

CHALLENGES TO THE CREATIVE MIND

Computers are everywhere—in the supermarket and the bank, under the hood of your new car, in your third-grader's math class, all over the office, and, much too often, on the other end of the phone. We have suffered these invasions of our daily lives with indifference, mild annoyance, even disdain. Computers, we say to ourselves, are only tools. They can follow the rules engraved in their programs; they cannot be creative.

Such is the uneasy truce that common sense has negotiated with the computer. On one side lies machine intelligence, a low mechanical cunning—repetitive, slavish, infallible, but inherently unimaginative. On the other, the human mind—error prone, inexact, wracked by distracting emotion, but capable of great and unpredictable leaps of intellect and fancy, capable also of commitment and care.

These caricatures of mind versus machine work well enough for our daily diplomacies with the automated bank teller, the word processor, and Pac Man. They grant a comforting conviction that the relation between man and machine is fixed and certain: We are the programmers, they are the programmed; we make the rules, computers follow. Whether or not this line holds, there is mounting reason to suspect that continued defensiveness toward computers is not in our best interest. The plain fact of the matter is that computers are now involved in complicated ways in many of the previously privileged areas of human creativity: fine art, music, philosophy, education. In this special issue of *Boston Review*, we take stock of these new encroachments on the line between machine and mind, and we re-examine the old metaphor of programmer and programmed.

Our report from the front is not, we think, without surprises for the skeptic or cautions for the true believer. We have found no sci-fi worlds where

computers create works of art or think thoughts on their own. But we have also found that the limits on computer creativity and consciousness may lie less in the nature of computers than in the immaturity of our own insights into artistic creativity and human intelligence. If we understood ourselves better, we might make machines more creative and smarter than they are.

In the current state of computer art, the relation between man and machine is best understood as a partnership. Computers remain the limited partner, providing a new medium for the artist, a new instrument for the scientist, and a new model of mind for the psychologist and the philosopher. The computer's role in these partnerships is to amplify human intellect, not to replace it. Still, such partnerships are not entirely innocent. As Joseph Weizenbaum and Daniel Dennett each make plain in their contributions to this issue, we pay a price for the intellectual prosthesis that computers can provide. As computers expand our powers, so may they limit our vision. As we program the computer, so may we be programmed. To expand the creative capacities of the computer without distorting our own—this may be the computer's real challenge to the creative mind.

The articles in this issue examine that challenge from the informed perspective of those who live with it in their art, their music, their philosophy, and their science. We hope this unique look at the creative partnerships now developing between mind and machine will amplify our own understanding of their dangers as well as their opportunities.

Eric Wanner

Eric Wanner is a psycholinguist whose next book, Language Acquisition: The State of the Art, will be out in the fall from Cambridge University Press.

Notes on Prosthetic Imagination

Daniel C. Dennett

According to Aristotle, the Pythagoreans thought that the heavenly bodies, being so large and moving so swiftly, should make a mighty—and harmonious—noise. "The music of the spheres" was not a metaphor at birth. Why can't we hear this music? The Pythagoreans' ingenious explanation was that "the noise exists in the very nature of things, so as not to be distinguishable from the opposite silence." What is truly pervasive, they realized, is virtually unnoticeable. It took the genius of Newton to "notice" gravity as a force, which might be stronger or weaker in unexplored parts of the universe. Such insightful strokes—the bogus as well as the real—arise more from the exercise of imagination than from either perception or calculation. Thanks to technology, science now has a panoply of prosthetic devices for extending our perceptive and calculative powers: We can see far and small and infra-red, and we can perform literally billions of mathematical calculations in a few seconds. If only there were a device that could prosthetically extend our powers of pure imagination! But there is: the computer.

Daniel C. Dennett is the author of Brainstorms and, with Douglas Hofstadter, of The Mind's I (Basic Books). He teaches philosophy at Tufts University.




Illustration: James Steinberg

What I have in mind is quite unlike the magnification of our intellectual range by earlier technological developments, though their benefits have certainly been great. The microscopist places some hunk of the world in the microtome and slices off a fragile, translucent cross-section, revealing to the light (and then, through the microscope, to the human eye) patterns previously outside our cognitive scope altogether. The photographer slices off a momentary cross-section of temporal process, suspending all animation, and permits the human eye for the first time to dally over the details, soaking up information from the still picture at the particular human rate that defines our epistemological horizons. That rate of perceptual intake has both an upper and a lower bound. Processes that occur too slowly relative to that rate are just as invisible to us. Time-lapse photography, by speeding up the relatively glacial changes, permits us to respond perceptually to the rhythms and patterns of plant growth, cloud formation, decay, and other slow processes that are otherwise invisible.

