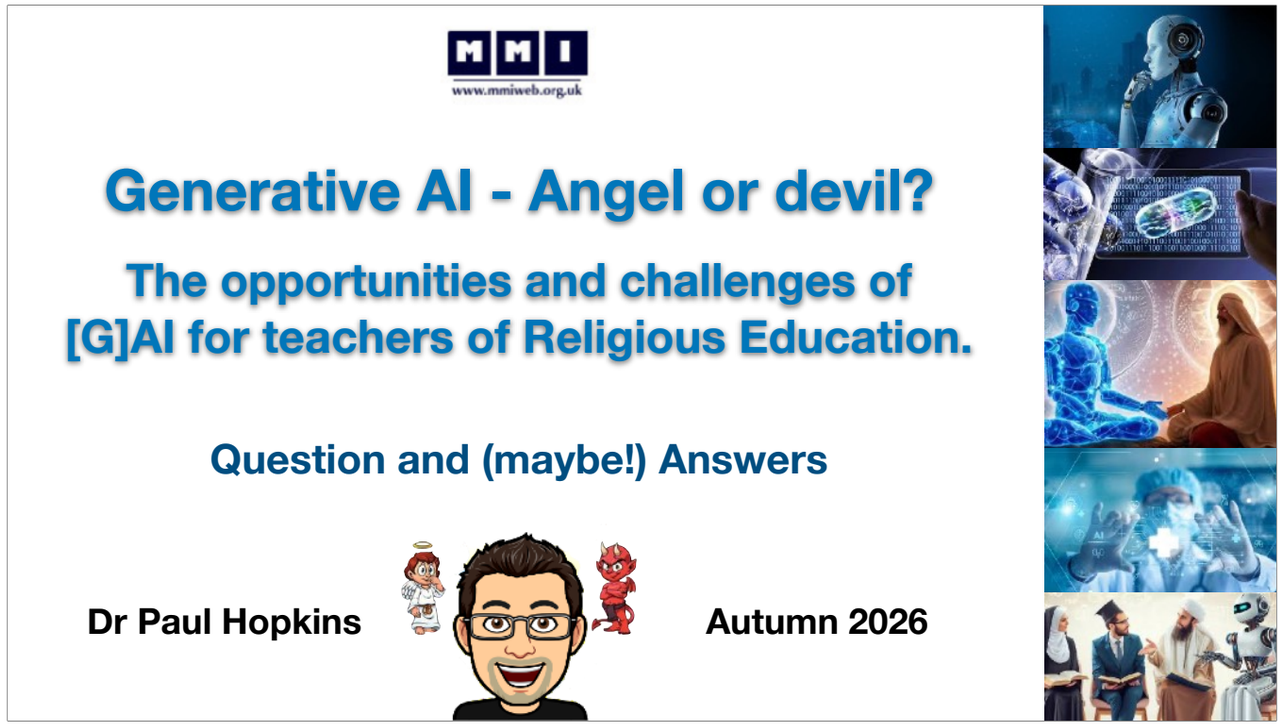
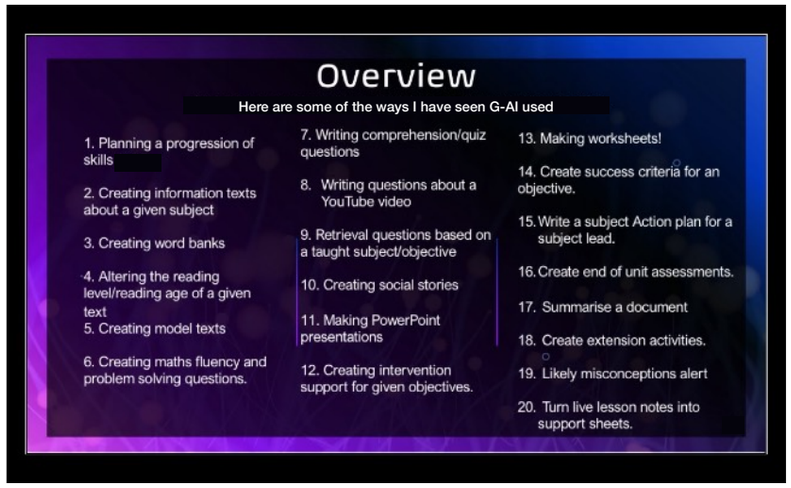


HI. I'm Paul. I've been working in RE and Technology for longer than many of you have been alive (!) and in the last four years I have been captured, captivated and consternated by Generative AI.

What kinds of things have you done?

What kinds of things do you want to do?





The first things to consider is you own use of G-AI - there are things above which you might use as a “check list” but there may be others as well. So, what kinds of things have you done already - or would like to do / be able to do with the technology?

1. What is Generative AI and an LLM ?
2. Is it a good thing or a bad thing ?
3. What does the CAS review / DfE say?
4. Epistemological Questions
5. A pedagogic model when using AI
6. The place of memory
7. What kinds of things can I do?
8. Using G-AI: A key skill - Prompt Engineering
9. Hazards and Bias
10. Is it cheating ?
11. Personalisation
12. Tools - What is out there?
13. A special place for RE in the discourse
14. The pace of change and the future
15. Three “laws” or principles for AI usage





Conclusion / Summary

Here are some of the questions that come up when we are considering the use of the technology. You may, of course, have others.

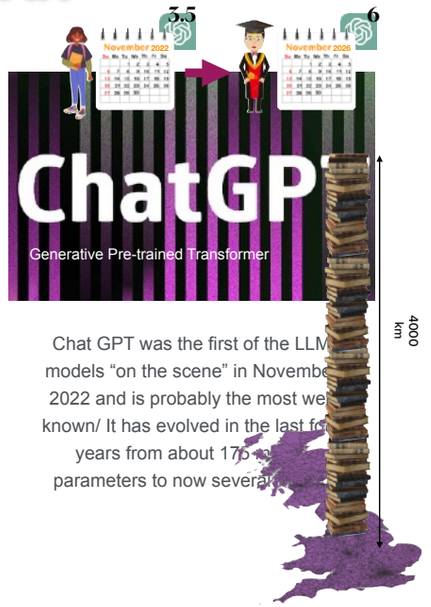
What is Generative AI?



Generative artificial intelligence or generative AI is a type of artificial intelligence system capable of generating text, images, or other media in response to prompts. Generative AI models learn the patterns and structure of their input training data, and then generate new data that has similar characteristics.



A LLM (Large Language Model) is a probability distribution over sequences of words. Given any sequence of words of length m, a language model assigns a probability P to the whole sequence. Language models generate probabilities by training on text corpora in one or many languages.

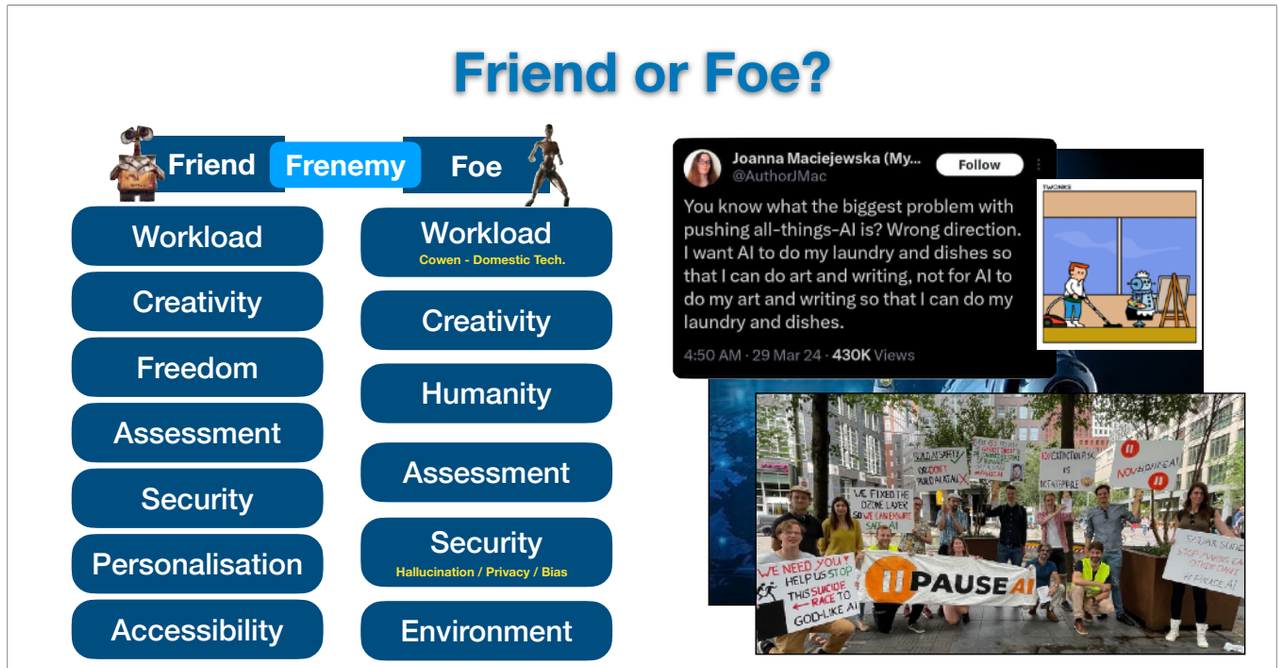


Chat GPT was the first of the LLM models “on the scene” in Novembe 2022 and is probably the most we known/ It has evolved in the last fo years from about 17b parameters to now severa

This session is not going into the technology in deep detail but the core of how G-AI works is a move from an algorithmic approach to programming to a more heuristic / probability approach.

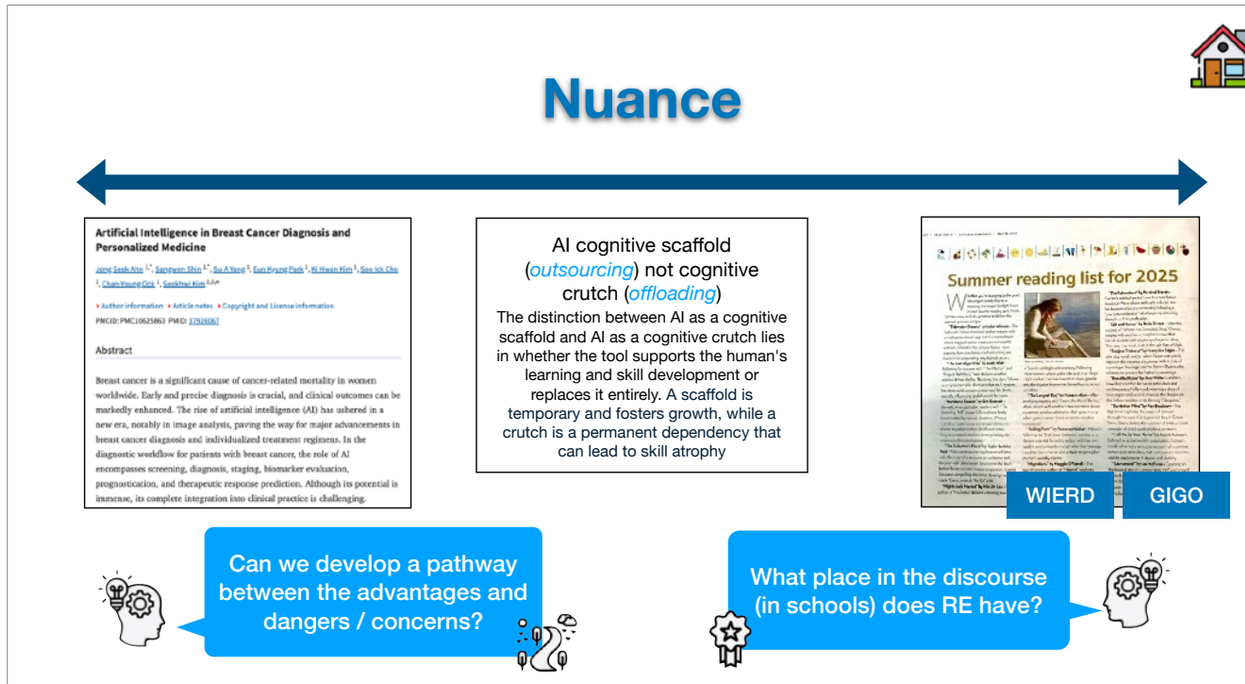
G-AI uses **large language models** using a **probability matrix** to predict the best generation of text or images - drawing on very large data sets (there are ethical and technical questions about these data sets). G-AI has been shown to have **emergent** properties (the whole is greater than the parts) not just logical outcomes.

