AI Explainer

A Short Guide for Teachers to Key Issues in the World of AI

The study of Artificial Intelligence is a wide-ranging issue which touches on many different elements of our society. With AI technologies like Large Language Models (LLMs) becoming more widely used every day, it can be difficult to keep up with the developments.

People come to conversations about AI from many different places, and with so much information out there about real AI technologies and science fiction approaches to AI, it can be useful to be able to navigate the wider issues which may be brought up when learning about AI.

This Explainer gives a top-level overview on some of the common discussion within the world of AI, and then explores some potential frequently asked questions that learners may ask during each lesson within this Teaching Pack.

Read on to find out more about:

- Defining “AI”
- Generative AI & Large Language Models
- Narrow AI & General AI
- The Environmental Impact of AI
- General FAQs

Defining “AI”

What does “AI” mean, and why is it difficult to define?

There is no single, universally agreed-upon definition of “Artificial Intelligence”.

The term “Artificial Intelligence” is broad, covering a vast range of different technologies, techniques and applications. Because of this, defining AI in a single, fixed way is challenging.

The term was originally coined in the 1950s, most notably in a 1955 proposal by John McCarthy and colleagues for the Dartmouth Summer Research Project on Artificial Intelligence. This foundational document stated:

> "*The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”*
>
> **John McCarthy et al., 1955, Proposal for the Dartmouth Summer Research Project on Artificial Intelligence**

Here is a breakdown of two different definitions of “AI”, exploring the pros and cons of each.

The first definition is more conceptual and popular, the second definition is more technical and is most often used in policy and governance contexts.

> *Definition A*
>
> **Artificial Intelligence is the field of study within computer science that focuses on designing computer systems that can perform tasks that we would normally consider to need “human intelligence”, such as making decisions and producing natural language.**
>
> **The technologies that result from this field of study are known as “AI technologies”.**

It can be unhelpful to define Artificial Intelligence like this because:

> ✗ It is a wide definition, so is not very exact
>
> ✗ It defines AI in comparison to “human intelligence”, which gives the false idea that AI technologies have any sort of agency, when in reality they are algorithmic processes designed to automate specific tasks.

It can be useful to define Artificial Intelligence like this because:

✓ It highlights that AI is not one thing, rather it is an entire scientific discipline, a whole field of concepts, problems, and methods for solving those problems. Thinking of “an AI” can be misleading as it suggests that AI is a thing that can be picked up and put down

✓ While this definition is evocative and helps us have a top-level understanding of the kind of tasks we use AI for, the comparison to human intelligence can lead to people misunderstanding the real-world AI technologies with conceptual “thinking” AI characters from science fiction

Put simply, this definition describes AI as systems that carry out tasks we'd typically expect from humans, but the systems themselves don’t think or understand like humans do.

Definition B

An AI system is a machine-based system that can, for a given set of human-defined explicit or implicit objectives, draw conclusions, from the input it receives, how to generate results such as make predictions, content, recommendations, or decisions that can influence physical, real, or virtual environments.

Different AI systems are designed to operate with varying levels of autonomy and adaptiveness after deployment.

- Based on the OECD Definition

Put simply, this definition describes AI systems as tools that process inputs and make decisions or recommendations based on programmed goals, without necessarily mimicking how humans think.

It can be unhelpful to define Artificial Intelligence like this because:

✗ It relies on the reader understanding some technological terms

✗ It is dense and abstract, and could miss the opportunity to explain AI to someone trying to understand AI in simple terms

It can be useful to define Artificial Intelligence like this because:

✓ It is precise in its descriptions of what is happening within an AI system

✓ It standardises a specific understanding of AI that is useful when discussing how to regulate AI technologies

✓ By not comparing AI to human intelligence, the description avoids attributing human-like qualities to AI, which can be misleading

Breaking It Down

The main difference in these definitions is whether "AI" is framed in the context of "human intelligence" or not. There are, however, common elements across both which may be helpful for you to decide on how you would like to define "AI".

