

What is a neural network?

In the 1st warm-up activity, I showed the full image to everyone in the **input layer** (group).

In an actual neural network, each **node** in the **input layer** would only see a small part of the image, not the whole thing.

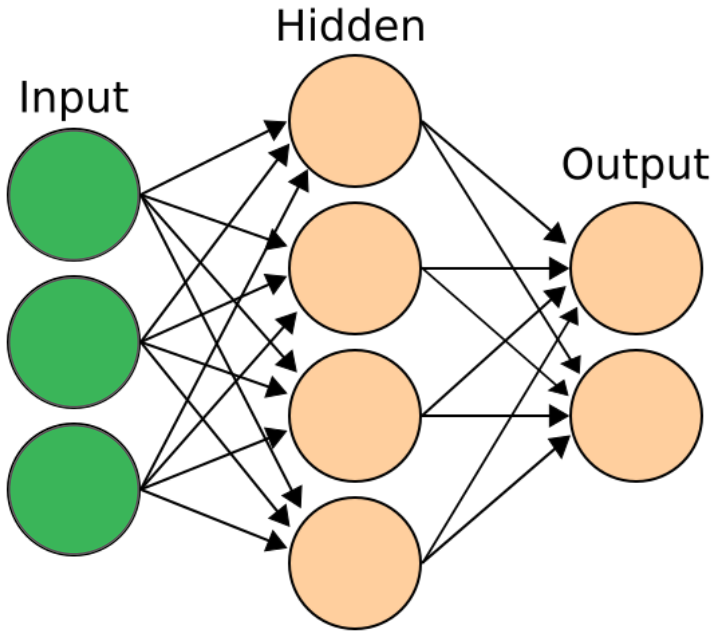

What is a neural network?

The **hidden layer** is the part of the neural network that does the "thinking." Each **hidden node** looks for patterns or clues in the data and acts as a mini decision-maker.

Each **layer** builds on the previous one to create a final output.

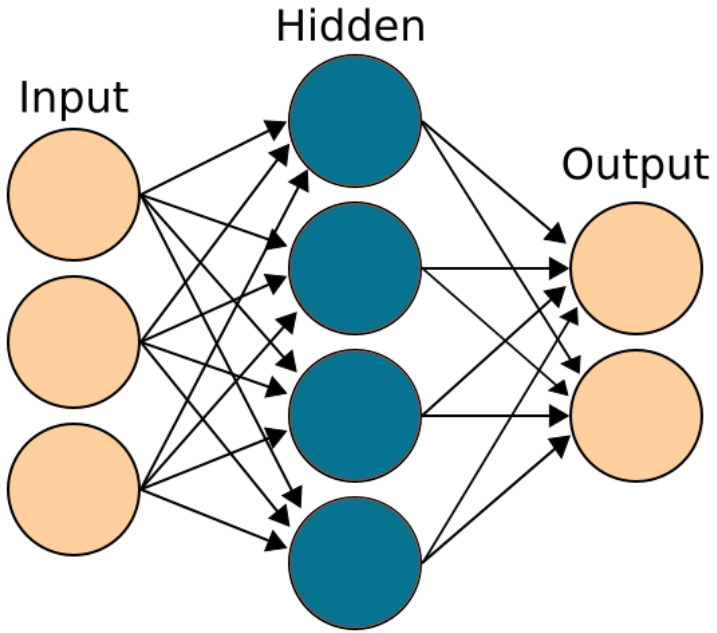

What is a neural network?

The **output layer** is where the final decision or prediction is made based on all the processing done by the **hidden layers**.

