Ray Kurzweil's Impossible Vision

Erik Larson

Ray Kurzweil, he really does say the craziest things. So crazy, in fact, that if he weren't Ray Kurzweil, we might stop listening. Kurzweil, a computer technology genius and head of Kurzweil Technologies, an R&D technology company that specializes in Artificial Intelligence applications that include computer vision and speech recognition systems, has made a fortune showing the world just what computers can do. As an author and visionary in the field of AI, he's also had no small success telling the world what computers will do. His record thus far isn't bad: in 1990 he predicted the year a computer would defeat the world chess champion, and as anticipated in 1997 Gary Kasparov waved the white flag to IBM's Deep Blue supercomputer. That, and Kurzweil's status as a pioneer in AI software, has given his more radical predictions an air of added credibility. With that in mind, try this one: *By 2029, machines will be as smart as us.*

That is Kurzweil's claim in *The Age of Spiritual Machines* (a successor to his well-known *The Age of Intelligent Machines* published in 1990). A prediction—one among many—in his book that has elicited an understandably strong response. 2029 is not, after all, that far from now. Many of us will be here, living in that world. What sort of world will it be? Kurzweil is just getting warmed up. We shall be surpassed by machines shortly after 2029. Why? Computers are naturally more efficient at certain things, even today. Their memories don't fail them. And they share knowledge easily. Computer A downloads a program for recognizing speech patterns from Computer B. Now A knows it, and the laborious process that B undertook is not necessary for A. In contrast, I learn Russian and it takes me years of hard work. I can't give that to you for free (you can't just download my knowledge and skill); you've got to do the work too. So computers, once they reach the "critical mass" point where they can think like humans, will "soar past them", as Kurzweil puts it. In fact they will soar so high that by the end of this century, computers will be vastly smarter than people, and humans will be, truly, inferior second-class citizens. *Ipse Dixit*.

This sort of weighty, mind-boggling talk has inspired discussion from some rather unlikely contributors. Bill Joy, chief scientist and all around tech guru for Sun

Microsystems shared a table with Kurzweil in the lounge at the Gilder-Forbes *Telecosm* '98 conference at Lake Tahoe, after missing Kurzweil's speech there. The bar chat that followed culminated in Joy's now-famous Wired magazine piece (April 2000), an oddly personal reflection on the socio-cultural implications of Kurzweil's vision. Unlike the cheery optimism of Kurzweil, who sees the coming of age of non-biological human-like intelligence as a fascinating twist in the general evolution of intelligence, Joy sees these possibilities as nightmarish. He cites Unabomber Ted Kaczynski, whose own vision of the future (agreeing in fact though certainly not in spirit with Kurzweil's) had humans reduced to the "pets" of machines.

To Joy, the truth of Kurzweil's prediction would mean a shift from the "Nuclear, Biological, Chemical"(NBC) paradigm of 20th century weapons of mass destruction to that of "Genetics, Nanotechnology, and Robotics" (GNR). Joy contemplates all manner of engineered weapons using GNR technology, each with global holocaust potential, and each more feasible than yesterday's catastrophes. At least with a nuclear bomb the design, acquisition, and deployment are effectively impossible for most everyone. With GNR, the techniques and tools to construct a nefarious germ or non-biological "replicating assembler" (tiny replicating germs that could spread across the globe in a great suffocating goo) will spring from corporations, the instructions will be available on the Internet, the cost within reach of thousands. To Joy, Kurzweil's vision of intelligent machines is no spiritual image, no Michelangelo. It's Munch.

Joy has been busy working out his views on "relinquishment", meaning essentially that we shouldn't pursue development of technologies that will make Kurzweil's vision come true. Noted environmental activist Amory Lovins has teamed with Joy, and the lineaments of the latest anti-progress, government regulated "save us from ourselves" movement has emerged. So the story goes. But now the Discovery Institute, a not-for-profit think tank based in Seattle, Washington has weighed in with "Are We Spiritual Machines?" (Richards, Jay W., ed. Are We Spiritual Machines? Ray Kurzweil vs. the Critics of Strong A.I.. Seattle: Discovery Institute, 2002.), a collection of responses to Kurzweil from several experts in philosophy of mind, biology, mathematics, and chemistry. The compendium is intended as a critique of the so-called "Strong Artificial Intelligence", or "Strong AI" position championed by well-known thinkers such as MIT's Marvin Minsky (The Society of Mind), Douglass Hoffstadter (Godel, Escher, Bach), Daniel Dennett (Consciousness Explained), and of course Ray Kurzweil. (Kurzweil writes the introductory remarks, and also a response to each of his critic's sections, making the cover-to-cover reading a lot like listening in on a debate, only in which each participant gets a good while to talk without fear of interruption.)