From the definitions above, AI is:

- More than one thing, one technology or one tool.
- A suite of different technologies used to perform complex tasks, such as making decisions, understanding language and making recommendations, based on the data they have been trained on.
- Sometimes described in relation to "human intelligence", sometimes not.
- A rapidly evolving and ever-changing concept, with definitions and understanding of "AI" shifting over time.

It is also important to note that the recent rise of one particular kind of AI technology, '**Generative AI**', has led to people considering "AI" to just be chatbots and image generation – but it's actually a much wider range of technologies, of which "Generative AI" is just one!

FAQs

Is AI smarter than people?	AI technologies can do some tasks really fast and well, like recognising patterns in lots of data, but it doesn't understand things the way people do. People can think creatively, feel emotions and make moral choices. AI doesn't have feelings or common sense; it only works within the rules it's been given.
Can AI be my friend or have feelings?	It might seem as if AI chatbots are friendly, but it's not really a friend. It doesn't have feelings or care about how you feel. It just follows instructions based on what it's been trained to do.

Generative AI

What do we mean by “Generative AI”?

Generative AI is a type of artificial intelligence that creates content such as text, images, music or video by analysing and synthesising patterns from large datasets.

These datasets often include copyrighted material collected from the internet, which raises important questions about consent and ownership. Although the outputs may seem original, they are fundamentally based on and influenced by the data used during training.

Large Language Models (LLMs)

One key form of Generative AI is the **Large Language Model (LLM)**.

These are trained on vast amounts of text to predict and generate words in sequence.

ChatGPT is an example of an LLM. When you type a question or prompt, it predicts what words should come next to form a meaningful reply.

LLMs don’t understand language the way humans do. They don’t have thoughts or feelings. Instead they are advanced pattern-matching systems, predicting likely sequences of words based on their training. This means they can produce fluent and well-structured writing but can also make things up or get facts wrong.

Some well-known LLMs include:

- **GPT-4** by OpenAI (used in ChatGPT)
- **Claude** by Anthropic
- **Gemini** by Google

These models are often built with billions of *parameters*, adjustable parts that help fine-tune how the model responds. The more parameters, the more powerful and complex the model tends to be.

Image Generators

Generative AI can also create visuals. Tools like **DALL·E**, **Midjourney** and **Stable Diffusion** use models trained on millions of images and their captions. When given a text prompt such as "a cat in a spacesuit on Mars" the AI tool builds a “new” image by piecing together learned visual patterns.

Image generators work differently from language models but follow the same principle. They generate content based on what they’ve learned from training data.

What Makes Generative AI Different?

Generative AI is distinct from other AI technologies & techniques because it:

- Creates new content rather than just analysing or recognising existing data
- Responds to prompts making it interactive and adaptable
- Can be used across media including text, sound and images

This makes it powerful and flexible but also raises questions about accuracy, bias, originality and ethical use. For example, the models may reproduce stereotypes from their training data or generate misleading content that looks real.

FAQs

Can I trust what Generative AI says?	Note completely. It can produce convincing and confident responses that are wrong or misleading. This includes making up facts, inventing sources or repeating bias from its training data. Always double-check important answers and use AI as a starting point, not a final authority.
Where do the answers come from in Generative AI?	Generative AI learns from large datasets made up of books, websites, articles and other online information. The answers it produces are shaped by patterns in that data, not by understanding. Some of the content it was trained on was used without the owner's permission, which some people think was wrong.

Scottish AI Alliance

Narrow AI & General AI

Are current AI technologies related to science fiction versions of AI?

In science fiction, AI is often portrayed as an all-powerful thinking machine that can take control of the world.

These portrayals, from omniscient computer overlords to sentient androids, have shaped public imagination for decades. However, they blur the line between speculative fiction and scientific reality and bring a lot of baggage into discussions around AI.