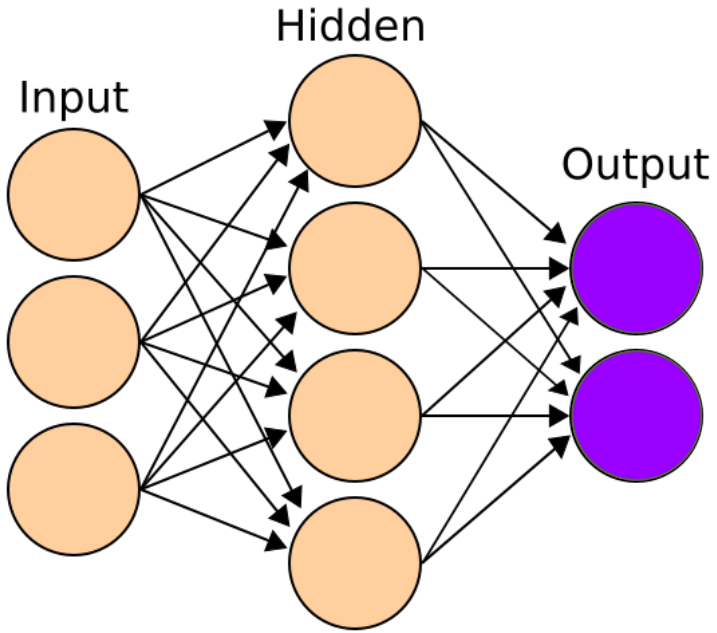

What is a neural network?

Think of the **input layer** like clues in a mystery, the **hidden layers** like the detectives putting those clues together, and the **output** as the solution to the case. Like our brains, **neural networks** can learn and improve over time as they're exposed to more data and experience.

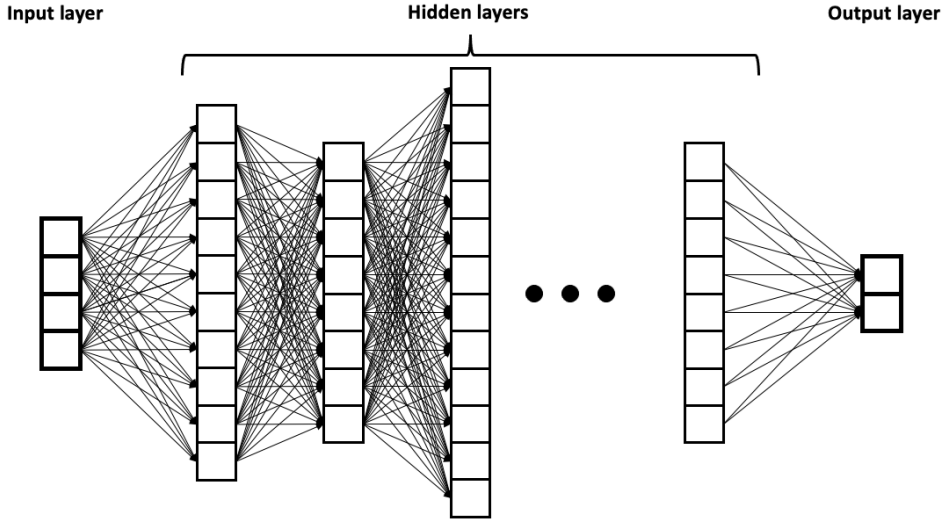

Neural networks can have:

- Multiple **input** nodes
- One or more **hidden** layers, each with multiple **nodes** (**neurons**)
- One or more **output** nodes, which produce **predictions** or **classifications**

Classroom Simulation Example	Actual Neural Network
input group	the **input layer**
hidden group	one or more **hidden layers**
output group	one or more **output** **nodes** in the **output layer**
each student	a **node** (also called a **neuron**)
passing index cards from one layer to the next…	…is known as **feedforward** (how data flows through the network)
reflecting on which index cards were most useful and adjusting connections for next time…	…is known as **backpropagation** (it's how the neural network learns from its mistakes)

Team Challenge

Our class has been tasked with figuring out how to teach a machine to identify a number (0-9) written in a 6x6 grid.

With a partner, explain how this could work using one of the three types of machine learning we've explored:

Supervised, **Unsupervised**, or **Reinforcement Learning**

Use what you've learned about how machines process data and make decisions. Think about how a **neural network** might help.

