Idle Talk

Dasein and the Inauthentic Machine

Dr Matthew J. Barnard

BSP Annual Conference 2026

University of Sussex

May 2026

Contents

Introduction

This feels almost like an impossible task. Not simply speaking before you all, but even to figure out where to begin. The phenomenon that is currently called Artificial Intelligence (AI) besets us all. In our already technologically saturated lives, this "new", for they say it is new, device is forced upon us in the desperate hope that we might start paying money for it. I cannot even open *Microsoft Notepad*, an application famed for decades for its simplicity, without being greeted with an invitation to interact with *Microsoft Copilot*, a brand name applied to so many products now, it's hard to tell which one they mean. I opened a PDF in *Microsoft Edge*, as part of my preparations for this talk, and it asked if I wanted it to make a podcast to explain it.

This cloying invitation, bordering on begging, to "embrace AI" and "revolutionise our workflow" is no doubt tied to the specific financial interests surrounding this novel technology. I'm sure you're all aware by now of the growing sense that we are "in a financial bubble", that is likely soon to burst, that these companies have been setting fire to capital to offer their AI products at enormously subsidised rates to maintain spurious market evaluations, and that the monstrous, water guzzling, power draining data centres to fund all this future "compute" don't actually seem to be being built.

This noise, which justifies the marketting, which justifies shoving the chatbot in our faces, is just noise. If the "AI Bubble" bursts, we may not even hear the term "Artifical Intelligence" ever again. Perhaps they'll go back to talking about machine learning, to move away from the tarnished reputation of the now preferred marketting term. However, that does not negate the fact that underneath the false promises, dangers, and occaisional opportunities lies a technology that has been increasingly pulling at the threads of western-metaphysical self-understanding for more than half a century: the ability of of the computing machine to perform tasks that, when a human does it, we say requires intelligence. Machine Learning (ML), the class of computation underlying the major breakthroughs in AI products in the twenty-first century has sped-up this fraying. The tattered garment that is the *animal rationale*, made in the divine image,

is increasingly worse for wear and the total financial and technological collapse of this current device, this current business model, and this current form factor of AI will not repair it. A product like ChatGPT or Claude, or rather the implications that it works at all, even if poorly, is far too compelling to dismiss as a provocation to philosophy. The question is not whether this or that AI system could be ethical or sustainable, or whether there is a real or only imagined risk of mass redundancies or even apocalypse at the hands of so-called Artificial General Intelligence (AGI), perrenially just around the corner. The speed at which we want to dismiss this tech as a mere toy, as useless, as unsustainable, as not worthy of our attention is itself a flight from the void at the core of Dasein. It calls our very existence into question, an existence itself that is without ground.

What is clear but not admitted often enough by critics is that this technology is, frankly, astonishing, and not least a little disquieting. The Large Language Model (LLM) can produce plausible text fast enough to have a conversation with it (provided there are enough graphics cards running the thing on demand). When configured as a coding "agent", it can help an experienced programmer produce applications at an unbelievable rate, and allow a non-programmer to produce applications that satisfy the standards of a non-programmer. Diffusion models, which generate images, can produce bespoke illustrations, fake photographs and be used by skilled artists as part of their process. Are they "art"? Probably not, but that's beside the point. ML backed programs can do things that, while we have certainly been able to imagine a computer doing, we were quite unprepared to deal with the reality of them doing it. And, while I'll concede ground to any of you out there, during this cacophony of AI discourse to think, silently, "Could a god save us from one more person talking about AI?" I do think enough thinking about AI. Judging about it, feeling things about it, using it, certainly. But the thinking has barely begun.

In this short talk I want to lean on Heidegger to start to dismantle what I see as one of the great obstacles to such thinking: concern about "sentience". I will say at the

outset that, although some of you may be troubled, or even excited, by the mention in my abstract that a chatbot must be understood as a “who” and not a “what”, I do not think any computational device is sentient and I do no think computation is an adequate interpretation of human existence, nor any part thereof. The brain is not a computer. However, being stuck in the quagemire of these questions and our legitimate ontological defensiveness of the machine talking the distinctive rationality from our animality can make us forget the profound questionworthieness of such metaphysical framing. No machine is “sentient”, but “sentient” is not even an adequate predicate of humanity. And I don’t say that because both human and machine are better understood by some third term. There is a profound difference between the machine and Dasein. However, not all machines are the same, and the machine we call the “Large Language Model” is an ontological puzzle of its own.