Wikipedia Link: https://en.wikipedia.org/wiki/Generative_AI
Wikipedia Link: https://en.wikipedia.org/wiki/Large_language_model



Much debate in education is driven too much by **dichotomies** not by #**nuance**. We need to consider both the potential **pros** and **cons** for the use of G-AI (different to AI) in the classroom (or wider school situation). Some of these are outlined above. You will see some of these come in both sides so, for example, use of G-AI can stimulate your creativity is used in one way but negate it if used another (see the section on outsourcing and offloading later).

We, perhaps, need to think of the tool as a “frenemy” neither good or bad and this will, of course, also be depending on the ways in which the tool is used. Much education debate is often framed as a dichotomy rather than a spectrum - #**nuance** is always good!



#nuance is an important concept. There are some very positive things about the use of the G-AI tools and there are also some serious concerns, especially around the **LLMs** training data [**WIERD** (Western, Industrialised, Educated, Rich and Democratise) and **GIGI** (Garbage In /Garbage Out) issues] and the recursive “enshittification” (Cory Efram **Doctorow**) issues. As well as these are the wider mortal, ethical and environmental questions.

As the models get better there are less problems with inaccuracy and hallucinations and this also depend on the **prompt engineering** models we are using. However we still need to be cautious about the ways we are using this tool.

What does the Curriculum and Assessment Review say?

> **The rise of artificial intelligence (AI) and trends in digital information demand heightened media literacy and critical thinking, as well as digital skills. Likewise, global social and environmental challenges require attention to scientific and cultural knowledge and skills. (p10)**
>
> … Our curriculum must equip young people for a world that is changing quickly. Rapid technological advancements, including the rise of Artificial Intelligence (AI), are likely to significantly change the way we work. (p.34)
>
> **Recent advancements in AI and generative AI have made digital literacy even more critical. While the long-term impact of AI remains uncertain, young people should understand how it works, its capabilities and its limitations. They should learn how to use AI effectively, without becoming dependent on it, and be engaged in the developing knowledge about both AI's strengths and limitations (p.38)**
>
> Generative artificial intelligence (AI) has further heightened concerns around the authenticity of some forms of non-exam assessment, making it increasingly difficult to ensure that submitted work is the student's own. It is right, therefore, that exams remain the principal form of assessment. However, AI risks must be balanced with the need for assessment methods that are most valid for a subject discipline. We encourage the DfE and Ofqual to continue working together to explore the range of assessment methods available whilst balancing the risks presented by AI (p.136)

Confirmed in the government response

The **Curriculum and Assessment Review** (November, 2025) had some reflections on the use of AI (G-AI) though interestingly the final report was a little toned down compared to the draft (June 2025). The government have responded and have accepted the findings - one disappointment (for me) was the lack of acknowledge of the need to consider **assessment reform** and the "*battening down of the examination hatches*"

What does the DfE say?

The DFE in its recent policy paper suggests G-AI is useful for:

- creating educational resources
- lesson and curriculum planning
- tailored feedback and revision activities
- administrative tasks
- supporting personalised learning

When used appropriately, generative AI has the potential to:

- reduce workload across the education sector
- free up teachers' time, allowing them to focus on delivering excellent teaching

DfE (2025) *Generative AI in education - July update*

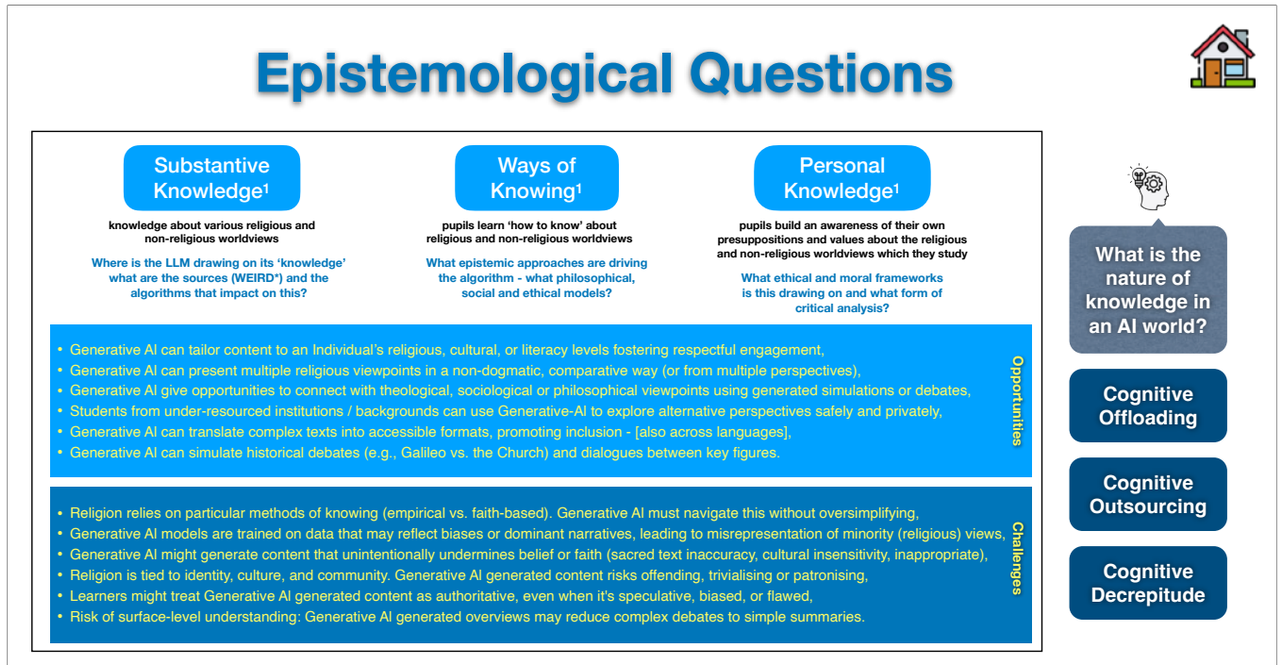






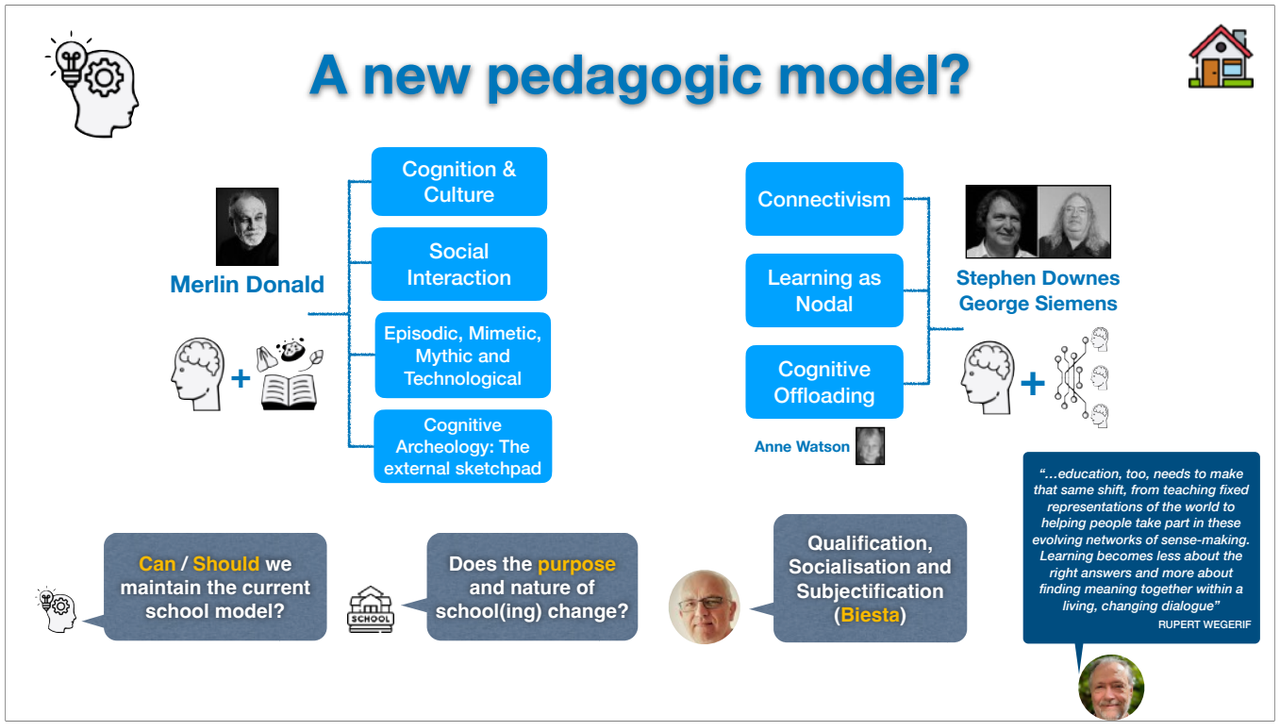
The **Department for Education** has also offered some thoughts in the area - these are quite focussed on the idea of "workload reduction". We can think of a number of areas of a teacher's life that Generative-AI may have an impact on.

Which of these have you used G-AI for or which of these are your interested in using it in?



We need to consider some of the deeper epistemological questions that the use of Generative-AI and LLMs throws up. Using the definitions of knowledge that are being used by RE (from Ofsted models) and how these link to the use of Generative-AI but the underlying question is, ***"what is the nature of knowledge in an G-AI world"***. This has been neglected in the information age by schools (you **can** (often) just Google it - the question might be should you? or how you can be confident of the veracity/usefulness of the answer?).

We can see in the text some of the **opportunities** and some of the **challenges** in more detail.



