Masterclass Session Script

The Royal Institution Science Lives Here
Masterclass network

 This icon refers to the supporting slide in the presentation.

 These icons indicate there is an activity to do, or a worksheet to complete.

 This icon indicates there is a video to watch.

Artificial Intelligence OTS Masterclass	Slides/ Activity
Introduction (15 minutes) Welcome to our Artificial Intelligence computer science masterclass! Today we are going to be looking at artificial intelligence, but what do we actually mean by that? Well let's start first by considering how we would describe intelligence. In a moment we will hear your ideas. • **What does intelligence mean to you?** • **What things do you consider to be intelligent, or how would you describe intelligence to an alien?** So now we know what intelligence means to us, let's start straight away by using an AI! This piece of paper I am holding is intelligent, probably even more intelligent than myself, and all of you! • **Who believes me?** • **How do you think it could be intelligent?** • **Has anyone here ever played a game of noughts and crosses?** This piece of paper is intelligent because it has never lost a game of noughts and crosses. It wins about half the games it plays and draws the rest. I need two volunteers to help test the piece of paper's intelligence – one to be the paper's hands, and one to play against the piece of paper. **To the first volunteer who is the piece of paper's hands* You will need to read aloud the piece of paper's instruction and then follow it by drawing an x in the location it specifies.* **To the second volunteer who is the paper's opponent* You will play against the paper, as noughts. Try not to lose!* • **So now that we can see that the piece of paper has won/drawn, and not lost, who believes me now – is the piece of paper intelligent?** The intelligence from this piece of paper has come from the person who wrote the instructions down. The volunteer who acted as the piece of paper's hands was really demonstrating what a computer does – follow instructions. Today we're going to investigate a bit more about what we mean when we call things artificial intelligence. Now that we've used a form of AI, what do we mean when we say that AI is artificial? We mean that it is made by humans, that it does not naturally occur. And broadly speaking, when we talk about intelligence, we refer to the process of being able to gain and apply knowledge and skills to new situations. It is important for us to understand that AI can learn by recognising patterns in data. AI can make decisions, based on its previous training. And it can also adapt to new information that it is shown. Crucially though, AI does not understand what it is doing, and it is only following the instructions given to it by a human.	 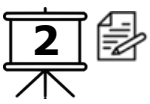 

AI or not (10 minutes)

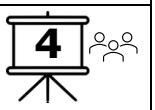

We're now going to put those criteria to the test and play a game of AI or not. For each of the following objects, I would like you to decide if it uses AI or not. First up we have a self-driving vehicle. Hands up if you think it uses AI. And hands up if you think it does not.
A self-driving vehicle uses sensors to detect roads, people, and obstacles, and then makes use of AI to decide what it should do: speed up, brake or change lanes for example.

Next, we have a toaster. Hands up if you think it uses AI. And hands up if you think it does not.
A toaster does not use AI, it just heats bread, it does not learn or make decisions. How could we add in artificial intelligence to a toaster? We could for example make a toaster that could learn how toasted each person likes their toast and stop the toaster at the correct time for each user.

Next, we have a robot vacuum cleaner. Hands up if you think it uses AI. And hands up if you think it does not.
A robot vacuum cleaner uses AI to map the layout of a room, avoid obstacles and find its way back to its charger.

Next, we have navigation systems and apps, like Google Maps or Waze. Hands up if you think they use AI. And hands up if you think they do not.
Navigation systems like Google Maps use AI to predict traffic, find the best route and even suggest alternative routes.

Next we have a standard calculator. Hands up if you think they typically use AI. And hands up if you think they do not. This is a tricky one, but a calculator is not traditionally considered artificial intelligence. It follows mathematical rules, but it does not learn or spot patterns that it can apply to new data. However, some AI's have been designed to solve more complex maths problems, or to suggest interesting problems that a mathematician could investigate.

What about voice assistants, like Alexa or Siri? Hands up if you think they typically use AI. And hands up if you think they do not. Voice assistants like Amazon's Alexa and Apple's Siri use AI to detect patterns in speech and respond with useful answers.