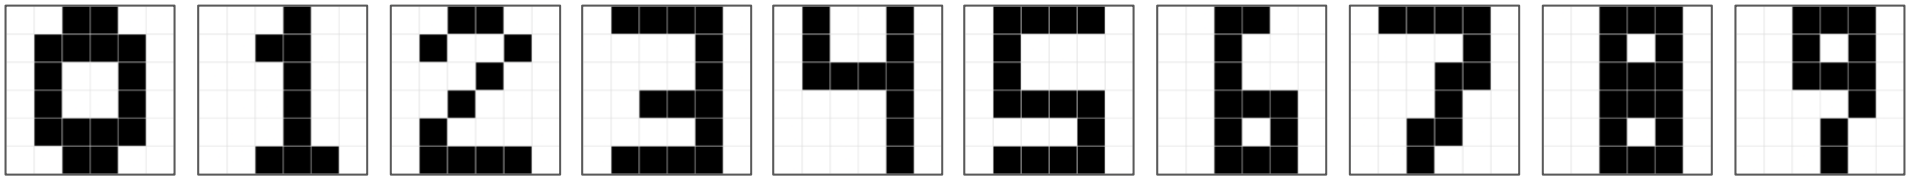

Hints

1. Recall that neural networks contain an **input layer**, **hidden layer(s)**, and an **output layer**. What would each layer's responsibilities be?
2. What kind of information does the machine need to see in order to recognize a number? Remember, it's a 6x6 grid. Think about what each **input** **node** might represent or focus on.
3. The machine won't just know the numbers - it must learn first!

Challenge Extension

Use **Teachable Machine** to train an AI to detect numbers drawn in a 6x6 grid of pixels.