In real-world research, the term **Artificial General Intelligence (AGI)** refers to a hypothetical kind of AI that could understand, learn and apply knowledge across a wide range of tasks at the level of a human or beyond. **AGI** does not exist, and whether it can or should be built remains a subject of active debate among scientists and researchers.

Most current AI systems are examples of **narrow AI**; they can perform specific tasks, like language translation or pattern recognition but they lack self-awareness, consciousness or general reasoning abilities. Importantly, these systems do not "understand" their outputs in the way a human or AGI would, nor do they possess desires or intentions.

Science fiction narratives can be useful thought experiments, raising questions about power, control and unintended consequences. These narratives can give a useful hook to explore complex issues with learners. But they can also distort public expectations and obscure urgent, real-world issues.

For example, current **narrow AI** systems already raise serious ethical concerns: bias in decision-making, lack of transparency, data privacy breaches and the potential for misuse in surveillance or automated warfare.

Are there better ways to talk about AI?

It can be helpful to always differentiate between the science fiction version of "AI" and the real-life Narrow AI Technologies, as it allows for a more realistic understanding of the current abilities of the technologies, helps prevent inflated expectations and avoids confusion.

One way we can do this is by using images to represent AI that are more grounded in the real-life way the technologies work.

Currently, in a lot of media, AI is represented with images of *glowing blue brains* and *men shaking hands with robots.*

If you are looking to find images to represent AI that resist that narrative, you can find a whole library via **BetterImagesofAI.com**

Traditional AI Imagery

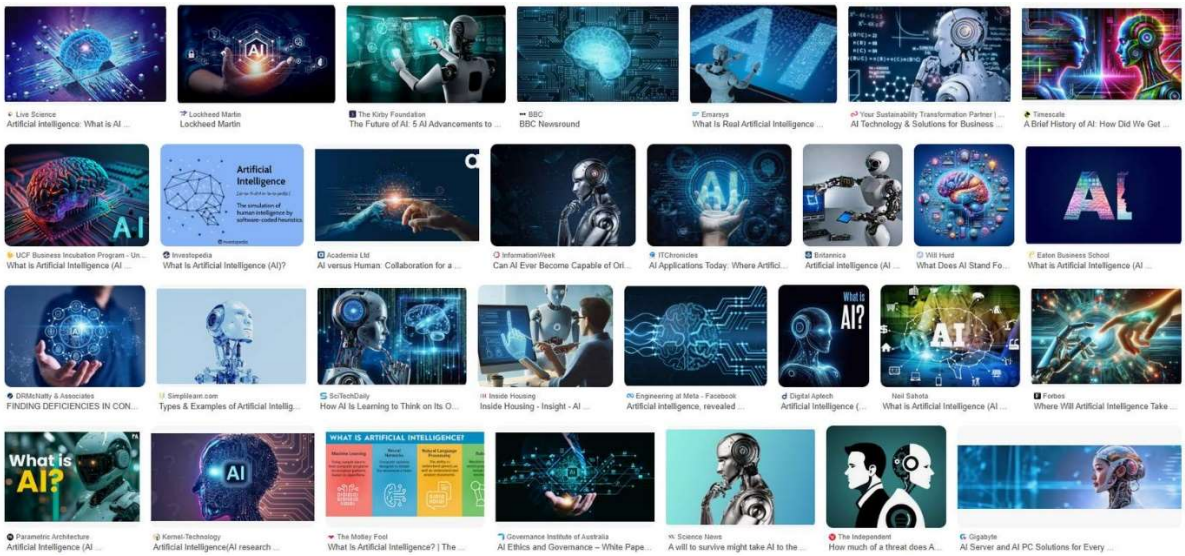

Better Images of AI

The images available via BetterImagesofAI.com aim to realistically portray the technology and the people behind it and point towards its strengths, weaknesses, context and applications