And how about YouTube recommendations? Hands up if you think they typically use AI. And hands up if you think they do not. YouTube can recommend videos they think you like. The algorithm they use to decide which videos to suggest for you, learns by looking at data from which videos you have interacted with before, and which you have had reactions to – for example liking it, or sharing it. This is a type of artificial intelligence.

And lastly, we have a dolphin. Hands up if you think they typically use AI. And hands up if you think they do not. A dolphin displays intelligence, but it is not an AI because it is not artificial. Its intelligence has not been 'created' by humans.

The Turing Test and testing for intelligence (10 minutes)

So how do we actually test whether something is intelligent or not? Well, one suggested way of doing just this is by putting your AI up against the Turing Test, named after computer scientist Alan Turing. Turing is most famous for his work during world war two, where he was able to work out how to decipher secret German messages that had been encoded using their infamous Enigma Machine. Some historians estimate that Turing solving the Enigma code shortened the war by up to two years, by allowing the Allied forces to anticipate the actions they would take before they were taken.

Later, in 1949, Turing was engaging in philosophical thought around what made something 'intelligent', or what it actually meant to 'think'. Through this, he designed what is now known as the Turing Test, initially called the Imitation Game, because it focuses on whether a human could be fooled into thinking that a machine was another human, based on some short conversations between the two.

The test can take many different forms, with different questions and scenarios given to the machine, but the main basis stays the same – if a human has a conversation with another human and a machine, and cannot correctly identify who the machine is, then the machine could be considered to be 'intelligent'.

- **What kind of questions do you think might have been asked in these conversations?**

These short conversations don't have to include questions that a human would find very tricky, an average human would probably not be able to answer a hard maths or science question straight away. Instead, the questions are designed to be easy and natural for humans to answer, like these questions here. Each of these questions would probably be quite easy for any of us to answer.

- **Does anyone have any thoughts or feelings towards the last question?**

If an AI answered whether they were a computer with 'no', then we would consider them to be lying. However, if an AI had been instructed and designed to answer that question falsely, then it is actually doing nothing wrong, and simply following the orders that it had been given.

- **Can you come up with any other questions that you would ask an AI to determine if it was human or not?**
- **Do you think any AI would pass the Turing test?**

Some AIs have reportedly passed different versions of the Turing Test. Back in March 2025, researchers looked at four different AI systems, including ChatGPT's GPT-4.5, and gave them two Turing tests. Human participants in the test had 5-minute conversations with another human participant and ChatGPT, then judged which conversational partner they believed to be human. GPT-4.5 was identified as the human 73% of the time, significantly more often than the actual human participants.

- **Do you think the Turing test is a good test for intelligence?**
- **How else would you want to test an AI's intelligence?**

Some AI researchers actually don't think that the Turing test is applicable for testing the usefulness and 'intelligence' of AI systems. One argument looks at plane design, and says that when planes are tested, they are tested to see how well it flies, not for whether it could be mistaken for a bird!
If we wanted an AI to be able to detect patterns in images for example, perhaps to confidently tell the difference between a bird, a plane, and superman, then we wouldn't necessarily need it to be a good conversationalist.

Neural Networks (15 minutes)

So now we have an idea of what 'intelligence' could look like in our AI. Let's say we are trying to make an AI that could, for example, pass the Turing Test, then it would make sense to base them on our own human brains, and that is exactly what a lot of computer scientists do! AIs are often trained by building something called a neural network – which is based on how neurones in our brain send signals to each other.

Sensory Neurones are the "detectors" of your body, found in your eyes, ears, nose, tongue, and skin. They detect things in the world around you, including noises, smells, and sights, and send this information to your brain.
Relay neurones are the "post office" inside your brain and spinal cord. They take the information from the sensory neurones and pass it on to other neurones, including motor neurones.

After receiving information from relay neurones, motor neurons send messages from your brain to your muscles, telling them what to do. They are responsible for movement, speech, and involuntary actions like swallowing and blinking.