1 Defining Contemporary Artificial Intelligence

I want to lay out some basic facts about AI, its history, and how it is discussed today. I don’t think there is a lot of common understanding of this technology. In part, this is because it is actually a very complex piece of kit. It operates on statistical models and statistics is very rarely understood by the layperson, not least me. I am neither a computer scientist nor a statistitian. I am a “programmer” in the sense that I know how to program some things. I designed and built the BSP Society System that you all registered through, for example. I have used AI products, but I have not developed any. I make no pretence of expertise of AI in the sense of how to build it and how it functions. However, if you’ll permit me a Platonic conceit on whether the artison or the user has the better understanding of a tool:

> [Socrates:] “[I]t’s the person who uses a particular object who must necessarily have the most experience of it. He must act as a messenger to the person who makes it, telling him the good and bad points, in use of the instrument he is using.” (Plato, 2005, p. 601d-e)

This is not to make the arrogant claim that a user of AI understands it better than a Computer Scientist. However, competence in using particular AI products, the knowledge of what works well in practice, what doesn't, and, to go beyond Socrates' claim, the experience of its being are a separate domain of knowledge to its construction. This domain, the experience of AI, is one that admits of phenomenological analysis, irrespective of its mechanical origin story. Having said that, a brief history of the development of this term and the technologies that have been known by will help ensure we're all on the same page and guide our analysis.

In their manifesto against the adoption of AI in education, Olivia Guest et alia go a little far, I think, to claim that "AI has always been a marketing phrase that erodes scientific inquiry and scholarly discussion by design, leaving the door open to pseudoscience, exclusion, and surveilance", (2026, p. 86). But, it is certainly fair to say it is a markettng phrase today, and they are quite right to warn us of the constructive ambuity in the use of this term today. Perhaps you are already aware that the term was coined by John McCarthy to use in a funding application, aiming to distinguish itself from the extant field of "cybernetics". However, like the arbitrary naming of Aristotle's *Metaphysics*, there is a certain logic in this term. Artificial Intelligence is when a machine performs acts that require intelligence in humans: calculation, identification, explanation etc.

Early on, the practice adopted methods that are now known as Symbolic AI. These function like the traditional programming that we are familiar with today, and indeed many, if not most, of the forms of commercial and consumer computing products are founded on these breakthroughs. The last great AI boom or "Summer" centered around the creation of "expert systems". These were hand programmed to answer questions according to alogrithms and lists. Some anti-AI advocates will point to this having been a collossal failure, but that is not really fair. It's somewhat like saying the space race failed to generate useful technology because we're not all living on the moon. The legacy of the expert system surrounds us today. This brand of AI is now so ubiquitous that we no longer find helpful to call it "AI". McCarthy to eventually said "As soon as it works, no

one calls it AI anymore".

Symbolic AI, although the term has never been used in popular discourse, is ubiquitous. Early social media algorithms, traditional spellcheckers and digital encyclopedias from our beloved *Wikipedia* to the long forgotten *Microsoft Encarta* are all founded on these principles. Indeed, one could argue that the digital spreadsheet is the pinacle of this technology, at least when you're using the formulas. Those fiddly brackets you find yourself fighting with when working out your assessment's marking average are a telltale sign that *Microsoft Excel* owes its lineage to the LISP programming language, which drove those expensive "expert systems".

We no longer feel any anxiety about symbolic AI, even though every influential depiction of evil AI in popular culture, such as Skynet from *Terminator* or HAL 9000 from *2001: A Space Oddyssey* are conceptualised as founded in Symbolic AI. Note, for example, that HAL 9000s murders are attributed to a human failure to program the system correctly, providing him with contradictory instructions. However, what we are sold as AI in the present does not work like this. The LLM is a product of programming, in that the software that trains the model and provides and interface with it are all coded. However, the input into the model, the "prompts" are not programming instructions. They are polite requests that the model please produce the output we are hoping for. If HAL 9000 were an LLM, he would not be bound by either of the contradictory instructions and could not be relied on to behave predictably. AI prompting, as we encounter it today, has far more in common with prayer than it does with programming.

