

Encounters with the Arch-Genius, David Gelernter

Jeff Brown for TIME

The artificial intelligence pioneer has some radical ideas about the supremacy of the human mind

Elections come and go. Markets rise and fall. Celebrities wax and wane. But now and then we meet a controversy of deep and lasting dimensions. Some of our leading engineers and most brilliant theorists say the future of artificial intelligence is such a matter. Will machines learn to think like humans—and then to outthink us? And if they do, what will become of us?

The topic, once a staple of science fiction, has become one of the defining facts of high tech. From Apple to Amazon, Facebook to Intel, Sergey Brin to Elon Musk, the titans of the 21st century are investing fortunes and countless hours in artificial intelligence (AI). Google's 2014 purchase of the British firm Deep Mind for something more than \$400 million produced a bonanza of publicity earlier this year, when its game-playing program whipped a human master of the ancient strategy game Go. IBM is pouring \$1 billion into building a business around Watson, the company's digital Jeopardy! champion that chats with Bob Dylan in its latest ad campaign. Amazon's new personal digital assistant, Alexa, dwells in the cloud and, like her cousins Siri (Apple), Cortana (Microsoft) and Google Now, will dispense instructions from speakers, smartphones, televisions and cars. It's remarkable how quickly we've adjusted to their presence.

Among the thundering vanguard, though, is a growing group of worried individuals, some of them doomstruck Cassandras, some machine-hating Luddites and a few who fit in neither group. They take in the rapid rise of superintelligent machines—which are already taking over jobs as factory workers, stock traders, data processors, even news reporters—and conclude they will eventually render us all obsolete. “The development of full artificial intelligence could spell the end of the human race,” warns Stephen Hawking, the renowned astrophysicist.

This is the topic that brings me through a snowy Connecticut forest to a house not far from Yale University. I am here to discuss the human mind and artificial intelligence with David Gelernter, artist, author, scientist, composer and stubbornly independent thinker. A conservative among mostly liberal Ivy League professors, a religious believer among the often disbelieving ranks of computer scientists, Gelernter is neither Cassandra nor Luddite. He is a computer virtuoso who happens to find human consciousness even more entrancing than the most amazing digital apparatus.

In his latest book, *The Tides of Mind: Uncovering the Spectrum of Consciousness*, Gelernter argues that the entire field of AI is off track, and dangerously so. A key question in the pursuit of intelligence has never been answered—indeed, it has never really been asked: Does it matter that your brain is part of your body?

Or put another way: What is the human mind without the human being?

This mind-body question has an odd place in the history of artificial intelligence. Alan Turing, one of the pioneers of the field, found it so daunting that he pushed it to one side. His seminal 1950 paper, “Computing Machinery and Intelligence,” drew “a fairly sharp line between the physical and the intellectual capacities of a man,” as he put it.

A similar attitude was struck a few years later by computer scientists at IBM. Reporting their breakthrough success in creating a computer capable of excelling at high school geometry, project leader Herbert Gelernter—David’s father—declined to say “whether our machine is indeed behaving intelligently.”

Over the half-century that followed, a funny thing happened on the dizzying road from electronic geometry whizbangs to the apocalyptic threat of Hawking’s direst thoughts. AI theorists stopped treating the human body as an overwhelming problem to be set aside and started treating it as an irrelevant matter to be ignored. Today the mainstream argues that there is no meaningful difference between the human brain, with its networks of neurons and axons—electrical and chemical on-off switches—and computers powered by 1s and 0s. And by the same analogy, computer scientists understand the human mind to be the equivalent of software running on the brain-computer.

Whatever differences exist between humans and machines, today’s gurus of artificial intelligence argue it will vanish in the not-too-distant future. Human minds, their memories and personalities, will be downloadable to computers. Human brains, meanwhile, will become almost infinitely upgradable, by installing faster hardware and the equivalent of better apps. The blending of human and machine, which Google’s Ray Kurzweil calls the Singularity, may be less than 30 years off, they theorize.