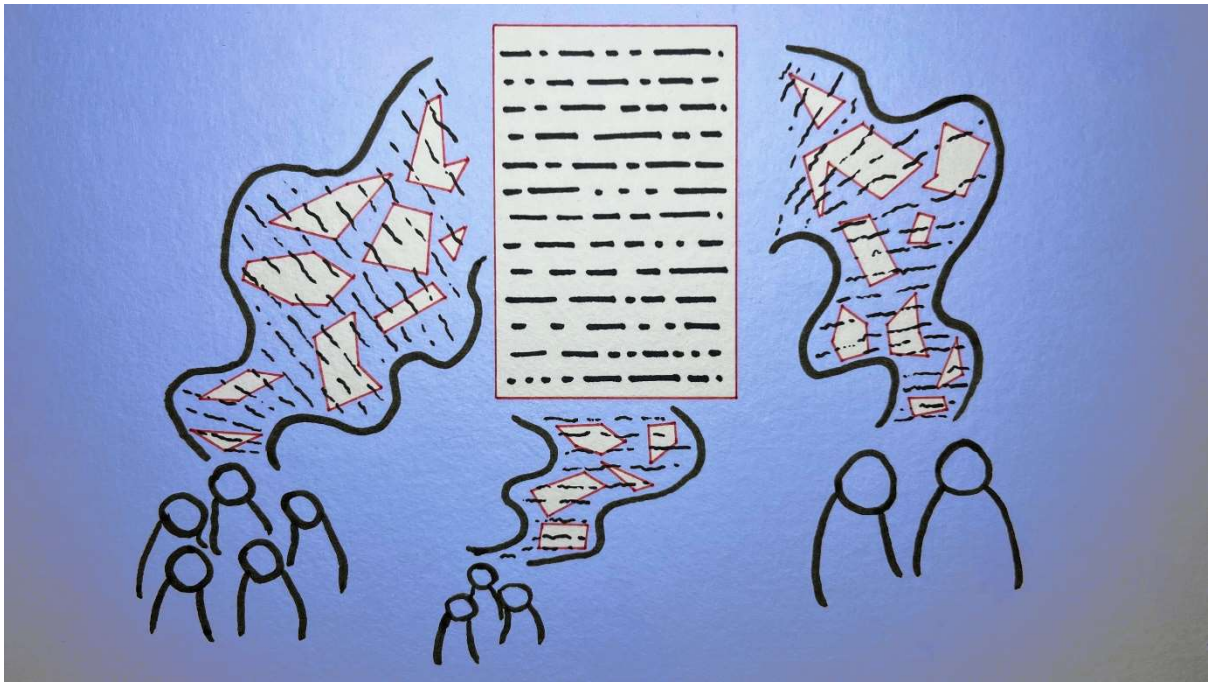

‘Data Processing’

Yasmine Boudiaf & LOTI / https://betterimagesofai.org / https://creativecommons.org/licenses/by/4.0/

Scottish AI Alliance

FAQs

Is AI the same as a robot?	No, not always. AI technologies are a kind of software or computer programme. Robots are machines, and some use AI technologies but not all of them do. Some robots are just remote controlled or follow simple instructions without using AI.
Can AI take over the world like in films?	Stories about robots taking over are fun and exciting, but in real life AI doesn't have plans or goals. AI technologies only do exactly what task they've been trained to do, and those tasks are quite limited. The people who build and use AI systems decide how they work. So, it's important to ask how people should use AI responsibly, and ask what people's plans or goals are with AI.

The Environmental Impact of AI

How do AI technologies use energy, carbon & water?

Training and using AI takes a lot of electricity. There are two main stages:

- **Training**, when the model learns from large datasets
- **Inference**, when it generates content or answers questions

Both happen in data centres, which are filled with high-powered computers that run non-stop. Training large models like GPT-3 requires thousands of specialist chips working for days or weeks. This process can use as much energy as several hundred UK homes in a year.

