

The Last Mile Is Human

Mohamed Abo Aziz

One night the pain in my chest was heavy enough that my grandmother took me to the hospital. They ran the troponin. It came back at 0.04 — detectable, not diagnostic. The gray zone, where the number alone decides nothing.

The protocol said: repeat in three hours, decide on the delta. A rise of fifty percent means admission. Less means discharge. Clean and inert. It cannot order itself, read itself, or decide. It cannot tell you what the number means for your life.

A doctor ordered the repeat. A doctor read the second result. A doctor signed the discharge when the delta was zero. The protocol required a human at every step — not because humans are wiser, but because without one, the protocol does nothing at all.

My grandmother's husband died of a heart attack. She knew what chest pain in the middle of the night could mean. She had already lived that night once. So when she looked at me in that chair, she was not looking at her grandson. She was looking at the man she had lost, and she was not going to lose him twice.

Three hours later, the delta was zero. My heart was fine.

The algorithm met its threshold. She met a different one.

This is when love becomes the diagnosis. Not the troponin, not the protocol, not the delta. Her love had already read the result before the first blood draw. It arrived as fear, and the fear became a decision, and the decision saved my life. No model was trained on that. No model ever will be.

I asked my brother, who works in computational biology, why we build AI to think when we already possess minds of our own. He gave me a practical answer. To save time. To process faster. To accomplish in seconds what would take a person years. His field is the art of translating living systems into mathematical models — taking the messy, wet, unpredictable machinery of biology and rendering it as code. He knows better than most what is lost in that translation, and what is preserved.

It was a good answer, and it was true, and it did not touch what I was asking. I was not asking why we build it. I was asking why we trust it.

The AI answers the question. The human questions the answer.

Philosophers have argued for centuries about what it means to know something, and one answer is that knowledge is a justified true belief, a fact you can defend with reasons. By that definition, AI holds more justified true beliefs than every doctor who has ever lived, combined.

But Aristotle drew a line we keep forgetting. He distinguished between *techne* and *phronesis*, technical knowledge and practical wisdom. *Techne* is knowing how to make something. *Phronesis* is knowing what to do in a situation that has no rules. AI excels at *techne*. It cannot reach *phronesis*, because *phronesis* is not stored. It is lived.

This world is full of intelligent systems. They have engineered revolutions in every field, from health to economics to politics. Data analytics dashboards. Language models upgrading by the day. Agents for every daily task. Automation spreading into every function. Together they have produced a world that runs at maximum velocity toward truth, clarity, and leadership.

Healthcare shows the gain most clearly. Cancer treatments once took decades of research. The process was expensive, exhausting, and often could not predict which compounds would work. Now the time has collapsed. The effort has shrunk to a fraction, liberating researchers to focus on the pure science instead of the machinery around it.

But stop here. Have we reached the truth? Does this prove superiority over us? Does AI possess the sensory and spiritual capacity for guidance, for judgment, for taking decisions and making change on its own?

Computational models like AlphaFold can now predict protein structures, revolutionizing drug discovery in a way the human mind has no ability to do. But they do not possess the human lens to judge whether a given protein is worth pursuing further in research and production, or whether curing the disease it is involved in is worth the cost and side effects. A diagnostic tool utilizing computational algorithms can scan tumors, even the smallest ones, with an accuracy, precision, and speed no human can cope with, yet it remains silent and useless in determining what this means for the person — whether to treat aggressively or wait. The human has the lens to see the world; its meaning should be accounted for and decided by that lens — it is a human advantage no tool can borrow.

These questions have simple answers. They are the same logic that already circulates in our minds and in the decisions we issue. We are the origin. The origin is always the truth, and the truth cannot be changed. This is reality, and it admits no argument.

Nobody says the machines in a factory, the ones that cut a plant, grind it into powder, and improve its shape and colour, are greater than the person who invented them. No one holds

that debate in his head. The same applies to abstract intellectual systems and digital technology. At their root, they are neural circuits imitating the human mind. The human mind brought them into our world. It borrowed from atomic physics: memory that does not end, speed of execution, and the capacity for learning and self-improvement.

