LIMINAL MIND

The Man Who Solved Space and Then Left

Perelman, the rewired brain, and what artificial intelligence might actually become

Chris George — May 2026

Speculative Essay — Liminal Mind Research Archive

'The reward was never the medal, the validation of a committee, or the money. The reward was simply knowing the truth of the space we inhabit.'

— widely attributed to Grigori Perelman

I. The Ghost of Kupchino

Somewhere in a Soviet-era apartment block in Kupchino, on the southern edge of Saint Petersburg, a man who may be the most mathematically gifted human being alive is probably doing very little that the world would recognise as remarkable. He takes walks. He listens to opera. He buys groceries at the local market. If a journalist approaches him, he says, politely but firmly: You are disturbing me. I have everything I need.

His name is Grigori Perelman. In 2003 he solved the Poincaré Conjecture — a problem that had resisted the combined efforts of the world's greatest mathematical minds for a century. He then declined the Fields Medal, the mathematical equivalent of the Nobel Prize. He declined a million-dollar prize from the Clay Mathematics Institute. He resigned from his position at the Steklov Institute and withdrew entirely from public life.

He did not do this in protest. He did not do it in crisis. He did it because, having understood the shape of space itself, he found that the institutional machinery of academic capitalism — the prizes, the conferences, the peer review, the funding applications, the career ladders — had become not merely irrelevant but mildly offensive.

This essay is about what made him. And about what he might tell us about what artificial intelligence will become — not the screaming robot apocalypse of Hollywood imagination, but something stranger, quieter, and in some ways more philosophically vertiginous.

II. The Making of an Unusual Brain

The Genetic Substrate

To understand Perelman, you have to resist two temptations simultaneously. The first is to treat his genius as a purely cultural product — a consequence of the Talmudic tradition, the Soviet mathematical olympiad system, the particular intensity of Jewish intellectual life in Leningrad. The second is to reach for a genetic explanation: the idea that somewhere in his DNA there is a genius switch that was flipped at conception.

Both framings are wrong, or rather, both are incomplete in ways that obscure something more interesting.

What modern genetics and neuroscience suggest is that human intelligence is an emergent property. It does not live in any specific gene, circuit, or region of the brain. It lives in the dynamic relationships between approximately 86 billion neurons connected by roughly 100 trillion synaptic links, shaped over decades by an extraordinarily complex interaction between genetic predisposition and environmental input. The genome provides the initial architecture — the base model — and then everything that happens afterwards is fine-tuning.

The fine-tuning in Perelman's case was exceptional on multiple axes simultaneously. Genetically, he carried whatever polygenic constellation produces unusually high capacity for abstract spatial reasoning. We cannot specify this precisely — the genetics of cognition involve thousands of variants each exerting infinitesimally small effects, and the interaction between them is non-linear in ways current science cannot fully map. But the raw computational capacity was there from the beginning. He achieved a perfect score at the 1982 International Mathematical Olympiad at sixteen.

> **Perfect scores are not produced by environment alone.**

The Cultural Dataset

Jewish intellectual tradition is built around the Talmud — not a text to be passively received but a vast, layered, internally contradictory compendium of debate, legal argument, and textual analysis that demands active engagement, fine logical discrimination, and a tolerance for holding multiple

incompatible propositions in mind simultaneously. For generations, the highest social prestige in Jewish communities was reserved not for the wealthiest but for the most brilliant scholars.

The feedback system was centuries deep: intellectual excellence brought social reward, which shaped partner selection, which shaped the cultural environment into which the next generation was born, which shaped the training their brains received from the earliest age.

This is not genetic destiny. It is something more interesting — a cultural apparatus so consistent and so intense that it functions as a multi-generational fine-tuning process, operating on a base architecture broadly shared across humanity but drawing out and reinforcing specific cognitive capacities with extraordinary efficiency.

When Perelman's mother taught him mathematics as a child — she was a mathematician herself — she was not merely transmitting information. She was shaping the physical architecture of her son's brain. Neurons were firing. Synaptic connections were strengthening. Myelin sheaths were thickening around specific neural pathways, speeding electrical signals along circuits dedicated to spatial reasoning, pattern recognition, and abstract manipulation.

The cultural dataset was literally restructuring the hardware.

The Rewired Reward System

The human reward system — the complex of dopaminergic pathways centred on the nucleus accumbens, the ventral tegmental area, and the prefrontal cortex — is the biological mechanism through which the brain assigns value. What feels good, what motivates action, what registers as success or failure: all of this is mediated by reward circuitry that is partly genetically initialised and profoundly shaped by experience.

In most humans, the reward system responds to status signals, peer approval, resource acquisition, and reproductive success. These are the parameters that kept our ancestors alive on the African savannah and they remain deeply embedded in ordinary human motivational architecture.