Once trained, models still use electricity every time they're used. For example, generating an image or answering a question with an AI technology like ChatGPT uses more power than a standard web search. Multiply that by millions of users and the energy use quickly adds up. The International Monetary Fund has warned that global AI electricity demand could more than triple by 2030.

Carbon Footprint - Electricity is only part of the story. What really matters for the climate is how that electricity is made. If data centres rely on fossil fuels, the carbon emissions are high.

Training GPT-3 produced an estimated 500 tonnes of CO_2, similar to driving a petrol car over a million miles. Another model, BLOOM, which used mostly nuclear energy, produced just 25 tonnes.

Generative AI is growing rapidly, and even major tech firms say it is making their climate goals harder to meet. Campaigners are calling for more transparency, cleaner power and better efficiency.

Water Use - AI also uses water, mainly to cool the servers in data centres. These systems often rely on chilled water or evaporation to keep machines from overheating.

Each 20 to 50 questions to a tool like ChatGPT can use about half a litre of water, mostly lost as steam. Training a model like GPT-3 used around 700,000 litres — roughly the amount needed to manufacture 300 cars.

Water use depends on climate, cooling method and time of day. Running models at night or in cooler seasons can save water, and some companies are exploring new cooling technologies. But as AI grows, so will its demand for fresh water.

Hardware - AI needs powerful chips and servers, which have environmental impacts of their own. These are made from materials like silicon, copper and rare earth metals, which require mining and refining.

Chip manufacturing is also resource-intensive, using large amounts of energy and water. Before an AI model is even trained, there is already a carbon and water cost from building the hardware.

The pace of AI development means that servers are upgraded regularly. If old hardware is not reused or recycled properly, it becomes electronic waste that can harm the environment.

FAQs

Can AI help fight the climate crisis?	Yes, it can. Some AI systems help scientists understand better ways to use energy or plan how to protect nature. But other types of AI, like the ones that make pictures and write stories, use a lot of power and water, which can harm the planet if we're not careful.
Is using AI bad for the planet?	It can be if we don't use it wisely. Some types of AI use lots of energy and water, which can add to pollution and the climate crisis. If people work on making it cleaner and better, AI could get more environmentally friendly.
Can I help reduce AI's impact on the environment?	Yes. You can ask fewer questions just for fun, using AI only when you really need it. You can also choose AI tools that try to be more eco-friendly, and encourage others to do so.

Lesson FAQs

What other questions might come up?

Lesson 1

Does AI know what it's doing?

No, AI technologies don't understand things like people do. They just follow patterns in data. They don't think or feel. They can only do what they've been trained to do, and they don't have any context they've not been given.

Can AI make mistakes?

Yes, AI technologies can make mistakes when they get given confusing data or if they haven't been given enough different examples. That's why it's important that people still check what AI technologies do, and don't just trust it to get everything right.

What do children's rights have to do with AI?

AI technologies are becoming a bigger part of life, and it can affect how people are treated. The UNCRC says that children have the right to privacy, fairness, education and being listened to. We need to make sure that the use of AI doesn't harm those rights, and helps them instead.

Lesson 2

Why does it matter what kind of data we give to an AI system?

AI technologies learn by spotting patterns in data, like how people can learn by seeing lots of examples. The better-quality data it has, the better it gets at doing its task. If we only show it one kind of example, like all the same kind of cat, it might think that's the only type of cat that exists. So it might miss or misjudge anything different.

What does "bias" mean in AI?

Bias means something is unfair or unbalanced. If an AI system has biased data, like mostly pictures of one type of cat, it won't work properly for everyone. It can lead to unfair results, even though the AI technology is not doing it on purpose.

Why do we need lots of different drawings in the catflap activity?

To help the AI technology learn all the different ways a cat might look – different colours, sizes, positions, types of pictures. If everyone draws the same type of cat in the same way, the AI technology might not recognise cats that look different.

Could this type of AI technology ever be used on people? What would happen?