David Gelernter isn’t buying it. The question of the body must be faced, and understood, he maintains. “As it now exists, the field of AI doesn’t have anything that speaks to emotions and the physical body, so they just refuse to talk about it,” he says. “But the question is so obvious, a child can understand it. I can run an app on any device, but can I run someone else’s mind on your brain? Obviously not.”

In Gelernter’s opinion, we already have a most singular form of intelligence available for study—the one that produced Bach and Shakespeare, Jane Austen and Gandhi—and we scarcely understand its workings. We’re blundering ahead in ignorance when we talk about replacing it.

Inside the house, evidence of the mind of Gelernter is everywhere. The towering walls of books—including his own works on computer science, religion, popular culture, history and psychology. His works of art—some abstract, some powerfully figurative, like the life-size evocations of the great kings of Israel inspired by Christian tomb art at the

Basilica of St. Denis outside Paris. Musical instruments fill the floor space. Flamboyantly colored birds survey the scene—a purple parrot in a cage near the kitchen and a multihued macaw named Ike that presides over the family room. Gelernter's conversation runs in torrents from the prophecies of Isaiah to the subtleties of Gothic engineering to the proper design of graphical user interfaces.

Indeed, the breadth of his interests and the range of his thoughts can be overwhelming. Gelernter expresses sympathy for the graduate students in his new seminar on software design at Yale, which has been his academic home for most of his life. The students are flummoxed, he reports, by their professor's extensive discussion of medieval architecture, specifically the so-called strainer arches at Wells Cathedral in Somerset, England. Wrought from massive stone blocks and installed nearly 700 years ago to prevent the church from collapsing, these arches are as strong as steel—and as lovely as butterfly wings. His reason for dwelling on the subject, he explains, is that all truly great designs are beautiful as well as functional. "I try to tell them, if you're going to do anything good in software, it will be beautiful."

Sun Microsystems co-founder Bill Joy has called Gelernter, who pioneered breakthroughs in parallel processing, "one of the most brilliant and visionary computer scientists of our time." Gelernter's 1991 book, *Mirror Worlds*, foretold with uncanny accuracy the ways the Internet would reshape modern life, and his innovative software to arrange computer files by timeline, rather than folder, foreshadowed similar efforts by several major Silicon Valley firms. (A patent lawsuit against Apple was ultimately decided in Apple's favor.) Yet Gelernter is not enthralled by the power of computer science, which he considers to be essentially a secular religion for its devoted disciples. His colleagues in computer science are so enamored of their own miraculous designs, he says, that they refuse to consider the limits of their machines.

Go back to that Gothic cathedral for a moment. How does it work its effects on the people who enter? In its scale and design, its vast weight and fortifying inspiration, its dark vaults and diffuse lights, in the ancient stories signaled through episodes of glass and carving, the church speaks to the mind of the engineer as well as the emotions of the pilgrim. The building can be measured and analyzed. But it is also felt. And how it feels depends on the time of day, the mental state of the visitor, the depth of the silence or the rumble of the organ. It smells of incense and age. It soars, and it terrifies.

The human mind, Gelernter asserts, is not just a creation of thoughts and data; it is also a product of feelings. The mind emerges from a particular person's experience of sensations, images and ideas. The memories of these sensations are worked and reworked over a lifetime—through conscious thinking and also in dreams. "The mind," he says, "is in a particular body, and consciousness is the work of the whole body."

Engineers may build sophisticated robots, but they can't build human bodies. And because the body—not just the brain—is part of consciousness, the mind alters with the body's changes. A baby's mind is different from a teenager's, which is not the same as an elderly person's. Feelings are involved: a lifetime of pain and elation go into the formation of a human mind. Loves, losses and longings. Visions. Scent—which was, to Proust, "the last vestige of the past, the best of it, the part which, after all our tears seem

to have dried, can make us weep again.” Music, “heard so deeply/That it is not heard at all, but you are the music/While the music lasts,” as T.S. Eliot wrote. These are all physical experiences, felt by the body.