Mr. Kurzweil, How Will It Happen?

The Strong AI thesis—the human mind can be reproduced by a machine—must be true if Kurzweil's prediction is to come true, in 2029 or whenever. An initial problem here is that the Strong AI thesis isn't just *obvious*; many people are convinced that it is downright false. In that case the timing of Kurzweil's prediction wouldn't be a suitable talking point (like arguing about how long it will take to reach the moon if one jumps from Earth); discussion of the viability of the prediction itself would be more germane. (A note: It is refreshing that "Are We Spiritual Machines?" addresses Kurzweil's position on Strong AI directly. Bill Joy and others who have assumed the truth of this thesis in order to move to discussions of its implications would do well to revisit or join these discussions.) Kurzweil is aware of the general skepticism about Strong AI and accordingly he spends a great deal of time in his opening remarks talking about developments in technology that will make the Strong AI position more appealing. Among these, he includes the rapid progress in computer chip design.

According to Kurzweil, gargantuan computing power (much more than currently available) is a necessary condition for realizing the Strong AI thesis. Blazing fast processors are necessary to run the massively parallel neural networks and "genetic" algorithms that will reproduce our minds. Indeed, a necessary condition for enabling GNR technologies generally will be fast computing, without which we cannot engineer complex non-biological things with nanotechnology nor manipulate huge amounts of genetic information to alter biological things. Kurzweil's vision therefore hinges on the plausibility of large, continuing increases of computing power tomorrow, even as we begin to push the limits of extant computer chip manufacturing techniques today.

Enter Moore's Law. Not a law at all, actually, but rather a prediction made by Gordon Moore, former chairman of Intel and one of the original designers of the integrated circuit. Moore suggested in the 1960s that computing power should double every 24 months, as new transistors were getting packed onto integrated circuits at that rate. Moore's "law" is in full effect; in fact in recent years the doubling rate has been every 18 months, give or take. Doubling every couple of years would produce magnificent computational power in just decades--if unabated. Problem is, we are approaching the physical limits of silicon-based techniques, after which without new methods progress will slow and eventually stall.

Kurzweil predicts the end of Moore's law as applied to current chip manufacturing techniques around 2019. But according to a broader principle of progress--the "Law of Accelerating Returns" as he has dubbed it--lithographic inscription of transistors on integrated circuits will be replaced by molecular electronics techniques that keep Moore's law rolling for another thirty years. By then, computational power will have doubled so many times that it will be possible to create humanly intelligent machines. (What is the "Law of Accelerating Returns"? Kurzweil's term of art to describe a process that does not run aground from lack of resources. Technology is supposedly governed by this principle, because as an evolutionary process technology grows more and more ordered while drawing on endless resources from the disorder in the broader environment.)

All well and good, but no one takes seriously the proposition that improvements in computer hardware alone will propel us to non-biological intelligence inside of three decades—certainly not any of the critics of Strong AI in "Are We Spiritual Machines?" (more on them soon), and not even Kurzweil himself. What is required is some roadmap—some blueprint—for intelligence that will give form and direction to all of this computational power. Such a roadmap is the "software" of intelligence—the algorithms and data structures embodied in code that convert various inputs into intelligent output. The Strong AI solution is to find a workable software design or architecture, and fill it in with the details until a fully functioning non-biological intelligence emerges. Kurzweil, a man who has made a successful business career out of designing computer vision and speech recognition software using, *inter alia*, parallel computing architectures such as “neural nets”—networks of nodes that are used to solve pattern matching problems by repeatedly computing and adjusting values assigned to the nodes—sees this approach as most promising for designing the software of human intelligence.

