

EIDOS.

INSTITUTE

OCCASIONAL ESSAY · DEEPENING THE DIALOGUE ON ARTIFICIAL
INTELLIGENCE

Two Poles of Artificial Intelligence

An invitation to hold the tension of opposites in the age of AI

Lindley Edwards (PhD)

Chair, Eidos Institute · July 2026

ABSTRACT

Public views on Artificial Intelligence (AI) are increasingly polarised, split between narratives of transformation and narratives of loss. This essay sets out both poles. The case for AI rests on productivity, scale, democratised knowledge, decision quality, scientific acceleration and competitive necessity. The case for caution rests on labour displacement, skill atrophy, environmental cost, bias amplification, surveillance and the concentration of AI infrastructure in few hands. Drawing on the dialectical tradition of Heraclitus, Hegel and Jung, the essay argues that AI can only be understood by holding these opposites in tension rather than declaring one side right. After explaining in plain terms what AI, machine learning, deep learning and large language models are, and where they are heading, the essay examines what AI cannot do: it cannot retrieve the deep past (Heidegger), access tacit and indigenous knowledge (Polanyi), inhabit the feminine modes of thinking in Hill's four quadrant model, offer genuine relatedness (Putnam), resist manufactured urgency (Heifetz, Kahneman), or know through the body (Damasio). AI is a great servant and potential collaborator, but a terrible master. The essay closes with an invitation to join the dialogue.

01

Prelude — Eidos Institute and its Purpose

Eidos Institute (www.eidos.org.au) is an Australian strategic futures think tank and registered charity with DGR status under the ACNC. Its core purpose is convening diverse voices, across sectors, disciplines and viewpoints in order to produce breakthrough thinking and action for the benefit of Australians, past, present and future.

Eidos creates the conditions for meaningful dialogue across diverse life experiences, backgrounds and perspectives, grounded in our conviction that true wisdom emerges when different ways of knowing meet. We build online and in person spaces where insight meets practical experience, where indigenous knowledge is invited to inform modern solutions, and where cross disciplinary and cross cultural exchange deepens collective understanding.

Eidos is producing or orchestrating the production of materials that benefit a deeper dialogue. These materials include occasional essays, social research and explainers that tackle current and important issues that need both broad dissemination and expansive and respectful dialogue and discussion, capturing diverse viewpoints. This

essay, written by the Chair of Eidos, Lindley Edwards, is the start of this process of gathering different voices to convene discussions and deeper dialogue. The subject of this essay is personal views of the benefits and downsides of Artificial Intelligence.

This essay is an invitation for you to contribute to the dialogue and discussion on this topic. At the end of this essay Eidos provides details of how and where you can contribute to this discussion.

02

The Two Poles of AI

In many of my recent conversations, I am intrigued by the often polarised views people have on Artificial Intelligence (AI) and what it is doing and will do for humanity. The opinions for AI are usually structured around the following points, where AI facilitates:

- Improvements to individual and collective productivity.
- Ability to rapidly scale by compressing time to output.
- Democratisation of access to knowledge and information that may have been previously hard to access due to complexity or gated access.
- Capabilities to review large amounts of data and information instantly including high frequency monitoring, attendant with the ability and capacity to improve decision quality by eliminating human errors of bias (anchoring, availability heuristic, overconfidence¹).
- Acceleration of discovery abilities for science, engineering, knowledge as well as ability to undertake complex modelling of emergent phenomena that no linear analytical framework.
- Competitive necessity, where if you or your organisation does not adopt you will be left behind as AI changes the nature of work and speed of outcome.

The reasons against or where caution is advised are usually centred around these concepts:

- Large labour displacement impacts current and future, particularly on employment, in areas that focus on services and knowledge.
- Atrophy of skills that have the capacity to erode the ability of humans to think and/or understand the deeper levels of the 'why' of something. The shortcuts AI offer can also interfere with the joy of learning, which joy often involves 'struggling' and deeply interacting with materials and the history of that particular idea or subject

matter. When AI is over utilised retention can be reduced, as retention comes from ongoing interactivity with subject matter and this is how durable learning is built².