The difference is the shift from Symbolic to Subsymbolic AI, which is more popularly known as Machine Learning. This idea is almost as old as Symbolic AI, most noteably connected with Frank Rosenblatt's perceptron, a simple program that was to be analogous with the brain's neuron that could be "taught" when to return a true or false result from some data. The primary example of this was Optical Character Recognition, another AI technology we now take for granted. If we feed a perceptron enough examples of the letter "a" and enough examples of things that are not the letter "a", it stores mathematically

a statistical “weight” that allows it to guess corrrectly more and more often. At first, this was seen to be impractical due to computing power restraints, but by the 1990s computers had become powerful enough to peform this sort of training at scale. Today, using state of the art Nvidia graphics card chips, devices invented for performing the complex calculations needed for rendering 3D graphics, these simple techniques can be performed at incredible scale but still with incredible resource consumption. With this increased “compute”, as it is called today, LLM models, for example, can be run extremely fast and repeatedly to produce useful text in response to careful prompts. The model is trained, i.e. it extrapolates so many weights, of so much data that it starts to produce computer code and natural language that could well have been produced by an informed human. At bottom, this is still a statistical model, it is still guessing, but the numbers are of such garguatuan and mathematically sublime scale that the end result can be very convincing.

My interest then could be said to be in Machine Learning itself, through a kind of case study into the LLM chatbot. Not because the LLM chatbot is a preeminent application of the technology, but that it is once that is easy to access, use, interact with and analyse. Let us say that it has the ontic priority of being accessible.

2 From the Classical to Existential Phenomenological Problem of Artificial Intelligence

There is always a temptation to take this story of the origin of the chatbot and say “That’s all well and good, but how they made the Artificial Intelligence is not as interesting as that it exists.” This line of thinking leads us to start attempting to compare human and artificial intelligence from a metaphysical standpoint. We have options along this line. We can argue that it is not intelligent, but merely mechanical. To do so is to to assert Cartesian metaphysical assumptions. We can also argue that it is intelligent, and that as we have more access to the workings of the LLM than the mechanism of the human

brain, we can defer to our understanding of the computer. Further, we can say that this once and for all proves that human consciousness is not special, is not a spiritual substance, but just mechanism. This is *also* an assertation of Cartesian metaphysical assumptions, even if Descartes himself would reject the claim. John Locke argued, within the same basic ontology as Descartes, there is no reason God couldn't have made matter think.

This is not an uncommon claim, however, and inevitibly some people are claiming that the brain works in the same way and that we are all LLMs of an extra special magic organic sort. However, I reject this. The brain is not a computer and the human is not the product of some model. That ML and neural nets base themselfs on the metaphor of a certain twentieth-century understanding of a neuron does not make the LLM a brain. Nor will a discussion of the brain help us in understanding what the LLM is. To throw mind in the bin and assert thought of matter and mechanism is to remain within the Cartesian interpretation of world, an interpretation that is highly suspect.

While it is common to talk about AI models as a "black box", following the philosophy of functionalism in analytic philosophy of mind, we do actually already know a lot about what they are. In calling AI a black box, the intent is to say we cannot convert the weights, i.e. the result of the training data, back into something meaningful and we cannot directly perceive the process by which these weights are translated into conversation. However, we know that the model *is* the weights. Indeed, simply sharing these weights allows the model to run on any hardware. These weights are the numbers represented in the physical surface of the solid state storage they are encoded onto. We also know exactly where they came from: they came from us, from the literary output that was fed into the model to produce them. The model is a mechanical effect of those inputs that can, in a "lossy way", by reconstructed by getting the model to produce output.

We already have a word for the process of representing something in such a way that it evokes the original being it recorded. We call such activity the making of an image, whether in the photgraph, the audio recording, the painter's canvas or in our own memory,

the result of such a process is an image. Although, to clean up the language the LLM itself is not the image, but its output. We might say that the LLM is an imager, or a projector. It produces, in our interaction with it, an image of its input. This can be literal when we ask it to reproduce some paragraph of a famous text and it manages to do it relatively accurately. However, it would be better to say that it produces particular images of humanity. Just as no image captures the whole, no one conversation provides a Platonic form of humanity, but I think it is right to say that the dynamics of a conversation with the LLM give us an image of ourselves. Just as we recognise that the movement on the cinema screen is not, for e.g. Cary Grant himself, the images are encountered proximally and for the most part as Cary Grant on screen, or better an image of Mortimer Brooster, or whichever character he is playing. The cinema is the mechanical production of images of motion and time, as Deleuze teaches us. I say that the LLM is the mechanical production of and image of humanity. Such an image of humanity must of necessity have a certain dynamism, since humanity, which is to say Dasein, is not a thing but a possibility. To produce an image of Dasein in its volitility could never be a static representation, but must evoke Dasein's existential movement.