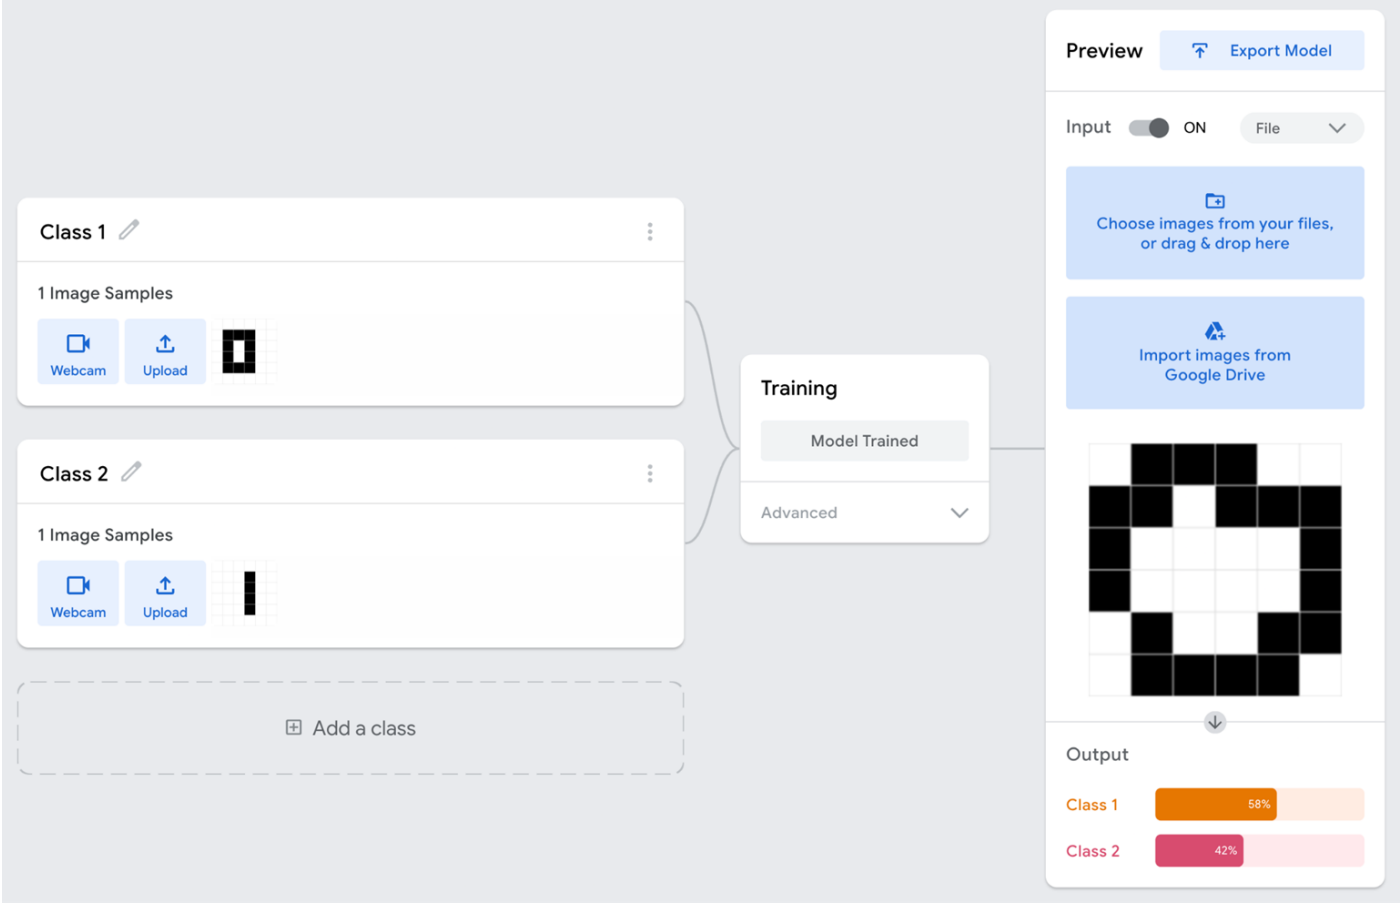

Lesson 3
Bias In Artificial Intelligence

In this lesson, we will...

- Define what bias in artificial intelligence means.
- Identify examples of how bias can enter an AI system.
- Begin to think about the implications of algorithmic bias and how we can minimize it.
- Reflect on the importance of fairness and perspectives in creating AI systems.
- Discuss how we can use AI ethically and responsibly.

Warm Up Activity: Predicting Outcomes

Read each scenario and predict what might happen as a result. Think about how the AI system's behavior could affect people.

1. An email filtering AI is trained on labeled spam and non-spam emails, but all of the examples are in English. **What might happen if it receives emails in other languages?**
2. A social media AI groups users after predicting their interests, then continues to recommend only similar content to the entire group. **How might this affect what users see or believe?**

