Department for Education

The safe and effective use of AI in education

Module 1 - Understanding AI in education video transcripts

June 2025

Contents

Video 1 – An introduction to AI in education

Welcome to module one, video one from the safe and effective use of AI in education resources.

Sufian Sadiq, Director of Teaching School, Chiltern Learning Trust:

“We are living through a time of rapid technological change, with artificial intelligence or AI transforming the way we work, communicate and learn, from automating repetitive tasks to supporting our planning of lessons. AI has the potential to support education in many ways, reducing workload and supporting teachers in powerful new ways. But as we embrace this technology, the humans are more important than ever, grounded in critical thinking, professional judgment, and ethical responsibility.

We know many educators feel uncertain about AI. Some are excited about its possibilities, while others are concerned about the risks or simply don't know where to start. That's why these materials are designed for everyone, whether you're brand new to AI or already experimenting with its capabilities.

In these resources, we will demystify AI, explore its implications for education, and consider how it can be used safely and effectively. We will highlight key risks, particularly in relation to students and provide practical guidance on responsible AI use. Our hope is that you feel inspired and empowered to explore AI confidently, use it to reduce workload and enhance learning while maintaining the values and professional judgment that define great teaching.”

These resources align with the Department of Education's position on AI in education. The resources have been developed through extensive consultation with educators from across the sector. We conducted a survey to understand teacher perspectives on AI, and their insights shaped our approach. In the first module we'll be exploring the question, what is AI and how does it work, with a particular focus on what we need to know as educators.

Bukky Yusuf, Deputy Headteacher and Science Lead, Edith Kay School:

“Some of the key benefits of generative AI for educators is the fact that you can create some really interesting, differentiated and personalised learning resources.”

Aaron Patching, Strategic Digital Leader and AI Network Lead, River Learning Trust

“So we found that image generators were a fantastic way to hook them in. And again, that we linked that to a specific purpose because we wanted to work on their oracy. We wanted to work on their descriptive writing.”

Chris Loveday, Vice Principal, Barton Peveril Sixth Form:

“So we had a leadership team meeting where we decided that we should explore how we can best leverage artificial intelligence, but to solve problems, not just because it was AI. We didn't chase AI.”

Presenter: Artificial intelligence or AI is a field of computer science dealing with the mimicking of human intelligence. The computer scientist, Alan Turing, spoke of AI as far back as 1947. He said in a public lecture, what we want is a machine that can learn from experience.

You'll have heard of generative AI applications, such as ChatGPT, Gemini or Copilot. These tools use a type of AI known as large language models, and many of you may well be using these types of tools regularly. Others may just be starting to think about how you might use AI tools.

We're now going to briefly cover some key definitions that we need to know, which we're going to return to later in the module. AI is an umbrella term covering a number of different types of AI. We'll look at generative AI in more detail later in the module, but there are some key definitions it's useful to know.

We're used to computer programs performing in a predictable way due to the way they're programmed. These traditional programs or apps are known as rule-based systems. Generative AI will generate or approach problems in a different way which makes it much less predictable or possible to understand how it works. This is because many of the newer AI tools are data driven rather than rule based.

Another important term is prompt. A prompt is the instruction that we give generative AI. It is often in the form of natural texts, but could also contain images or files.

Narrow AI, this is the type of AI designed for specific tasks and can't be easily adapted to do other things. It works using carefully selected data to complete a particular job. Some examples include facial recognition, such as that used to unlock phones, and the AI in spell checkers or adaptive learning tools.

General purpose AI. General purpose AI can perform many different tasks rather than just being limited to one job. It's trained on large amounts of data and can be adapted for different uses.

Generative AI is a type of AI that creates new contents, such as text, images, videos, or even music, based on a user's input or prompt. This is the type of AI we'll be focusing on throughout this toolkit.

Large language models are a type of general purpose AI designed to understand and generate human-like texts. These models are trained on vast amounts of text and can answer questions, summarise information and even translate languages. Some will also deal with video or images.

Machine learning is a key subset of AI where computers learn from data instead of being directly programmed for each task. There are also different ways in which generative AI systems learn from the data. There's what's known as supervised learning, where the AI system learns from examples with clear labels, such as teaching an AI system to recognise apples by showing it labelled images of apples and then other fruits

Unsupervised learning is where the AI system finds patterns in data without being given labels, such as identifying emails as spam based on recognising suspicious patterns without being explicitly shown labelled data.

