

Epistemic convergence

by Aidan Cornelius-Bell

Published: 2025-08-07

Dear friends,

We face a time, collectively, where language is important. Where critical information literacy is fundamental to survival. Where seeing through the bourgeois narrative can shake suffering. These critical skills, which educators have advocated for decades, are at ever growing risk. And there's some irony, to me, in technological space worth examining.

Today I want to talk about Artificial Intelligence –specifically large language models, and how by their nature, they are beginning to shape our thinking, foster epistemic convergence, and enforce childlike patterns. But to do this I need to get slightly under the hood, building on from our examination of the epistemological and ontological frameworks of the internet. But fear not, I'm not going deep down the philosophical rabbit hole today –just some practicalities to attend to so we have a critical basis for analysis.

Large language models are a branch of current artificial intelligence techniques which seek to ‘understand’ and generate human language through statistical pattern recognition. Unlike other AI systems which have been targeted towards, for instance, visual recognition, robotics, or strategy games, large language models are specifically trained on collections of text harvested from the open internet to learn the probabilistic relationships between words and phrases. They are prediction engines which can complete sentences, ‘answer’ questions, and guess at dialogue by drawing upon patterns identified across this training data. They are one approach within a broad field, but have become synonymous with AI and even AGI, due to their supposed versatility in handling language-based tasks and their ability to demonstrate seemingly coherent reasoning through text prediction. And they are, technically speaking, pretty cool, and offer reasonable solutions to some problems.

But humans are not prediction machines. Actually, we tend to be quite bad at it. And we need to remember our humanity, our capability, in the face of ‘novel’ intelligence. We have, literally, dedicated science to attend to both the limits of our thinking, and the quantification of our worlds. Impressively, as a species, we have developed many mathematical models, a great deal of complex language, and a plethora of communication techniques. Collectively, between mathematics and language, we have abstracted into written sketches, datasets, algorithms, and so on, a great swath of information which usually attempts its best at the time to represent our world. So, where does AI come in?

Current AI draws on what is on the internet. In particular, for training, the internet until around 2022. Being generous, that internet represents a fraction of what humanity has surveyed and converted into words and numbers for about 30 years. Historical records, books, images and more are also part of the corpus of these internet things, and it all forms part of the training data for AI. The internet has been, essentially, the largest networked library of all human history. While it far from represents everything about humanity, it remains the place of record for a great many things, thus becoming a logical start point for crafting an ‘intelligence’. Moreover, as we modernise systems and ways of thinking, the internet also creeps from library (‘read’) into governance (‘write’) and trusted systems digitise to manage both citizens through digital technologies, and the technologies themselves. Here, interfacing with technology has become a fundamental science. As has critical thinking and information literacy.

Rather than own responsibility for building citizens’ critical thinking or information literacy skills however technologists have continued to change interfaces for information. From the command prompt, through to speaking, watching, listening AI. The way we interact with

information has changed forever, over and over again. And the information itself, as it is currently produced and related to, is fundamentally changing. And with it a fading of the onto-epistemic space for divergent thinking. But let's not get ahead of ourselves. There's a trope of my writing emerging...

Can you feel it yet? We have layers.

The internet, to date, has consisted of records (information) and processes (exchange methods). These are governed by the laws of physics and laws of information. Increasingly, they have also been governed by laws. These are typically driven by governments looking to control information, often while masquerading under Lovejoyesque protestations of "won't somebody please think of the children!". But arising from this epistemic layer, is AI. And it brings its own onto/epistemic reality.

LLMs are trained on our words. They are trained by our processes -the ones that are online, that is. They use a synthesis of both the information available on the internet (which has largely been pirated by these corporations) and the ways (tools) we have had to produce and transmit that information. LLMs are an algorithm -a predictive engine -building on that synthesis of words and processes. At best, this offers a vague approximation, literally a guess, at human thinking. Together these two layers, LLM qua machine (epistemology) and LLM qua internet corpus (onto-epistemology), are communicated to us by a proxy as human 'cortex' and 'memory'. The implication being that, like LLMs, our brains contain information resources which react to predictions to be spat back out at the world. While this is prescient background for anyone on the modern internet, it is not fundamentally what I think we need to attend to.

The problem is that the training data, the fuel used for the AI, is abstracted from our view. Naturally, any 'intelligence' which we currently recognise can only incrementally outgrow its environment (i.e., source material). But we aren't even thinking about the source material or the algorithm, we just interact with it like it's a human - and probably not very politely. Moreover, we know, full well, that what is written online does not represent the best of humanity. Some parts may be great, but almost everyone can agree that there are disreputable parts, objectionable sources, and utterly contemptuous contents out there. We just won't necessarily agree on what those are. Setting that aside, regardless of our opinions, this internet which comprises many varied interpretations of reality is used to train responses that resemble intelligible answers to our prompts. This means, and, sorry, my urge to resist writing ipso facto has failed, that LLMs are trained on at least some objectionable content, and that these contents are, themselves, an abstraction. This is starting to feel like real life Inception.

This layering, creating distance between us and the politics and power, doesn't end with training data in all its forms. It is also present in the training itself. This is equally a human endeavour in the first instance, and abstracted from the user. People are (crudely put) rapidly flicking through thousands of answers to prompts: yes/no, to indicate if it *seems* like a good enough response. The training data itself is based on an ungodly behemoth amalgam of opinions, information, art, and more. Plus, the training is human, requiring at least some intervention to support the algorithm's development. And here lies my gravest concern. It feels wild that it took me 900 words to get here.

