AOSIS

Roland reflects on the implications of generative artificial intelligence on homo-digitalis

Check for updates

Author:
Dieter de Bruin[1] ⓘ

Affiliation:
[1]Department of Practical Theology and Mission Studies, Faculty of Theology and Religion, University of Pretoria, Pretoria, South Africa

Corresponding author:
Dieter Bruin,
dieter.debruin@up.ac.za

The rapid development of artificial intelligence (AI) has sparked theological debates about its potential to achieve consciousness and personhood. This article critically evaluates these claims, arguing that consciousness cannot emerge from computational systems as an intrinsic quality of living beings. Through philosophical analysis and theological reflection, it examines the fallacies in attributing agency or personhood to AI.

Contribution: These reflections are offered in a playful reflective tone, drawing on the form and content from two creative works by David Bentley Hart, where persons other than human beings are given voice to critique an algorithmic and mechanistic development of consciousness.

Keywords: generative AI; consciousness; animal theology; eco theology; David Bentley Hart.

Introduction

Dear Leon

We begin, as one cannot, but when reflecting on the implications of generative artificial intelligence (AI), with a portion of the Singer-Songwriter Nick Cave's letter (2023) in response to a question about using generative AI in songwriting. Ideally, one should listen to it read by the magnificent Stephen Fry (Letters Live 2023). The question was posed by a certain 'Leon' in Los Angeles: He wrote to Cave that he was employed in the music industry and that there was a buzz about ChatGPT. He mentioned that he had recently spoken with a songwriter from a band who had told him that he was using ChatGPT to write his lyrics, explaining that it was 'faster and easier'. Leon admitted that he didn't have a counterargument for that.

In one way, we must agree with Leon. There is no question that AI can do things as well as, if not better than, human beings, scoring in the bar exams of lawyers in the top 10%, for example (Barnard & Otte 2024:290). It certainly can do it 'faster and easier'. Leaving aside what it would mean for AI to write better songs or at least the same quality of songs as humans would do, even though it would be faster and easier, Cave begins his answer with the story of creation involving God working in the act of creation and having to rest on the seventh day. In contrast, according to Cave, ChatGPT eschews all:

> … notions of creative struggle that our endeavours animate and nurture our lives, giving them depth and meaning. It rejects that there is a collective, essential and unconscious human spirit underpinning our existence, connecting us all through our mutual striving.

Cave (2023) further, in no uncertain terms, denounces Leon's friend if they use generative AI to 'write' 'their' lyrics because the songwriter:

> … is participating in this erosion of the world's soul and the spirit of humanity itself and, to put it politely, should fucking desist if he wants to continue calling himself a songwriter.

This article argues that Cave's inclination is, theologically and philosophically, in the main, entirely sound and necessary. To wit: that the world *has* a soul, or perhaps more appropriately, as I shall endeavour to argue further on, that Soul has a world, and all reality is underpinned, not so much by an *unconscious* spirit but by Consciousness itself. Furthermore, in an extension perhaps of Cave's denouncement of Songwriters using AI to write their songs – I will argue here that theologians should desist, not so much, from calling themselves theologians if they use generative artificial intelligence (GAI) to write their theology – perhaps an argument to be had on another day – but that theologians should desist from entertaining the idea that if AI could write theology, or songs, or code or pass the Turing test, among others from calling GAI *Theologians*. In other words, I will contend that the part of the 'all' connected by the world's soul is not, and *will not* be, AI beings that have come to consciousness or achieved sentience. To come to the

defence of this thesis, I will call not on AI-generated or 'generating' beings but on a canine theologian-philosopher called Roland and some of the Greek gods who had a leisurely conversation about the small matter of the nature of consciousness. This could be another way of saying that David Bentley Hart will be my main interlocutor.