Yes, some AI systems are used to recognise people's faces. If those systems are used, and they're only trained with certain types of faces, they might not work well for everyone. This could mean that some people are wrongly identified by the AI technology.

Lesson 3

Can AI keep my secrets safe?

AI chatbots collect all the information that you share with it, and the company that runs it will now have that information. You shouldn't share things that are personal or private with an AI chatbot, it's much better to talk to a trusted adult.

Will AI take over jobs or stop people from working when I grow up?

AI technologies are changing jobs by doing certain tasks faster, but people still need to be there to oversee the task and make decisions. As AI technologies are used more and more, some jobs could be different or disappear in the future, but new jobs dealing working with AI might appear.

Does AI remember everything I tell it?

Some AI chatbots can store some information, but it all depends on how it is designed. Sometimes it will remember and store what you have told it, sometimes it will not. You should be able to look into the specific AI chatbot you are using to find out what it does with the data you share with it.

Lesson 4

How can AI be used to help me learn?

AI can be used in different ways to support learning, like using AI chatbots to give you extra practice, explain things in new ways or creating fun learning activities. But AI is just a tool, it doesn't replace teachers or know the way you learn best. Sometimes it works well, sometimes it doesn't work well, and it's important to think carefully about how it's used and what you use it for.

Can AI make mistakes when helping me with my work?

Yes, AI can make mistakes. AI chatbots are programmed to give you the most likely answer based on the data it has been trained on, it is not programmed to give you the truth or the correct answer. If an AI chatbot has been trained on information that it is wrong, it may give you the wrong answer. It's always good to double-check answers from AI chatbots with a teacher or a trusted adult.

Will AI replace my teacher?

AI can help teachers with some tasks, but it shouldn't replace them. Teachers understand how you feel, help with tricky questions that an AI technology can't handle and know you personally. AI can't do those things. Teachers are very important for learning and supporting you.

What should I do if I don't like how AI is used at my school?

If you feel worried or uncomfortable about AI in your school, it's important to tell a trusted adult like your teacher or a parent. You can share your feelings and ask questions. Everyone should listen to your ideas about what's right and what's fair when using AI, as children's voices matter a lot!

Lesson 5

Why do some people want to keep certain information private?

Some information about us is personal because it tells other people things about who we are or how we feel. People want to keep some information private to feel safe, protect their identity and to choose who they share it with themselves. Privacy helps us control our own life and stop others from using that information about us in ways we don't like or expect.

Why do some apps and websites ask for permission to use cookies?

Cookies are small bits of information that apps and websites save on your device to remember things like your preferences or what you looked at on the site. Apps and websites ask for your permission to use cookies because the law says they must be clear about what information they know about you. This helps protect your privacy by letting you decide if you want to allow the app or website to remember that information about you or not.

How can I decide what information is okay to share with an AI chatbot, and what should stay private?

Deciding what to share is about thinking who will see your information and why. If it's something like your favourite colour, it might be okay to share. But things like your address, passwords or medical details should usually stay private. When you use an AI chatbot you are sending information away to a company. It's good to ask yourself "Would I be okay if anyone in the world knew this information?" If the answer is no, it's probably better to not share it with an AI system.

Why do apps and companies collect data?

Companies collect data for different reasons. Sometimes, they use it to make their websites or apps work better and more fun for users. But they also collect data because it helps them make money. For example, by knowing what you like, companies can show adverts you might be interested in, and advertisers pay companies for this. It's important that companies do this in a way that respects people's privacy and keeps their information safe.

How do I know if a picture or news I see is real?

Sometimes AI can make pictures or news stories that look real but aren't. To check if something is true, you can do a few things.

- Look for the source: is it from a trusted website or person?
- Check with an adult: ask a parent or teacher what they think.
- Search the image or story: see if it shows up on other trusted sites.
- Trust your instincts: if something seems strange or too good to be true, it might be fake.

Even grown-ups get tricked sometimes, so it's always OK to ask questions.