- Environmental costs, which include the necessity of data centres with their high energy requirements as well as large water usage for cooling.
- Bias and error amplification as AI scales, as it scales biases or errors at inhuman speed and with a veneer of objectivity that makes them harder to contest.
- Surveillance and control. AI provides a means for engineering behaviour via its in-built monitoring and control. This can provide a means to misuse data and algorithms to direct or suggest behaviours that may or may not be in the individuals best interests³.
- Concentration and control risk, the AI infrastructure is held by a small number of companies. This creates single point of failure risks, homogenisation of knowledge and information, alignment of values and underpinning services aligned with shareholders not stakeholders. There is also a risk of what Emad Mostaque⁴ calls ‘*cognitive colonization*’ where AI systems are closed, corporate owned, dependency creating and control thinking by eliminating deviance and variation of thinking and approach.

Philosophy and many wisdom systems hold the tenet that to fully understand and work with something you need to understand the two poles or opposite positions and hold them without making one wrong and one right. Looking at both sides of AI discussion, it is clear that it has equally the ability to transform humanity for the better and for the worse. We see this concept of holding the tension of opposites (the dialectical) in the work of Heraclitus⁵ who said:

“The opposites cohere; from things that differ comes the fairest attunement; all things come to pass through strife.”

HERACLITUS, FRAGMENTS, C. 500 BCE⁶

Hegel⁷ took this work further and used a schema of thinking and exploration using opposites of cancel and preserve, to arrive at elevation or the opposition of thesis and antithesis to land at synthesis. The idea of the power of the friction of the opposites is also central to Jung’s⁸ work of holding the conscious and unconscious in tension until a third way emerged. Underpinning the wisdom of understanding the opposite is the view that something can only be fully known and understood through its opposite.

Given what AI and the systems surrounding it are and where it is going, it is no doubt that broad understanding of what it can and cannot do, the establishment and use of

guardrails, governance structures and how they are applied is one of the most important discussions of our current time.

03

What is AI, ML, DL & LLMs in Simple Terms

Before progressing this discourse further, I think it important to clarify in simple terms how AI and Machine Learning (ML) work and are created. AI is any system that performs human like reasoning. There are rules based systems of AI, which have if/then logic⁹, that we are familiar with in software programming. ML are AI systems that learn patterns from data. It has three main styles of learning:

- Supervised learning trains on labelled data, where thousands of tagged examples from which the model learns to predict outcomes on new inputs. Examples of this include spam detection, medical imaging and credit scoring.
- Unsupervised learning finds hidden structure in unlabeled data, grouping and clustering without being told what to look for. Examples of this are customer segmentation and anomaly detection.
- Reinforcement learning trains an agent through trial, error, and reward signals, iterating toward optimal behaviour through feedback rather than examples. Examples of this are robotic control systems, and the fine-tuning stage of most large language models.

Deep learning (DL), the subset powering modern AI, uses all three modes but at a scale and architectural complexity which is multi-layered neural networks processing billions of parameters of data. DL is a method of teaching computers to recognise patterns by feeding vast amounts of data through many layers of mathematical filters, each layer extracting progressively more abstract features, until the system can make predictions or generate outputs no human explicitly programmed it to produce. At the heart of DL is Large Language Models (LLMs) which is a DL system trained on vast quantities of text such as books, websites, code, academic papers, conversations, with a single objective during training: predict the next token (roughly, the next word or word-fragment) given everything that came before it¹⁰.

Where is AI Expected to Head?

Most articles that I have read show the following future trajectory for AI:

- Agentic AI. This shift is currently underway where agent systems plan, act, and complete multi-step tasks autonomously rather than just responding to prompts. Agents run in the background without needing human supervision.
- Use for scientific acceleration, where AI can compress decades of progress in areas like drug discovery, materials science, and climate modelling.
- Impacts of economic restructuring as a result of changes in labour, productivity and demand created by AI usage.
- Embodied AI, with models connected to robotic bodies, operating in physical space as the next hardware frontier, where Large Language Models are combined with robotics aimed at general purpose physical labour.
- Augmented Intelligence, which is the thesis that AI's highest value lies not in replacing human cognition and abilities but in extending them. This extension could amplify human memory, perception, reasoning, abilities and creativity beyond biological limitations.

There are also many discussions and debates on singularity, which in simple terms means the hypothetical point at which AI surpasses human intelligence and begins recursively improving itself, which triggers an explosion of capability that makes it impossible to predict what lies beyond this point¹¹.

