appg christians in parliament

Artificial Intelligence and the future of humanity

Implications for Policymakers

Introduction

AI is now in use everywhere, in the home, at work and in public services, sometimes working in the background without our knowing. Recent developments in Generative AI like ChatGPT have captured the imagination of citizens, government and business, but also resulted in warnings of an existential threat to humanity posed by Frontier AI. Numerous concerns have been raised by civil society and experts about some of the current uses of AI, including ChatGPT. The UK hosted a global AI Summit to address safety concerns of the most advanced AI systems, launching an AI Safety Institute. A pro innovation approach to regulation is in danger of focussing on a too narrow set of risks, and failing to address the reality of the broader ways in which AI is already harming our humanity.

Christians are in a unique position to understand the true risks that AI poses to humanity because the Bible provides a theological foundation for understanding what it means to be a human being, whether we are a Christian or not. This paper provides a brief overview of what AI is, sets out the Biblical foundation for evaluating the risks to humanity and concludes with some policy implications.

What is AI?

Chat GPT (the letters mean Generative Pre-Trained Transformer) is an example of what has come to be known as Generative AI. These computer algorithms produce artificial content such as texts, computer programmes, images, audio or video in response to a user prompt like – 'write a short children's story about a family of rabbits'. The output is statistically generated from Large Language Models (LLMs) trained on a vast amount of data scraped from the internet, much of which is copyrighted. Frontier AI refers to these more advanced AI systems that use LLMs.

Discriminative AI is designed to classify something, such as identifying a person from a database of faces, or making a prediction. These systems also use training data, like images of people, animals or MRI scans, from which the probability that a new input matches something in the trained model is calculated. For example, the system outputs the likelihood that a new MRI scan shows a cancerous tumour or a CCTV image contains the face of a 'person of interest'.

NO REAL INTELLIGENCE

For many people, the response of Generative AI chat bots and image generators seems completely human, even intelligent. In reality they have no intelligence at all. They are stochastic parrots, that is statistical pattern matching engines that generate output based on an assemblage of what people have already written or artists created, including old masters. In that sense the I in AI is a misnomer, there is no intelligence at all.

Care is needed in giving too much credence to the claims of the companies driving Frontier AI and Artificial General Intelligence. Whilst such systems may appear impressive if and when they are developed, they will have no understanding as a human does. Many are buying into the idea that AI will solve the world's problems, like health and climate change. The reality is that the technology is a long way from being able to meet those claims.

SHAPING SOCIETY

Many believe that technology is neutral, it is how we use it that counts, but this is a fallacy. An evaluation of the development and use of technology since the dawn of time shows that it shapes society and society in turn shapes the values of those that create

new technologies and their application. Human sinfulness also has a bearing on design values and how we use technology – the use of AI to create fake news and videos is one obvious example of this. The desire for companies to grow big and for investors to want a return on investment influences both technology design choices, marketing and business models, as the issues with social media illustrate.

In order to understand the way in which AI, or indeed any technology, shapes humanity for good or bad, we first need to consider what it means to be human.

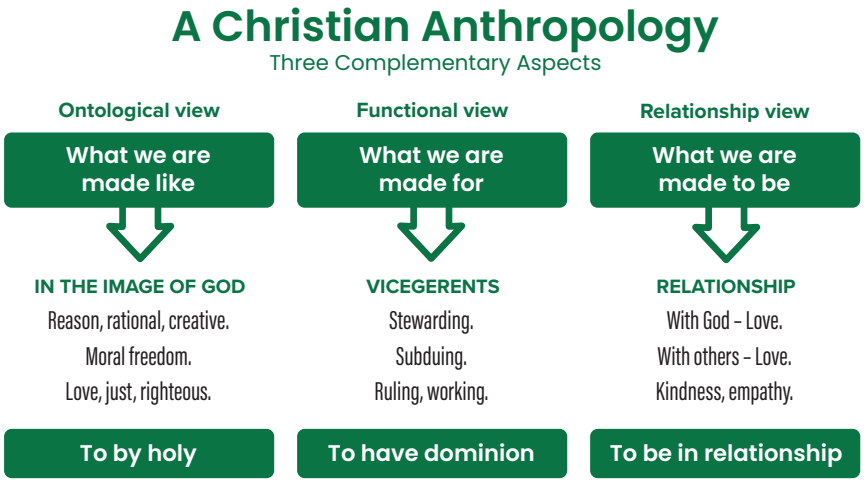


Figure 1: Three complementary aspects of what it means to be made after God's likeness or in his image.