I chose David Bentley Hart for a few reasons to be my main theological interlocutor. Firstly, I take a leaf from the late South African theologian and philosopher of science, Wentzel van Huyssteen (2006), when he states that:

> … it has become impossible, and certainly implausible, event to talk about 'theology and science' in any generic, abstract, or a-contextual sense. This increasing awareness of the radical *social and historical contextually* of all our rational reflection should make it abundantly clear that in interdisciplinary dialogue the rather vague terms 'theology and science' should be replaced by a focus on *specific theologians* who are trying to do *very specific kinds of theologies* and are attempting to enter into *interdisciplinary dialogue* with very *specific scientists*, working within specified sciences on *clearly defined, shared problems*. (p. 4 [*author's own emphasis*])

It would be tautological almost to claim that in radically contingent positioning I resonate with Hart (a specific theologian) and his theology (as far as I understand it!). To a great extent, he has formed my worldview and theology. This theology and philosophy (in this case, the philosophy of mind) is, among other things, directly directed against reductionistic materialist or functionalistic metaphysics, as one would find, for example, in Daniel Dennett (Hart 2010, 2013) or David Chalmers (Hart 2024:56).

Secondly, in line with his non-material metaphysics, Hart is very open towards other forms of intelligence and sentience other than human beings. And in continuation of that point. his writing, especially his book *Roland in Moonlight*, but also 'embodies' this conviction. His book, *Roland in Moonlight*, can be seen as a kind of Socratic dialogue between his dog Roland and himself. It would be easy to say that this conversation is merely fiction or a clever rhetorical device; but in some sense, that would be too simplistic. In his most recent book, *All Things are Full of Gods*, Hart is explicit that his use of a symposium between the gods, Psyche, Hephaistos and Hermes to convey this idea is a rhetorical device.

The main foil for my argument is an article written by the distinguished theologian Dion Forster (2023):

In this article, the theological implications of artificial intelligence (AI) achieving consciousness are analyzed, with a focus on the feasibility of this development and its impact on concepts of personhood and spirituality. Initially, Dion Forster's view of AI as a potential theological entity is presented as a contrasting perspective. Subsequently, the philosophical arguments of David Bentley Hart, are utilized to assert that machine consciousness is unattainable due to the pleonastic fallacy, the category error, and the narcissistic fallacy. Lastly, the ethical risks and potential degradation of human-centred spiritual principles resulting from the pursuit of AI consciousness are evaluated.

Forster as foil

But before we can start dismissing the idea of artificially intelligent sentient beings, we should reflect on what we are talking about when we invoke the idea of AI. And here already, Forster is of great help.

On the one hand, Forster delineates what you would call utilitarian AI, which consists of symbolic AI and sub-symbolic AI.

Symbolic AI explicitly represents human knowledge in a declarative form through rules and facts encoded in computer programmes. It operates by mimicking human reasoning within predefined structures. Here, one could think of a computer programme that can play chess, such as Deep Blue, or grammar and spelling checkers such as Grammarly. In theology, this technology can be used to analyse big data sets texts for example to find common themes.

Sub-symbolic AI advances beyond rule-based systems, focusing on machines with cognitive capabilities – reasoning, learning, perceiving and communicating – often self-adapting through experience. Here, we have less of a closed rulebound system and more of an inductive system that sifts through data to find patterns that can help it make predictions in the stock market; it can, for example, analyse a large data set of demographic data in an area, together with theological resources on missional theology, and church council minutes, to come up with a strategy for community outreach. Sub-symbolic AI learns from mistakes and can reprogramme itself for better results.

These two utilitarian uses of AI are not what is at issue here. Rather, Forster (2023:5) foresees instances of AI where it '… is not only a utility used by the theologian, but rather, a fellow theologian, a conversation partner, a phenomenon that invites reflection, conversation, and mutual engagement'.

Foster is speaking about extending organic theologies and organic theologians to non-human, specifically non-organically intelligent sentient beings who we not only reflect upon but reflect *with* and have their reflective capabilities. For Forster, it is perfectly possible that robots could be '… contributing theologically from its existential perspective'.

Marcel Barnard and Wim Otte (Barnard & Otte 2024) are *ad idem* with Forster on this point. They posit that if language, creativity, grounding and meta-representation are aspects that make humans unique (human exceptionalism), AI's excelling in these domains question the unique status of humans. Barnard and Otto (224:306) speculate that AI is in the children's shoes as it were, but that the child can grow up to reflect on itself and escape human influence, whereas AI can become a religious call on God and become an independent agent.