Deep learning is a more advanced form of machine learning, inspired by how the human brain works. We'll be learning more about this in the next video.

What makes generative AI so accessible is that it is a type of AI that can generate content based on an instruction given by the user. This instruction is known as a prompt. Interaction with an AI system can be conversational, using everyday language, and can also involve images, audio, or video.

While generative AI systems can appear convincingly human, it's important to always remember that it's only ever a very effective technology.

Miles Berry, Professor of Computing Education, University of Roehampton:

“In thinking about what we mean by artificial intelligence, it's worth taking a step back and thinking about what we mean by intelligence. Think perhaps of your pupils who demonstrate something that you see as intelligence. What does that mean? It means that they have some sort of understanding of the world in which they find themselves? It perhaps means that they're reasonably articulate, they're able to express their ideas well. Perhaps you would say it means that they have some sense of agency. They can choose what they want to do. They're able to make decisions for themselves.

“You would also, I hope, acknowledge that it includes an element of being able to learn from their experience to be able to get better at doing things. When we talk about artificial intelligence, we might be thinking of many of those things about being able to make sense of experience, being able to learn from the environment, learn from data and spot patterns and determine rules from that, much like an intelligent person would.

It perhaps also includes some more general sense of all round understanding of knowing things in some sense and being able to express ideas and perhaps even to have some degree of agency to be able to affect things. So artificial intelligence covers a whole multitude of things, but I think the easiest way to think of it is when a machine can do something which perhaps previously we'd have assumed for a human being to do requires a degree of intelligence.”

Presenter: As a user of a generative AI tool, you are responsible for what you put in and how you use the output. AI can be used with varying degrees of success to complete tasks such as writing emails or newsletters, generating ideas for lessons, translating articles, simplifying content, and much more.

Bukky Yusuf, Deputy Headteacher and Science Lead, Edith Kay School:

“Resources can be in the text formats. It could be audio, even video now as well. And I think it really allows you to just do different things in a way that we haven't done that can really engage young people, but also allows young people to learn in different ways, both within and beyond school.”

Presenter: AI can be very effective and useful for educators by saving them time and reducing workloads. A recent Education Endowment Foundation study found that AI can reduce some areas of work by more than a third. The Department for Education have created clear guidance in the policy paper, Generative AI in Education. This clearly sets out the opportunities and risks associated with generative AI use in education. And this course is designed to support you in exploring these opportunities and risks.

If we're going to use generative AI safely and effectively in education, we need to make sure that we're following best practice whilst being aware of the limitations and potential biases. For example, whilst UNESCO has said it's committed to harnessing the potential of AI technologies for education, it also emphasised that the use must be guided by the core principles of inclusion and equity at a human centred approach to AI.

The Department for Education has set out some basic actions all educators should take if we're using generative AI. For example, we need to critically evaluate the outputs of generative AI to make sure it's accurate, free of bias and appropriate for the context. There has to be human oversight and accountability. We should adopt good data practices and avoid actions which would infringe data protection or intellectual property law.

It's important to follow the relevant policy in your setting. You should also refer to the generative AI policy paper and Product Safety Framework from the Department for Education. We need to understand and follow the process for reporting any issues or concerns about generative AI at school or college, and, given that it's such a rapidly developing field, it's good to stay up to date on best practice for the use of generative AI in education, and this support will help you with that.

AI is already all around us. According to Ofcom's 2024 media use and attitudes report, nearly half of children, aged eight to 17, had used generative AI technology – 46% of them had. As educators, we have a responsibility to understand the rapidly changing digital landscape they are experiencing.

In these materials, we'll explore some of the use cases in education as well as limitations. This module, understanding AI in education, asks the question, what is generative AI and how does it work? Module two looks at interacting with generative AI in education, module three looks at developing the safe use of generative AI in education and module four looks at some of the use cases.

Accompanying the teacher toolkit, there is a leadership toolkit, and this is specifically designed for education leaders looking to integrate the use of AI safely and effectively in their settings.