Neural networks. Neural nets were developed in the 1960s under the name “perceptrons”, and after gaining some quick recognition by initial successes they then sank into obscurity for a couple of decades in the wake of faltering progress, and a devastating critique (which the entire AI community apparently read) by AI pioneers Minsky and Seymour Papert in 1969, titled “Perceptrons”. But now the emergence of ‘back propagation’ techniques and other improvements in neural nets have pushed them back into popularity for solving certain classes of problems, particularly (or exclusively) pattern matching ones. And success in areas well suited for pattern matching solutions—computer vision for instance—has led many researchers to hope that node networks might provide answers to larger questions. Noted philosopher Paul Churchland, for instance, has adopted neural networks as a way of empirically developing and testing possible solutions to philosophic problems about the mind (the so-called mind-body problem) by accounting for its properties entirely by the mathematical functions of firing neurons—of which nodes are crude approximations.

Yet there is almost certainly more to human thinking than pattern matching (a point which Kurzweil does not apparently appreciate), and at any rate one doesn’t have to go swimming in deep philosophical waters to find limitations with neural nets—even in their revised, modern form. Minsky himself, in an interview with AI publication Generation[5], has pointed out that his original 1969 critique would still apply to the neural nets formulated today, and concludes: “Generally, the neural-network community has now silently recognized that more complex architectures are needed for harder problems, but they don't like much to talk about this.” None of this bothers Kurzweil in the least, by the way, because in suggesting neural network architectures will create human intelligence in machines, he assures us he means tomorrow’s neural networks that have been suitably improved—mostly by taking the investigation of real human neurons as a guide.

You Say You Want A Revolution…

One of Ray Kurzweil's charms is his ability to present an idea as if it is brand new, as if he had just pulled together the pieces days ago, and now it is a discovery, worthy of getting announced more than just argued about. That is, indeed, what helps make Kurzweil an engaging writer, a brilliant thinker, and a persuasive proponent of AI. These qualities are not unique to Kurzweil, of course, and are shared by many innovators and visionaries. But perhaps lesser appreciated is that his unbridled optimism about the future progress of the field of AI in particular is hardly original either. In fact, if one sees his position in historical context, Kurzweil is just the latest voice in a long chorus of wide-eyed optimism marking the strange, forty-year history of Artificial Intelligence.

Rewind. Forty years ago Nobel laureate and AI expert Herbert Simon won the prestigious A.M. Turing Award as well as the National Medal of Science (Kurzweil himself received the National Medal of Technology from President Clinton in 1999). That didn't stop him from declaring in 1957 that "there are now in the world machines that think, that learn and that create", and later to prognosticate in 1965 that: "[By 1985], machines will be capable of doing any work Man can do." In 1967, Marvin Minsky himself predicted that 'within a generation, the problem of creating "artificial intelligence" will be substantially solved' (he obviously did not think that neural nets were the solution!).

Kurzweil joined this chorus in 1990 with the publication of his *The Age of Intelligent Machines*, but by this time many of AI's strongest proponents had already abandoned the more visionary aspirations of the field, and had begun refocusing their research toward more modest ends, such as improving existing software systems and applications (an interesting field in its own right). Nonetheless, Kurzweil's books have been well received, and there is considerable interest in his ideas and their implications (witness Bill Joy's ruminations), which is (or ought to be) no small feat.

No small feat because the press--whatever press still covers the field of AI--got wind of its tribulations long ago. Although the arrival of the year 2001 caused a brief splash (symbolic because of Kubrick's depiction of the lip-reading super-smart computer

HAL 9000 in "2001, A Space Odyssey"), along with Spielberg's quixotic film "Artificial Intelligence" the same year, most of the coverage of AI has come to include the obligatory mention of the challenges the field faces, and the difficulties it has in making substantive progress. Kurzweil remains unfazed (manifestly unfazed!), buoyed perhaps by the success of his own forays into computer vision and speech recognition software, or perhaps because he somehow knows something the rest of the world does not. Does he?

Kurzweil and His Critics

One thing that "Are We Spiritual Machines?" does, invaluable to this entire discussion, is to bring fundamental philosophic and scientific questions to bear on Kurzweil's claims. That is apparent in the panel of Strong AI critics: Berkeley professor and renowned philosopher of mind John Searle, biologist and human geneticist Michael Denton, zoologist Thomas Ray (also designer of popular evolution software Tierra), and philosopher and mathematician Bill Dembski.

Take Searle. A key player in philosophical criticism of Strong AI since his 1980 paper "Minds, Brains and Programs" (*Behavioral and Brain Sciences*, Vol. 3), Searle argues in "Are We Spiritual Machines?" that Kurzweil's prediction will fail. Kurzweil fails to grasp some conceptual distinctions that are fundamental to assessing the possibilities for pure machine intelligence.