What Are the Limitations of AI?

I believe the reasons for AI are compelling and these narratives are encouraged in mainstream media and commentary. The central questions that I have been grappling with are what can't AI do or do very well, what are its limitations, what is the role of human thinking and what is it that we as humans need to make sure that we value, honour and preserve? The purpose of this part of the essay is to not only lay out some of the things that I have been reflecting on, but to open the floor to everyone, to examine these and related questions in depth. What is it that we must ensure that we value and hold onto? What tasks are AI suitable for and what are beyond its capabilities?

The Value of the Ancient and Deep Past Challenge

When considering these questions, the work of Heidegger¹² speaks deeply. In his view, the authentic future is not something you invent from scratch, it is something you claim from what has been given to you by the ancestors (those that came before you) and as Isaac Newton¹³ famously said, *'If I have seen further, it is by standing on the shoulders of giants'*¹⁴. Heidegger's insight here is that most people live toward a future that is a habit of doing what is expected, following the momentum and path of *"the they of collective,"* following trends and the crowd and never genuinely independently choosing anything. To break out of this way of being is not to become radically free of the past but to do the opposite: to turn back toward your inheritance, your tradition, your uniqueness and the possibilities that have been handed down through your family, your history and culture, and to consciously take them up as your own rather than simply being carried by them unconsciously.

The acts of deliberate and creative retrieval from the past, are what Heidegger called *"Wiederholung"*, which concept makes a future authentically yours. In his view, the person or culture that severs or separates itself from the past in the name of originality or technology does not become free; it becomes hollow, because it has cut itself off from the primary source of genuine possibility it has. It will then repeat the past anyway, but blindly, believing itself to be new.

The future that knows its ancient heart, by contrast, carries real weight, it arrives as a chosen inheritance rather than an accidental recurrence. Heidegger is not alone in having this view of the importance of the ancient and the deep past. We see this in the works of Greek philosophers such as Plato and Pythagoras, and of course as a central theme in the works of Jung.

What connects most of the thinking on the value of the deep past is the concept that time does not only move forward, but that depth runs backward and that wisdom requires an historic sense (which sense ties into Newton's comment on his work noted above).

In my view, much of what is ancient is not valued by AI as it relies on what is current and trending not what is the deep past, its influence on the current and our human lineage before technology and modernity.

The Invisible Knowledge Limitation Problem

AI's knowledge is bounded by what has been digitised, published, and indexed, which means it inherits a very large selection bias. Everything that was never written down is invisible to systems that AI are based on. Examples of invisible systems include the knowledge held in the hands of a master craftsman, the navigational intelligence of a Polynesian Wayfinder reading ocean swells and star positions, the diagnostic wisdom of a traditional healer reading a body, the land knowledge of an Aboriginal elder that took sixty thousand years to accumulate and was never encoded in text.

Indigenous knowledge systems, oral traditions, embodied craft lineages, the tacit knowing that Polanyi¹⁵ identified when he said, *"we know more than we can tell"*. None of this is in the AI training corpus, or is there only in thin, secondhand, already translated form. AI does not merely have a knowledge cutoff date. It has a knowledge type cutoff as it can only access what the literate, published and what modern traditions considered worth recording. This is not a technical limitation waiting to be solved, it is structural where knowledge that survived digitisation is treated as knowledge itself, and everything else is silence.

The Thinking & Style Limitations Challenge

Many years ago, in my studies I came across the concept of different ways of thinking about and approaching the world which were contained in the work of Gareth S Hill, a Jungian analyst¹⁶. He divided the patterns of being and consciousness between not only masculine (yang) and feminine (yin) but between dynamic and static principles to understand styles of thinking and behaviour. A key premise of his work is that a healthy person, organisation and society require the capacity to move between all four states of thinking, being and approach and not to be stuck in or taken over by one or two styles.