The thrust of Forster's (2023:11) article is also to question human exceptionalism. He connects to a tradition where the

capacity for agency and theologising is extended from humans to non-human creation. At stake for Forster is that 'personhood' could be extended to non-human beings, or to put it differently, that being human and being a person are not necessarily interchangeable concepts. For Forster, in line with Barnard and Otte's perspective, the fact that AI is 'created to represent our image' in that it takes on some of the qualities that make us human, Forster takes the realist pragmatist view that AI can become moral agents. Or, in line with Cave, and quoting Warehem, 'AI can become genuine subjects of harmonious communal relationships, exhibiting solidarity and identity'.

It seems that Forster is arguing in line with that the possibility for AI to attain subjecthood will be actualised when humans can no longer detect when a 'syntactical machine agent' fools (his word) 'us into the mistaken belief that it genuinely experiences empathy and cares for us'. When the barrier is breached where humans can no longer 'detect the intersubjective inadequacies of AI technologies' (Forster 2023:15), it is then the moment for the reconsideration of the boundaries of moral personhood.

Forster (2023:8) connects with the impulse of an African theological perspective on intersubjective identity that is praiseworthy. It is an anthropology that refuses to absolutise human uniqueness. It is an anthropology that embraces humanity's interweavement with each other but also with other non-human creatures. Forster connects to a theological trajectory that rejects the absolute claim that humans are in every way better than the rest of creation. Creatures such as rivers and trees prossessing sentience and consciousness thought should be embraced; it is even true that it is not only human beings who can be theologians. It is, however, that one of the best thinkers of the non-human world is of the firm conviction that artificially intelligent beings can never do so; he will be backed up by some of the Greek gods.

As we have seen, GAI, (Bajohr 2024; Neuman, Danesi & Vilenchik 2022) such as large language models, operates by analysing massive datasets of text to identify statistical patterns, enabling it to predict and generate coherent responses to prompts. While this process yields impressive results – often mimicking human-like conversation, My critique, however, is not aimed at what AI can or cannot achieve technically. Instead, it is a philosophical exploration of whether such systems, no matter how advanced, can possess the subjective experience and intentionality central to consciousness and personhood. This distinction is key: my concern lies not with technological limits but with the deeper question of what it means to truly "be".

Forster's vision of AI as a theological companion suggests machines might one day join us in contemplating the divine. Yet, this optimistic outlook invites critical examination. Can artificial systems truly possess the consciousness necessary for theological reflection? David Bentley Hart and his dog Roland argue that they cannot, offering a rebuttal that challenges the core of Forster's position.

Roland rebuffs

Roland was sadly David Bentley Hart's wise and trusted friend and collaborator. He was a dog with unwavering theological instincts and philosophical prowess. Again, as referred to earlier, I will not pronounce whether Dr Hart's conversations were imagined or fictitious – It seems wrong to do so. I choose to receive the conversations between Dr Hart and his companion as it is presented to us. Perhaps an endorsement of the book by Henry Weinield, can frame the role of Roland in the reflections of Dr Hart's book:

> David Bentley Hart's Roland in Moonlight, part work of fiction, is a brilliantly ambitious and deeply moving book – lyrical and philosophical by turns – that centres on conversations between the narrator and his dog, the mighty-souled Roland, who, in differing from humans (he is at once more dispassionate and more compassionate), at the same time presents us with a model of what humans are capable of being and becoming. The conversations, ranging far and wide over a host of topics, from the mundane to the arcane, are interspersed with poems, some of them very beautiful, written by the narrator's Uncle Aloysius, a brilliant creation in his own right. Charting the years of the narrator's illness and partial recovery, Roland in Midnight is at the same time a palimpsest that sketches a symbolic movement leading from the loss to the partial recovery of Eden (Hart 2021:loc.24–32).

I cannot recount the remarkable and poignant tale of how Roland entered David Hart's life and how they started conversing, chiefly nocturnally. Still, I can say that I have gained the insights I purported to have on this topic from him and the putative 'owner' student he so ably taught.