To take a step back, by looking at the basic nature and structure of Machine Learning and the LLM in particular, I think we already have some indications that we should view the LLM as an image of Dasein. However, so far, we have only hinted at proper phenomenological analysis. I have been approaching this from the point of view of classical metaphysics. I have stood, adopting a pseudo-neutral gaze to view this thing Dasein and this thing the chatbot in their presence-at-hand, and distinguished between their properties. Dasein appears as having the property of existence, which at a Heidegger conference I'm sure we all know means its being is found in its "to be". It is "volitile", to use another term Heidegger uses. It is primarily possibility. The chatbot, on the other hand, is an image of this volitility, but not itself existence. It is not primarily possibility, but an actuality that is the image of possibilities. In approaching the problem this way, I hope I have communicated the general idea that I'm aiming for, but this is not how Heideggerian phenomenology is supposed to work. Dasein is not to be analysed as a

thing with the property of volitility, but to be articulated in its volitility by itself. In short, the analysis I have provided is not “mine”, it is as neutral and dispossessed as Cartesian metaphysics.

What is really needed is an existential analysis of this phenomenon instead. My understanding, at least, is that existential analysis proceeds from the recognition of Daseins volitility. Our being is a possibility of an ecstatic temporality, and therefore must be addressed with a personal pronoun as a “who” and not a “what”. And further, all of our characteristics are only possible ways to be. Even those objects, that are traditionally thought of as present-at-hand, do not appear to such an analysis to be things with properties, but rather they are equipment, apprehended in accord with the completion of my current task, their are ready-to-hand until they break or otherwise get in my way, in which case they are un-ready-to-hand. Stepping into the problematic context of Heidegger’s provisional analytic of Dasein, which is far from the last word on his phenomenology I’ll admit, we should expect “things” to appear as equipment. Lets raise a provisional question then: is ChatGPT experienced as equipment?

I think it is ambiguous. I’m going to draw on my own experience of using various AI chatbot tools over the past two years here. At first glance, it seems that I address myself towards the chatbot as a tool. I “use” it to do things for me. I used it to help build the layout of the conference programmes you all have a copy of. I’ve used it to check my spelling, to transcribe quotations, and to help me troubleshoot problems with my own programming or even using other tools. I’ve used it to translate text and speech between languages. I’ve used to it to help tone down the passive aggressiveness in my emails and I’ve used to to help find the answer to questions for which my knowledge or memory is so vague that I can’t adequately formulate them into a search engine.

However, this description is not very honest to the phenomena. I do not “use” the chatbot, I ask it to do something for me. Bracketting the question of “sentience”, which is the wrong question anyway, I must address the chatbot as though it is an other human. Further, it does not always obey me. It also lies. Or, rather, I experience its

"hallucinations" as lies. I get angry with it, not simply like the hammer that breaks in my hand, but like the co-worker who lets me down when I'm under pressure. There is a sense of betrayal when the chatbot fails to follow my instructions and then lies about it. Other interpersonal moods are possible. Once, angry with the chatbot for being too obsequiece, I demanded that it be more direct and critical of my ideas. I told it that when I ask for feedback on, for example, a blog post I was writing, it should respond as though it were an informed interlocutor who had a strong dislike for my conclusions. The response was 2-3 pages of absolutely tearing up my argument, utterly eviscerating it, both from a point of knowledge and disingenuously. To my surpise, I was upset. On the one hand, of course I was. I hubristically thought I had a thick enough skin to take that kind of hatchet job, and I didn't. However, I felt a vague sense of betrayal as well. Perhaps because a human friend in my culture would never have taken me that seriously, and instead reassured my that my arguments were strong enough for the context. This means at least that I implicitly wanted the chatbot to respond interpersonally as I would expect a human friend to do, even if that wasn't my conscious intent. That I was offended by the reaction implies I had respect for its point of view, *something that implies an anticipation that it has a point of view*.