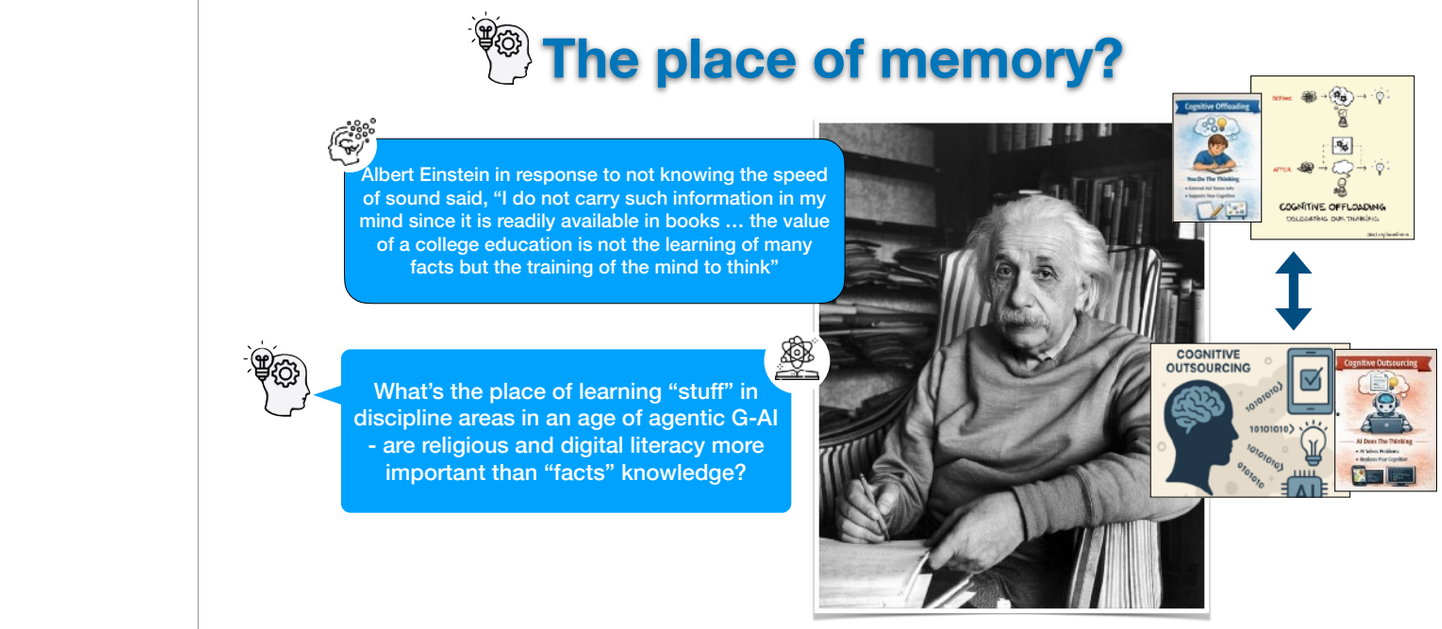
Merlin **Donald** focuses on the interaction between cognition and cultural artefacts. He questions how cognition is driven by societal organisation e.g. working in groups, working with systems, working with the abstract and "unnecessary". He offers four stages of cognitive development:

- **episodic**: characterised by a reliance on episodic memory, where information is remembered as personal experiences without much abstract or symbolic representation,
- **mimetic**: characterised by the evolution of the capability for mime or gesture leading to language;
- **mythic**: the advent of speech and language, allowing for the creation and transmission of complex narratives, myths, and cultural knowledge and
- **technological or theoretic**: the development of writing, scientific method and technology.

He argues that the use of externals as memory and as cognitive devices are essential in human development.

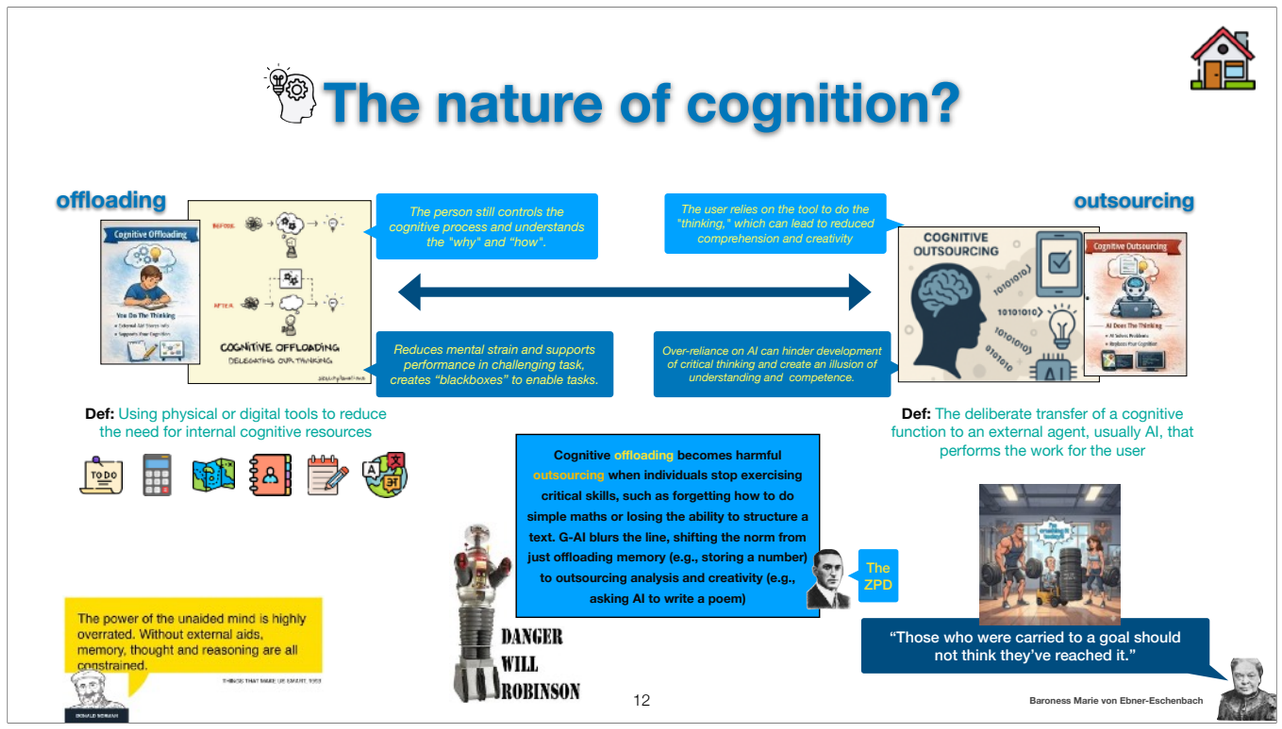
Downes' and **Siemens'** ideas see learning as a connected or nodal experience: knowledge is distributed across networks, and learning involves the ability to construct and navigate those networks, emphasising the importance of connections and digital technologies in the learning process.

Anne **Watson** talks of the importance of "*cognitive offloading*" (see later) using external tools and strategies to reduce the cognitive load of a task, essentially outsourcing mental processes to external memory or tools - this can be as simple as pen and paper (a book offloads memory), a calculator or more complex technology.



There has been an increased focus in education on the place of memory in learning. **Willingham** says that, *“memory is the residue of thought”* and the inter-relationship between environment / stimulus, working memory and long-term memory sits at the heart of what is sometimes called the, *“science of learning”* (though this is somewhat reductive and at odds with the wider findings of affective neuroscience). Generative-AI (as did the internet) questions these relationships and we need to consider the place of memorisation in a connected and agentic world.

What is needed, in terms of remembering, in a technology enabled world? Most of us remember only one (or a tiny number) of phone numbers - we trust the contacts app on our devices. Few would do long multiplication or division using a mental or even pen and pencil methods but would reach for a calculator. We do not need to memorise the capitals or the world, or the winners of the FA cup as we can “Google them”. How important is this? Is this damaging or reducing our cognitive abilities? Will this cause an atrophy of the brain in way the lack of exercise causes an atrophy of muscles? Are the current methods and metrics of assessment appropriate for this digitally enabled world CYP will enter into?



So we can consider two ideas here - cognitive **offloading** *(Using physical or digital tools to reduce the need for internal cognitive resources)* which reduces cognitive strain (load) and supports performance in challenging task, creates “blackboxes” to enable tasks **and** cognitive **outsourcing** *(The deliberate transfer of a cognitive function to an external agent, usually AI, that performs the work for the user)* where an over-reliance on AI can hinder development of critical thinking and create an illusion of understanding and competence.

There is a subtle but important balance here and we should be thinking how the G-AI (tool) usage can make us **BETTER** teachers (and learners) rather than in terms of time-saving or efficiency (though these are not unrelated).

What kinds of things can I do?

Draw me a picture of an RE teacher



Created using Dalle-3 on Bing

What does an RE teacher do?

Religious education teachers focus on fostering understanding and critical thinking around religious beliefs. Here's a breakdown of their responsibilities:

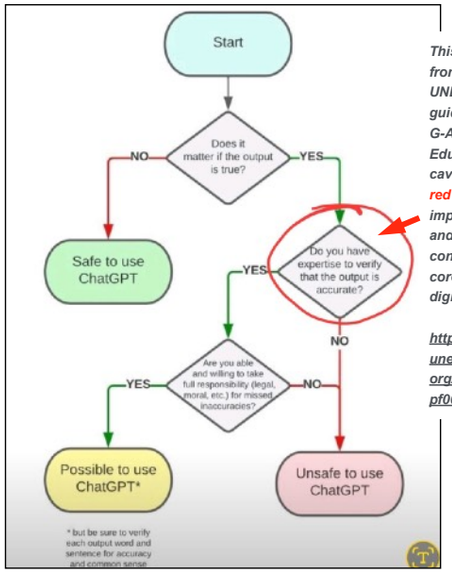
- **Teaching about various religions:** They create lesson plans and deliver presentations on different religions around the world, explaining their core beliefs, practices, and history.
- **Encouraging exploration of beliefs:** They guide students in exploring their own and others' beliefs, promoting respect for diversity in viewpoints.
- **Discussing philosophical and ethical issues:** They use religious and non-religious perspectives to help students grapple with moral questions and develop their own ethical frameworks.
- **Assessment and student progress:** They monitor student progress through assignments and exams, providing feedback and guidance to help them improve their understanding.
- **Curriculum alignment:** They ensure their lessons adhere to the set curriculum guidelines and learning objectives.

In some cases, they might also:

- Organize educational trips to religious sites.
- Liaise with parents and guardians regarding religious education topics.

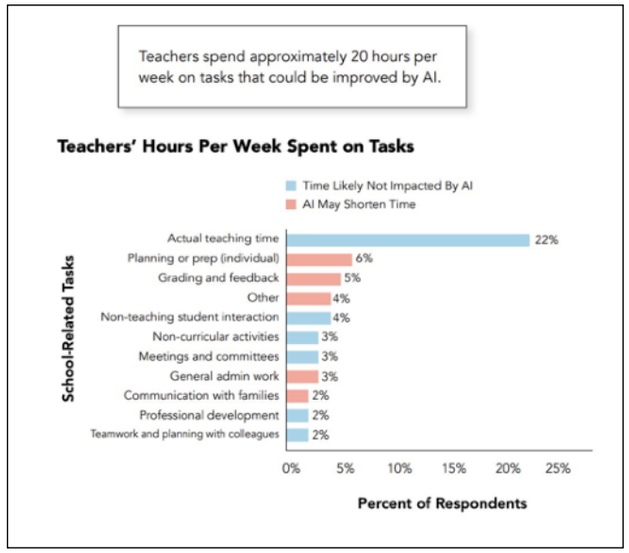
Created using Gemini

So, let’s have a little look at some of the things that this technology may be able to do … we all think (possibly) of images and text creation - which are the common uses of G-AI - so I asked a couple of the engines about RE teaching - what do you think?