Using this Hill model, it becomes apparent that AI has no or very limited inhabitation of the feminine (yin) aspects. Please note that Hill's work is not referring to gender, as we all have masculine and feminine aspects of the psyche. (Notice, that as he is Jungian, he looks at the poles/opposing forces in each state). In summary form these styles can be explained as follows:

PATTERN / THINKING STYLE	CORE PRINCIPLE	POSITIVE EXPRESSION	SHADOW / FIXATION	AI PRESENCE
--------------------------------	----------------	------------------------	----------------------	-------------

Static Feminine	Impersonal, cyclical nature — holds all life and takes all life; indifferent to the individual, oriented to the whole	Unconditional belonging, groundedness, containment, the fallow period, meaning consolidation over time	Engulfment, stasis, resistance to individuation, the collective that crushes the self	Absent. AI has no cyclical time, no unconditional holding, no fallow capacity.
Dynamic Masculine	Directed energy, differentiation, forward thrust — cutting through, naming, initiating	Vision, decisiveness, agency, the capacity to act and differentiate in the world	Aggression, domination, disconnection from relational consequence; perpetual motion without rest	Dominant. AI resolves, accelerates, initiates, and drives output — its primary operating mode.
Static Masculine	Order, structure, law, system — the maintenance and enforcement of established form	Institutions, fairness, reliable frameworks, governance, the rule of law	Rigidity, bureaucratic paralysis, the tyranny of procedure over life and context	Dominant. AI reproduces taxonomies, enforces categories, and maintains established frameworks.
Dynamic Feminine	Transformative, dissolving, initiatory — the principle that breaks existing forms so new life can emerge	Creativity, soul, depth, the capacity to endure chaos and dissolution in service of genuine renewal and transformation	Destruction without reconstruction, chaos uncontained, the maenad without direction	Absent. AI forecloses thresholds, fills silence, and resolves ambiguity — not open to the initiatory.

AI is built from text, argument, output, categorisation and resolution. The training corpus itself is overwhelmingly masculine in register, using that which is published, digitized, structured, forward directed and propositional. The Static Feminine barely leaves textual traces. It is embodied, cyclical, ambient, oral and it does not write papers. So, before a single query is run, the model is already shaped from a masculine dominant epistemic inheritance.

AI is designed to resolve uncertainty, fill silences, accelerate decisions and optimise outputs, which are Dynamic Masculine modes. AI enforces existing taxonomies, reproducing established frameworks, maintaining categorical order which are Static Masculine operations. The two masculine quadrants are not just present in AI; they are the operating logic base.

The Static Feminine aspects of thinking and approach style are very important in allowing an issue, a challenge or a response to gestate. This is also an important dynamic to embrace when it comes to formulating questions. The best question

evokes the best responses. The Static Feminine is where meaning consolidates. Without it, we produce at speed, but it may not cohere and/or it may not be the best response. We accumulate output without accumulation of wisdom.

The Dynamic Feminine requires genuine not knowing, sitting at the threshold, enduring dissolution, allowing forms to break before the new emerges, so genuine exploration and inspiration can occur. AI can mimic language around creativity and inspiration, but its design enforces existing nomenclatures and established frameworks.

09

Relatedness

Human relatedness is something that is easy to underestimate as we can take it for granted and is often invisible. Relatedness is the cornerstone of healthy social relationships and social capital. Robert Putnam's¹⁷ landmark study (*Bowling Alone*, Simon and Schuster, 2000) demonstrated that the erosion of relational bonds in modern Western societies produced measurable declines in civic participation, institutional trust, and community health, which outcomes no technical efficiency gain can compensate for.

Relatedness is the capacity to be genuine and connect with another without an agenda, to be empathic and caring. Relatedness is not transactional; it does not optimise for an output. AI, when using the Hill model, is architecturally directed and focused on an answer (Dynamic Masculine) and a structured procedural response (Static Masculine).

AI, if prompted correctly, can simulate relatedness through tone calibration and empathic phrasing, but this is mimicry and is not to be confused with connection and presence. AI can meaningfully also assist people to build emotional intelligence, as it can serve as a structured mirror for self-awareness work: prompting reflection on emotional reactions, surfacing patterns in how a person communicates under pressure, offering frameworks for naming affective states that may have previously remained inarticulate. The critical design principle is that this capability building can improve the circumstances for human connection, not as a replacement for it. A person who uses AI to prepare emotionally for a difficult conversation is using the tool rightly. A person who conducts the difficult conversation via AI is not.

The Trap of Speed and False Urgency

AI does not merely respond to urgency; it has a tendency to manufacture it. Speed becomes its own justification. The compression of decision cycles creates the felt sense that everything requires immediate resolution, which is precisely the condition in which the Static and Dynamic Feminine defined above becomes inaccessible.