And if I had understood anything about Roland and Dr Hart's interchanges, Roland would have heartily agreed that personhood should not be confined to human beings and that creatures other than human beings can be theologians. In fact, he does not think very highly of our species' ability in that regard.

He would, however, short-circuit our speculation that so-called strong AI would ever develop agency, rationality or consciousness. In short, though it is perfectly laudable to reflect on AI or, God forbid, use GAI to do theology, (they) will never be able to become theologians.

Neither would Roland deny the staggering ability and utility for good or ill – in fact, he has already done so.

In his own words, speaking in this case of programmes such as Deep Blue, which already in the 1990s, beat Kasparov:

> What's interesting is how beguiled even some programmers were by the illusion of intelligent agency. Of course, theirs were the only mental agencies present in the actions that occurred through the software they'd devised. (Hart 2021:28)

For Roland, the only true minds at work are those of the human creators whose ingenuity animates these lifeless systems. The machines' victories, he insists, belong to their makers, achieved not through any spark of thought but by its

stark absence – a void where affectivity, intentionality and the pulse of consciousness should reside. This, he argues, is the perennial story of technology: tools fashioned to expedite our labours by stripping away the very conditions of soul and sinew that define us. Machines heft burdens not merely with brute strength beyond our own but because they lack will, feel no strain or tedium, and harbour no purpose of their own. Their power, Roland asserts, stems from a profound emptiness – no effort of reflection, no whisper of desire, no grasp of time's unfolding. Their ceaseless algorithmic dance lacks the bridge between an intending spirit and the living moment, devoid of that sensuous awareness of the here and now that marks consciousness as agency. To him, these programmes are no more thinkers than Archimedes' lever was a toiler of flesh and blood.

The key phrase that I want to latch on to here is the fact that if we were to follow the reasoning of Forster, Barnard and Otte that AI can become sentient, conscious beings capable of their own agency – we would indeed be fooled by the appearance of intelligent agency in the form of language, creativity problem-solving ability among others.

David Bentley Hart's argument rests on a broader starting point: that consciousness cannot 'emerge' from AI or that AI can 'develop' into conscious beings. The reason is that consciousness does not *develop* or *emerge* from anything; everything comes from consciousness, and consciousness, language and life are but other names for God himself (Hart 2013, 2024). As Hart (2024) states himself:

> … life is itself the pervasive 'organic' logic of the material order from the first, not emerging from that order but instead creating, governing, forming, and quickening it from within. I believe, moreover, not only that mind and life are both irreducible; I believe that they are one and the same irreducibility. And, in fact, I would add language to that combination as well, as yet another aspect of one and the same irreducible phenomenon, ultimately inexplicable in mechanistic terms. (p. 56)

Consciousness cannot be reduced to something like information, as Tononi (2008) might argue, but is 'is not a machine; it is an act that, like any act, can be understood only in its intrinsic unity' (Hart 2024:56).

To think that consciousness and sense can develop out of material causes, such as computational processes, rests on at least three fallacies. We will enlist the God Psyche to explain these fantasies to us.

The Pleonastic fallacy

Psyche, who speaks for Hart (2024:22) as to his 'convictions regarding the structure of life and mind', rejects the notion that the mind somehow emerges from non-mind material. And to think that mind somehow 'emerges' from matter or information would be to fall prey to the pleonastic fallacy where the mistake is made, 'an infinite qualitative distance can be crossed, or even simply diminished, by a sufficient number of finite quantitative steps'. For Hart (2024) – the:

> [D]istinction between objective physical events and subjective phenomenal episodes is, I submit, just such an infinite, untraversable distance. No accumulation of mindless physical proficiencies by themselves would ever be enough to add up to even the most elementary of mental powers. (p. 56)

This is the fallacy at play when the argument is made that although AI technology is not 'there' yet (there being the critical threshold when AI will 'cross over' the minimal threshold of personhood or consciousness or agency), it is still in its children's shoes and will indubitably grow exponentially – that the objective 'third person' can cross over and develop into the subjective first-person awareness.