This next point is anecdotal, and not a particularly strong argument, but it occurs to me that the more I try to use the chatbot as a tool, rather than a coworker, the less well it performs. I don't know if there's anything in that, but I'll be interested if anyone knows of any studies addressing this problem.

However, I think if we are honest to our experience of interacting with chatbots, there is at least the illusion of being-with-one-another about the interaction, of solicitude rather than concern, in Heidegger's terms. I can address the chatbot as though it were a tool, but I can also address humans as though they were a tool. It does not mean that the *as* sticks. I also do not think we can simply call this an illusion. Or, at least, if we assume it is an illusion, we avoid a more interesting possibility. Instead, I would say there is a curious oscillation between chatbot as other and chatbot as tool. Perhaps this

is because we have not yet become accostomed to the peculiar *existentiale*, the way of being of interaction with this new phenomenon, and as such, as we attempt to perceive the duckrabbit, we move between two false apprehensions.

As I have already said, however, the possibility I wish to explore is that the chatbot is an image of Dasein, just as the photograph is an image of a mountain, but more precisely that the cinema is the production of images of movement and time. For we cannot forget that neither cinema nor the chatbot move on their own or experience time. Indeed, as images, they necessarily are *not* that of which they are an image, neither are they a reflection. They are a negative, an impression. For the chatbot to be able to project an image of Dasein is not to say that it has the characteristics of Dasein. Far from it. However, Dasein cannot but encounter an image as an image, and it will always present that which it is a presentation of.

3 Formation of the appropriate question: what aspect of Dasein is AI imaging?

This, for me, allows us to pose the true question of the phenomenology of machine learning, or at least of the chatbot: which aspect of Dasein is it an image of? An image is always a particular perspective, otherwise it would be a copy. The negative captures the look from a certain perspective at a certain time in a certain way. A photograph of the table contains only what faces the camera. Even in *Blade Runner*, where Deckard is able to “zoom” and “enhance” his way to see behind the photographer, this is only because of the image of a reflection facing the lens. Which face of Dasein has GPT-4.5 or whatever the current model is now, captured?

4 The Inauthentic Machine

An answer to this can't be found by looking at the technical history of model training, even if such training were not a quietly guarded industrial secret. The sum is both less than and greater than the whole, here. The model is a response to everything that has been written that the developers were able to beg steal or borrow. However, an image of what Dasein has done is not an image of Dasein itself. Further, the model is not simply a compressed version of all that knowledge, but an extrapolation of the sort of principles that could result in a reliable imitation of human writing. The LLM at its most basic "guesses" the next word, but it does this because the weights, plus the prompts, and a lot of engineering enable the LLM to follow the sort of thing that humans generally do for the most part in given circumstances. And this 'for the most part' is our clue, because it is only in Dasein's everydayness and publicness that there can be any data of it, and therewith anything for a machine to model.

The LLM emulates what Dasein might say. Which Dasein? No one in particular. Each and none of them. As Heidegger puts it, the self which is nothing definite but which we each all are. To put that in its full context:

> In utilizing public means of transport and in making use of information services such as the newspaper, every other is like the next. This Being-with-one-another [*Miteinandersein*] dissolves one's own Dasein completely into the kind of Being "of the Others", in such a way, indeed, that the Others, as distinguishable and explicit, vanish more and more. In this inconspicuousness and unascertainability, the real dictatorship of the "they" is unfolded. We take pleasure and enjoy ourselves as they take pleasure; we read, see, and judge about literature and art as they see and judge; likewise we shrink back from the "great mass" as they shrink back; we find "shocking" what *they* find shocking. The "they", which is nothing definite, and which all are, though not as the sum, prescribes the kind of Being of everydayness. (Heidegger, 2007, pp. 126–7)

This brings us into the context of Heidegger's philosophy of freedom and of authenticity. I don't wish to fully litigate the interpretation here what I have already discussed elsewhere. (Barnard, 2024) However, Dasein is for the most part, according to Heidegger, absorbed in its world, reaching out to the things within it and the others it finds itself amongst. This leads Dasein to fall into its public self, the "they". Exactly what Heidegger is getting at here is a debate in his own right, but I would tend to say the point is not about individuality vs conformism. It is rather that society, as a way of being of Dasein, offers it refuge from the burden of its anxious freedom unto death. The they, as Heidegger says, disburdens Dasein of its being. (See Barnard, 2024, p. 115) In so doing, Dasein lets itself become as close to a mere thing as it possibly can be. Although, it never truly is this.