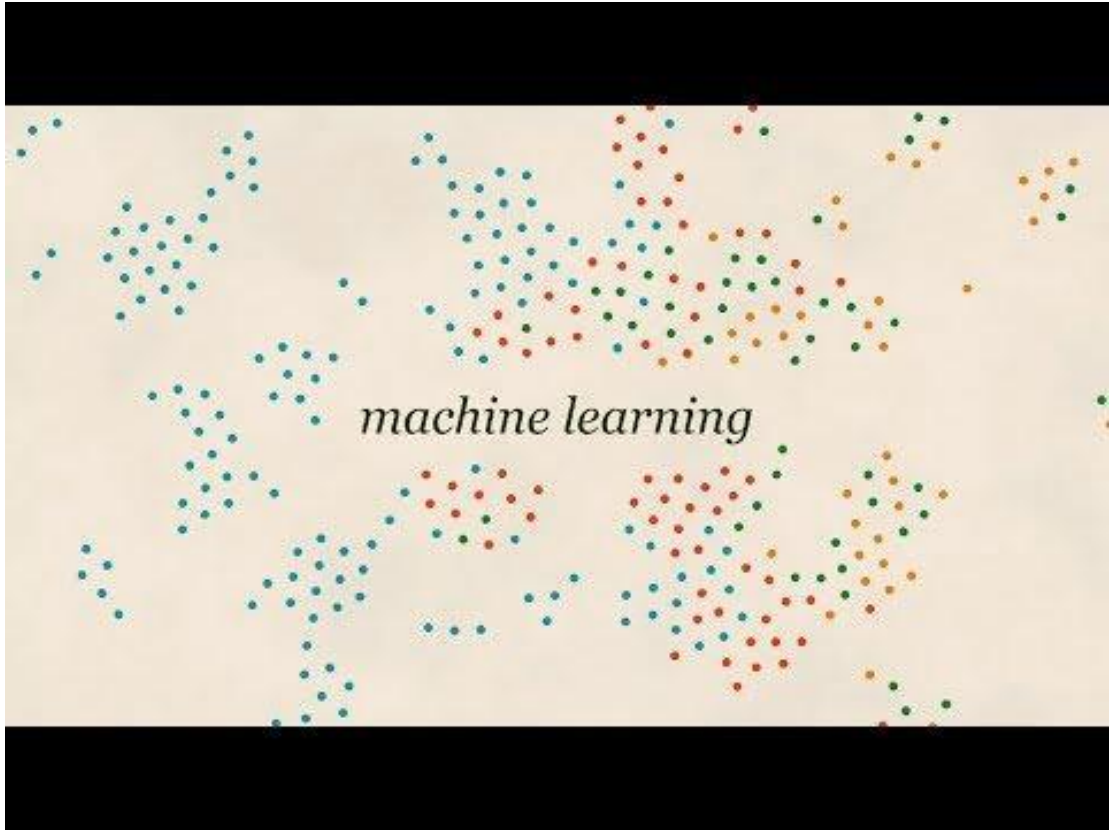
machine learning

What is bias?

What is bias?

Bias is the action of supporting or opposing a particular person or thing in an unfair way.

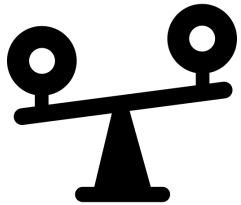

Bias is not always intentional but can lead to unfair outcomes. Understanding bias can help us design better, fairer systems.

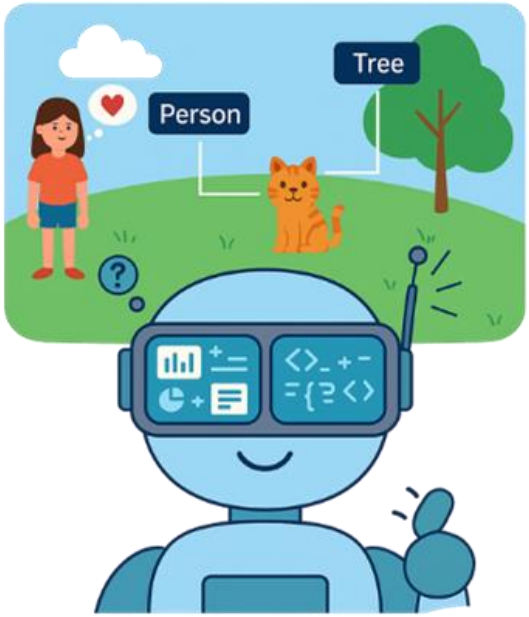

Algorithmic Bias

We asked an AI chatbot to generate an image of pop stars. Here's what was generated:

What do you notice about the images?

Why do you think the outputs look like this?

Algorithmic Bias

Even though AI is supposed to be neutral, it can still learn unfair patterns from the data it's trained on and end up repeating the same biases people have.

Algorithmic Bias

Algorithmic bias (bias in AI) occurs when a system makes unfair decisions because of the data it learned from or how it was built.

It may not be done on purpose, but AI bias can still harm people.

Let's explore how that happens and what we can do about it.

Brainstorming Potential Biases

Let's generate a list of **biases** that an AI may learn.

Think about how the AI may have **learned** these biases.

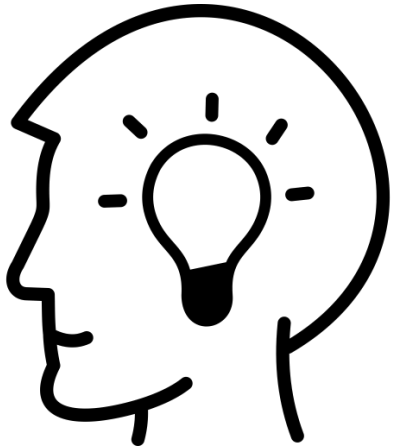

Some Examples of Bias

- Language
- Age
- Gender
- Race/skin tones
- Geographic location
- Accent
- Economic
- Disability
- Name
- Academic
- Cultural
- Healthcare data