To think that AI can somehow achieve consciousness of moral agency is to believe in a physicalist register that 'consciousness must be explicable entirely in terms of the interaction between our neural constitution and the concrete world around us' (Hart 2013:153). Or to think that AI can develop consciousness from distinct functions such as pattern recognition, problem-solving or learning and adaptation. Whereas Hart (2013:153) 'is undoubtedly correct that neither stimulus nor response is, by itself, a mental phenomenon; neither, as a purely physical reality, possesses conceptual content or personal awareness'. By extension, we could say that AI's inputs and outputs lack personal awareness. Just as the physicalist attempts to think that consciousness is a result of magnificently complex interactions within the brain's neurons, the approaching infinite complexity of large language models cannot 'become' conscious. Consciousness is not a bottom-up phenomenon, according to Psyche (Hart 2024), but rather that:

> consciousness, uniquely, is first-person in its phenomenal structure, all the way down. In fact, it's the very phenomenon of the 'first person' as such, the sole act whereby someone is anyone at all … (p. 35)

Category error

The above-stated relates directly to the classic maxim that 'correlation is not causation'. The fact that Large Language Models (LLMs), as Barnard and Otte (2024) would argue, can generate an infinite variety of sentences in the same way that humans can create sentences, or create poetry as humans do, or can pass 'theory of mind' tests, even better than humans would or can extract real word knowledge in a grounded way, does not mean that these qualities that would correlate with human personhood when present would necessarily indicate or cause personhood or moral agency. In the same way, to ascribe AI beings traits such as the simulation of human behaviours such as empathy or solidarity in the face of humans forming emotional bonds with AI beings or robots producing theological statements or outputs would be evidence of personhood, would be to fall prey to the same category error. Just as humans enact behaviour that signifies empathy or theological reflection, it does not mean that when one sees text or movement that indicates theological beliefs, one could infer personhood or moral agency from such movements. John Searle's *Chinese Room* thought experiment

vividly illustrates this limitation: just as a person manipulating Chinese symbols according to syntactic rules does not truly understand the language, so too does AI, no matter how advanced, remains confined to computational processes that lack the semantic depth and intentionality inherent to human consciousness (Searle 1980).

Narcissistic fallacy

This leads directly to another fallacy that pertains to our ascribing agency and sentience to computers, robots, among others. According to Psyche (Hart 2024:70), it is 'the error we make when we project our own mental agency onto some inanimate medium or device we're employing, as though it were the source of that agency …' See, for example, in AI's production of poetry or Nic Cave-like songs, evidence of creative agency on the part of the AI would be to mistake the true agency of the programmers that created this tool to generate poetry, songs translations, among others.

In his prescient review of Hart's book, Charles Carman (2025) summarises the folly of ascribing attention, meaning and agency to the abacus:

> We are not so delusional as to ascribe to the wooden beams and rods and beads any awareness of what it means for the beads to be placed to the right or to the left. Only the student attending to the tool knows what their placement entails. Likewise, AI language models move linguistic items around but without any recognition that they mean anything. They needn't mean anything, in fact, for the algorithm to analyze them. (p.1)

To put it in another way, meaning does not 'emerge' from tokens but symbols, words are used to convey supervenient meaning.

To use an analogy, Tolstoy's mind and his intention and meaning with his novel *Anna Karenina* do not somehow emerge from the markings and pages of a book. Instead, the markings and pages of the book are the result of the intention of the author, and it is impossible for the book and the markings to read and interpret itself. It is necessary for an interpreter to use spectacles, brain operations and all other chemical and biological processes. Still, it needs an interpreter to decode that meaning, as it were, and for Tolstoy to have encoded that language's new paragraph, the same would go even for so-called genetic coding. It is not the organism's material, genes, or whatever that 'reads' in quotation marks, are the genetic material that needs to be interpreted. A camera lens does not see – it requires an interpreter to see, and judgement is necessary; all of those things depend on certain ends, according to which it would be interpreted, according to which forms would be discernible at all; another way of putting it would be that agency cannot emerge from the information.