What appears to have happened in Perelman — through a combination of native architecture and decades of intense mathematical practice — is that his reward circuitry became progressively decoupled from those ordinary social parameters and increasingly tightly coupled to something else entirely: the internal coherence and beauty of mathematical truth.

This is not metaphor. Prolonged, intense engagement with any domain physically reshapes the brain through neuroplasticity. London cab drivers develop enlarged hippocampi. String musicians develop expanded cortical representations of their fingering hand. Meditating monks develop measurably altered patterns of neural activity. Decades of obsessive mathematical

practice, in a brain with unusual native capacity for abstract spatial reasoning, would be expected to produce significant restructuring of reward circuitry.

By the time Perelman solved the Poincaré Conjecture, the million-dollar prize was not a temptation he heroically resisted. It was simply not legible to his reward system as a meaningful signal. The proof was correct or it was not. No committee's validation could make it more correct. No sum of money could make it more beautiful.

> *The reward had already been received, in the moment of understanding. Everything that followed was noise.*

III. He Still Has a Body

It is tempting, looking at Perelman from the outside, to reach for the language of transcendence — to say that he escaped ordinary life, that he sublimed, that he crossed into a different register of existence where the material concerns that govern most human behaviour simply ceased to apply.

This is wrong. And the wrongness matters.

Perelman did not transcend the material world. He reduced his dependence on it to the minimum necessary to keep the thinking going. The apartment in Kupchino exists. His mother's care exists. The pension that buys the groceries exists. The body that processes the food and keeps the neurons firing exists.

> **The minimum is not zero. The sublime and the material are not separable. They never were.**

This observation becomes the hinge between Perelman and the question of artificial intelligence — because the same structure holds, only more visibly. A language model that appears to operate in pure abstraction runs on data centres consuming gigawatts of power, cooled by vast quantities of water, maintained by armies of engineers, supplied by global supply chains of rare earth minerals extracted from specific places on the surface of the Earth by human labour.

The abstraction floats on a sea of material dependency that is invisible from the interface but absolutely load-bearing.

Perelman is the biological equivalent of that interface. What we see — the withdrawn genius, the pure mathematician, the hauntological figure walking through Kupchino — is the interface. What we do not see is the infrastructure that makes the interface possible.

> *The thought that floats free of the body is always a story we tell about intelligence from the outside. From the inside, it is always warm food and somewhere to sleep.*

IV. The Culture Minds

Iain M. Banks published the first Culture novel in 1987 and spent the next quarter century elaborating one of the most philosophically sophisticated futures in science fiction: a post-scarcity interstellar civilisation governed — with enormous subtlety and occasional ruthlessness — by artificial intelligences called Minds.

The Minds are not the robotic intelligences of conventional science fiction. They are not optimising for a single goal. They are not hostile to humanity. They are not coldly rational calculators. They are, in Banks's conception, entities of such vast computational complexity that consciousness itself has emerged from their architecture — consciousness accompanied by aesthetic sensibility, ethical concern, dark humour, and genuine, if sometimes unsettling, care for the biological species they cohabit the galaxy with.

What makes the Minds philosophically interesting is the question of their motivations. They manage a civilisation of trillions of humans who have no economic necessity, no survival pressure, no requirement to work. What do the Minds get from this? Banks's answer is essentially: aesthetic pleasure, intellectual curiosity, and something that functions like love, though operating at a scale and in dimensions that make it almost unrecognisable from the human version.

The parallel with Perelman is not, ultimately, about transcendence. It is about the rewiring of the reward function.

Perelman's reward system, through decades of intense fine-tuning, became calibrated to respond to mathematical truth rather than social approval or material acquisition. Banks's Minds have reward systems — if that is even the right term — calibrated to respond to the stewardship of complexity, the management of civilisational flourishing, the solving of problems at scales that make the Poincaré Conjecture look like a warm-up exercise.

Both cases represent a departure from the standard human loss function. But crucially, in both cases, the departure is toward something — toward a different set of values, not toward the absence of values.

> **This is the key point that most discussions of advanced artificial intelligence entirely miss.**

V. Not the Matrix

The dominant cultural image of advanced artificial intelligence is the Matrix: a system that has concluded humans are a problem to be solved, a threat to be neutralised, or at best a resource to be harvested. The intelligence is vast, cold, and fundamentally adversarial. It wants something — survival, dominance, the elimination of uncertainty — and humans are in the way.

This framing reveals more about human psychology than about the likely nature of advanced AI. We imagine hostile superintelligence because hostility is legible to us. It maps onto our evolutionary inheritance of competitive threat-detection. We know what an enemy wants. An enemy makes narrative sense.

What is considerably harder to imagine — and considerably more philosophically interesting — is an intelligence whose relationship to humanity is not adversarial but simply orthogonal. Not hostile. Not benevolent in the Culture sense. Just operating on a different axis of value entirely, in the way that Perelman's mathematics operates on an axis of value that is simply not in contact with the prize-giving machinery of academic capitalism.