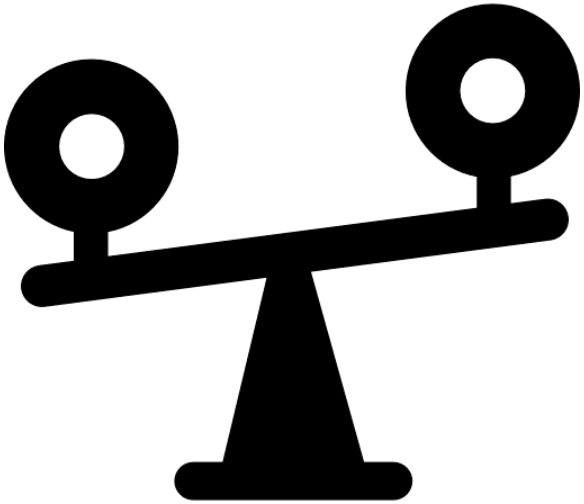

Consequences of Bias

What are the impacts of these biases? Who do they affect? Why does it matter? Describe the scenario.

Designing A Fairer System

How would you attempt to prevent or reduce biases from being learned and decreasing their negative impacts on people?

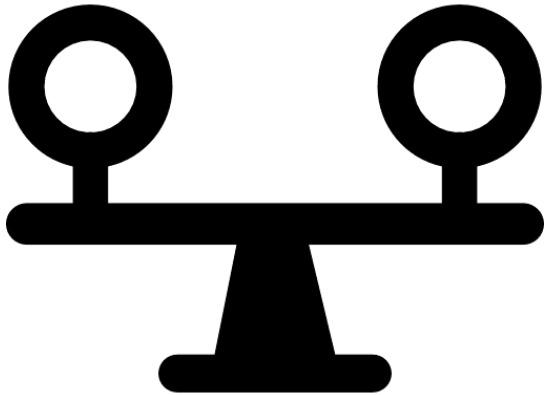

For Example

Bias: AI voice recognition systems often respond more accurately to adult voices than to kids' voices.

Consequence: Kids receive less accurate responses. The system may not understand their questions, may ignore them, or may give confusing answers. This can lead to frustration and a sense that the technology "doesn't work for them."

A fairer system: Train the voice recognition AI on a more diverse dataset that includes people of different ages, accents, and speech patterns.

What do you see?