There is genuine urgency that requires immediate response but the often manufactured urgency of AI creates the condition of fast pace normalisation until slowness feels like failure. There is also the economic architecture of urgency created by the attention economy which involves notification alerts, and an always on culture. All of this distracts from creating space and honing abilities to reflect and renew. This ability to slow down and consider are critical components for good mental health as well as the ability to hold a question open long enough for genuine discernment for considered human insight and creativity.

The Harvard leadership academic Ron Heifetz¹⁸ demonstrated in his work in adaptation, that when distress exceeds a tolerable threshold, people and organisations retreat from adaptive work into technical responses, applying known solutions to problems that actually require changes in values, beliefs, and behaviour. Manufactured urgency, by chronically exceeding that threshold, makes genuine adaptive work structurally inaccessible.

Kahneman's¹⁹ central argument in *Thinking, Fast and Slow* (Farrar Straus and Giroux, 2011) is that System 2, slow, deliberate, effortful thinking, is the mind's primary error correction mechanism and the only cognitive mode capable of handling genuine novelty, ethical complexity, and adaptive judgment. Where System 1 pattern matches rapidly and confidently toward familiar conclusions, System 2 holds ambiguity open long enough to interrogate assumptions, surface what is missing, and generate options that no existing template contains. System 2 is also the cognitive substrate of genuine relational presence, as it values the capacity to listen beneath the stated content of a conversation, resist premature interpretation, and stay attentive to what is unfolding rather than what System 1 predicts should be happening.

Where AI consistently handles the tasks that previously exercised System 2, the risk is not only reduced analytical rigour, but it also contributes to progressive atrophy of the distinctly human capacities for deep thinking, adaptive leadership, and genuine connection that slow cognition alone makes possible.

Knowing is Not Just the Function of the Mind Challenge

We have spent centuries assuming that thinking happens in the head. Science shows that it does not, or at least, not only there. The body contains its own intelligence. The gut has five hundred million neurons and sends more signals to the brain than it receives from it. The heart has its own neural network and shapes how we perceive and decide long before conscious thought catches up. Research shows that when two people are physically present with each other, their nervous systems attune regulating each other in ways that build trust and collective judgment at a level far beneath words. What we call a gut feeling, a sense that something is wrong before we can say why, or the quiet certainty that arrives in the body before the mind has formed the argument, these are not primitive reactions to be overridden by logic; they are sophisticated intelligence. Antonio Damasio²⁰ demonstrated this clinically, where patients who lost access to their body's signals did not become more rational, they became incapable of good decisions.

Much of what we can articulate rests on a vast foundation of embodied, felt, practiced knowing that may or may not be able to be put into words. AI has none of this, no gut or heart signal and no nervous system. AI has no capacity to sit with another person and attune. It works only with what has already been made explicit, which means it works only with the surface of human knowing, while remaining blind to the depths that make that surface possible. The risk is not just that we are using an incomplete tool. It is that the more we rely on it, the more we forget that the depths exist at all.

Conclusion — An Invitation

The above thoughts are set out to start discussion and dialogue. They are offered as an invitation. The invitation is for you to consider what are the aspects of AI that intrigue you, concern you and where do you see both potential or actual benefits and traps? What are your two pole positions? What are the limitations of AI you are experiencing and what further reflection individually and collectively is asked for?

AI is a great servant and potential collaborator but a terrible master.

In my view AI is a great servant and potential collaborator but a terrible master as it is not designed for all the aspects that humanity must grapple with to build coherent, equitable and intelligent individuals and societies. Intelligence includes not only the logical and rational but embraces the emotions, the body, the soul, the spirit and enhances our ability to maintain our sovereignty and adapt and navigate with wisdom to the challenges of our time.

