


How Far Can Human Knowledge Go?

by [Christopher Ormell](#) (January 2024)



Bounds of the Intellect– Paul Klee, 1927

There is a long-standing branch of philosophy–called

'epistemology' –which studies the limits of human knowledge. It addresses the question how far future knowledge can go. It uses logic, reasoning and realistic imagination (thought experiments) to try to envisage and conceptualise the boundaries which knowledge might reach. It was much discussed and mulled-over by the best thinkers from Ancient Greece onwards. However in recent times it has fallen out of favour. It seemed to hit a brick wall after the collapse of faith in general reasoning-and-theory in the 1970s: the post-modern sinkhole. As a result it has been (a) neglected by most of today's philosophers, (b) trampled-on by the now-dominant digital computerists, and (c) ignored–insofar as it poses constraints–by some rash late 20th century scientists who started speculating that a 'theory of everything' might be 'just round the corner'.

We know the leading anglophone philosophers of the early and middle 20th century–Russell, Whitehead, Ramsey, Wittgenstein, Ayer, Popper, Quine, Lakatos–paid their full attention to epistemology. But, in spite of their strenuous efforts, they were unable to reach any kind of satisfying conclusion. Since the 1970s, it has almost disappeared. Most of the academic 'philosophy' published since then has been in semantics, ethics, wellbeing, or preoccupied with social issues like cultural identity, personal identity and racism.

Does it matter? Yes, epistemologists have recently been sorely missed–because they were needed to defend commonsense against unwitting assaults. These assaults were mounted by computerists–who also, incidentally, virtually ticked themselves into a 'New Age' box, because they didn't claim, or even pretend, to have any acquaintance with epistemology. Instead, they brashly proclaimed a new subject, 'knowledge engineering,' and set-up previously unknown quests like 'creating knowledge,' 'getting computers to talk' and 'training computers to show "artificial intelligence."' (It is ironic that they assume that "intelligence" can be "trained,"

because for hundreds of years it was agreed that humans needed education, *not* training, to grow intelligence.)

There is a discontinuity here, one which provokes sceptical questions, such as whether knowledge *can be* 'engineered,' whether it *can be* glibly 'created,' whether it makes sense to call the audible outputs of computers any kind of 'speech.' Finally, there is the moot question whether computers could ever operate on an intellectual level where attaching the praiseworthy, high-status, personable word 'intelligence' makes proper sense.

Here is a thought experiment: if a student produces a body of work on a difficult conceptual area which happens—by sheer accident—to be almost identical with the output of ChatGPT on the same problem, should we immediately say that the student had shown *intelligence*? There is a significant minority of sceptical digital experts who will say that the answer is "No!" They are aware of a vein of overblown hype in the groupspeak emanating from Silicon Valley. They are aware that AI is a selective *copying* operation. It can certainly cobble together a semblance of prose which sounds initially quite credible, but closer inspection reveals the shocking fact that it has no genuine *understanding* of what it is "saying." It is a concoction put together using a vast amount of computer power ... cleverly copying and fielding a melange of OK-phrases, idioms and hitherto unnoticed usages, but doing so in a copycat, randomised, objective-neutral way which blandly ignores the self-questioning and emotive dimension of human intelligence. (In other words, it ignores the very things we tend to look for, when we are minded to attribute the high-status word 'intelligent' to someone.)

When challenged about their AI outputs, apologists for AI tend to give examples of things we have always known computers can do well, like scanning huge areas of technical information looking for unnoticed patterns and hidden pitfalls.

But what about the absence of genuine human feeling, goodwill, moral values or emotion in AI outputs? AI apologists are apt to reply that they are moving towards incorporating simulations and concentrations of human emotion into their neural networks.

Oh dear. They seem to be unaware that the very concept of 'simulated human emotion' is dodgy. 'Simulated human emotion' is precisely what sensible, responsible people abhor, and which they try to avoid. It is the stock-in-trade of con-artists, scammers and swindlers. And it won't do the public image of AI a lot of good, if AI is seen to be actively trying to join this unsavoury band.