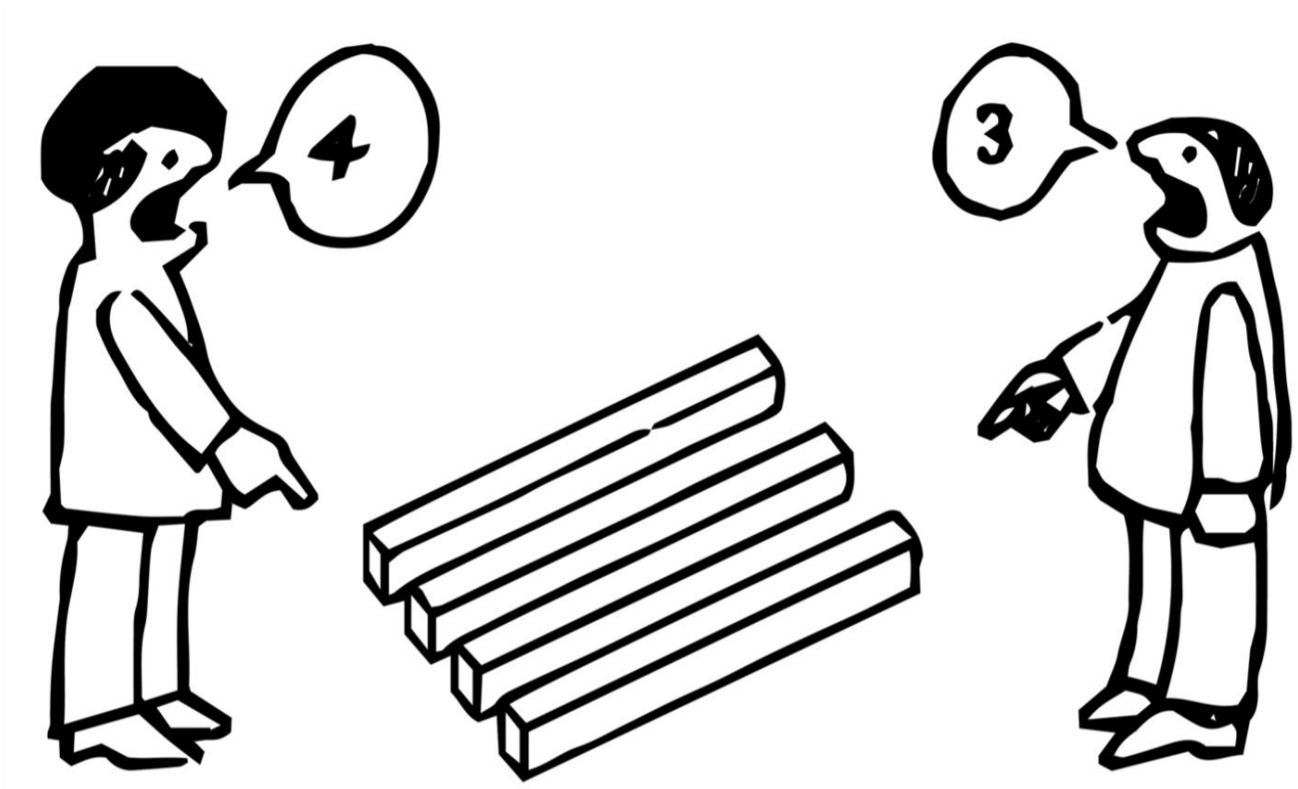

What do you see?

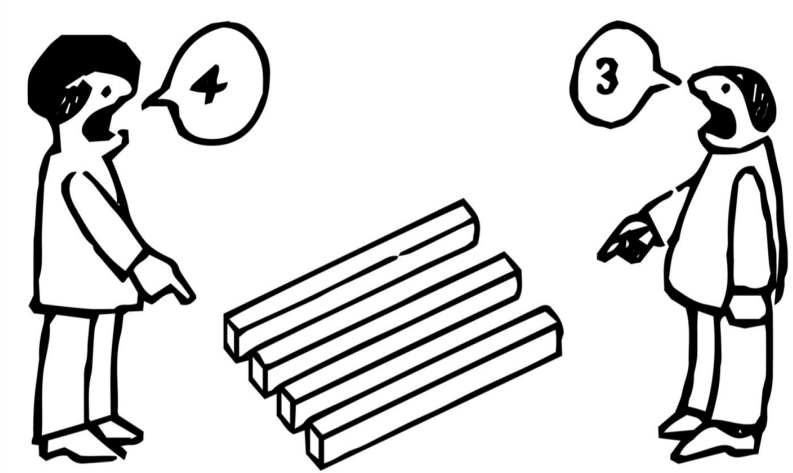

Perspective - a particular way of looking at something; a point of view

If some classmates see something differently, does that mean they're wrong?

Why is it important to have people with different backgrounds, experiences, and perspectives involved in creating AI?

Algorithmic Bias

Now that we've seen how computers can learn unfair things if their data isn't balanced, let's look at a real-world example.

This video clip is about a researcher named **Joy Buolamwini**. She was using a computer to recognize faces, and something surprising happened.

The computer worked well for some people… but not for others.