Through computer graphics, these perceptual horizons can be further expanded. Since any feature of the world can now be "color-coded" and represented on a screen, changing at whatever rate suits the human observer, we can use technology to see, quite literally, labor migrations, changing geographical concentrations of Baptists or adolescents or millionaire women executives, heat loss from private homes,

and whatever else interests us. But what should interest us? That is a question posed to the imagination.

Using computer graphics to spoon-

sternation, a sort of plaid began to appear, growing in resolution as thousands and thousands of color dots were added. There were hot spots and dark

Finding a pattern in a superficially patternless sequence can be extraordinarily difficult, but the human eye does it effortlessly under the right con-


wonderful plaything for scientists. What about using visual exploration to study literary style? Computer-aided analyses of word choice and style to date have been rather clanking and statistical affairs, but that could change. The delicate, subliminal effects of rhythm and other stylistic devices—often quite beneath the conscious threshold of their creators—can perhaps be magnified and rendered visible or audible. This would not be like looking at a painting under a microscope, for the features that would be intensified and raised to human ken would not have to be physical features

"Purely abstract phenomena can be turned into objects of human perception by the computer."

feed information to our idiosyncratic visual systems is not restricted to information about physical processes and events. Purely abstract phenomena can

regions that repeated themselves with perfect regularity across the field. These had to be the visible effects of a probabilistic bias in the random number

ditions—especially if those conditions involve symmetry, motion at the right rates, or color contrasts of the right sort. Could such a graphics system



be turned into objects of human perception by the computer. Late one night several years ago I was visiting with a few friends at MIT's Artificial Intelligence Lab. We were playing around with a color graphics system that projected its displays onto a large Advent TV screen, and someone had the bright idea of using the computer's random number generator to create a visual display. A different color was assigned to each digit from zero to nine, and a simple program randomly sprayed these different colored dots in an array over the field, growing uniformly in density as time elapsed.

As the first few hundred dots appeared, they looked just the way we had expected them to look: multicolored snowflakes, with no discernible pattern. As the density increased, we should have seen first more or less uniform confetti (with a few random blotches), and finally, a sort of *pointilliste* beige (or whatever color should emerge from the uniform mixing of those ten colors). Instead, to our con-

generator, or in our use of it to cover the field. There were regions in the field where (red) sevens were more frequent than the other digits, and hence were slightly more predictable. Anyone "betting on red" who bet only during periods when those regions were being peppered would have an edge. So either

the random number generating system itself or at least our apparently careful and innocent way of exploiting it was seriously flawed. (Were it to be used thus to provide a random background against which to search for some pattern in nature, its own spurious pattern could be mistaken for a pattern in nature.)

become an investigative tool in a new sort of quasi-empirical mathematics? Several years ago it was discovered that prime numbers fell into remarkably regular patterns if the sequence of numbers was displayed in a tight, rectangular coil. Exploring this fact visually was made possible by a computer

of the work, but abstract patterns, biases of connotation and evocation—of meaning, not matter. Many people would just as soon not learn whatever secrets lie hidden in this domain, for it smacks of murderous dissection of art. I think they underestimate the power of the art they wish to protect. (Some

"Most Artificial Intelligence programs look like preposterous Rube Goldberg machines when you study their details in slow motion, as it were, but when they run at lightning speed they often reveal a dexterity and grace that appears quite natural."

display. Are there hundreds of other telling patterns in the number system waiting to be discovered as soon as features of large portions of that system are made visible?

What other apparently featureless or complex systems will yield their secrets to this new variety of visual search? There is no reason to reserve this

chess enthusiasts have feared that computer chess programs would somehow "crack" chess by discovering a foolproof strategy for chess, and turn it into a game of no more interest than ticktacktoe. It hasn't happened. Even if a computer program soon becomes world champion in chess—a distinct possibility—the games it will play will

no doubt exhibit the same sorts of beauty as the great chess games of the past. Perhaps we will even learn to see deeper into the astronomical vastness of the game.)



Since our powers of imagination depend critically on our powers of perception, these novel perceptual opportunities are one way of pushing the imagination into new places. The front-door approach, one might say. But

effect; what we call common sense is composed largely of popularized borrowings from ancient and recent science. Computers, however, can lead their users not only to strikingly new

through the antechamber of the word processor, noting where I am by the familiar billboard—the “no file menu,” it’s called. When I get to the system level, I can recognize my location by

portions of a logical space that has become as vivid a terrain of activity for me as the corridors and offices of the building in which I work. Moving around in this logical space involves performing

“Computers can lead their users not only to strikingly new concepts but to new ways of using concepts—playing with them, tinkering with them, perceiving and ‘handling’ them.”

concepts but to new ways of using concepts—playing with them, tinkering with them, perceiving and “handling”

another landmark, the system prompt “A>” that appears on the screen. Without such landmarks I would soon

a new variety of abstract locomotive acts, and just as the pianist becomes oblivious to the individual finger mo-




Illustration: James Steinberg

what we can imagine or experience also depends on what concepts we have. As Kant said, intuitions without concepts are empty. My hunch is that in the long run, computer-borne concepts will expand our imaginative powers even more than computer-aided perception.

them. A sort of fluency and familiarity can be achieved that *feels* quite novel.