Moreover, Gelernter observes, the mind operates in different ways through the course of each given day. It works one way if the body is on high alert, another on the edge of sleep. Then, as the body slumbers, the mind slips entirely free to wander dreamscapes that are barely remembered, much less understood.

All of these physical conditions go into the formation and operation of a human mind, Gelernter says, adding, “Until you understand this, you don’t have a chance of building a fake mind.” Or to put it more provocatively (as Gelernter is prone to do): “We can’t have artificial intelligence until a computer can hallucinate.”

Gelernter’s new book is the fruit of a lifetime’s reflection on such matters. Rejecting the analogy of brain to computer and mind to software as “childishly superficial,” he describes a variable human consciousness that operates along a spectrum from “high-focus” to “low-focus”—up and down, back and forth, many times each day.

At high focus, the mind works exactly like a computer. It identifies specific problems and tasks. It calls on the memory for data and patterns and instructions necessary to answer the questions and perform the jobs at hand. High focus finds the mind thinking about thinking; that is, thinking on purpose.

At low focus, the mind may drift, even seem to go blank. Notions and daydreams pop up without being consciously summoned. At the lowest focus, when the body is asleep, the dreaming mind churns up images and memories and patches them together—not according to a rational blueprint, Gelernter argues, but according to some sensation or emotion that they share.

“As we move down-spectrum,” he writes, “mental activity changes—from largely under control to out of control, from thinking on purpose to thought wandering off on its own. Up-spectrum, the mind pursues meaning by using logic. Moving down-spectrum, it tends to pursue meaning by inventing stories—as we try to do when we dream. A logical argument and a story are two ways of putting fragments in proper relationship and guessing where the whole sequence leads and how it gets there.”

Inevitably to modern, logical readers, this description suggests a hierarchy. “Up-spectrum” sounds superior to “down-spectrum,” “high-focus” better than “low-focus.” We might ask—even if Gelernter is correct about the workings of the mind—why artificial intelligence should not operate solely at high focus and up-spectrum? Leaving the lower range of consciousness behind might be progress, right?

No, Gelernter contends. The full expression of the human mind requires the entire spectrum. His book, like his conversation, is a celebration of the full span. He quotes not only scientists and psychologists but also poets and novelists. A mathematical proof or scientific discovery is no greater sign of intelligence than is the “Ode to a Nightingale” by John Keats, who ends his masterpiece by wondering where on the spectrum of

consciousness he was: “Was it a vision, or a waking dream?/Fled is that music:—Do I wake or sleep?”

For that matter, not all logical breakthroughs come from minds operating at high focus. Consider the story of pioneering neuroscientist Otto Loewi. A century ago, Loewi tried to devise an experiment that could test his theory that the brain transmits some signals chemically. When he finally grasped the answer, it was at low focus, in a sequence of dreams. The experiment that Loewi envisioned while sleeping in 1921 eventually led him to a Nobel Prize.

David Gelernter was born in 1955 with a front-row seat on the computer age. On the dedication page of *Tides of Mind*, he hails his father as “one of the six men who invented AI.” After earning his bachelor’s degree at Yale, where he majored in religious studies and pursued an interest in neurobiology, Gelernter did his Ph.D. studies at Stony Brook University. (His father was a professor there.) He joined the computer-science faculty at Yale and pitched into the vital problem of parallel processing—in rough terms, how to make computers perform more than one task at a time.

His breakthrough in that field cemented his reputation for brilliance—and it came, he says, courtesy of a down-spectrum moment. After thinking at high focus about the problem of gridlocked signals, Gelernter daydreamed a vision of Grand Central Station so crowded that no one could move. The escalators were in motion, though, churning people from one level of the station to another. “To be conscious of a thought does not mean we know where it came from,” Gelernter observes in *Tides*. Whatever its origin, the image freed Gelernter’s mind to unstick the flow of signals in his software.