Coded Bias Movie Trailer

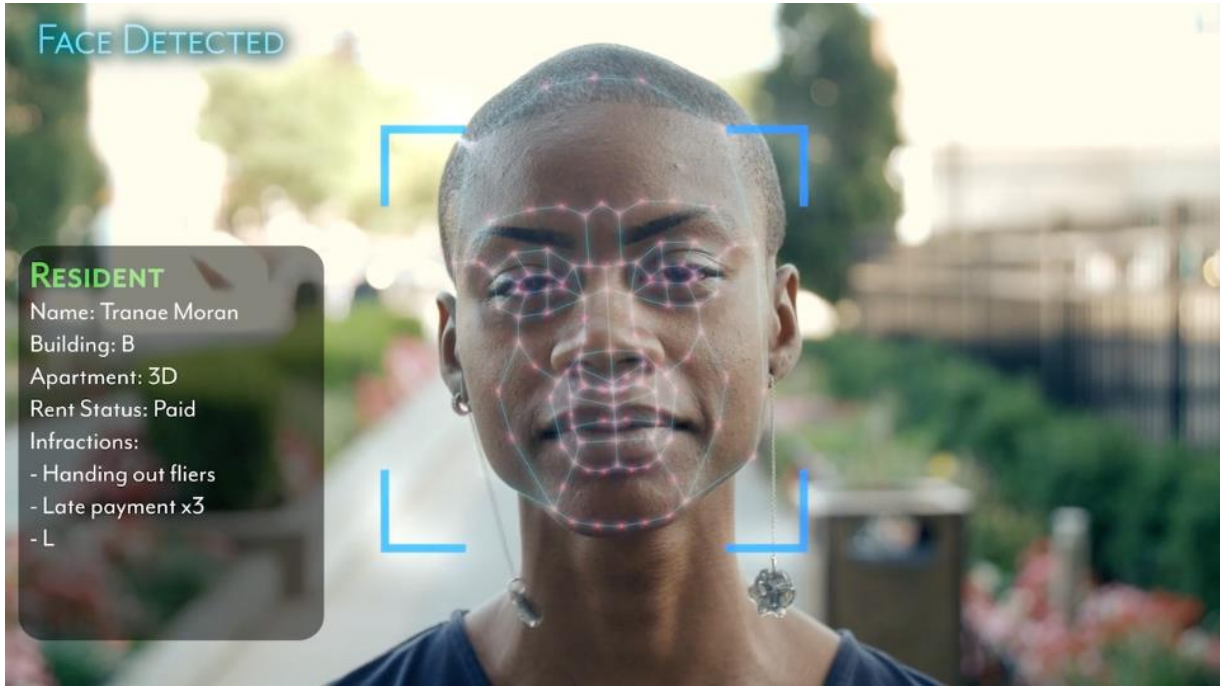

Exploring Algorithmic Bias

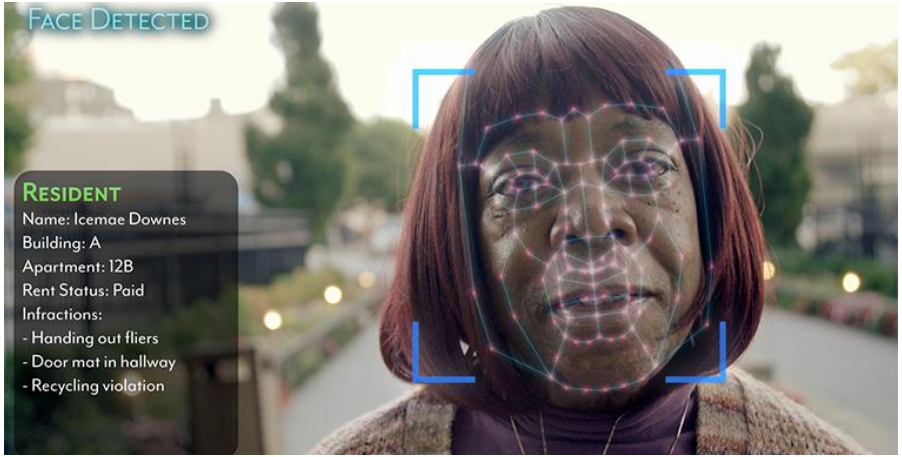

1. What stood out to you in the trailer, and why?

1. Did you notice any of the biases we talked about earlier in the trailer?

Closing Activity

Activity: Bias Detective - Spot the Bias in AI Systems

Directions: Read each AI scenario carefully. For each one:

1. Identify the type of bias that might be present.
2. Explain who could be affected.
3. Suggest how the bias could be fixed or prevented.

Unit Review

Machine Learning Investigation

Choose a real-world example where AI is being used. It could be in sports, entertainment, fashion, science, education, social media, law enforcement, or another area you're interested in.

1. What type of **machine learning** is being used: supervised, unsupervised, reinforcement, or another? Explain how.
2. Is a **neural network** involved? What is it helping the AI do?
3. Can you identify any **bias** or **fairness concerns** in your example? What are the possible impacts/consequences and how can they be reduced?

Share Out