What is AI? Understanding Artificial Intelligence in your daily life

A brief introduction to AI

Everybody seems to be talking about AI – Artificial Intelligence - at the moment, and how it will change the way we live and work. Many of us are already using AI systems at home – on our phones, streaming platforms, or perhaps with voice-activated smart speakers - sometimes without even being aware of it.

For example, when we're using an internet search engine to find out about something, AI is often operating in the background to generate results and offer a short explanation of the information we're looking for.

In the workplace, we may see it as an option in the software our email uses, in Teams or Zoom for collaborative working or calls, or when using specific software to perform particular tasks like taking notes of meetings or transcribing recordings. You may have seen Microsoft CoPilot and other AI software pop up offering to help you with drafting documents or completing tasks.

But what exactly is AI? What might it be helpful to know about it? And how do you decide if you want to use it, and what you might use it for?

Let's start with the basics. The truth is, you've probably been using AI for a while without realising it. Let's look at some examples you might recognise:

So, at home, when a TV streaming service like Netflix, Amazon or Disney suggests what you might want to watch next, that's AI analysing your viewing patterns.   When your phone's camera automatically focuses on faces or suggests people to tag in photos, that's AI.

Even your smartphone's predictive text that finishes your sentences – that's AI too.   And, if you use one, your voice assistant – whether it's Siri, Alexa, or Google Assistant – will use AI to understand what you're saying and respond appropriately. And the spam filter in your email that catches unwanted messages?   Yes, also AI.

And AI is also present in the workplace. Those meeting transcripts that automatically convert speech to text? – That's AI in action.   Or when Microsoft Outlook suggests replies to emails or CoPilot offers to help draft a document - all AI, integrated already into everyday work tools. And AI might also be working in the background, ensuring that your files and documents are safe and secure.   So, what is Artificial Intelligence?

Artificial Intelligence involves teaching computer systems to perform tasks that typically require human intelligence - for example, like recognising patterns, understanding language, making predictions, or solving problems.

What makes AI different from regular computer software is that traditional software follows exact, pre-written instructions. This limits how useful a computer system can be as circumstances change around it. The system can't think for itself and if we want it to do something different, we need to upload a new programme.

However, AI software has the ability to adapt. Instead of following rigid instructions, it learns from examples and sets of data, and is able to use its learning from that information to make new decisions and adapt in new situations. Because it is bringing "fresh eyes" to a set of data, it might even discover new patterns or clusters that give users a new way of approaching a problem or issue.

This learning ability makes AI incredibly powerful, but it also introduces new challenges and considerations that we all need to understand, especially when the system is directly affecting, or interacting with, real people.

Understanding how Artificial Intelligence learns is crucial, especially if you work in a sector that uses large systems as this could impact on people or communities. And while AI has huge potential, there are some possible pitfalls for us all to be aware of and to look out for too.

The most common type of AI which is used at home and in the workplace is Machine Learning - we will talk more about machine learning in our later videos - but for now it is useful to understand that Machine Learning is a type of AI that learns from large amounts of data. This data could be published information, academic works, creative content, websites, or more specific documents or files. Large language models, such as ChatGPT and Google's AI answers are a type of machine learning model that you might be familiar with.

Through processing large amounts of data, the system produces a model - essentially a set of learned patterns - that allows it to process new information. This training allows the AI system to interact with new information, help users find the answers they are looking for, and possibly identify patterns that humans might miss or take much longer to spot.

An AI system might analyse many thousands of pieces of data from a list of stakeholders, customers or clients and learn to identify which needs, requests, services or products require specific attention - which people might need to be prioritised and what services might support them best. AI can work 24/7, processing information much faster than humans, and spot subtle patterns across vast amounts of data.

This suggests potential efficiencies and could be beneficial to enhancing services. But there are significant pitfalls to be aware of too.

Information security is a major concern. When data is inputted into an AI system, it is not always certain where that information goes, or if it is stored securely. This could mean, for example, that sensitive information might end up being used to train AI systems elsewhere, and for different purposes.

Most organisations have internal systems will have security measures in place – in the UK there is also the General Data Protection Regulation which governs how organisations must handle data - but AI itself isn't necessarily going to warn you about data security.

But perhaps the biggest risk in AI is bias. AI systems learn from data, and if that data reflects discrimination, prejudice and/or out of date thinking, language or practice, there is a very real danger that the AI model will perpetuate this. For example, if a system is trained on historical data where certain ethnic communities were over- or under-represented due to systemic bias, the AI might learn to identify people from those communities as more or less in need, at risk or engaged, when they're not.

So, the system may appear to be objective, but it's actually amplifying human prejudices. That is why it is so important that there is human oversight of AI – so that errors or biases can be identified and challenged.

Also, in the same way as humans can, AI systems can also make mistakes. AI systems are eager to please, and that means that sometimes they can make confident-sounding recommendations that are

completely wrong or miss important context. But humans can use other experiences and knowledge to understand the context. AI doesn't truly 'understand' situations the way that we do – it is recognising patterns but not bringing personal experience or exercising professional judgment.

This is why it is absolutely essential that human beings have the final say when AI systems are being used. AI simply can't replicate the human insight, ethical reasoning, and contextual understanding that comes with real-life experience. This is especially critical to recognise if AI is being used in the context of public services including social care, health, social security, and in decision-making.

AI isn't some distant future technology – it's already woven into our daily lives and for many of us, our work too. From the predictive text on our mobile phones to the information management systems our workplaces might be using, AI is quietly in the background processing information and making suggestions all day, every day.

In our next videos we will be looking at AI in more depth.