Or in the words of (Hart 2024) himself:

> philosophy of mind should liberate itself from the nonsensical notion that mental agency can be understood mechanically, as a combination of extrinsically allied parts and discrete functions. Whatever the mind may be, it is not a machine; it is an act that, like any act, can be understood only in its intrinsic unity. (p. 20)

I have argued, with the aid of David Bentley Hart and his canine collaborator Roland, that consciousness – far from emerging from the machinations of GAI – is the primal ground of all reality, an infinite qualitative realm no algorithm can breach. Against Dion Forster's hopeful vision of AI as a 'fellow theologian', capable of reflecting from its own existential stance, I have countered that such a notion falls prey to fallacies of emergence and projection, mistaking computational prowess for the soul's depth. Roland's wisdom affirms that while rivers and trees might whisper truths of the divine, machines remain mute, their outputs a shadow of human agency. So why Cave's strident tone in the beginning? Is ChatGPT generative AI an existential threat to be fought for the sake of the soul of the world? Perhaps not as such, but entertaining AI somehow obtaining consciousness and agency presents at least three dangers.

Missing the real enchantment and the real threats

The first is the danger of having an extravagant adventure that misses the point, wasting research time and money, and, importantly, diverting attention from the very worthy goal of engaging in an expansion of the personhood of theology and theologians that can be found without the human world. Focusing on mechanistic intelligence could be understandable if lonely men could enter into a relationship with a holographic being, such as in the movie Bladerunner 2049, or robots eliciting a sense of purpose and agency, such as in the movie WALL-E (2008). These artistic impulses could say something about our longing to be in a broader and deeper relationship with the world. As Psyche (Hart 2024:556–557) would have it:

> The proper habitat of a living soul is an enchanted world, charged with *mana* or filled with fairies or *kami*…and in the absence of those numinous or genial presences, human beings feel abandoned and very much alone. (pp. 556–557)

In another context, in an interview with Oxford Scholar and author of the book *A Brief History of Artificial Intelligence: What it is, Where we are, and Where we are going*, Michael Wooldridge, Jonathan Bi (2025) makes a prescient observation. Rather than the benign intent to include AI beings with agency in communities of care, the question revolved around the existential threat that it might pose if AI beings can develop agency and pursue goals that are inimical to human interest. Apart from the fact that Wooldridge places such an eventuality without the realm of plausibility, Bi (2025) makes the observation he and his friends came up with:

> [T]his half-joking term called 'existential risk', risk, which is the risk upon a society that focuses too much on existential risk and away from other risks that we could actually be facing because of AI today … (09:27-09:41)

Or as Wooldrige in his conversation with Bi (2025), puts it when we fixate on the narrative of machines becoming as intelligent as human beings:

> It tends to suck all the oxygen out of the room … And it tends to dominate the conversation and distract us from things that we should really be talking about. (03:38-03:50)

These risks are real, and one of them relates directly to our need to rediscover an enchanted world.

Missing the opportunity for real ubuntu

To reiterate, I concur wholeheartedly with Forster (2023):

> We know that the human body is not the only location for theology – eco-theologians, animal theologians and a host of others invite us to reflect theologically from locations outside of ourselves. (p. 7)

Including dogs, trees, and mountains in our communities of care is what we should do. However, including humanoid robots within our church communities as persons would be a grave mistake because it would detract from our care and inclusion of non-human persons. Although Forster (2023:7) would concur with the statement that 'virtually all Christian theologies are 'organic theologies' (McBride 2019:669). The problem is that only organic beings pay the price for the creation and/or AI technologies. Focusing on so-called moral and sentient robots can obscure what Gray and Suri (2019) 'ghost work', where the 'human labor powering many mobile phone apps, websites, and artificial intelligence systems can be hard to see – in fact, it's often intentionally hidden'. Paradoxically, to include AI beings within the community of human persons could have the effect of subordinating the needs of the environment to sustain these beings in their 'lives'. It already takes an enormous amount of energy to train these AI models; it takes huge data centres to store the data. It generates vast amounts of e-waste – and returns to the point again – that this environmental desegregation has a disproportional effect on marginal communities (Brevini 2022) – to persons with whom we embody ubuntu.