Consider what this actually means. A biological neuron is a cell — wet, organic, powered by chemistry, shaped by evolution over billions of years. It communicates through electrical impulses and chemical gradients, and it is embedded in a body that feels hunger, fear, exhaustion, and joy. A silicon transistor is a crystal — dry, inorganic, powered by electrons moving through a lattice of atoms. It has no body, no chemistry, no stakes in anything. When we say AI is “rewritten with atomic specifications,” we mean that scientists have taken the mathematical logic of how neurons process information — the weighted sums, the thresholds, the learning rules — and transferred that logic onto silicon. They have kept the architecture of thought and discarded the architecture of life. The result is extraordinary. It is also incomplete. The silicon can hold the pattern. It cannot hold the experience that gave the pattern its meaning.

And yet these models lack the one thing we cannot feed them and cannot program: perception. Feeling. Awareness. This is the human feature that is not programmable and cannot be imitated, because it originates as one of the components of the human mind itself, in its mechanism, in what triggers it, in how it moves. It is what mixes creative mental thinking with the human, sensory, emotional, existential side of us. That mixture produces results built on real experience, results that stay grounded.

Consider what happens when you look at a colour. Light leaves a surface at a specific wavelength. It enters the eye, strikes the retina, and becomes an electrical signal. A machine can measure that wavelength with perfect precision. It can name it, classify it, reproduce it on a screen. But it cannot see the colour. It cannot know what red means at a funeral, or what blue means in a hospital room at 3 a.m. The photon carries information. It does not carry meaning. Meaning is what the human mind adds when the light arrives.

And if that is true of something as simple as light, it is even more true of the deepest layer of reality. Quantum mechanics tells us that a particle exists in every possible state at once — in superposition — until it is measured. The act of observation collapses the wave function into a single reality. AI can calculate every possible outcome. It can hold every probability in its weights. But it cannot collapse them. It cannot decide which one becomes real. That requires a party. A conscious observer. A human being who chooses to look.

The universe, at its deepest level, waits for someone to choose. AI can describe every possibility. It cannot make one true.

This is not an abstract physics lesson. It is the same gap I saw in the hospital, and the same gap I see in every algorithm that claims to know what a person should do. The machine can hold every probability. It cannot be the one who chooses. That choice is ours, and it is the only thing that has ever been ours.

Closer to home, I once asked a language model which faculty to choose after secondary school. It answered well. It listed every option, ranked them by employment, and compared salaries, mapping where each path led. Everything it said was true. It told me which choice was reasonable. It could not tell me which choice would be mine. It did not know what my grandmother's illness had already decided for me. I am in pharmacy now. The decision was mine, and that is why it carries weight.

When my grandmother was sick, the system saw a diagnosis. It saw a prescription. It saw a cost. It did not see her fear, or mine, or the way our family rearranged itself around her. The data was accurate. It was also beside the point. And that gap between accuracy and meaning is not a small one. It is the same gap that opens in every crisis, in every market, in every moment where the numbers say one thing and the human situation demands another.

The most dangerous version of that gap appears when the smartest models meet the most chaotic reality. When the assumptions built into the mathematics stop holding, and the model keeps calculating as if nothing has changed. That is not a hypothetical. It happened in 1998, to the most sophisticated financial minds in the world, and it almost took down the global economy.

In 1998, Long-Term Capital Management was a hedge fund run by Nobel laureates and built on mathematical models of risk. These models are trained to match past historical relationships, assuming the market remains stable and calm. When the crisis hit, human judgment stepped in: the Federal Reserve Bank of New York, under President William McDonough, took action alongside other financial institutions by bringing the major banks together. Fourteen financial firms eventually agreed to provide about \$3.625 billion and take a 90 percent stake in LTCM, resolving the fire sale and the systemic collapse. The models could not see the panic coming. The humans in the room could. That is the only reason the system survived.