A biblical foundation

The creation account in Genesis and a variety of New Testament texts reveal three aspects of what it means to be made in God's likeness or image. These are often referred to as ontological (the essence or nature of being), functional and relational (Figure 1). Understanding these different aspects is crucial to evaluating the impact of AI on humanity.[1]

MADE IN GOD'S LIKENESS

The opening chapter of Genesis tells us that we're made in God's likeness and therefore bear something of his nature or essence, albeit tarnished by sin. The Bible, especially the Psalms, teach us about God's moral excellence and his moral freedom, from which love, truth, reason, justice and righteousness flow. Although marred by the fall, all humanity still reflects something of His image (Gen. 9:6) and especially believers, who are being restored to the likeness of Christ (2 Cor. 3:18), who is the exact representation of God (Heb 1:3). These attributes are unique to humanity and distinguish us from animal life.

MADE TO BE GOD'S REPRESENTATIVES

God commissioned all human beings to be stewards over his creation and made us his representatives, his vicegerents. In particular, he calls his redeemed people to be salt and light, his ambassadors, given responsibility to show the world God's character. We are to image him in the world through who we are and how we act. God also made us to work in his creation and even though it's made more difficult by the fall, there's dignity in work.

MADE FOR RELATIONSHIP

The God of the Bible reveals himself to be three persons in eternal relationship. To be created in his image therefore means that human beings are unavoidably relational by nature ("Then God said, "Let us make mankind in our image, in our likeness.." Gen 1:26). The Genesis account reveals two main

1 For an in depth discussion see: Jeremy Peckham, *Masters or slaves, AI and the Future of Humanity*, IVP, 2021. See also mastersorslaves.com for an online bible based course that can be used in small group or individual study.

implications of this creation by a relational God. First, humans are created for relationship with God – God walks and talks with the first human beings and gives them instructions for how to flourish. Second, he creates human beings for relationships with each other because it is not good for them to be alone (Gen 2:18). And from the Genesis account all the way to the end of the Bible, we learn that God exists in eternal loving relationship within himself (Mt 3:17) and loves the world and people he has made (Jn 3:16).

MIRRORING GOD'S IMAGE

In sum, the Bible describes humans made in the likeness of God, with special moral dignity and responsibility, with an innate calling to work in stewarding creation and to do this in loving relationship with God and each other. Any account of human flourishing in a world of AI must take these foundations into account.

We are called to imitate God (Eph. 5:1), that is to mirror his nature, having 'put on the new self which is being renewed in knowledge after the image of its creator' (Col. 3:10). This requires us to ensure that as technology and AI mediates our world, we do not allow it to diminish that calling. Loving our neighbour means that we should also be concerned for the ways in which it diminishes the image of God in them.