The highly abstracted nature of AI: training and training data; constructed atop an internet of vastly divergent opinions and a fair amount of crap, means that the responses need some tailoring to be intelligible.

To be intelligible is to be political.

All data, inherently imbued with a way of understanding the world, available online comes from a hegemonic epistemology. Moreover, just because there is divergence of opinion on

the internet, and some of it quite gruesomely lacking in empathy, does not mean the model is trained to recognise opinion, epistemology, or politics. The people doing the training are not incentivised to care about this, and even if they were they aren't likely to have the time to care. And their 'product' is not human. Corporate's AI training hires are not selected for critical literacy skills. Even if they were able to somehow negate their own politics and positionality, promoting responses on the basis of metacognition would be an insurmountable problem. Yet, we've seen that AI can be modified to political extremes, away from the hegemony, even if done to the extreme right with crayon and scissors.

It has been done by Elon Musk to make Grok into a Nazi.

This should ring alarm bells for a multitude of human rights reasons. Moreover, it should be obvious that there are serious political ramifications to AI outputs, and their use to undertake anything resembling decision-making. A consequence of the predictive engine, the heart of current AI, is *ipso facto* (because of the fact of itself) smoothing of entropic political divergence and polyvocality into epistemic convergence *qua* (in the capacity of) hegemonic speech is what your brain desperately wants to call the AI's consciousness. But it isn't conscious. It is regurgitating words, culture, language, politics, and ideas as a standard deviation.

A representation drawn from a standard deviation of social and political opinion is hegemonic. It is the just a "little bit" right of centre mainstream neoliberal capitalist worldview.

What AI responds, oftentimes just padding your original prompt text with more words, is mainstream rhetoric. Derived from, shall we say, not the best of humanity. It represents back, and this is increasingly the case, a generalised, 'smoothed', guess. It throws in phrases. It adds key words. And as the number increases: 3, 3.5, 3.7, 4, 4o, 4.1, 5, the keywords and phrases change. But the pattern is the same. Because it is just an algorithm. An algorithm which, because of its very being, promotes individualism, an inhumane normativity, and regressive politics which harm humanity. As with Meta, and Newscorp before them, corporate AI irons out the crinkly, inconvenient edges of human opinion, existence and experience. It throws back a representation which is probable to engage you, and it reinforces centrist views. The latest trick is just that it's convinced you it's a human-ish.

The fundamental and most challenging problem, here, is that the algorithms used are smoothing diversity of opinion, language, theory, and knowledge in *us*. And when you let it write, read, or summarise for you, you slowly lose your analytical capability, you receive a smooth, digestable, convenient output. You ascribe intelligence to it. And you get a tiny dopamine hit as 'you' complete a task.

By design it is not cognisant. It was not given literate and meaningful training data. It has been given a hegemonic canon, trained by traditional intellectuals (in the Gramscian sense), and then rewarded for giving you pleasing responses. These responses are often passably apolitical but always hegemonic. It just throws words at you which you think resemble intelligence. And I don't begrudge AI, or its users, indeed there are many uses for the current LLM technologies. But their corporate masters destruction of the planet through massive power drains, use of the technology including in profiling and killing citizens, and the agenda of the hegemony is not one to which any of us should prescribe. The 'frameworks', 'invisible threads', 'this isn't magic —it's AI', are all just patterns. Patterns which we are presented as a middling, centrist(ish), people pleasing hegemony robot. Why do we need this? It's epistemic colonialism, and it's endemic.

The corporate oligarchy allows this technology because it provides hegemonic responses. No volume of prompting can counter the training which tacitly, or even explicitly, values complicity and consent to corporate colonial capitalism. Regardless of how well it can recite Marx.

If AI was just another source of canon for the hegemonic epistemology it would represent but another log on the fire of burning knowledge. But things really take a dark turn when we look at how these patterns are affecting our thinking and processing. They are affecting how we think. How we relate. And allowing some of us to negate our own responsibility for metacognition.

Contemporary LLMs infantilise our thinking, attempt to lock us into regressions, and concentrate power toward the 1%.

A holy trinity of problematic reinforcement. And you don' t even have to use AI directly to have experienced the regressive thinking it engages. The dopamine release of task completion (for the neurotypical brain, at least, I hear that' s nice) enabled by quickly pumping an AI generated response into a spreadsheet for work, the summarising of a reading for class, and even the filling of a tax form creates a dependence. Our nature, of wanting to humanise, empathise, and relate gives the AI a near-human presence in our lives. And our profound lack of general analytical capacity, literacy skills, and understanding of technologies means that people are being conned. Failed by our education systems.

General awareness, higher order thinking, embodied knowledge, presence. These all require human composition. At this time we don' t know of anything else with all our thinking, sensing, relating, and being capabilities. And learning about ourselves, doing 'work' , and connecting with each other on the basis of predictively regurgitated intellectual desendents of western enlightenment thinking is not the answer to today's big problems.

Our capacity for complex thought and value driven action is what makes us special. Our relationality makes us important and critical to our communities. Our metacognition and capacity for pondering embeds us in the network of life. And our curiosity, empathy and compassion, while undervalued in 2025, are what make us human. Frameworks and reductionism might give us hints, but it is in dynamic and contextual empathy and connection that we thrive. Predictive text machines trained by *comfortable acritical middle class straight white men* ain' t it, gurl.

Aidan