When reality breaks the assumptions, human judgment becomes more valuable than computation.

It is not about whether humans have responsibility and feelings while AI does not. It is simpler and harder than that. A dashboard cannot take responsibility. It cannot take a decision. It is a tool that produces evidence — quantitative or categorical, a result no human could match in speed, storage, or efficiency. It offers additional evidence, free of

any interest of its own. Nothing more, and nothing less. But evidence describes what is. It does not say what ought to be done. It cannot be asked to answer for anything. The accountability and the decision are ours.

The doctor signs the discharge. The diplomat answers for the room they misread. Richard Hakluyt, a sixteenth-century English clergyman and geographer, put his own name on the collection and answered for its choices. An AI is never there afterwards to be asked why.

Hakluyt understood this four centuries before the AI arrived. When he compiled his voyages — gathering accounts of English exploration for a work he published in 1589 — he was doing what a database does now: gathering, sorting, comparing. The difference lies in what he refused to do. He refused to flatten the accounts into one clean summary. He let translated Spanish, Portuguese, and French voyages sit beside the English ones in the same collection, and he trusted his reader to weigh them. That trust was not a technical choice; it was a human one. The algorithm sorts. The editor judges.

There is a second lesson in his work. Why did he compile those accounts at all? Not because the data was waiting to be collected, but because England was falling behind Spain and Portugal in the race for the New World, and he wanted his country to compete. The motivation came first. The collection followed. Every algorithm we build today follows the same order. A human being decides the world needs curing, and then builds the AI to help. The AI does not decide. The human decides, and the AI serves.

None of this is a case for turning the AI off. That would be absurd, and it would not work. The case is narrower and harder: it is for knowing which parts of the work are ours to keep.

My brother gave me the right answer to the wrong question. Speed, he said. Time. Efficiency. All true. But speed was never what I was asking about, and I think he knew that. He answered the question he could answer.

You can feed an AI every scan, every trial, every result. It will answer in seconds. It will never once have to walk into the room and say the words out loud.

That part is ours. It always was. It always will be.

My grandmother did not wait for the delta. She only knew the shape of a man she had lost, and she saw it in me, and she moved.

Her love became fear. Her fear became a decision. It arrived before the first blood draw.

It saved my life. It is not in the data.

Creative Contribution Statement

I conceived, researched, and wrote this essay. The personal narrative — the troponin test, my grandmother, my brother, and my decision to study pharmacy — is my own lived experience. I used AI tools for grammar-based corrections, to help generate possible global examples through prompting, for brainstorming, and to suggest strategies for presenting my ideas and essay effectively. I selected which examples and strategies to use, integrated them into my own argument, and made all final creative decisions about structure, voice, and wording. The essay is not AI-generated. The creative contribution lies in connecting a personal medical moment to the distinction between techne and phronesis, and in using historical and contemporary examples to show that AI produces evidence but cannot take responsibility, judge, or decide.

AI Use Disclosure

I used AI in the following ways: grammar correction, brainstorming, generating candidate global examples via prompts, and suggesting presentation strategies. All final choices, edits, and writing are mine.

References

Aristotle. *Nicomachean Ethics*. Trans. W. D. Ross. Book VI.

Fleming, Michael, and Weiling Liu. “Near Failure of Long-Term Capital Management.” *Federal Reserve History*, November 22, 2013.

<https://www.federalreservehistory.org/essays/lcm-near-failure>.

Fuller, Mary C. *Lines Drawn across the Globe: Reading Richard Hakluyt’s Principal Navigations*. Montreal: McGill-Queen’s University Press, 2023.

Hakluyt, Richard. *The Principall Navigations, Voiages and Discoveries of the English Nation*. London: George Bishop and Ralph Newberie, 1589.

Lowenstein, Roger. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. New York: Random House, 2000.