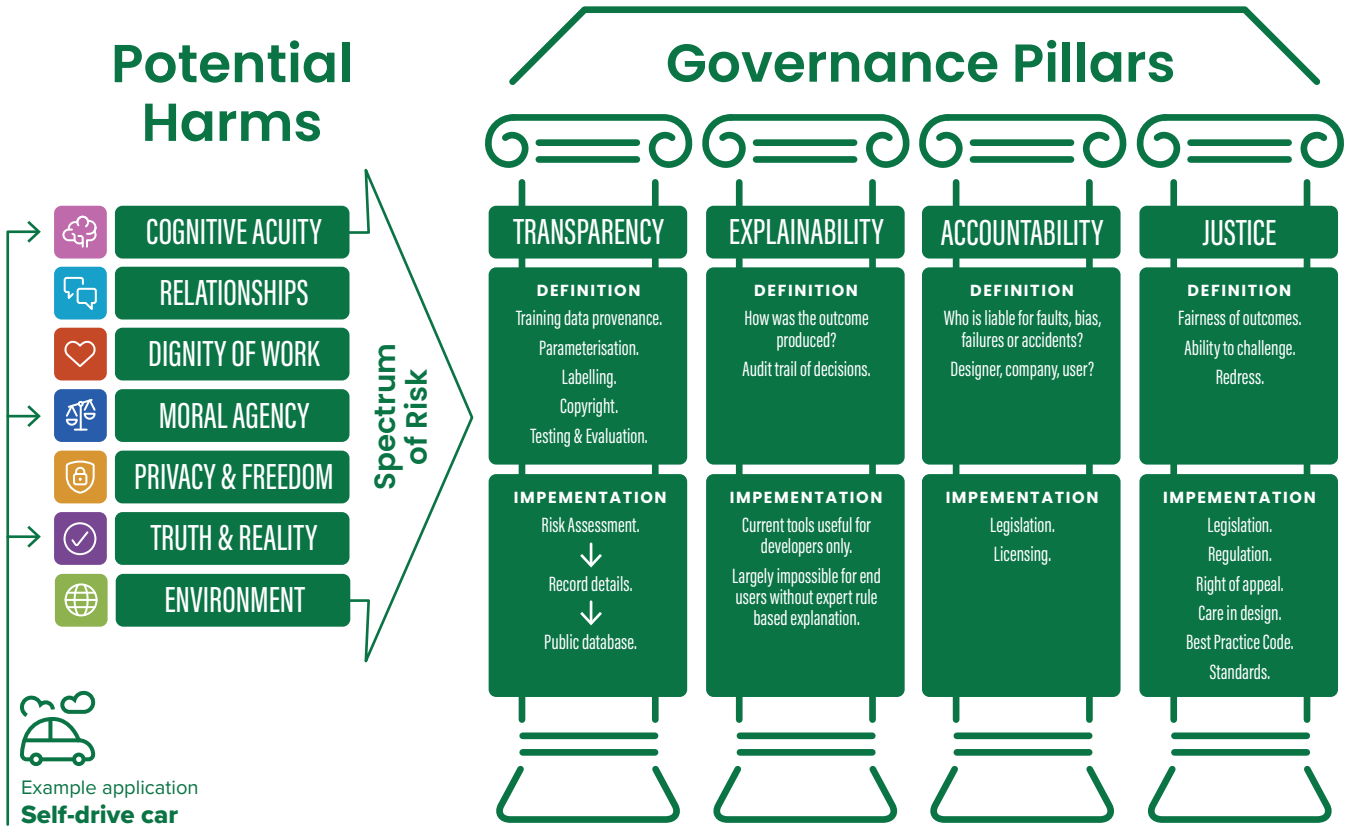


Figure 2: Mapping ethical principles into governance issues for AI applications and how they are implemented. Potential harms to humanity each lie on a spectrum of risk and are set orthogonal to these governance pillars. An application may have more than one type of harm. Illustration from: Jeremy Peckham, A harms and Governance Framework for Trustworthy AI, IEEE Computer, March 2024.

Policy Implications

The key question surrounding the use AI is: what sort of society do we want? There is a danger that following the techno-optimists' vision of human flourishing will mostly benefit the elite in society.[2] Technology can of course benefit society but there is a balancing act between what is gained and what is lost.