Video 2 – Focus on generative AI

Presenter: Welcome to module one, video two from the safe and effective use of AI in education resources.

In the first video in this module, we learned about some of the uses of AI in education and had an overview of how generative AI works. Generative AI has many uses for us as educators. We'll explore some of these in module four, such as generating lesson resources, supporting lesson planning, and generating a podcast with virtual presenters to support revision on a specific topic.

Let's now explore in simple terms how AI works. Most of us use AI tools every day without even realising it. Narrow AI, such as email spam filters, predictive texts, and autocorrect, all rely on AI to analyse patterns and make predictions based on large data sets. There are also other applications for AI beyond generative AI, such as adaptive testing platforms that personalise learning and image recognition systems to describe something shown to a camera for someone who is visually impaired or blind.

Generative AI focuses on creating new content such as text, images, audio, video, and code. This area of AI is evolving rapidly and could support education in a number of ways. It is all possible because of deep learning, which is a type of machine learning. Machine learning is a field of study within AI where a computer program is trained on data and uses this training to make predictions based on that data.

For example, in medicine, machine learning has been used to analyse large data sets and assess the likelihood of a patient developing certain conditions based on their medical history and risk factors. Deep learning is a subset of machine learning, and a key element of deep learning technology is neural networks. Neural networks consist of interconnected nodes designed to replicate how our brain works. These neural networks allow a deep learning model to learn from uncategorised data efficiently, making it easier to use much more data to train the model.

Deep learning is used by large language models, or LLMs. You may have heard of and perhaps used large language models, such as ChatGPT, Gemini, or Copilot. LLMs are trained on a huge amount of information, including written text, images, or videos. Depending on the LLM you're using, it may also be trained on the instructions or prompts that you give it, although this doesn't happen in real time.

This means we need to be particularly careful with the data and information we provide in our prompts, avoiding inclusion of personal data such as name, address, phone number, protected characteristics, and business-sensitive information such as financial details. We must also avoid using intellectual property where we don't have permissions in place, and this could include pupil or student work.

To help us with this, it's important to refer to the DFE's generative artificial intelligence and education paper, filtering monitoring standards, and product safety expectations. There could also be issues in the data that the large language model is trained on, which could cause it to be biased or inaccurate, or the people designing the system may have built-in features that we are not aware of.

We'll learn more about this later in the resources in module three. Large language models, such as ChatGPT, Gemini, or Copilot, have natural language interfaces, which means we can use everyday text and speech patterns to give instructions to the model. These instructions are the prompts that we mentioned earlier. In some models, it is also possible for a prompt to contain images or code.

To summarise this process, you type in your prompt, the large language model breaks down your prompt and analyses it for context and meaning based on the patterns that it has learnt. Then it comes up with a predicted response based on its prior learning. It gives this back to you as an output, and the output could be text, but it could also be images, audio, video, or even code.

The output often feels natural and human-like. It's important to remember, though, that an AI model isn't human. It doesn't think, feel, or have the same contextual understanding of the world that humans have. But the process doesn't stop once you've received this response. You can fine-tune it by giving further prompts. You might not be happy with the level of detail, or it may be too complicated, too long, or even in the wrong language.

Your follow-up prompts will improve your outputs, just like with an email spam filter where we mark things as "not being spam", or when we tell our messaging apps not to auto-correct certain words.

We need to avoid using free personal accounts or public large language models that offer free accounts. This is due to the data protection risks associated with inputting personal or sensitive data. These models are typically hosted and managed by external providers, and any data submitted may be processed or stored outside of your organisation's control, potentially breaching data protection regulations.

They may also use your data to develop their AI model further. Under UK GDPR, organisations must ensure that personal data is handled lawfully, transparently, and securely. Inputting such data into a public large language model risks unauthorised access, cross-border data transfers, and loss of control over data retention.

Educators should only use tools approved for use in their education setting. These enterprise accounts offer added protections for data. If you need to use public large language models, ensure that you do not enter any identifiable information to protect personal data.

Department for Education

About this publication:

enquiries https://www.gov.uk/contact-dfe
download www.gov.uk/government/publications

Follow us on X: @educationgovuk
Connect with us on Facebook: facebook.com/educationgovuk