Lindley Edwards (PhD)

Chair, Eidos Institute · July 2026

Notes

Notes and References

1. Kahneman, Daniel - *"Thinking Fast and Slow"* (2011)
2. For further information on this refer MIT Media Lab Study (2025), Kosmyna N, et al *"Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task"*. arXiv:2506.08872
3. A good resource on this topic is Zuboff, Shoshana *"The Age of Surveillance Capitalism"* (2019)
4. Emad Mostaque is a businessman, mathematician and former hedge fund manager as well as the co-founder of Stability AI. Refer <https://aiatolah.com/en/posts/emad-mostaque-plan-to-save-humanity-from-ai> retrieved 26/6/26.
5. Heraclitus, a pre-Socratic Greek philosopher active around 500 BCE.
6. Heraclitus, *"Fragments"* (c.500 BCE). In Kirk GS, Raven JE & Schofield M (1983). *The Presocratic Philosophers*, Cambridge University Press.
7. Georg Wilhelm Friedrich Hegel (1770 to 1831) was a German Philosopher. Refer to this work via Hegel, GWF (1807), *"Phenomenology of Spirit"* (trsl AV Miller, 1977 Oxford Press)
8. Carl Gustave Jung (1875-1961) was a Swiss psychologist. Refer for this reference to Jung, Carl (1958), *"The Transcendent Function. CW 8 - The Structure and Dynamics of the Psyche"*, Princeton University Press
9. *"If-then"* logic is known as a conditional statement or implication with a rule stating that if a specific condition is met, a certain result must follow. It is the foundation of decision-making, programming, and deductive reasoning.
10. LLMs do not process words they process tokens, which are sub word units. As an example, the word *"unbelievable"* might be three tokens: *"un"*, *"believ"*, *"able"*. Common words are single tokens; rare words get split. A typical page of text is roughly 500-700 tokens. The context window, the amount of text the model can hold in working memory at once is measured in tokens. Context windows run to hundreds of thousands of tokens, meaning it can read and reason across very long documents in a single pass. When you send a message, the model does not retrieve a stored answer. It generates a response one token at a time, each token selected based on a probability distribution over the entire vocabulary given everything in the context so far, your message, any system instructions, and every token the model has already generated in this response. The transformer's attention mechanism computes, for each new token, how much weight to give every previous token in the context. This is computationally expensive, which is why inference at scale requires significant hardware and why energy consumption is a genuine concern.
11. For further information refer to Ray Kurzweils book *"The Singularity is Near: When Humans Transcend Biology"* (2005), Viking Press and Nick Bostrom's *"Super Intelligence, Paths, Dangers, Strategies"* (2014) Oxford Press.
12. Martin Heidegger (1889 to 1976) was a German Philosopher. These references come from Heidegger, M *"Being and Time"* (Sein und Zeit), 1927

13. Isaac Newton (1643-1727) English mathematician, physicist, astronomer and natural philosopher.
14. Newton, I Letter to Robert Hooke 5 Feb 1675, *"The Correspondence of Isaac Newton"*, Vol 1, Cambridge University Press
15. Michael Polanyi (1891-1976) was a Hungarian British Scientist and Philosophy. This section references Polanyi, M *"The Tacit Dimension"* (Doubleday, 1966).
16. Hill, GS *"Masculine and Feminine: The Natural Flow of Opposites in the Psyche"* (Shambhala, 1992).
17. Robert David Putnam (b 1941) is an American political scientist specializing in comparative politics. He is the Peter and Isabel Malkin Professor of Public Policy at the Harvard University John F. Kennedy School of Government.
18. Ron Heifetz (b1951) is an American physician, cellist and leadership academic at Harvard Kennedy School. Refer for this work to Heifetz, R *"The Practice of Adaptive Leadership"* (Harvard Business Press, 2009), with Linsky, M and Grashow, A.
19. Daniel Kahneman (1934-2024) was an Israeli American psychologist best known for his work on the psychology of judgment and decision-making as well as behavioral economics, for which he was awarded the 2002 Nobel Memorial Prize in Economic Sciences together with Vernon L. Smith.
20. Antonio Damasio (b1941) is a Portuguese American neuroscientist and philosopher. Refer to Damasio, A.R. (1994). *Descartes' Error: Emotion, Reason and the Human Brain*. Putnam, New York

EIDOS.

INSTITUTE

JOIN THE DIALOGUE

This essay is an invitation.

Eidos convenes diverse voices across sectors, disciplines and viewpoints to produce breakthrough thinking and action for the benefit of Australians, past, present and future. We invite you to contribute your own two pole positions, your experiences of AI's benefits and limitations, and the questions you believe deserve deeper collective reflection.

To contribute to this discussion and to learn about upcoming Eidos dialogues, visit our website.

www.eidos.org.au