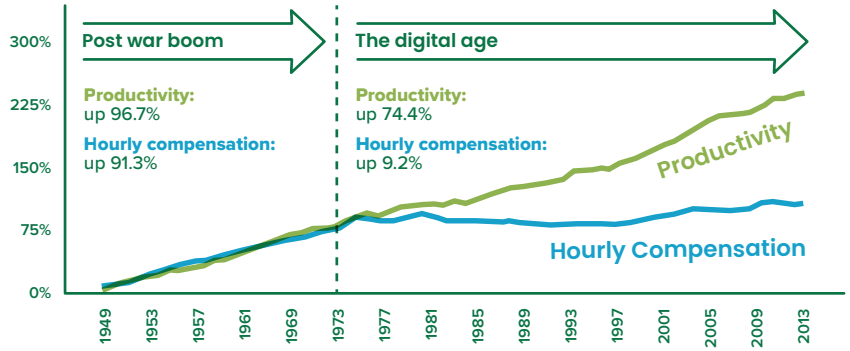


Figure 3: The relationship of productivity growth to wages post war to the digital age. Productivity hasn't benefitted workers in the digital age. Source: US Bureau of Labor Statistics.

There are six key areas of our humanity, our imaging of God, that are at risk of harm when we engage with applications of AI (Figure 2). Applications lie on a spectrum of risk, some are benign whilst others have a high risk of impacting our humanity. These risk areas, along with four governance pillars, provide a framework for considering the policy implications of AI applications.

1. PRIVACY & FREEDOM

Old Testament (OT) law lays a foundation for property law that could be seen as legitimately extending to the records and data that we hold about ourselves and that belong to us. The OT is also clear about theft. 'You shall not steal; you shall not deal falsely; you shall not lie to one another.' (Lev. 19:11)

Unfortunately data is an intangible asset and has limited protection apart from GDPR and copyright law. The development of LLMs has been at the expense of a huge infringement of people's copyright material. Policy in the area of data should be guided by the biblical requirement not to steal, and to be transparent, a concept held by many societies. These principles extend to areas such as uninformed 'one click' contract consent to provide data in exchange for services.

Policies should respect human autonomy and freedom where they are impacted by company or state surveillance using personal data and imaging. Examples include identifying individuals of interest, working patterns or personalised adverts.

2. DIGNITY OF WORK

The act of creation itself shows us that God works and creates. Work was also part of God's design for humankind and clearly a part of the natural order when he created Adam and Eve putting them to till and keep the garden. Policies should consider preserving the dignity of work rather than capitulating to something like Universal Basic Income. The evidence suggests that the digital age has not been as economically inclusive as the post war years (see Figure 3). Policy will also consider how productivity

2 Daron Acemoglu and Simon Johnson,Power and Progress, Our Thousand-Year Struggle Over Technology and Prosperity, Basic Books UK, 2023

improvements can be more inclusive and re-skilling provided. It will need sensitivity to the fact that some workers may not be able to transition to other forms of work that, for example, require soft skills.



3. TRUTH & REALITY

The Hebrew root 'emet' is used in various ways in the OT in terms related to truth and conveys the ideas of stability, support, faithfulness and conformity to fact. God, speaking through the prophet Isaiah declares 'I the Lord speak the truth; I declare what is right' (Is 45:19). The word is also used to convey the idea of conformity to reality as opposed to something erroneous as for example the response of witnesses in Isaiah 45:9 'Let them bring their witnesses to prove them right, and let them hear and say, it is true'. The New Testament (NT) usage of the term for truth accords with these OT concepts and also implies the absoluteness of truth, such as when Jesus prays for his disciples, 'Sanctify them in the truth; your word is truth' (Jn 17:17).

Various injustices result from data and human bias in AI systems that cannot be completely mitigated technically and policy will need to be cognisant of that. The untrustworthiness of Generative AI applications like ChatGPT, coupled with the statistical nature of their output and their use to generate fake content, runs counter to the Biblical view of truth that affirms its correspondence with reality. Coupled with the previously mentioned copyright infringement, this should shape policy on whether and how LLMs should be used in public services, business and free services.

As the Apostle Paul writes to the Ephesian Christians, 'having put away falsehood, let each one of you speak the truth with his neighbour,'. In using the likes of ChatGPT we must be careful to avoid what the prophet Isaiah laments, that 'truth has stumbled in the public squares' (Is 59:14).