Later, Gelernter and a colleague attempted to program a computer to mimic low-focus consciousness. He imagined a sort of dial on the device that would move the machine up and down the spectrum, from Spock-like logic to loopy hallucination. Though the attempt did not achieve quite the results he had hoped for, the program did show a degree of suppleness that they eventually used to advance the role of computers in diagnosing diseases.

Gelernter has no doubts that huge strides can be made in expanding the spectrum of artificial intelligence. “Computers already have more than enough capacity to mimic low-focus thought,” he says. With sufficient resources, a huge database can be compiled from human subjects connecting myriad images and sensations with matching emotions, he says, describing the work to be done. From there, machines equipped for “deep learning” could eventually become adept at faking the feelings that give structure to down-spectrum consciousness.

But his name for such machines conveys his concern and contempt: “zombies.” They might be made to look like humans, and even to react like humans. But they would not have genuine human feelings. They wouldn’t know the fear and exhilaration of riding a roller coaster, much less the racing heart and flip-floppy stomach of young love, nor even the depressed exhaustion of grief.

Perhaps most important, the computer won't feel the existential dread or weird magnetism of death. Admittedly, this is exactly why the Singularity is so appealing to Kurzweil and his followers. By merging the human with the machine, the software mind is freed from its wet mortality and crosses to eternal life. But what is human consciousness without the shadow of death? "The meaning of life," wrote Franz Kafka, "is that it ends."

"Kurzweil," says Gelernter, "is a good man and very bright, very capable. It's just that some of what he says doesn't make sense. He's going to upload his mind to the cloud and live forever—what does that even mean? If my mind is running on another computer, it is no longer me."

Issues of mortality and limitations are not abstractions to Gelernter. His range of human feelings, including his familiarity with death, expanded horrifically one early summer day in 1993. In his office at Yale, Gelernter opened a package that had come in the mail. A pipe bomb, prepared by Ted Kaczynski—the so-called Unabomber—destroyed his right hand and left him fighting for his life. More than two decades later, Gelernter still deals daily with the pain and disability.

On some days, he feels estranged from a world that has little patience for "long-term consequences," he says. "It is hard for people to sustain their attention to chronic conditions and permanent injuries." He seems reluctant to say this, because he doesn't like complainers. He quickly adds, "In the final analysis, there is an insulating layer of kindness for which one thanks God."

Gelernter is not the only dissenter from the AI orthodoxy. Silicon Valley entrepreneur Roman Ormandy, for example, has criticized the brain-as-processor model. "The more neural research progresses, the clearer it becomes that brain is vastly more complex than we thought just a few decades ago," Ormandy has noted.

But Gelernter is vastly outnumbered—so much so that he worries that his ideas might simply be ignored. "There has never been more arrogance and smugness" than in today's self-congratulatory scientific culture, he asserts. "The spectrum of our consciousness is such a part of who we are and how we live in the world. But we make such a virtue of ignoring it. We have a fundamental cultural prejudice that high-focus thought is better, when if we would just examine our own lives we would see that we all hallucinate every day as we dream and see visions as we're falling asleep."

Computers are going to grow much more powerful, and they will be relied upon to complete far more tasks than they do today. Scientists won't stop in their pursuit of better programs on faster processors—nor should they. Gelernter fully appreciates that this progress will take machines deeper and deeper into the spaces previously reserved for human intelligence. Their memories will be bigger than ours and more rapidly accessible. Their importance will grow as they do more, tirelessly and cheaply. They will change the nature of work, of learning, of relationships.

This is precisely why we need to understand what computers are not, and can never be. For that, we must know ourselves. "We've turned away from exploring the human mind

just when it was getting interesting,” Gelernter says. But that’s not entirely true. He hasn’t.

This appears in the March 07, 2016 issue of TIME.