We are going to demonstrate how a neural network works by playing a game of snap. I need seven volunteers. Four students will face towards the board and be the sensory neurones in the brain, who can see the colour of the two squares. Each are given a rule for when they should tap the relay neurone behind them. Person one will tap the relay neurone on the shoulder when square 1 is red, student 2 will tap when square 2 is red, student 3 will tap when square 1 is black, and student 4 will tap when square 2 is black.

Students 5 and 6 are the relay neurones, responsible for delivering messages from the brain to the body. When they receive two taps on the shoulder, they should then tap student 7, the motor neurone. Student 7 is the motor neurone, responsible for delivering an action in the body. When they receive any tap on the shoulder, they should shout Snap. Students 5-7 should face away from the board, so they cannot anticipate the response. Let's see our brains in action!

If we want to build artificial intelligence, what better thing to use as an example than the human brain?! AIs are often built in ways that mimic, or copy, the way the human brain works, using artificial neural networks. Computer scientists programme the decisions that the AI model should take when they are given different inputs.

- **For example, can you remember how you learnt the difference between a cat and a dog?** Normally, when humans are very young, we are shown images of dogs and cats, and told "this is a cat, can you say cat?". You might also be asked "what noise does a cat make?", and guided by your caregivers to get to the correct answer. Over time, as you see more and more different types of cats
- **How might we go about doing this with an AI model, so that it can correctly identify whether an animal it is shown is a dog or a cat?** It is pretty much exactly the same! We must give the AI model lots of different images of cats and dogs as input, and classify these images as a cat or a dog. The AI model will look for patterns between the cat images, and try to find those same patterns in new images it is shown.
- **What if we only showed the AI black cats, and ginger dogs?** It might not be able to identify that a ginger cat is a cat, or a black dog is a dog. We need to give it lots of rich inputs that show a whole diversity of images and patterns it may come across.

AI Dance Trainer (10 minutes)

Let's imagine I am now an AI dancer. It is your job to train me to dance– whatever moves you show me I will try to copy. I want you all to do whatever dance move comes into your head! Now the AI dancer will try to copy.

- **What went wrong here, why does our AI dancer look a little bit silly?** Because we gave it too much messy data to start with.
- **What could we do to try to get the AI to do clearer dance moves?** Give the dance moves one by one!

Let's pick one dance move to teach the AI dancer first. The AI dancer's moves have improved, because we give it much more clean data. • **Why is it a good thing that we got lots of people to do the same move?** Because now the AI knows how the dance move looks on different people, it could copy and recognise that dance move no matter who in the class is doing it. Let's keep giving the AI dancer more moves.	
Training an AI (10 minutes) We now have an activity for you to practice your skills on how to correctly train an AI model. This worksheet describes Ash's attempt to make an AI bot that gives prizes when people win at a funfair game, however the bot seems unfair and does not give prizes to everyone. For each question you will need to identify why the prize has not been given, and how the AI bot could be improved.	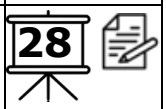
Training our own AI (30 minutes) You are now going to train your very own AI model, using Google's Teachable Machine. You will need to open this up on your browser, and then decide on what it is that your AI is going to classify. You can either choose to make an AI that can tell the difference between different pictures that you can find online, different poses that you can make using a webcam, or different sounds you can record using the laptop's microphone. Remember to also make a class that shows a 'neutral' pose or contains background noises. Follow the examples we have here, or get creative and come up with your own. When you are ready, we can share our creations with each other.	
My AI Model (10 minutes) Now it's time to describe what you did when creating your own AI model. Use the worksheet prompts to explain the steps you took when training your AI model, and what you could do to improve it.	
End of session – recap Thank you very much for joining this masterclass today, we hope you enjoyed it and learnt a lot about artificial intelligence! If you would like to have a go at some extension activities, there are these activities available on similar topics. If you have any questions, comments or thoughts we would be happy to hear them now. You can also ask the Ri, by emailing any questions to masterclasses@ri.ac.uk.