

Old McDonald had a server farm

Paul Hopkins is an independent consultant and academic, as well as being Chair of the Shap Calendar Group.

— AI-AI-oh!

Religious education has the tools to lead the conversation on AI in schools – tackling ethics, humanness, truth and diversity in ways other subjects rarely can.

It has been five years since the pandemic – which challenged the way that schooling was carried out in the UK, and around the world, and led to an increase in the use of technology and changes in pedagogical practice. Since then, schooling, and pedagogies, have mostly ‘bounced back’ to how things were before the pandemic, though there are some remnants of changed practice.

Just under three years ago, in November 2022, OpenAI launched ChatGPT, and the ripple from this continues to expand. Since November 2022 there have been a plethora of artificial-intelligence-enabled (AI-enabled) apps, chatbots, websites and more.

In this article I want to think not so much about the practical aspects of using AI (you can see my article from the Spring 2025 edition of *REtoday* for some practical ideas, or go to www.mmiweb.org.uk/gaiandre/), but rather about the wider discourse around the use of generative AI (G-AI) in education and how we go about thinking about its use, and the use of it by the children in our classrooms. Schools are pretty conservative places, and the response to challenges and changes to established structures is often to ignore or to ban (think about the current debates around smartphones) or to revert to the comfortable status quo – think about the move back from blended learning post the pandemic. I do not think we can do either of these with G-AI. Pandora’s box is well and truly open.

However, I want to argue that religions and worldviews education (RWE) and RWE leaders and teachers are in a strong position, not only to consider the discourse but to lead it; for the questions that are being raised when thinking about G-AI are the kinds of questions that RWE has been tackling in its curriculum for time immemorial – or for a good while! These fall into two areas:

1. RWE has the disciplinary tools needed when thinking about the application of G-AI
2. The substantive ‘stuff’ that is included within the RE classroom is at least analogous to the kinds of stuff that need to be explored in relation to G-AI

1. RE is good at dealing with controversial issues

Dealing with controversial issues is at the heart of RWE and has been for a long time. This might involve the über-issue of belief itself or questions around particular issues, but RWE has the toolkit to approach these issues and to explore and examine the claims made by each side (see also point 4 below). At the centre of the discussion on G-AI is the question 'Is this a positive force that will offer reductions in workload and drive new and exciting discoveries and innovations, or is this a harbinger of doom increasing unemployment and ultimately leading to human destruction?' As ever, the truth lies in the context, in the application and the actions of the humans developing and using the technology. We can see parallels with the debate on nuclear energy and nuclear weapons, for example. RWE is good at both dealing with controversial issues and also teaching about them to children.

2. RE is good at asking questions about human-ness

One of the growing concerns about G-AI generally and about artificial general intelligence (AGI) in particular is the question of what it means to be human. Again, RWE regularly tackles this question as it lies at the heart of what many religions and philosophies teach. This concept of humanness is also challenged as AI moves from 'simple AI' to AGI and onto the concept of the conscious machine. RWE has long tackled this question of *deus ex machina* and again has procedures and pedagogies for so doing.

In RWE we can explore the ideas of humanness as something that can be related to our ability to be rational, thinking beings. If we reduce this to calculative intelligence, then AI has overtaken us. Calculators are much better than people at long division, but if we have an idea of humanness as a fundamental and reciprocal relationship with and recipient of love (divine or other humans'), then it is hard to see any way in which technological advancement can threaten that.

3. RE is good at exploring moral and ethical dilemmas

There are a significant number of moral and ethical questions that are thrown up by the use of G-AI. These include questions around cheating, copyright,

intellectual ownership, the distribution of harmful content, privacy, and the removal of meaningful and enjoyable roles, as well as the future of teacher employment and the environmental footprint of AI systems.

RWE has a range of pedagogies and models for dealing with ethical questions. It is able to explore ideas from different standpoints and understand the concerns and beliefs of the other. These pedagogies and models might be narrative approaches drawing on the wisdom from sacred or philosophical texts, or ethical pedagogies such as virtue ethics, utilitarianism or consequentialism.

There is obvious overlap between this and the next point.

4. RE is used to dealing with truth claims

There are significant questions in the use of truth claims when using G-AI. Is the information that G-AI is creating 'true' or not? RWE is used to dealing with the idea of complex truth claims where truth is not a simple 'true' or 'false'. We have seen many examples of where G-AI has either distorted truth because of the datasets (large language models [LLMs]) on which it draws or where intentionally deceitful or malicious content has been shared (e.g. the recent image of Donald Trump as Pope). In Finland, for a while schoolchildren have been learning the difference between misinformation (inaccurate or erroneous), disinformation (intended to deceive or propagandise) and malformation (divorced from context in order to deceive).

There are important and significant concerns about the bias of LLM datasets, about the transparency of the algorithmic processes used in these models and about AI hallucinations where the AI generates made-up information, and where incorrect information is amplified. RWE has tools and experience in dealing with these kinds of issues.

5. RE is good at using stories

Merlin Donald talks about four stages of human cognitive development: the episodic, the memetic, the mythical, and the technological or theoretic – the use of external thinking systems.

G-AI creates narratives from a range of experiences, and RE has experience in the use of narrative as truth and different kinds of narrative to explore different kinds of truths: mythological, analogical, religious, scientific and so on ... RE has an epistemological toolkit that is unique.

Conclusion

Finally, this issue's theme is diversity. There are serious, significant and relevant concerns about the LLMs that the G-AI algorithms and programmers are drawing on. For example, are these too WEIRD (Western, educated, industrial, rich and democratic)? RWE has always been at the centre of exploring and promoting diversity in schools (and in wider society). This makes it even more vital that RWE teachers, leaders and champions are not just involved, but at the centre of the discourse around G-AI.

Pandora's box is well and truly open – we can't ignore or ban generative AI in schools.

Religions and religious practices have a history of embracing technological change, from radio sermons to televised eucharists, and from digital apps to virtual places of worship and even a G-AI prayer writer.

There is not a dichotomous answer. G-AI is not 'good' or 'bad', a 'friend' or an 'enemy', and we need to walk a careful path between these extremes. RWE has a hinterland of walking on eggshells, and RWE teachers have a toolkit for doing so. So, I am throwing down a gauntlet to RWE teachers to lead the discussion about the use of G-AI in schools. ■

NOTE

This is a synopsis of a longer article that will hopefully appear in the *British Journal of Religious Education* later in the year.

FOOTNOTES

Types of AI Mentioned in this article:

Artificial Intelligence (AI) is a broad term for machines or software that mimic human intelligence. The article refers to several kinds:

Simple AI performs narrow, specific tasks (like recommending products online).

Generative AI (G-AI) refers to AI that can create content such as text, images, or music – e.g., ChatGPT or DALL-E.

Artificial General Intelligence (AGI) is a hypothetical type of AI with human-level reasoning across many areas, not yet achieved.

Large Language Models (LLMs) like ChatGPT are trained on vast amounts of text to predict and generate human-like language. They raise deep questions about ethics, truth, and what it means to be human – questions central to RE.

Merlin Donald (born 1939) is a Canadian cognitive neuroscientist and psychologist best known for his theory of how human cognition has evolved over time.