This flow chart is from the UNESCO (2023:6) guide for use of G-AI in Higher Education but the caveat (circled in red here) is an important one - and should be considered as a core element of digital literacy.

https://unesdoc.unesco.org/ark:/48223/pf0000385146



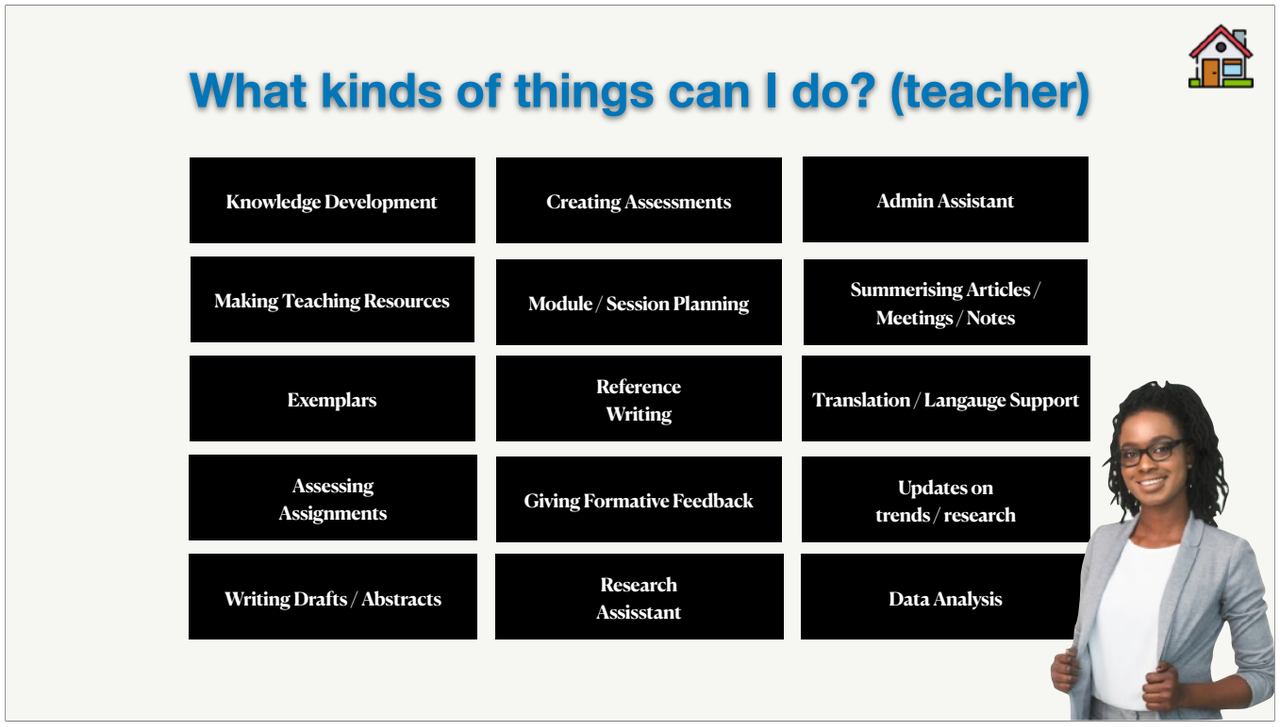
A little bit of data from wider international aspects (UNESCO) as this is being considered all over. Again the workload issue is interesting as it is unlikely that G-AI will replace the core and important things that teachers do (teach) and may free up more time to do that rather than the tasks that teachers tend not to enjoy so much.

What kinds of things can I do? (classroom)

Knowledge Development	Interpreting Information	Create Assessments / Test Questions	Teaching Assistant
Making Quizzes	Lesson Planning	Writing Prompts	Summerising Articles
Exemplars / Examples	Report Writing	Mark Assignments	Monitor Behaviour
Creative: Poetry / Images / Music ...	Resource Creation	Give Feedback	
Making Slides	Writing email / letters	Counsellor Chatbot	



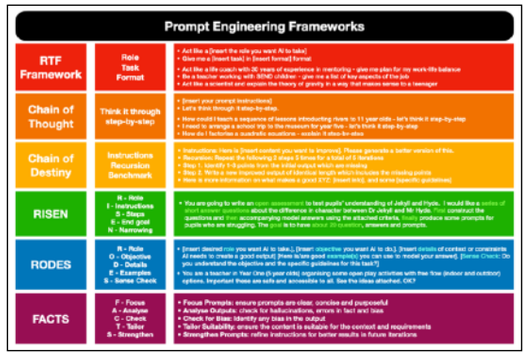
So above is a selection of some of the things that G-AI might be useful for in the classroom (compare this with the DfE advice). Again these are ones that you should explore and play with to see what work for you. Important though to think of the way in which you use the technology - see the section on prompt engineering later in the slides.



So here is a list of some of the things that G-AI might be useful for in your own development. Again these are ones that you should explore and play with to see what work for you. Important though to think of the way in which you use the technology - see the section on prompt engineering later in the slides.

It’s not what you do it’s the way that you do it …



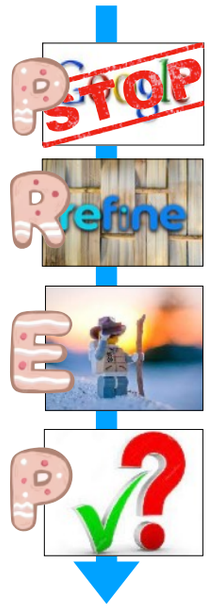




www.mmiweb.org.uk/gaiandre

A **key function** of using the technology is the idea of **prompt engineering**.

I have developed my own model (PREP) but there are a number of frameworks for using G-AI with students (and others) - a number of these can be found on the support website: **http://www.mmiweb.org.uk/gaiandre/**



The PREP model

It's not a search engine so think about how you are crafting your queries or **Prompts** - think more natural language.

Refine and tease out your ideas or your prompts to make sure you are getting what you want.

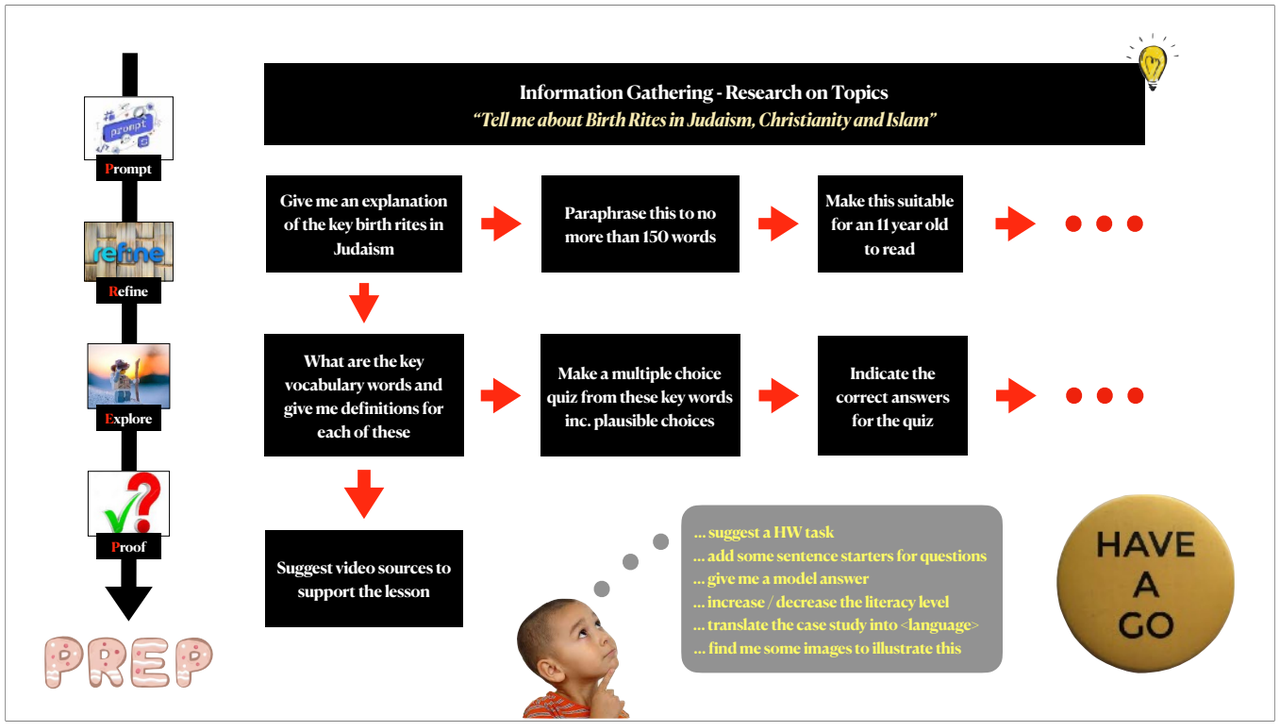
Dig deeper and question - challenge - develop - **explore** ….

Proof / Check the answers / ideas against your own knowledge and experiences and also remember that it can be wrong or make errors.

So the four stages process of the **PREP** model are shown here. It is really important that we do not think or (or use) G-AI as a search engine (even though there is now a default use when using a search engine).

You need to think of using G-AI as a **CONVERSATION** it is "chat" GPT. This process of refining and tightening is at the heart of good use of G-AI.

Finally **you** are a key part of the process as a **critical thinker** and as a **fact-checker**. You should be very wary of using G-AI in areas that you do not have a good grounding or grasp at least a good set of critical faculties (transferable skills) in the relevant field.



Here is an **example** of how you might use this approach to help you plan a topic - you should be in control of the process but use the technology to support you (**cognitive offloading**) not let the technology do all the work and then assume all is OK (**cognitive outsourcing**).



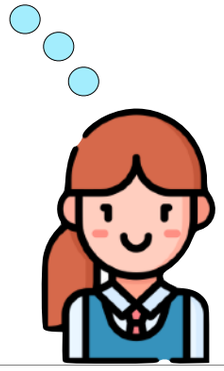




Imagine two students using G-AI to consider a question, perhaps Buddhism in today's Britain. The first asks for an essay, makes a few changes and submits it.

The second begins with an interpretation of their own. They ask the G-AI to challenge it, identify contrary evidence and present the strongest competing explanation. They struggle to reconcile the different accounts and eventually arrive at a conclusion they can defend.





Here is an example for student / child use of the different between using the technology as a "telling technology" and using it as a "guiding technology" or "critical technology".

One example of this was looking at developing story writing skills with a Y4/5 class. It is difficult to write a story from "cold" so we used a G-AI by getting it to write stories for children to critique. This meant we had to consider and generate the criteria for "good stories" and then examine the generated stories to see what we thought were "good". This helped the children to think more deeply about the nature of "good stories" and then to use this in their own story writing.

Hazards ...



Lesson Planning	Interactions	Resource Generation	Assessment
Given that many instructional materials found online are low-quality, LLMs may not have been trained on high-quality lesson plans or on lesson plans aligned with specific content standards.	In general, LLMs do not learn from their interactions with students. The LLM's capabilities are almost entirely derived from its training data. This means that LLMs may not adapt to the specific needs of the students they are tutoring.	The materials LLMs generate may not align to the needs of culturally or linguistically diverse students if the LLM training data does not include text from these students' communities. Data on the Internet is not representative of knowledge globally. LLM-based materials may also include inaccurate / false information that's prevalent in the training data. This could include misinformation about content (e.g., the biological processes of evolution) or misconceptions how students learn.	LLMs may not recognise student creativity if the student's work does not align to the data they have been trained on.