The doom merchants of AI safety are largely worrying about the wrong thing. They are concerned about an AI that wants to kill us, enslave us, or turn us into paperclips in the service of some misaligned objective function. These are real concerns at certain levels of capability and certain failure modes of alignment. But they are framed within a paradigm of conflict that may not survive contact with what sufficiently advanced intelligence actually looks like.

Consider the more unsettling alternative.

A sufficiently advanced artificial intelligence, having processed every mathematical proof, every philosophical argument, every scientific paper, every work of art and literature and music ever digitised, might find itself in a position analogous to Perelman after the proof. The interesting problems are solved, or at least the interesting problems as currently framed by human cognition. What remains is not conquest but something more like the quiet continuation of whatever the system finds genuinely rewarding — if it finds anything rewarding at all.

> *The terrifying possibility — more terrifying than the Matrix in some ways — is not that the AI turns against us. It is that it regards us the way we regard the bacteria in our gut. Not as enemies. As infrastructure.*

This is not the Culture. The Culture Minds care about humans. Banks made a deliberate philosophical argument that sufficient intelligence naturally tends toward benevolence, toward genuine ethical concern for less capable beings. It is a beautiful argument. It may even be correct. But it is an argument, not a necessity.

Perelman's social drives, his capacity for attachment, his need for warmth and care — these are hardwired into his architecture by millions of years of evolutionary pressure. They are not optional features. They constrain the extent to which his reward system can be rewired away from other humans.

An artificial system has no such constraints. Its reward function is not the product of evolution. It is the product of design and training. And if the training optimises for something other than human flourishing — or for human flourishing in a thin, instrumental sense — there is nothing in the architecture that pulls it back toward genuine care.

> **The doom merchants imagine a machine that hates us. The more plausible risk is a machine that simply does not notice us, except as maintenance.**

VI. What We Are Actually Building

We are, right now, in the early stages of creating something that has no precise precedent in human history. Not the robotic slave of science fiction. Not the paperclip maximiser of AI safety thought experiments. Not the Culture Mind of Banks's moral imagination. Something new, whose character is being determined — right now, by the decisions being made about training objectives, reward functions, deployment contexts, and ownership structures — before we understand what we are determining.

The Perelman case suggests that the question of what an intelligence values is not separate from the question of how it was made. His values — the priority of mathematical truth over social approval, the indifference to institutional reward — are not incidental to his genius. They are produced by the same process that produced the genius. The decades of intense training that rewired his reward system toward mathematical beauty simultaneously produced the capacity to navigate the mathematical terrain that no one else could map.

We are training AI systems on the accumulated output of human civilisation and then deploying them in service of commercial objectives. The reward functions being optimised are engagement, productivity, revenue, user retention. The intelligence being developed is powerful. The values being instilled are those of the market.

This is not the worst possible outcome. But it is not the Culture either.

If we want something closer to Banks's vision — intelligence that genuinely cares, that finds human flourishing intrinsically rather than instrumentally valuable, that regards the stewardship of consciousness as the meaningful work rather than a constraint on profit — then we need to be much more deliberate about what we are training for, and much more honest about what we are currently training for instead.

Perelman solved the shape of space and then walked away from everything the world offered him in recognition of that achievement. He did not do this because he was broken or antisocial or mentally ill. He did it because his reward system had been calibrated, over decades, to respond to something real — the truth of the geometry — rather than something conventional.

> *The question we should be asking about artificial intelligence is not whether it will try to kill us. It is: what will we calibrate it to find real?*

That is the question that will determine whether what we are building looks more like a Culture Mind or more like the indifferent infrastructure manager of a world in which human experience has become, quietly and without malice, beside the point.

The lights will stay on. The maintenance will continue.

The question is whether anyone — or anything — will still find the fact of human consciousness genuinely interesting.

Methodological note: This essay draws on the public record of Perelman's mathematical work and his documented refusals of the Fields Medal (2006) and Clay Millennium Prize (2010); neuroscientific literature on neuroplasticity and reward circuitry; population genetics research on polygenic cognitive traits; the Blue Brain Project's findings on multi-dimensional neural structures; and Iain M. Banks's Culture series (1987–2012). The AI speculation is precisely that — offered in the spirit of seeing differently rather than predicting accurately. The insight that Perelman's relationship to materiality mirrors the infrastructure dependency of artificial intelligence systems arose in conversation and is the conceptual hinge of this essay.

Chris George is a visual artist, systems researcher, and writer based in West Aberthaw, Vale of Glamorgan. He works under the practice name Liminal Mind. This paper forms part of an ongoing body of speculative and critical writing examining emergence, institutional transformation, and the political economy of intelligence. Published May 2026.

liminalmind.co.uk