4. AUTHENTIC RELATIONSHIPS

The relationship aspect of being created in the image of God helps us to think about whether our use of AI technology might affect the way in which we relate to people – does it diminish authentic, embodied and situated soul to soul relationships with others and, of course, ultimately my relationship with God? Various studies have shown how the manipulative algorithms in social media platforms have impacted the digital generation's relationships and attitudes. This will be exacerbated as the interaction with computers becomes even more frictionless through Generative AI and Virtual Reality or 'spatial computing'. Policy implications in this area might be to require service providers to offer the option of human interaction rather than just a bot, particularly in high stakes areas like medicine and finance.



5. COGNITIVE ACUITY & CREATIVITY

When we use applications like ChatGPT that can produce a story, provide medical information, write an email or even a prayer, we are giving up our own creative activity and imagination, all for the sake of convenience. The author is not human and we rob God of glory when we fail to image him through human creativity and the use of our minds. The argument that AI democratises creativity by allowing anyone to produce a piece of art is a lie because humans are not doing the creating. These concerns are in addition to the 'truth and reality' concerns already described. Policy will need to consider how or even whether such applications should be used in the public and business sphere given their impact on creativity, cognitive acuity and truthfulness. In private use, a labelling system could be required to warn users of potential mental health risks from using certain applications.

6. MORAL AUTONOMY

When God made Adam and Eve he gave them freedom to choose - moral freedom or agency. They chose, against God's command, to eat the fruit of the tree of the knowledge of good and evil when they were tempted in the garden of Eden. This choice ended with disastrous consequences for creation and humankind. However, free will is foundational to moral responsibility, regardless of whether one thinks morality has external agency or not. As a concept it's therefore important in the debate about AI potentialities, as well as in framing ethical questions in the use of AI and our relationship to it. It's also important in relation to the moral responsibility at the heart of the ethical debate about self-drive vehicles, autonomous weapons and systems that act as agents. Effectively, when we step into a self-drive taxi or our own self-drive vehicle, we are still moral agents. Humans uniquely have been given moral agency by God and this cannot be transferred to an artefact, it is only proxy agency.

Moral agency is also an issue in the use of many Generative AI applications where users effectively give proxy agency to the artefact, for example when getting it to generate an essay or story, in place of doing it themselves. Policy will need to ensure that humans remain responsible for their actions, even those of autonomous vehicles, robots and and the output of Generative AI. It should avoid granting 'rights' to robots in order to preserve the concept of human moral agency.

Potential harms from AI applications often fall into several of the areas described above. This means that policy considerations will need to be holistic rather than piecemeal and balance a pro innovation approach with protecting society from the broader range of harms discussed here, rather than a narrower set of concerns and supposed existential threats.

Conclusion

The techno-optimist's vision of a productivity boom and AI solving our biggest problems is driving the development and deployment of AI faster than governments can consider the risks and appropriate regulation. All this is at the expense of people's private data and copyright material. Yet the same people are warning of a future existential threat whilst development continues apace and they set the agenda on safety, perhaps to deflect from the impact AI is already having on society.

There are many benign applications for AI technology where it enhances or extends human capabilities rather than replaces them, such as in precise microsurgery and prosthetics or in hazardous environments like mine disposal. However, many of the mainstream applications are replacing human cognitive capabilities, shaping society in ways that will ultimately be detrimental to human flourishing. It also remains to be seen whether the promised productivity boom will be inclusive or not, the evidence so far suggests not.

Jeremy Peckham FRSA is a technology entrepreneur who has spent much of his career in the field of Artificial Intelligence, firstly with the MoD and then the software consultancy firm, Logica. He founded his first company in 1993 and later listed it on the London Stock Exchange. He is passionate about applying the Bible to the ethics of technology and AI.

This is not an official publication of the House of Commons or the House of Lords. It has not been approved by either House or its committees. All-Party Parliamentary Groups are informal groups of Members of both Houses with a common interest in particular issues. The views expressed in this report are those of the author and do not necessarily represent the opinions of the APPG.