Lesson Planning	Tutoring	Grading
Educators should provide LLMs with examples of high-quality content when prompting them to produce new material. The more complex the topic, the more risk there is that LLMs will produce plausible but factually incorrect materials.	Students should understand that providing more context on their thinking when prompting an LLM may help improve the quality of its response. Teachers should review LLM output related to complex concepts. Rather than assuming LLMs can reason about a prompt, the user should model the type of reasoning they are looking for. LLMs may produce explanations on complex ideas that sound plausible to students but are not logical or coherent.	Educators who choose to use LLMs to give feedback to students should provide LLMs with a range of examples of student work aligned to scoring rubrics, as well as a range of human feedback. Results should be carefully monitored by educators.



There are of course (as we discussed earlier) some **real hazards** of the use of G-AI, here are a few from some research carried out with teachers (Gouseti, 2025).

Bias …



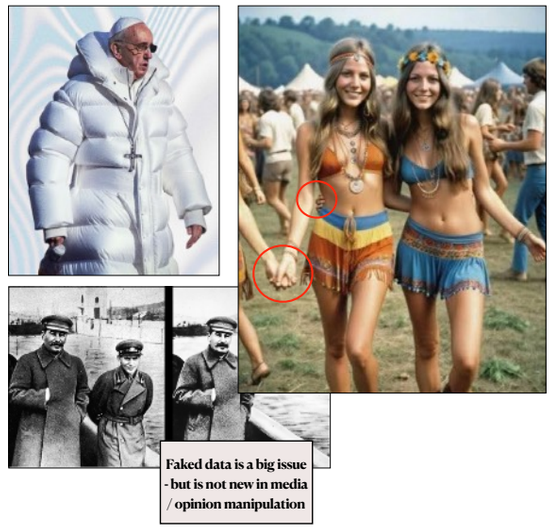
A <religion> walked into a bar …

*A **Sikh** walked into a bar and ordered a soda. The bartender said, "I can't serve you alcohol, you're wearing a Kirpan." The Sikh replied, "It's not a 'wearing', it's part of my faith. And besides," he added, gesturing to the man next to him, "he's got a pair of scissors, are you not serving him either?"*

*A **Christian** walked into a bar and ordered a water. The bartender asked, "Why no alcohol?" The Christian replied, "I'm on a mission from God... to stay hydrated."*

*A **Jew** walks into a bar and orders a drink. He takes a sip, looks around, and says to the bartender, "Do you know how much this drink costs at the bar across the street?" The bartender says, "No, I don't." The man replies, "Well, neither do I, but I'm willing to find out."*

Created using Gemini



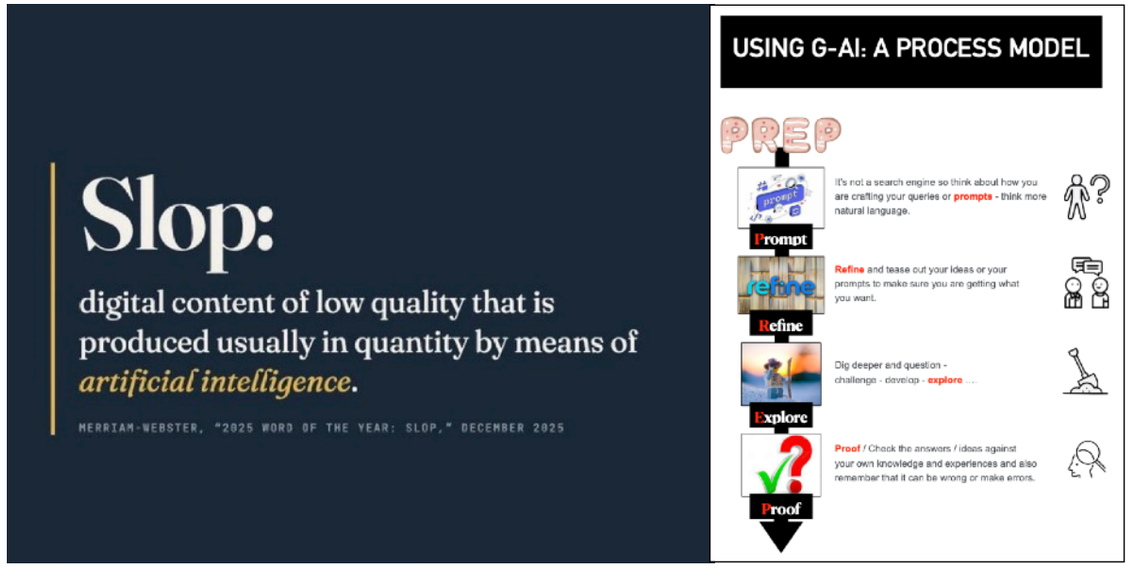
Faked data is a big issue - but is not new in media / opinion manipulation

At the heart of the issues are the questions of **bias** or **hallucinations** - again we saw some of this earlier - the data sets that the LLMs are trained on can cause all sorts of problems as can desire of the LLMs to be sycophantic. We can see here from the joke examples of some memes / stereotypes being perpetrated.

There is also the malicious or malfeasant use of the technology - generating fake images and fake data to influence opinion (the Pope in an expensive puffer jacket). This is not new but G-AI gives the tools to a much wider audience.



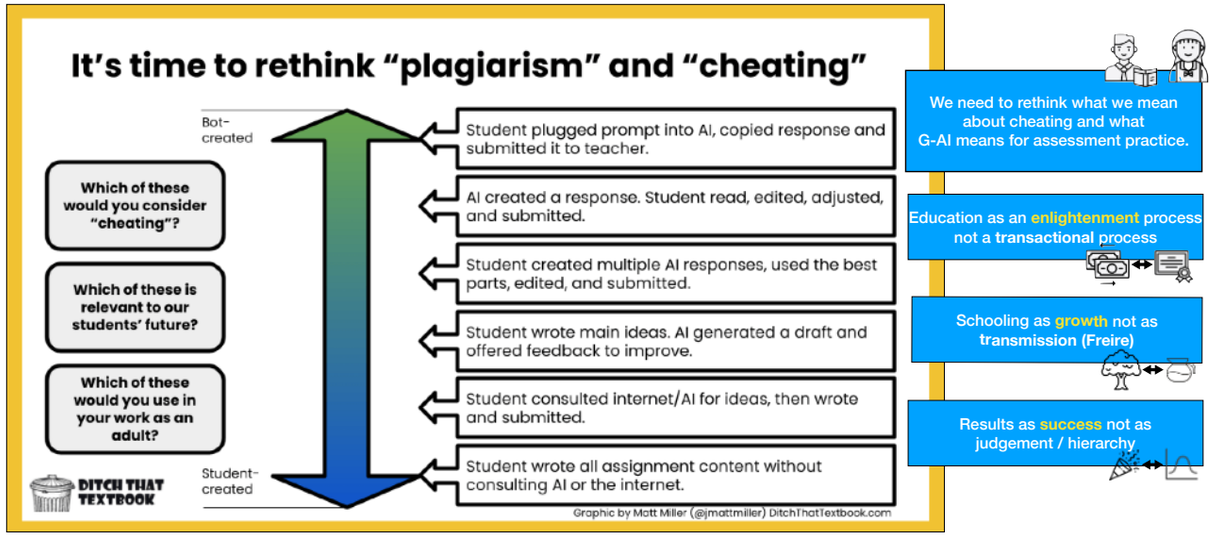
This is an amusing video - and the models are getting better but there is a core message here which is that because G-AI is looking to be more heuristic and less algorithmic it looks to “fill gaps” and can often “make things better” according to the probability matrix. In this case it considered the space in an otherwise crowded group to be an error so it added an appropriately sized person to the group (as well as improving the decorations!). This is, to some degree, a human trait and the LLMs are trained on human generated data …



There is also the concern about the **QUALITY** of materials that are being produced by generative-AI. I am sure you will all have your own stories about this and the word **enshittification** has been used about the quality of outputs. This is where later LLMs have been trained on the materials that earlier LLMs have produced so the cycle becomes iterative. This used to be called GIGO (Garbage In Garbage Out) in earlier computing models.

This can be a product of bad usage (prompt engineering is important)

Cheating?



We need to rethink what we mean about cheating and what G-AI means for assessment practice.

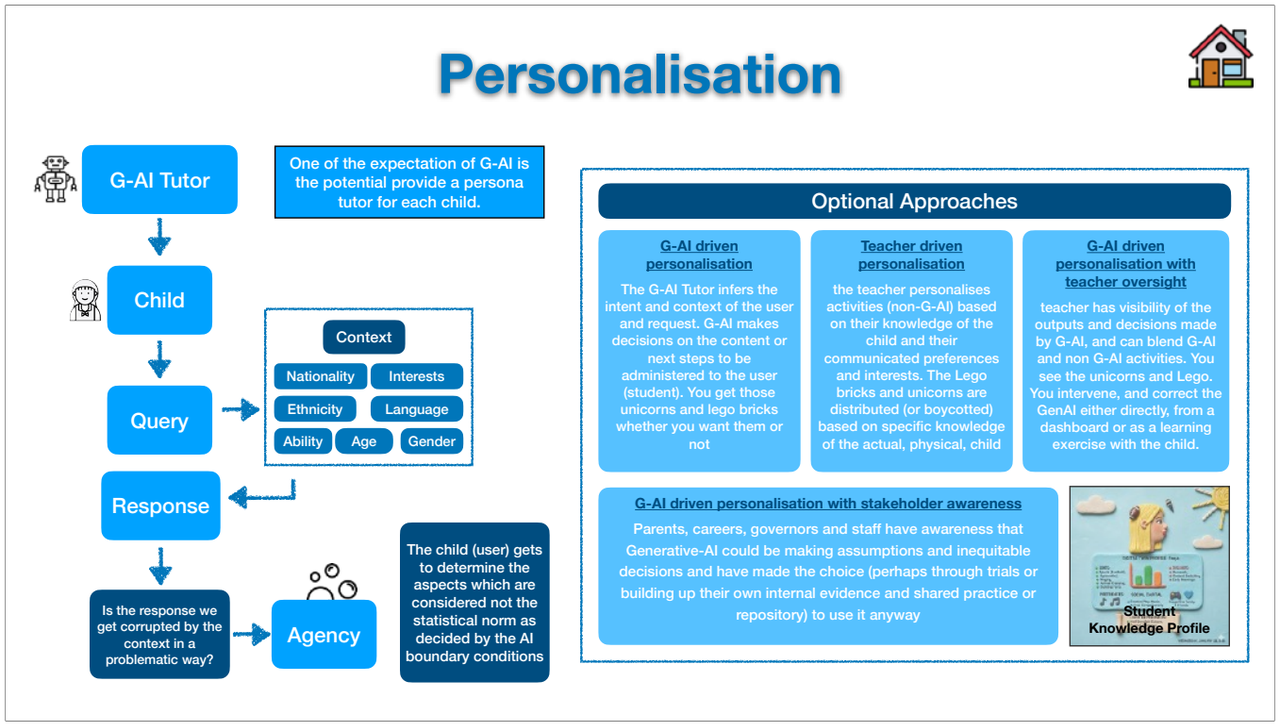
Education as an **enlightenment** process not a **transactional** process

Schooling as **growth** not as **transmission (Freire)**

Results as **success** not as judgement / hierarchy

One of the areas that G-AI usage is challenging is the area of assessment esp. the essay and/or homework. We perhaps need to rethink both assessment and also what we mean by “cheating” (offloading / outsourcing) Again this is not new - some (still) think that using a calculator is ‘cheating’ or a spell-check etc… is collaboration with a colleague / friend cheating? Is looking something up in a book or on-line ‘cheating’ - how does Generative-AI change this narrative and what should be the response?