That the latest AI *can* simulate what looks initially like meaningful human speech is, of course, an amazing technical feat. Trying to do this used to be considered, until quite recently, to be a hopelessly over-ambitious goal. It is astonishing that such plausible outputs can result from a nexus of inter-connected microchips. But the devil is in the detail, and the detail is, unfortunately, that the level of *really good judgment* expected by responsible human users, depends on the very element (honest feeling, responsibility, shared emotion, self criticism) which a mixture of deterministic algorithms and randomised neural connections don't, can't, and never will, deliver.

Too many otherwise savvy experts have let themselves be seriously dazzled by the technical software feat involved in getting computers to *appear to speak sense*. They have been bamboozled by their own propaganda, and have let this amazement blind them to the devilish detail ... that reaching the level of 'appearing to make sense' is not good enough.

The human brain is said to have billions of neural links in place—mysteriously organised in some baffling, obscure, currently un-understood way. These billions may be apparently up-staged in the future by vastly expensive mega-clouds filled

with *trillions* of microchips! But the sheer size of a neural network is not, actually, the main factor which will determine whether or not it is going to be trusted as the “last meaningful word” by ordinary human beings. The *feeling, sense of relevance* and *emotion* associated with genuine human speech is evidently a direct consequence of its essential biological, organic, personal origin. This isn't going to be replicated by a mixture of deterministic and randomly accessed microchips, however many trillions are pressed into service. Apologists for AI seem to assume—with very little reflection—that it *is* eventually going to be replicated. They are evidently convinced that “in the last analysis” what a human being says *must—of course—be a product of biochemical and neuro-electric micro-processes in her or his brain ... and hence ultimately be a mathematical construct.*

This demeaning assumption about what it means to be human appears to be a quite widely, even casually, swallowed, notion among software apparatchiks. It cannot claim, however, to have any kind of logical basis. What it takes for granted lacks common credibility. We have no reason whatever to think that mathematics, which is an inherently regimented, clunky, predictable, inert, timeless medium—and which has been used to discover representations of similar regimented, clunky, ordered, inorganic systems throughout history—can be expected to be equally good at unlocking infuriatingly fluid, defiantly-unpredictable, elusive, purposive, holistic, personable, emotional living systems. The human brains of cosmologists have also somehow managed to achieve a partial picture (a fragmented, semi- understanding) of a vast distant universe out there, which also, crucially, includes us.

So the real “devil in the detail” is that so-called AI lacks the key, vital element—genuine humanity—which is the essential *sine qua non* if the new agency is to contribute to the cultural wellbeing of the human race. In terms of factual content, word-counts, grammar, subject framework, and

associations ... AI does well. Overall its utterances might score 98% for believability. But 98% is not enough in many sensitive areas of human endeavour. In aviation, space travel, finance, medicine, weapon performance ... we expect the experts to get their answers *absolutely right*. NASA famously expects 0% error.

We have a mantra which applies when this level is not achieved: *A miss is as good as a mile*. There have already been various would-be public demonstrations of the new "artificial intelligence" which have ended in embarrassment.

It seems likely, therefore, that the inevitably unobvious vein of unreliability present in the new AI outputs, will gradually loom larger and larger with the passage of time. We can be sure rigorous thinkers are prudently going to avoid taking the risk it packs. It is, I'm afraid, fairly likely that the current AI craze will eventually turn into a bubble ... which will then burst... much like the ill-fated Japanese project to invest heavily in early, premature versions of AI in the 1990s.

So the casual neglect and dismissal of more than two thousand years of hard thinking about the nature of knowledge (and its limits) comes at a price. It involves going out on a risky limb, which comes about when a 'precise result' is offered which is known *not to be based on understanding*.