What I have in mind can be illustrated with even such a nontechnical device as a word processor. Here I am typing along on my microcomputer keyboard, watching the letters appear on the screen. I am “in” the word processor

be lost in the maze, and knowing where you are is very important; in this magical world, simple actions have different effects in different places.

From the system level I can open a trap door of sorts, and descend, like a nosy passenger on an ocean liner, into the engine room. I can see by pressing a

tions making up a well-practiced arpeggio, so leaping about in the logical space “inside” the computer comes to be done by acts that are decomposable into their constituent keystrokes only with conscious attention and effort.

The reason these places are just logical places and not—at least not *necessarily*—real, physically locatable places in the hardware, is that when I interact with the computer, I am permitted by the design of the system to conceive of my interactions at a high level: the level of a “virtual machine” that is much “friendlier” than the actual hardware.

What is a “virtual machine”? Imagine a man with a broken arm in a cast. The motions his arm can go through are strictly limited by the physical rigidity of the plaster cast. Now imagine Marcel Marceau miming the man with the broken arm. His arm *could* move in ways it now *won't* move, given his temporary intentions; he is mimicking the behavior of the first man by deliberately adopting a certain set of rules of

“Imagine Marcel Marceau miming the man with the broken arm. His arm could move in ways it now won't move, given his temporary intentions; he is mimicking the behavior of the first man by deliberately adopting a certain set of rules of motion for his arm. He has a virtual cast on his arm—a cast ‘made of rules.’”

Computer science and its allied fields have added a host of concepts to the working vocabulary of almost everyone. Shallow notions of programming (and “deprogramming”), software, and information storage and retrieval have already established themselves at the lowest, broadest level of popular culture. There is nothing novel about that

(that’s the program I’m interacting with). More specifically, I am “in” a file I have called “BOSREV,” which contains the draft of my notes for the *Boston Review*. But now I am going to “leave” for a moment, and go back out to the system level, a sort of vestibule from which I can enter other programs. On my way to the system level, I pass

few keys that the “A” disk drive is running 740 microseconds fast. This is perplexing territory for me, which I will leave to the crew, so I quickly return to BOSREV where I can compose these recollections of my journey.

These places are not different physical locations inside the computer on my desk. They are *abstract* places,

motion for his arm. He has a virtual cast on his arm—a cast “made of rules.”

Suppose you wanted to put a seesaw

raising of the other side at just the right rate. You’d have a virtual hole in the wall and a virtual solid plank. A virtual hole doesn’t have all the causal proper-

ties of a real hole. You can’t shine light through it or squirt water through it, for instance. It has only all the properties

of-girders, pulleys with pulley-descriptions—together with some compensatory machinery at the edges (our tipper-uppers and their kin).

expectations, and emotional reactions of the reader, as these contribute to the aiming of her eyes. Imagine a poem that would be invisible to anyone who did not read it at the right speed. Imagine a poem that was many different poems depending on how the reader dwelled on the words. The hyperbole of those critics who insist that the poem is a different text for every reader can be made the literal truth if we wish. Consider how a reader might react to a text that absolutely could not be re-read: The pages dissolve into nothingness for that reader once they are turned; any attempt to go back leaves the reader with a somewhat different version of the story—like trying to catch up with the shadow of the past in recollection. The encounter-

“If I tell you Smith ate the last piece of cake, I don’t have to add that Jones didn’t eat it because you know that only one person can eat the same piece of cake (unless he is sharing, etc.). But that is just the sort of thing a computer stumbles over.”

through a wall, so that a child in one room could sit on one end and a child in the other room could sit on the other. One way of doing it—the obvious, physical way—would be to cut a horizontal slit through the wall, and place a stiff plank halfway through the slit. But suppose the wall was impenetrable. You could achieve the same effect by cutting the plank in half, and attaching each half to an electrically powered hydraulic tipper-upper device butted up against the wall, and connected to its mate on the other side of the wall by a purely informational “feedback” link. As one side went down, information describing its course would control the

needed for a certain application. If you know exactly which causal powers you want to preserve, you can replace hardware with software, girders with rules-

ties of a real hole. You can’t shine light through it or squirt water through it, for instance. It has only all the properties