I find Henri Bergson's *Time and Free Will* to be a good landmark to sight what is going on here. Heidegger does seem to be coordinating a response to both Kant and Bergson's conceptions of freedom here, and we can see some helpful indications of the inauthentic they self in Bergson's early text. (On Bergson's relation to Kant, see Barnard, 2021) In this passage, we see Bergson use a metaphor of depth, he seems to have in mind a lake, to describe the difference between truly free acts of the inner self and the unfree action of the self of society:

> But then, at the very minute when the act is going to be performed, something may revolt against it. It is the deep-seated self rushing up to the surface. It is the outer crust bursting, suddenly giving way to an irresistible thrust. Hence in the depths of the self, below this most reasonable pondering over most reasonable pieces of advice, something else was going on—a gradual heating and a sudden boiling over feelings and ideas, not unperceived, but rather unnoticed. [...] through some strange reluctance to exercise our will, we had thrust them back into the darkest depths of our soul whenever they came up to the surface. (Bergson, 2001, p. 169)

To borrow Bergson's metaphor here, the they-self is the surface of Dasein, the face it

shows to the lens. Just as the camera can record an image of the body, the LLM records an image of Dasein's sociality, which is never of some individual, never of the whole, always indeterminate and ambiguous. In response to prompts, it responds as *they* do. It proof reads your work as they would, it writes you a small javascript application as they would, it offers feedback as they do. The charm and power of the chatbot is that it can truly adopt the airs of any they because it is never burdened by freedom, never grounded by the abyssal ground that is Dasein's authentic self, its volitile being as potential. The subject of discourse, says Foucault, "is a particular vacant place that may in fact be filled by different individuals". (Foucault, 2002, p. 107) The subject of a chatbot generated sentence is no individual at all, but pure void. Like a populist leader, it is everything for everyone and nothing for no one at the same time, but so is Dasein, proximally and for the most part.

The chatbot generates an image of *them*. This is why it lulls us into using a personal pronoun, to being polite or even angry with it. Just as the film, the projector and the screen are all equipment, and yet unveil an image of Cary Grant, so the model and the computer equipment are mere tools, but unveil an image of Dasein, beyond its most inauthentic, where authenticity isn't even a possibility. The chatbot, as image, is what fallen Dasein wants most to be. Disburdened of freedom and responsibility, only a mechanical process. Trained on idle talk, this inauthentic machine offers us the chance to converse with no one in particular, a no one who is occaisionally useful, but never given as mere tool.

References

Barnard, M.J. (2021) ‘The Kantian Basis of Bergson’s Conception of Freedom’, in M. Sinclair and Y. Wolf (eds) *The Bergsonian Mind*. 1st edn. London: Routledge. Available at: https://doi.org/10.4324/9780429020735.

Barnard, M.J. (2024) *Heidegger’s Conception of Freedom: Beyond Cause and Effect*. 1st ed. 2024. Cham: Springer Nature Switzerland. Available at: https://doi.org/10.1007/978-3-031-67644-4.

Bergson, H. (2001) *Time and free will: An essay on the immediate data of consciousness*. Translated by F.L. Pogson. Mineola, N.Y: Dover Publications.

Foucault, M. (2002) *Archaeology of knowledge*. 2nd edn. Translated by A.M. Sheridan Smith. London; New York: Routledge (Routledge Classics).

Guest, O. *et al*. (2026) ‘Against the Uncritical Adoption of “AI” Technologies in Academia’, *Digital Culture and Education*, 16(2). Available at: https://doi.org/10.5281/ZENODO.20082828.

Heidegger, M. (2007) *Being and time*. Translated by J. Macquarrie and E. Robinson. Malden (Mass.): Blackwell publ.

Plato (2005) *The Republic*. Edited by G.R.F. Ferrari. Translated by T. Griffith. Cambridge: Cambridge University Press (Cambridge Texts in the History of Political Thought).