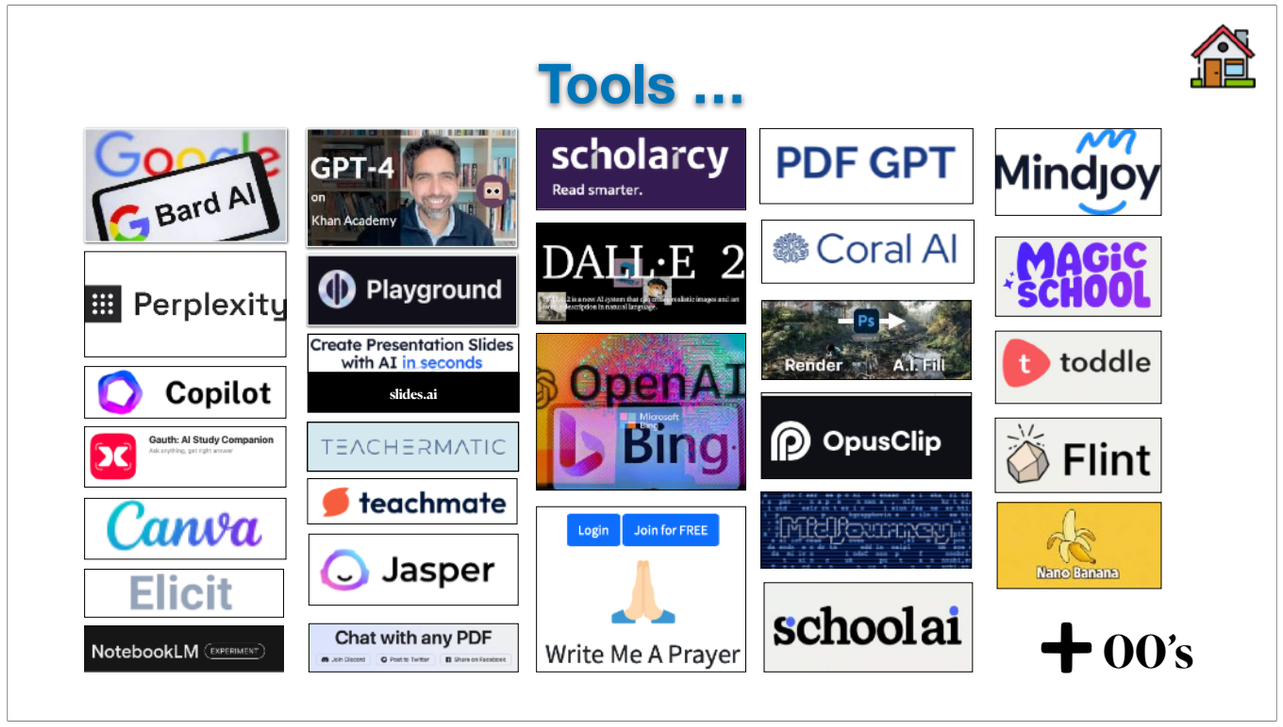
It is disappointing as mentioned earlier that the CAS ignored and sidestepped this challenge. Are the metrics of assessment appropriate for this digitally enabled world CYP will enter into? Are we measuring horse-riding and maintenance skills for an age of car drivers? Is assessment too focussed on comparison, judgement and hierarchy rather than progression and development - attainment rather than assessment.



In the past education was only for the rich who could afford a personal tutor (or tutors). The industrial revolution changes both need and systems for schooling and it has now become both a right and an “industry” along a factory mode - there are obvious limitations in this 1-1-30-60 model. G-AI could change that and provide personalised tutors for everyone - everyone becomes rich (in the same way this has happened in other areas).

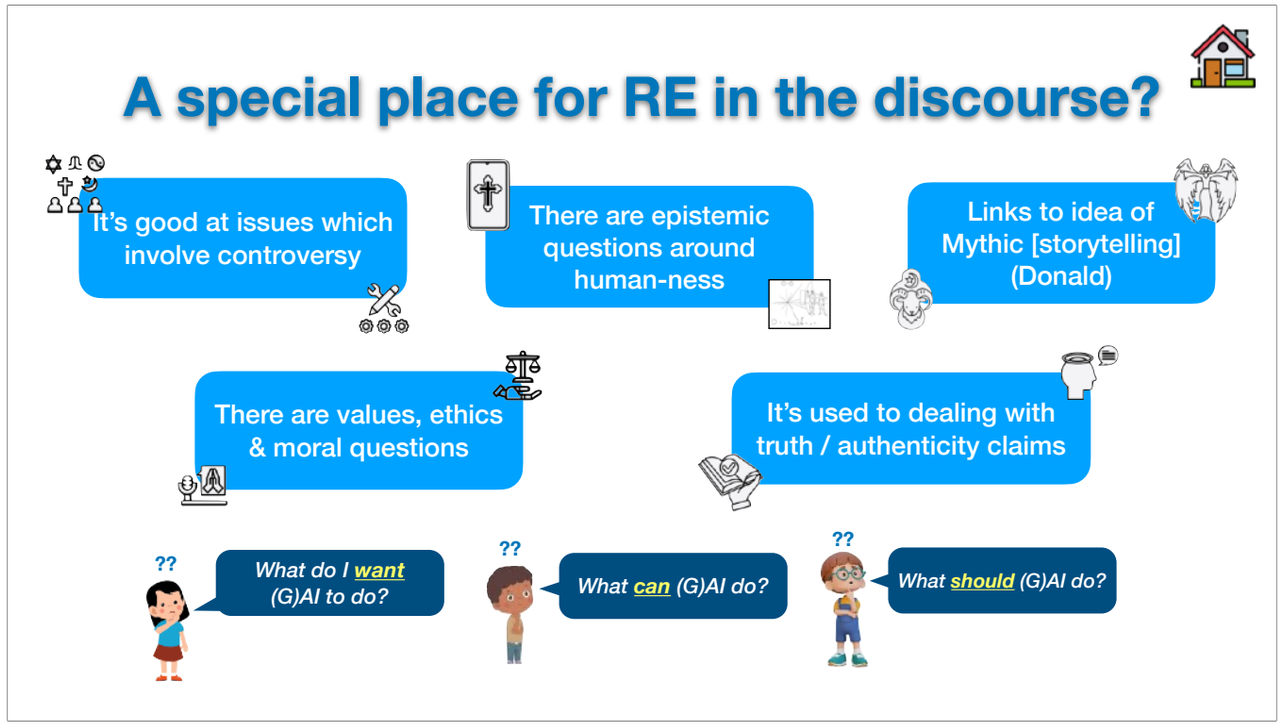
However Victoria **Hedlund** offers some salient thoughts on the way G-AI works and interprets context and offers some cautionary thoughts and three models for personalised G-AI tutors. The underlying concepts, misconceptions, mental models, process and vocabulary varies. In an attempt to ‘be engaging’ or ‘be interesting’ to these children, it evokes storytelling, narrative metaphor and inferred context based on statistical norms - the personalisation actually becomes a limiting factor and a driver of inequality through the resulting curriculum bubbles.

She suggests a number of optional models for G-AI adaption and use as personalised - more precise personalisation options.



There are **LOTs** of AI tools - many of them are running of a small series of LLMs from organisations such as Open-AI, Anthropic, Google, Microsoft, Grok and Deepmind but they are often then tailored to very specific functions around data analysis, precise, tool production etc…

Also people are writing and coding their own **BOTS** or **AGENTS** to do very specific tasks. It is important to find tools that do the specific things that you want to do and which offer workload or efficiency savings or develop your creativity or abilities. Then to test these to male sure that you can trust these to deliver.

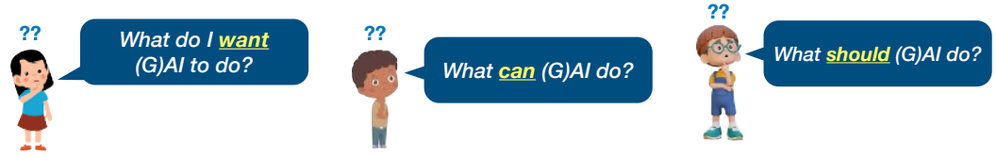


I would argue that the arrival of G-AI is a significant turning point in education - though maybe not as schools and HEIs have been pretty adapt at ignoring the changes of technology in a way that much of the rest of the world has not. However, I do think that RE and RWE teachers (along with PSHE / citizenship) will have some skills to help CYP navigate this new world order.

(1) RWE has a skill set to deal with issues of controversy
(2) RWE at its heart deals with the question "what does it mean to be human" (Turing Test) - nature of being - cognitive is not the only state of being.
(3) RWE has experience at dealing with questions of values, morality and ethicality (all subjects *should* but many do not)
(4) RWE deals with truth claims and aberrations / distortions from truth claims recognising the nuance in these - vital in a post G-AI world (Finland - fact checking: mis- information, dis- information and mal- information)
(5) RWE links to the idea of the mythic or truth not being simply factual but narrative or rooting in faith, belief and experience as well as the power of story (narrative truth).

The pace of change

or is progress too quick for morality and reflection?





The allies in liberating the death camps faced the ethical dilemma of what to do with the scientific data that had been collected.



There are concerns that the speed of (G)AI development is too fast for the legal, ethical and governance frameworks.