There have been plenty of warnings about what can go wrong. The phrase 'creating knowledge', for example, brazenly contradicts the traditional sense that discovering new knowledge often involves much hard, self-critical, submissive effort. New unexpected knowledge is discovered when careful observers humbly observe, not when they are imposing their preconceptions onto the empirical scene. It also requires things like: achieving a clear conceptualisation, valid mathematisations, rigorous checking, re-checking, more re-checking and finally clear-cut acceptance (of the final

results) among acknowledged independent trustworthy intellectual opinion leaders.

So, getting computers to say or print words is one thing: getting them to do this in a way which meets the best acceptable human levels of relevance, point, substance and responsibility, is another. When a person speaks we can hold them responsible for their message. And if it turns out badly, we can *reprove* the speaker. But modern computers don't wince when they get their prognostications wrong. (Of course Apple, Dell or HP could easily build a "visible wince response" into their machines—one which would cause the motherboard to shake briefly if it detected that it had been caught-out by a glitch. But it wouldn't carry much appeal for potential buyers, because they would know perfectly well that computers lack psyches, and they are certainly not going to feel the pain of the rebuff. Also, any future wince-like mechanical response might *look* clever, but everyone will know it is a fake.)

We have no good reason to suppose that what underpins the intrinsic freedom of the human mind is a brain "which can be modelled by mathematics." In the past some unwary highly intelligent people have jumped to the conclusion that it "must be" describable mathematically, because they think that this is all it can be. (Mrs Thatcher's phrase was 'TINA' = 'there is no alternative'.) For more than 2,000 years it has been a virtually unquestioned tenet in science that any evidently elaborate, but not-understood, complex structure *must have a mathematical explanation*. It is only in the last few years that anti-math has been around, offering a radically alternative explanation. Math can be used to great effect to describe structures composed of timeless, rigidly-ordered elements. But not all elements (maybe fewer than we think) are ultimately composed of timeless, wooden, infinitely static, absolutely indestructible elements. Anti-math is the new 100% abstract, 100% rational, lucid discipline which studies the

logic of transient realities.

So now we have a huge conundrum for epistemology: is the human brain more likely to be a mathematic or an anti-mathematic structure? The answer is obvious: the human brain is inherently transient, because we are transient, mortal beings.

Actually there are two major arguments which tell us that the human brain and the universe cannot expect to be describable using math.

They are: (1) because mathematical modelling can only happen (operate) with the help of an infusion of ordinary human imagination. It is like a food which consists of a powder, which can only be digested after it has been activated by mixing it with water. In math modelling, we have to interpret specific math configurations which contain—and are controlled by—the variable t : this may be t minutes or t microseconds. We have to *imagine* these configurations changing as t changes. No one with poor envisioning powers can hope to be effective as a math modeller, because they will be unable to “see” (in their mind’s eye), the way the model is changing with time. (Those dimensions which go up and down, those shapes which shrivel, burgeon or invert.)

Now a so-called “math model of everything” can’t call-in the human imagination to infuse it with activity, because the math is supposed to have *already* incorporated everything there is, into its structure, including human imagination. There is no way that the existence of this special activator—imagination—*operating onto the model from the outside* could be (ever) explained. (It also incidentally needs an external clock, again outside the math.)

(2) Math modelling depends on axioms. The first math modelling which sported process and changing structures came in the 17th century. It was introduced into a scientific culture already committed (since the renaissance, which had turned its back on

Aristotelian holistic thinking) to piecemeal methodology. Local axioms were needed to provide a sound basis for the math. They were based on specific empirical knowledge which had been around for some time, and which was known to be thoroughly trustworthy. But, as items of generalised empirical knowledge, these axioms also needed to be explained. This need not present any problems when the scientific culture is piecemeal, because a different bit of piecemeal knowledge can pull the trick. But a 'math theory of everything' would have to empirically explain its own axioms, which is a mathematical impossibility.

This is the nub of the granitelike enigma which Western science hit with devastating consequences around 1900. It was at the time, a sickening realisation. When Einstein discovered that there is no such thing as objective simultaneity, he discovered—though he didn't want to admit it—that 'the objective cosmos' was itself an illusion. The math modelling, on which the human race was accustomed to rely, was evidently not suited to modelling such an un-rigid, unobjective universe.