Searle asks us to imagine him locked in a room, receiving questions in Chinese through a slit in the wall. Searle does not speak Chinese, so the questions are meaningless to him. But Searle has a computer program for manipulating the symbols; when he receives a bunch of symbols on a slip of paper, he shuffles the symbols according to the rules he has (using the computer program to speed things up), and scribbles the results onto a piece of paper, handing it back through the slit. His shuffled symbols are answers in Chinese. According to Searle, since he was acting as a computer in the Chinese room, and didn't understand Chinese, so too the computer cannot be said to understand the symbols that it shuffles, though it may output correct answers in response to its input.

Searle's thought experiment is a popular philosophical conundrum, and it underscores a more general point: since symbols in a computer program don't literally "mean" anything to a computer, how can a computer be said to "know" anything at all? Isn't a computer *necessarily* mindless?

Searle suggests that Kurzweil is confused about other conceptual issues, such as the distinction between simulating and replicating a process. Suppose I believe Kurzweil and decide to download the entire informational contents of my brain onto a suitable computer storage media (in 30 or 40 years, say). I expect that when they hit the switch, if my brain was just a computational system anyway, the lights would "come on", and I would be conscious and functioning again, correct? Wrong. That is a confusion of simulation versus replication. A computer might have all my information—the neuronal connections, connection strengths, etc—and might perfectly simulate my brain's processes from this new medium (suppose), but that would not suffice to reproduce the actual causal powers of my biological brain. As Searle puts it, a meteorological program that simulates hurricanes by computing wind speed, direction, and so on, is not an actual hurricane. Or, a computer that simulates human digestion cannot be fed a hot dog—the simulation does not causally reproduce the actual digestive process.

There are opposing views on Searle's arguments, and some smart people have argued that the cases of cognition (thinking) and that of digestion (or what have you) are not analogous (see for instance David Chalmers' discussion of Searle's argument in his "A Computational Foundation for the Study of Cognition", http://www.u.arizona.edu/~chalmers/ai-papers.html). Nevertheless, the price of wading into deeper philosophical waters with ideas from engineering disciplines like computer science and AI is that you assume some degree of responsibility for addressing these issues on philosophical grounds. Kurzweil isn't quite up to this task at times, it would seem, and resorts to accusing Searle of assuming what he is trying to prove in his thought experiments. Then, in a particularly peculiar move, he supposes that the Chinese Room argument concedes the Strong AI point by allowing that a computer program inside the room could carry on a conversation in Chinese—thus demonstrating it had human-like intelligence. That, of course, was Searle's point—that a computer could be as effective as you wish and it would still not "know" anything! There is a real sense in which the computer scientist Ray Kurzweil and the philosopher John Searle are talking past each other.

There is no dearth of thought provoking discussion in “Are We Spiritual Machines?”, as each contributor brings his particular interests and expertise to the table. Among the most interesting (credit the Discovery Institute for including biologists here) are the depictions of the amazing properties of biological systems by Thomas Denton. Denton’s interest is in the organism versus machine comparison, and he brings to light some fascinating differences that make Kurzweil’s easy optimism about the quick convergence of man and machine seem a bit naïve. For instance, the ability to replicate through reproduction is commonplace in nature, and yet no machines come close to possessing this ability. And then there is “morphing”—fluidly changing from one shape to another. In discussing the capabilities of living things to morph themselves into different forms—take egg cells changing into “amoebic” cells, then into specialized cells such as blood cells—Denton asks us to imagine the machine equivalent: computers and appliances that could morph into wildly different things effortlessly, a computer into a helicopter for instance. It might as well be magic.

Thomas Ray in equally eye-opening discussion argues that the so-called Turing Test—a test for human intelligence in a machine based on the criteria that a human judge posing questions to a computer or person concealed in a room could not distinguish between them—is “one of the biggest red herrings in science” because machine intelligence will evolve (if it does) distinctly from human intelligence, in its own digital medium (perhaps on something like the World Wide Web). The “Turing Fallacy” as Ray puts it, is thinking that as computers progress, they will become more like us. Instead, they will become more like them, whatever that will be.