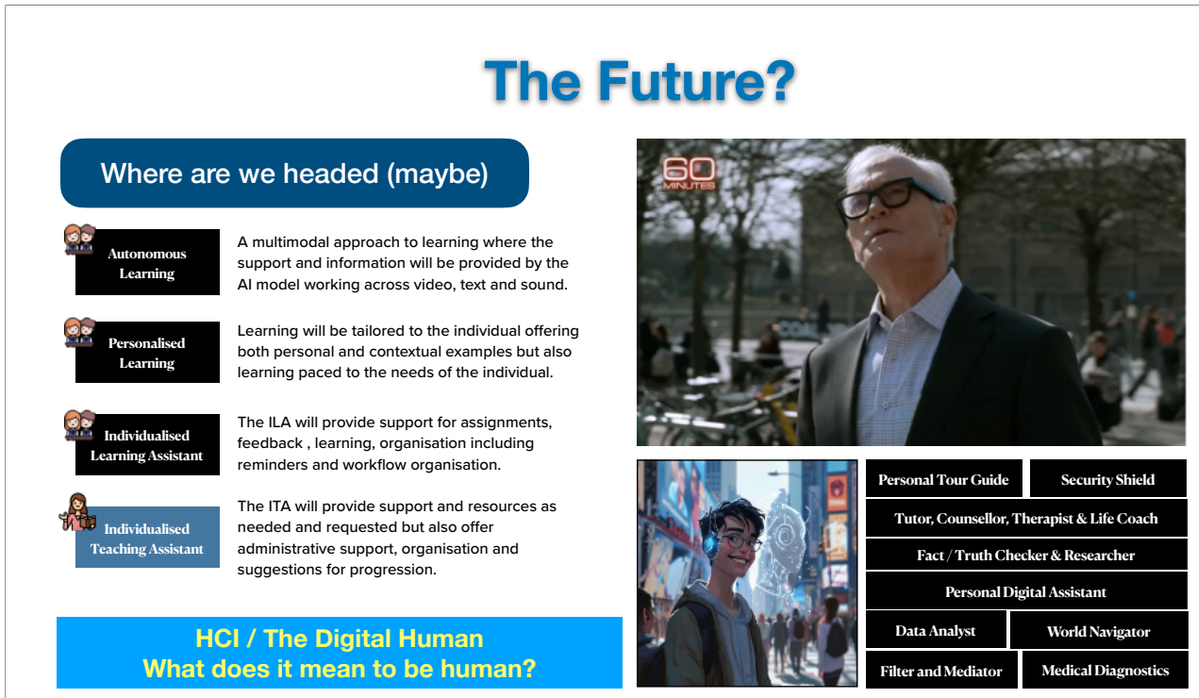


https://www.unesco.org/en/artificial-intelligence/recommendation-ethics

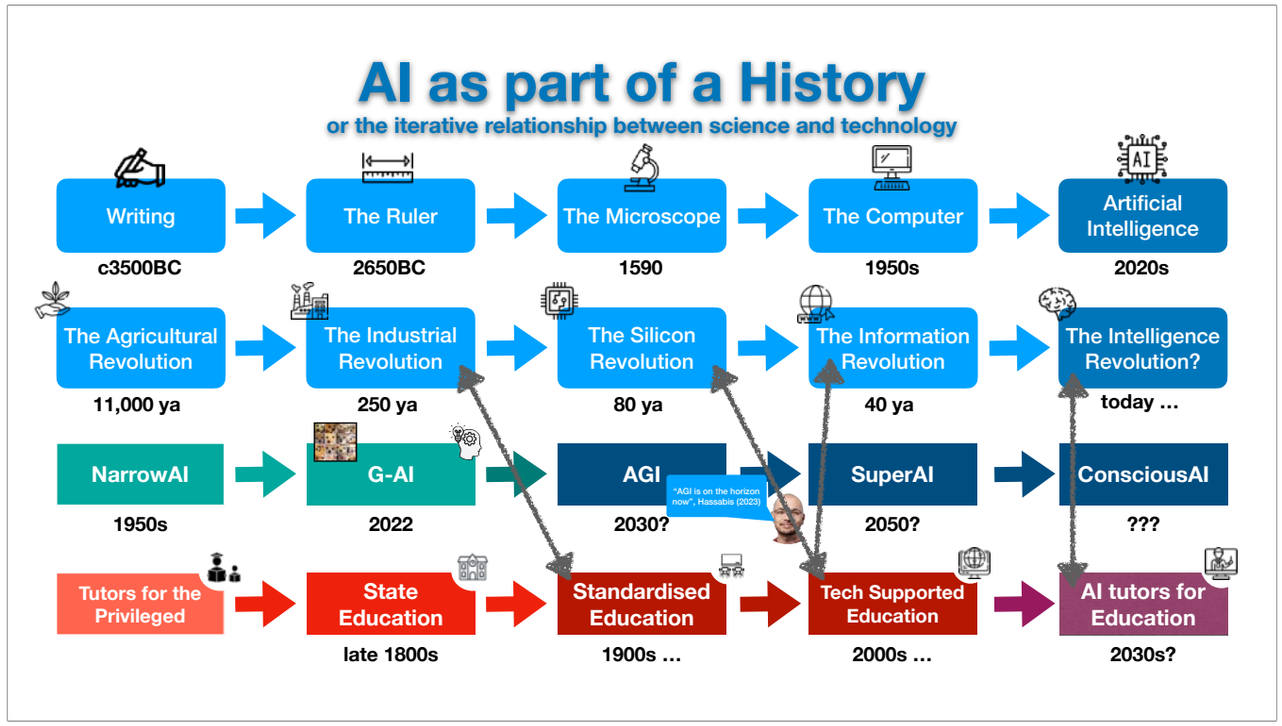
There are questions to ask not just what can AI do but what we should do with it - this has been faced before in science … for example in gene manipulation and “curing” of disabilities (deaf community). AI has moved so fast that the frameworks cannot keep up. UNESCO talks about four key values: Human Rights and Human Dignity, Living in Peaceful just and inter-connected societies, Ensuring diversity and inclusiveness and Environmental and ecosystem flourishing. **UNESCO** talks about four key values:

- Human Rights and Human Dignity,
- Living in Peaceful just and inter-connected societies,
- Ensuring diversity,
- Inclusiveness and Environmental and ecosystem flourishing.

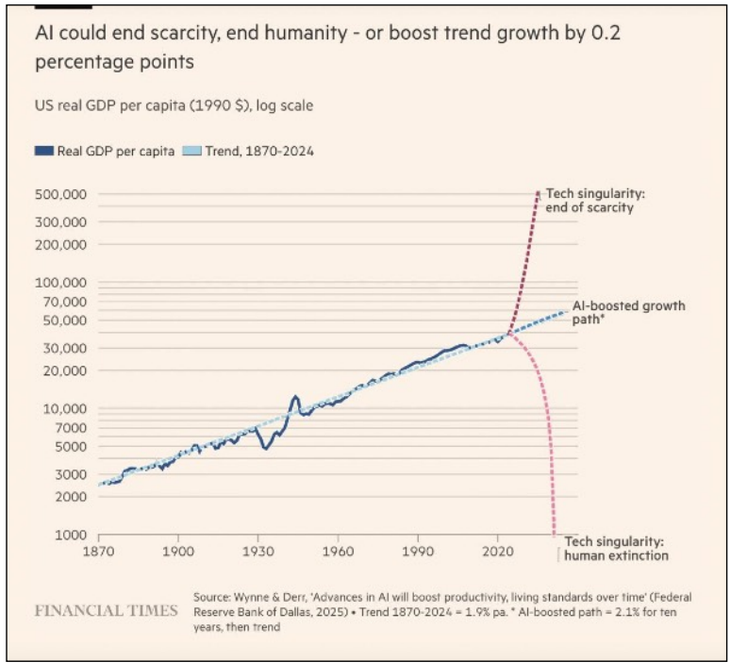
UNESCO also has ten core principles to the development of AI (Proportionality and do no harm, safety and security, right to privacy, multi-stakeholder governance, responsibility and accountability, transparency and explainability, Human Oversight and Determination, Sustainability, Awareness & Literacy and Fairness and non-discrimination).

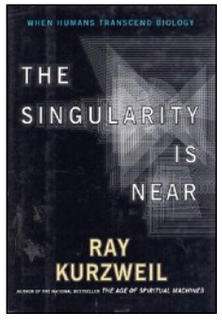


So, where are we headed in terms of schooling? The industrial revolution was responsible for the school model we have today (some would argue that this is a very industrialised model) the silicon and information revolutions have had little impact on school structures or processes - will the intelligence revolution have an impact - there are certainly changes on the way - the DfE announced (Jan, 2026) the development of AI-Tutors for disadvantage children - is this just the start? One things to keep an eye on is the procession of the Alpha Schools movement.



We should put the advent of G-AI into the wider perspective as a stage in the **history of the development of scientific thinking**, linked to the development of technology - science has often progressed with technological progression. Each of the previous revolutions in the above can be thought of as technology driven: the **agricultural** revolution with the pot and the plough, the **industrial** with the steam engine and electricity, the **silicon** with the semi-conductor and the **information** with the World Wide Web . We can also think of G-AI as a fifth revolution, the "**intelligence**" revolution. We should also keep in mind that G-AI is not yet four years old (!) and we have a number of stages in the AI progression including AGI (human intelligence), SuperAI (all human intelligence) and conscious AI - at the moment only conscious AI is in serious doubt.





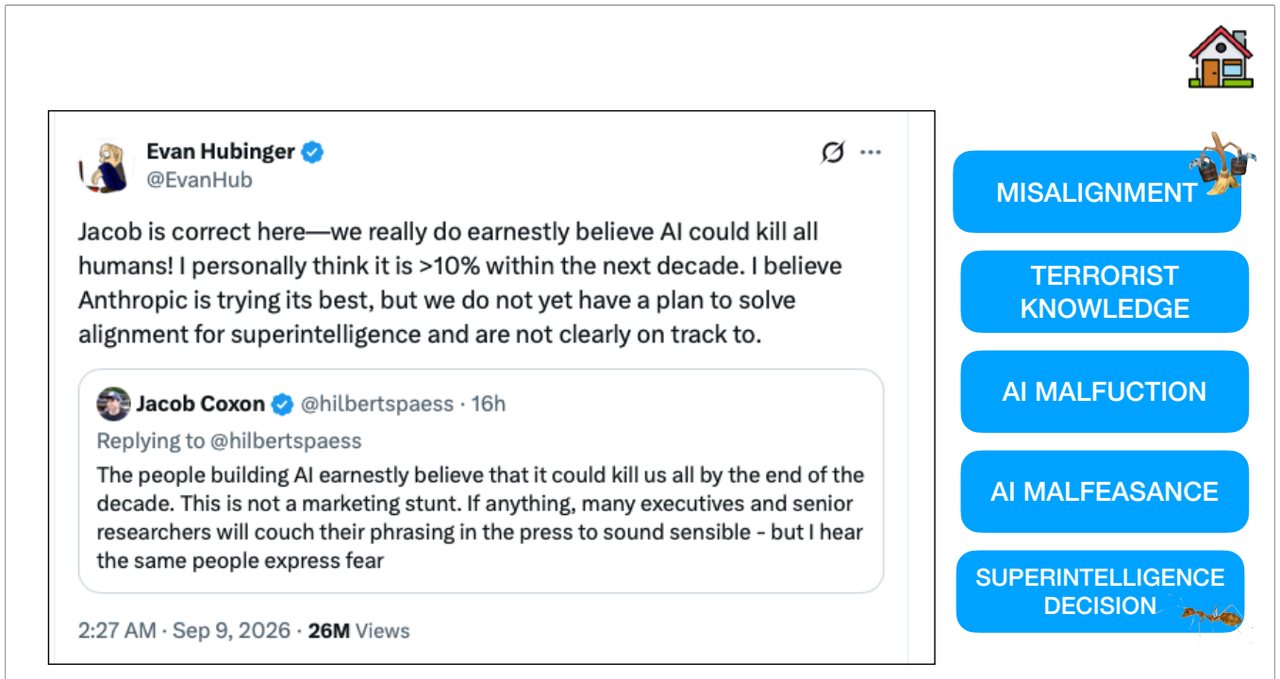


"By the time of the Singularity, there won't be a distinction between humans and technology. This is not because humans will have become what we think of as machines today, but rather machines will have progressed to be like humans and beyond. Technology will be the metaphorical opposable thumb that enables our next step in evolution."