We have been in a state of severe cognitive shock ever since. Confidence in the capacity of human brainpower to ever understand the deepest secrets of the universe has plunged alarmingly: today it is beginning to sap confidence on every level. Fortunately a solution—anti-math—has emerged recently, and it is anti-math-based scientific modelling which now offers our best bet to unlock the secrets of the cosmos.

Anti-math begins with a picture of astronomic numbers of jumping random tally sequences, treated as reflections of jumping random shadows "out there" —real, but possessing no deeper substance.

It lets us off the hook of the impossibilities (a) and (b) above because: (a) it is self-activating (and energising) as a result of the random vitality of the vast substratum of

chaotic jumping shadows on which it rests. And (b), it can be a source of essentially transient reliable structures—objects, processes and conditions which last for quite a long time—and we have every reason to believe that it can also support (=provide models for) agents with cybernetic capabilities, the pinnacle being human capabilities. So it is reasonable to believe that it can support agents (us) which are capable of establishing their own final anti-axioms. ('Final anti-axioms' = anti-math axioms capable of creating representations of us.) Our being and intelligence can thus be understood as the final result of these immensely structured, unconscious, creative anti-axioms. This is a new, explicit, modern version of what Kant was saying more than 200 years ago. It might seem to be unlikely, but what can be potentially achieved via modelling with anti-axioms, is an as yet un-opened openbook. And there is every reason to think that it is more suited to producing fluent cybernetic outcomes than modelling with inherently rigid math.

Let's admit it, we *do* have a machinelike side. We are only semi-miraculous beings, because the material structures of our brains are dependent on the anti-axioms we tacitly (unconsciously) fully respect (probably through layer after layer of complexity built into our DNA) to secure our existence. But we do have an ocean of freedom to think and act, because we are the source of our own interpretations: we are not the victims of some strange, unknowable, omniscient, invisible, alien devilbrain. We need to be thoroughly responsible, though, on all the levels where we have an initial freedom to think and act, because there is no other agency guaranteeing the stability of our world. The total amount of goodness around is that embodied in individuals of goodwill. Here, too, Kant showed amazing percipience.

The human brain is obviously the most sophisticated, most complex structure present in the physical universe. It is also the most powerful, because without the mass of anti-math

axioms it unconsciously imposes onto chaos, there would be no structures “out there” (or “in here”), only a jumping random shadow wasteland. So it is likely that the physical cosmos is a necessary by-product of the key final-anti-axioms which support our being and intelligence.

And if it takes an immense, vast, awesome physical universe to serve as the by-product of the anti-axioms which are needed to support our being and intelligence, it is rather unlikely that a similar result could be pulled-off with a relatively few microchips.

So the (very unexpected) answer to the \$64 question *How far can human knowledge go?* is *All the way*. The epistemology of the future can aspire to be ‘Total Epistemology’, an epistemology which leaves nothing hanging as an unknowable unknown. (But because anti-math is based on absolute randomness, most future answers can only be indications of the *kind* of things that are there.) This conclusion may seem to be (and is) completely at odds with today’s unsure, weary, demoralised, despairing outlook. It is completely at odds with common expectations *precisely because* it isn’t any kind of product of today’s seriously muddled public world. It has emerged instead—unexpectedly under the radar—by rigorously following the trail of epistemology, the long forgotten, searching, traditional quest. It may not be the easiest trail of reasoning to follow, but it is, in the last analysis, thoroughly checkable. We also know that *hope springs eternal* and that hope, especially determined hope, is capable of snatching success out of the jaws of defeat.

[Table of Contents](#)

Christopher Ormell is an older philosopher whose ideas spring from a thoroughly de-mystified interpretation of math. He set

out a thought experiment which was the mirror image of Descartes' *I think therefore I am* in six articles in the journal *Cogito* 1992-94. It showed that absolute randomness was both a logical possibility and an unavoidable fact of life.

Follow NER on Twitter [@NERIconoclast](https://twitter.com/NERIconoclast)