Apart from Ray’s qualms with Kurzweil’s human-centric vision of machine intelligence, he also strongly opposes Kurzweil’s suggestion that a human brain can be scanned and “copied” in a non-biological form. According to Ray, the structure and function of the brain cannot be separated, and the interdependence of the information processing capacity of the brain along with its essential life-support role in the overall maintenance of a human person make re-creation in Kurzweil’s sense an utter illusion. This proposal is “doomed to failure” according to Ray.

One of the most fundamental philosophic contributions comes from mathematician, philosopher and all around smart guy Bill Dembski, an associate professor at Baylor University who is best known as the gadfly of evolutionary theory for insisting that biological systems found in nature provide clear, scientifically justified evidence of "Intelligent Design." Dembski is unabashedly a Christian thinker, and his interest in the current debate is in Kurzweil's metaphysical presuppositions. Consider: we could not copy ourselves completely, including our conscious minds and souls, by copying the brain if the mind is not just the brain. Dembski suggests that Kurzweil is assuming without argument a metaphysical position of "scientific materialism", and only on this assumption does Kurzweil's claim that machines will be "spiritual"—everything about us will appear in our machines—have any grounds.

Perhaps the most profound suggestion in "Are We Spiritual Machines?" comes at the end of Dembski's section, where he suggests that the question of whether a machine can be just like a person (have our minds, our souls) is not a scientific question at all. It is a question of "ontological commitment", as Dembski puts it, referring to the philosophical study of what really exists, Ontology. If our souls are a different "stuff" than material and its operations, then purely material constructions of a person are not possible. Dembski hints here at a solution to the centuries old mind-body problem—the problem of how minds and brains interact. To Dembski, the problem is ill posed. At root is a model of science that seeks reduction of the contents of mind—purpose, intention, consciousness—to material and its processes. That view is myopic precisely because the ontological commitments of its adherents have closed off better alternatives.

One suggestion Dembski offers: instead of taking smaller and smaller stuff as primary (molecules to atoms, to…), we take the objects in our world as primary. Then we see that a person is not merely a collection of parts but also something in the world, a person. Then we see that persons (as we know) have real properties such as consciousness, and souls. The task then would be to integrate, or rather rebuild, a scientific methodology on top of this new metaphysics.

Kurzweil's Missing Link

Despite the perspicacity of these contributors, one can't escape the feeling at the end of "Are We Spiritual Machines?" that Kurzweil, in his responses, manages to sidestep his critics. A crafty arguer himself, Kurzweil manages to deflect much of this discussion as either irrelevant because advances in computing techniques will tame the thorny complexity issues (Denton and Ray's talk of the magnificent complexity and functional interdependence of biological organisms), or short-sighted because talk of limitations with today's technology is not tomorrow's problem. *Everything you say today is true. It won't be true tomorrow.* One feels that everything, and everyone, is getting swept up in the "Law of Accelerating Returns"; the ineluctable march of technological progress.

For many of us, this is not a satisfying result. After all, Kurzweil really has said something crazy: Machines will be as smart as we are by 2029. He means this full stop: our greatest art, our science, our thoughts and feelings, our souls will all be reproduced in a non-biological form by 2029. That certainly seems a bit too quick, doesn't it? Or perhaps we should say, unequivocally, that it is simply false. Let's do that. I won't bother telling my son in four or five years how he should prepare for the robot civilization after college. There will be more important things to worry about (perhaps, to take up a small piece of Bill Joy's concern, dangers involving genetic manufacturing of germs and viruses). And here I wish that the manifest falsity of Kurzweil's claim would have seen more of the light of day in "Are We Spiritual Machines?" The problem, mind you, does not lie in what was said by the likes of Searle, Denton, Ray, and Dembski, but rather what was left unsaid.

Kurzweil's research area in AI is in pattern matching, or slightly more specifically, in describing physical or "psychophysical" processes such as hearing and vision in the language of mathematical algorithms. There is an impressive body of work here (see Lloyd Watt's work on auditory processing (http://www.lloydwatts.com), or CalTech researcher Carver Mead's work on visual processing), and indeed, one can readily accede that real medical applications (e.g., hearing implants) are produced that rely on this research. That is the (mostly welcome) march of technology.