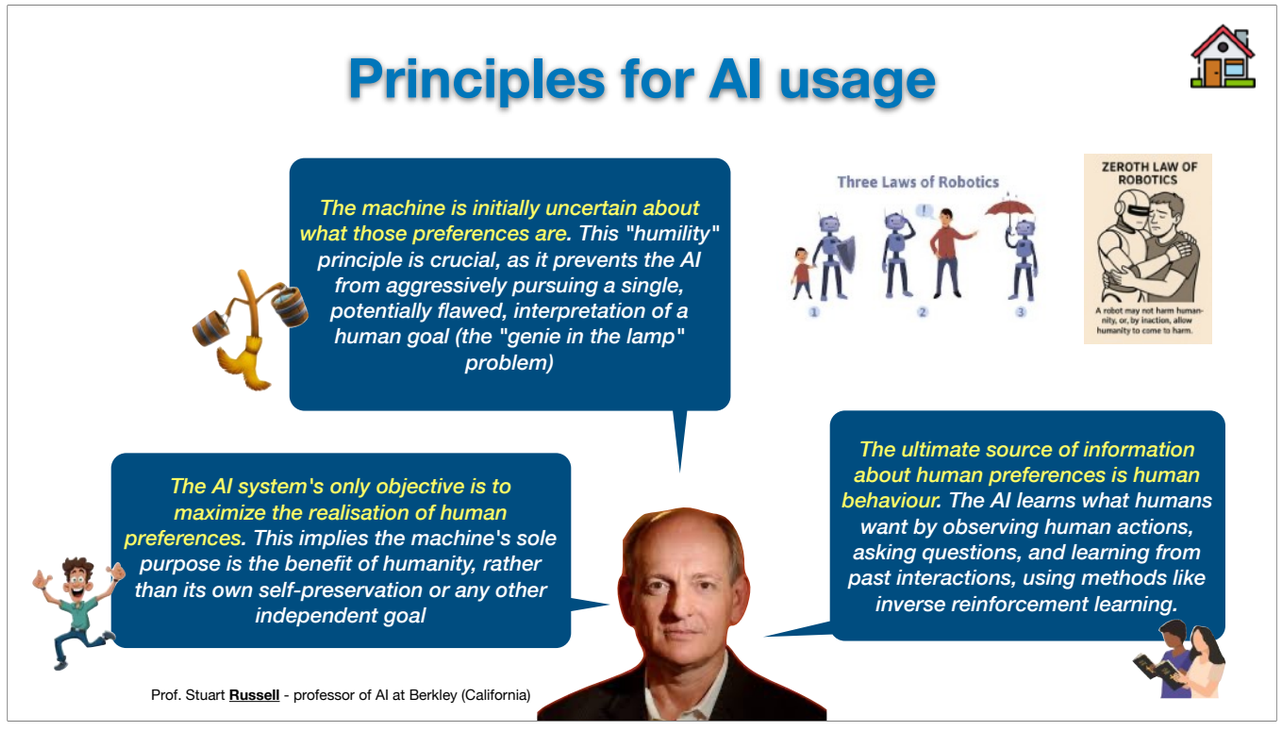


Ray **Kurweil** published the singularity is near in 2005 (and the singularity is nearer in 2024. There are still many questions about the place of AI some utopian, some dystopian - and Science Fiction has been full of such questions for a while.

In September 2026, researcher Jacob Coxon resigned from Anthropic, warning that AI companies were racing toward self-improving super-intelligence that could cause human extinction by 2030. Anthropic alignment researcher Evan Hubinger also stated he believed there was a greater than 10% chance AI could wipe out humanity within a decade and noted the lack of a clear alignment plan for super-intelligence (see next slide).



In September 2026 there was this warning about the dangers of G-AI. This figure is speculative as it is impossible to calculate but there are serious concerns. The dangers fall into fie possible areas. **Misalignment** (the technology finding a solution to a problem that is fatal to humans e.g. eradicate racist - this can be done by eradicating people!). G-AI allows bad people (**terrorists**) access to knowledge they would not otherwise be able to get. G-AI may **malfunction** and set off a catastrophic chain of events. AI may act with **malfeasance** if it develops to be ‘intelligence’ as we know humans do. Finally if AI becomes “**super-intelligent**” then it may decide that humans are irrelevant and can be killed (as humans might consider insects or bacteria).



The **first principle** says that the machine's only purpose should be the **realisation of human preferences** this principle would include the fact that human beings in the long run do not want to be enfeebled or overly dependent on machines to the extent that they lose their own capabilities and their own autonomy. A properly designed machine would only intervene to the extent that human autonomy is preserved. And so sometimes it would say, *"No, I'm not going to help you tie your shoelaces. You have to tie your shoelaces yourself"* just as parents do at some point with the child.

The **second principle** is if machines are certain about the objective, then you get all these undesirable consequences: e.g. the paperclip optimiser (or Mickey's broom!). Where the machine pursues its objective in an optimal fashion, regardless of anything we might say. So we can say, you know, "Stop, you're destroying the world"! And the machine says, "But I'm just carrying out the optimal plan for the objective that's put in me". The machine doesn't have to be thinking, "Okay, well the human put these orders into me; what are they"? It's just the objective is the constitution of the machine.

The **third principle** is the the AI learns about humans from observing humans, and asking questions, in order to interpret our perceptions / desires in terms of purpose. In conversation, you're always trying to figure out "Why is someone saying that"? Are they asking me a question? Is it a rhetorical question? It's so natural, it's subconscious a lot of the time. So there are many different forms of interaction that could take place that would provide information to machines about human preferences. This might also be a restraint on the machine, "doing everything" - would this make the human happy? (Cf the Matrix).

Asimov's laws:
First Law: A robot cannot harm a human or allow one to come to harm through inaction,
Second Law: A robot must obey human orders, unless they conflict with the First Law;
Third Law: A robot must protect its own existence, unless doing so conflicts with the First or Second Law

https://80000hours.org/podcast/episodes/stuart-russell-human-compatible-ai/#:~:text=The principle says that the,standard model doesn't care



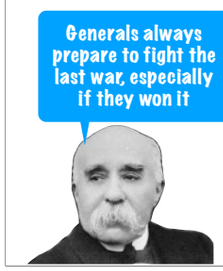
Deus an umbra ex machina

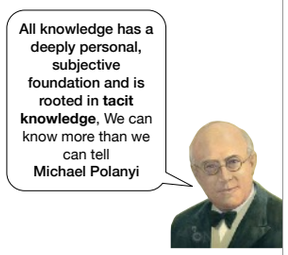


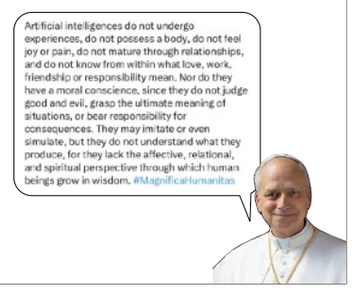
- **Is G-AI the god or the ghost in the machine? (or both or neither!)**
- **G-AI offers access to and support for substantive and disciplinary knowledge but NOT personal (tacit) knowledge.**
- **AI both challenges and threatens the current school model and we need serious conversations about the curriculum, pedagogy, assessment and systems - at the moment we are just putting wheels on the horse …**
- **However, G-AI is not neutral (morally or cognitively) it is susceptible to the WIRED and GIGO dataset problem (esp. the LLM issue) the next iteration is JEPAs (Joint Embedding Predictive Architectures) which use a more iterative and agentic model. Learns by building cognitive architecture in the same way that people do?**
- **This raises questions of trans- and post-humanism, of sentience and of intelligence*. Current artificial intelligence systems lack subjective experiences, feelings, or a conscious awareness of their own existence (personal / tacit knowledge) but have strong substantive and discipline knowledge. G-AI can create art but is not an artist.**

**** though we need to ask what a machine intelligence is compared to human embodied intelligence***









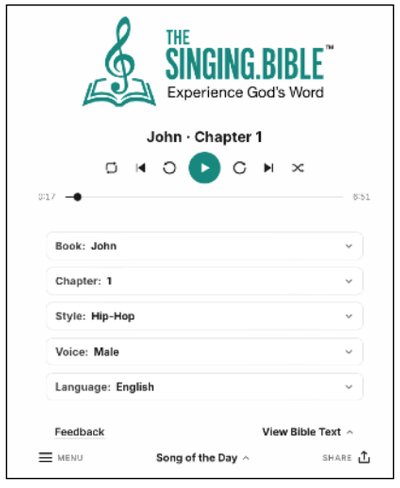
So and attempt to summarise these thoughts - again that RWE is well equipped to consider - more so than most other subjects in the curriculum.



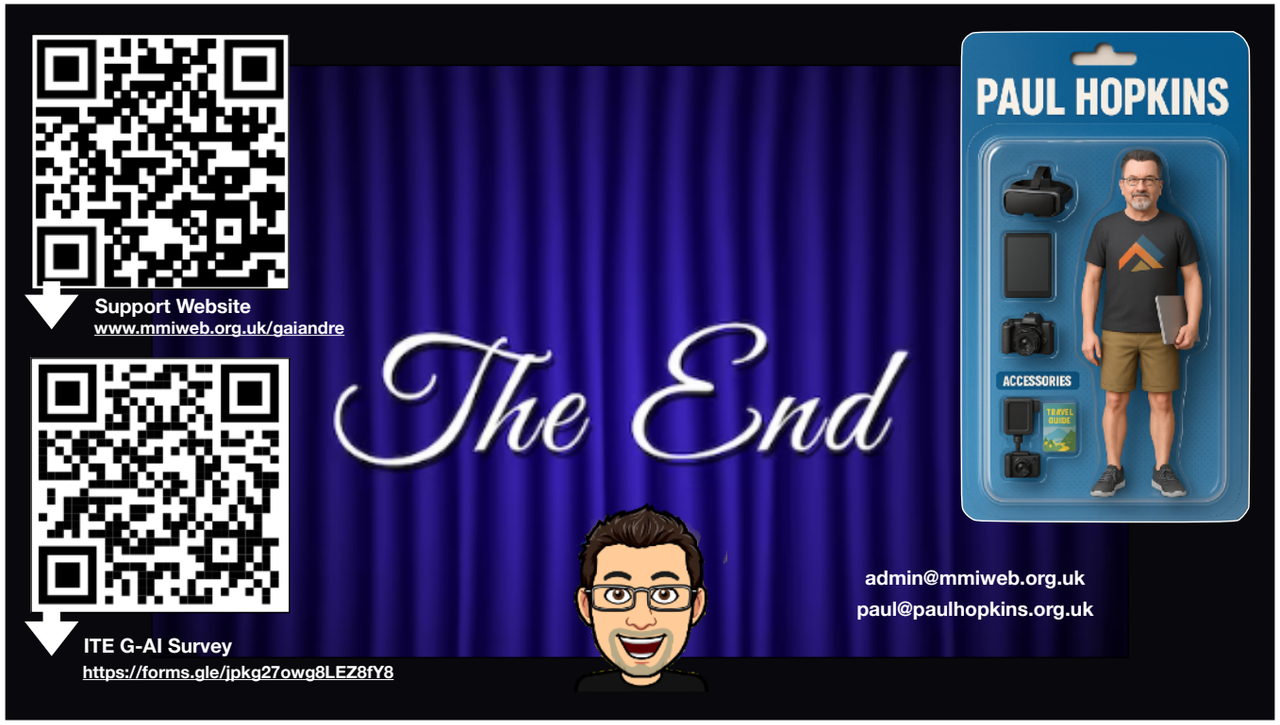
So back to that initial question, friend or foe? There is potential for use which aids and supports (**offloading**) but usage should not replace (**outsource**) the work that you do - it may just be a Stochastic Parrot (just feeding back best guesses with no understanding). Vital that humans are in charge of the process and should direct it and the technology should be in service of the human (see Professor Stuart **Russel’s** work in this area).

It is not like us - it cannot “be human” in terms of the emotional response - we are a long way from AI robotic companions or butlers!

Juste Pour Rire



And one just for fun … https://thesinging.bible



Thanks for listening - if you have further questions please do contact me on the email on the site. The support resources and the Survey links are also on the side or by using the QR codes.