Conclusion
The problem with Narcissus

Professor Forster references an article written by Hart about the danger of Narcissus, where we 'create these technologies in our image' (2023:18) with our own prejudices and biases such as racism and sexism, for example. The danger is far deeper than that. Perhaps the problem is not primarily that we can consciously create AI in our image but that we will fall in love with AI and think of ourselves in AI terms, furthering our self-conception as computational, algorithmic, mechanistic beings.

As Carman (2025) would have it – this would be the main concern of David Bentley Hart's book *All Things are Full of Gods*. A world where life, language and mind are conceptualised in AI terms will be 'a network of inputs and outputs'. And what's 'left of life will grow quiet and illegible, replaced with the noise of the machines'.

Ultimately, we should begin to think of the soul of the world not as love but as a manipulative algorithm.

Acknowledgements

References

Bajohr, H., 2024, 'Dumb meaning: Machine learning and artificial semantics', in E. Voigts, R.M. Auer, D. Elflein, S. Kunas, J. Röhnert & C. Seelinger (eds.), *Artificial intelligence – Intelligent art?: Human-machine interaction and creative practice*, Digitale Gesellschaft., 1st edn., vol. 64, pp. 45–57, transcript Verlag, Bielefeld.

Barnard, M. & Otte, W., 2024, 'Is the machine surpassing humans?: Large language models, structuralism, and liturgical ritual: A position paper', *International Journal of Practical Theology* 28(2), 289–306. https://doi.org/10.1515/ijpt-2023-0078

Bi, J., 2025, *Don't believe AI hype, this is where it's actually headed | Oxford's Michael Wooldridge | AI History*. Viewed n.d., from https://www.youtube.com/watch?v=Zf-T3XdD9Z8

Brevini, B., 2022, *Is AI good for the planet?*, Polity Press, Cambridge.

Carman, C., 2025, 'One to zero how to destroy the future', *The New Atlantis*, viewed 17 June 2025, from https://www.thenewatlantis.com/publications/one-to-zero.

Cave, N., 2023, *The red hand files*, viewed 17 June 2025, from https://www.theredhandfiles.com/chatgpt-making-things-faster-and-easier/.

Forster, D.A., 2023, 'African theological perspectives on intersubjective identity: In conversation with developments in Strong Artificial Intelligence', *Stellenbosch Theological Journal* 9(1), 1–23. https://doi.org/10.17570/stj.2023.v9n1.at4

Gray, M.L. & Suri, S., 2019, *Ghost work: How to stop Silicon Valley from building a new global underclass*, Houghton Mifflin Harcourt, Boston, MA.

Hart, D.B., 2010, *Atheist delusions: The Christian revolution and its fashionable enemies*, Yale University Press, New Haven, CT.

Hart, D.B., 2013, *The experience of God: being, consciousness, bliss*, Kindle edn., Yale University Press, New Haven, CT.

Hart, D.B., 2024, *All things are full of Gods: The mysteries of mind and life*, Yale University Press, New Haven, CT.

Letters Live, 2023, Stephen Fry reads Nick Cave's stirring letter about ChatGPT and human creativity. Viewed n.d. from https://www.youtube.com/watch?v=iGJcF4bLKd4&t=2s

McBride, J., 2019, 'Robotic Bodies and the Kairos of Humanoid Theologies', *Sophia* 58(4), 663–676.

Neuman, Y., Danesi, M. & Vilenchik, D., 2022, *Using AI for dialoguing with texts: From psychology to cinema and literature*, 1st edn., Routledge, New York, NY.

Searle, J.R., 1980, 'Minds, brains, and programs', *Behavioral and Brain Sciences* 3(3), 417–424.

Tononi, G., 2008, 'Consciousness as integrated information: A provisional manifesto', *Biological Bulletin* 215(3), 216–242. https://doi.org/10.2307/25470707

Van Huyssteen, W., 2006, A*lone in the world?*, Eerdmans Pub Co., Grand Rapids, MI.

WALL-E., 2008, Directed by Andrew Stanton [Feature film], Walt Disney Studios Motion Pictures, Burbank, CA.

Wooldridge, M.J., 2021, *A brief history of artificial intelligence: what it is, where we are, and where we are going*, First U.S. edition, Flatiron Books, New York, NY.