But, and this is the central Achilles heal of Kurzweil's broader AI fantasies—there is no known theoretical bridge from simple input-driven perceptual modules (hearing and vision) to human-level cognition such as thinking and reasoning. I'll repeat: *no theoretical bridge is known to exist between simple perceptual modules and human thought*. Indeed a growing body of literature coming out of cognitive science in the last decade suggests that such a bridge may not be computationally viable at all (see for instance, Fodor's "The Mind Can't Work That Way", MIT Press, 2000). Add to this some broader philosophical criticism on how we could ever get a computational reduction of the contents of mind—such as our beliefs that refer to things in the world—given some deep problems with developing algorithmic (read: computational) methods for resolving meanings (i.e. of words in a language), and one would think that there is much more than 27 years of research left on the Strong AI question! (This last point about the deep problems with reducing mind to computation has been made by the famous logician and philosopher Hilary Putnam. Ironically, Putnam helped launch the current "Computational Theory of Mind" by arguing for a mind-as-computer theory called "functionalism" in the 1960s).

What is ironic here—from the standpoint of Kurzweil's "law of accelerating returns" anyway—is that much of this head scratching and nay-saying has arisen in the last decade, when progress is supposed to really be picking up. But instead of solving the problems, one by one, many researchers "in the know" about these sorts of things are beginning to categorize: easy (solvable) problems go here, hard (unsolvable) problems over there. ("Psst. Make sure CNN only hears about progress on the solvable problems.").

Kurzweil's answer to this—to the extent that he has recognized it as a problem—is that we will keep extending research on local brain functions ("modules" in the language of philosophers and cognitive scientists) to include more and more of the brain, until it is understood mathematically in totality. But that simple additive approach to solving the types of problems encountered by cognitive scientists pondering how we reason with a set of changing beliefs (i.e., how we think like people), seems woefully inadequate. It doesn't even begin to take seriously the conceptual distinction between perceptual processing in local brain regions and the sort of "global" context-sensitive inferences we make in deliberate thought. But these global, context-sensitive inferences are really the essence of human intelligence, and they (not perceptual processes such as vision or hearing) are precisely what must be demystified by Kurzweil to give any substance to his claims.

Kurzweil has not offered any plan for demystifying the "hard" problems of how we think, and in fact he skirts the serious philosophic discussions (e.g., from Fodor, Putnam and others) about reducing mind to computation. More specifically, he skirts what might be called the problem of "epistemic context"—grasping what is relevant in a situation such that the entire space of one's beliefs need not be searched in order to make (sometimes) simple inferences. This philosophical conundrum of how we understand context is cross-referenced in AI research as the "frame problem." Call it what one wishes, it is a core problem that is almost certainly not solvable by increases in computing power alone, nor by simply scaling mathematical treatments of local brain processes.

The reasons for this are involved and would require lengthy explanation, but here is a synopsis: any procedure that is locally specified (read: written out in a computational format) seems, *ipso facto*, an impossible candidate for performing global belief-based reasoning (in other words, for performing everyday human thinking!). To put it another way, the construction of a computer program—using any architecture we choose—becomes a candidate for interesting cognition only by "knowing" many facts about its world. But by the very fact of its having such deep, rich knowledge, it apparently won't have any means of picking out the small, relevant subsets of those facts needed for particular instances of reasoning. It is not simply that this problem is "hard", and will take a while to solve, but rather that it seems to lack a solution entirely, given our concept of computation. And since we really don't understand at all how a computer could ever reason in context, *a fortiori* we have no clue how a computer could ever reason like a person. So not only is there no known theoretical bridge to reproducing genuine human thinking, there is no such *imagined* bridge either! And from all of the engaging discussion in "Are We Spiritual Machines?", I wish that this problem had been put more directly to Kurzweil. I would love to hear the answer.

Still, what "Are We Spiritual Machines?" accomplishes as a compendium of thought on Kurzweil and Strong AI is a commitment to viewing such ostensibly technological endeavors through the broader lens of philosophical criticism. These discussions are welcome and indeed unavoidable, because in the end what is at issue is not the merely the truth of Kurzweil's prediction but indeed the reality of ourselves, our very natures. Kurzweil, for his part, remains the true technology visionary, insisting all the while that his critics suffer from myopia with regard to the future of technology. Perhaps. But perhaps, as Dembski suggests, the whole question of a computational solution to the mind runs aground on an imaginative failure of a different sort. That failure, as he put it so well, may be a failure to see the world itself as it really is.