Virtual machines are of great utility to the scientist, but who else might want to use them? Since they are as

“Learning a new system of concepts, like learning a foreign language, changes one’s habits of thought—breaking some bad habits, but only by creating others.”

needed for a certain application. If you know exactly which causal powers you want to preserve, you can replace hardware with software, girders with rules-

versatile as our powers of imagination, they can be used for anything imaginable. We can make new musical instruments that depend on “physical principles” that we make up for the occasion. We can make fantastical virtual environments populated with virtual creatures of our imagination. (Yes, I know, the early fruits of electronic music and “adventure” computer games are not in general inspiring works of artistic genius, but then the first tunes played by dragging a bow across a string were probably not all that prepossessing either.)

There are other possibilities still to be explored. Consider the strange promise of the device I call the Berkeley machine, after Bishop Berkeley, the philosopher. This actually exists at the University of Illinois, where it was developed to study the psychological processes involved in reading. When a person reads—or just looks at the world—her eyes dart quickly about in a “ballistic” leap known as a *saccade*. While the eyes are in motion between resting points, nothing is seen. When a person sits in front of the Berkeley machine reading the text displayed on the video screen, a device detects her eye movements so accurately and swiftly that it can predict where her eyes will alight and focus *before the saccade is completed*. Quick as a magician, the machine erases the word that has been sitting at ground zero for the saccade and replaces it with a different word! What does the reader see? Just the new word—and with no sense of anything having changed. As the reader peruses the page, it seems to her for all the world as stable as if the words were carved in marble, but to another person reading the same page over her shoulder (and saccading to a different drummer), the page is aquiver with changes. The apparently stable world of persisting words is created to order for the wandering eye of the reader; the words’ *esse is percipi*. The resulting world can be an intimate function of the attitudes,

machines one might create with such techniques would not be novels, or poems, or plays, but something new. There is not even any point in trying to name the genre until we see what works emerge.

Viewed one way, virtual machines are just a special case of what happens in every computer simulation: The “physics” of the simulation is exactly and only what you declare it to be. As the psychologist Zenon Pylyshyn has said, you don’t get any of the physics for free. It is this feature, more than any other, that turns computers into imagination-machines, for it forces one to notice the normally unnoticeable. The dog that *doesn’t* bark is often the important clue.

This comes out particularly vividly in Artificial Intelligence (AI), where the task is to design a “smart” system of one sort or another. Since one starts with a blank slate—the totally ignorant agent—every logically independent bit of requisite knowledge must somehow be inserted. What does the artificial agent need to know to do its job? Well, what does a person need to know to do that job? The rather brief list one first comes up with almost inevitably falls far short, for among the things a person needs to know are things so colossally banal as to be entirely beneath our notice. If I tell you Smith ate the last piece of cake, I don’t have to add that Jones didn’t eat it, because *you* know that only one person can eat the same piece of cake (unless he is sharing, etc.). But that is just the sort of thing a computer stumbles over, and in stumbling, it draws our attention to features of the world beneath our normal notice, jogging our mind-sets in ways that are often—but not always—fruitful.

Another gift to the imagination from

computer simulations is their sheer speed. Before computers came along the theoretician (in psychology, for instance) was strongly constrained to ignore the possibility of truly massive and complex processes because it was hard to see how such processes could fail to appear at worst mechanical and cumbersome and at best vegetatively slow. Of course a hallmark of thought is its swiftness. Most AI programs look like preposterous Rube Goldberg machines when you study their details in slow motion, as it were, but when they run at lightning speed they often reveal a dexterity and grace that appear quite natural. This can often be a dangerous illusion of verisimilitude, but at least elbow room has been created for the imaginative exploration of gigantic models of thought and perception.

No one is more in need of this elbow room than philosophers. It is a sobering but fascinating exercise to canvass the history of philosophy for occasions on which philosophers have mistaken their incapacity to imagine or conceive this or that as an insight into *a priori* necessity. It has often been remarked that inconceivability is our mark—really our only mark—of impossibility, and philosophers often claim that something is necessarily nonexistent by observing that they cannot conceive of it. There is no way out of the *hubris* implicit in such a claim. What we can't conceive, we can't conceive, and what better grounds could we have for deeming something impossible? But at least we can improve our position by trying harder. We can make a deliberate effort to find ways around the obstacles that our current conceptual scheme has placed like walls around our view of the possible.

A nice case in point from the history of science (and philosophy) is the debate that raged over the mechanism of heredity before the advent of modern microbiology. Did the sperm contain a tiny person, which in turn contained still tinier sperm which contained still tinier persons and so forth *ad infinitum*? That was surely an impossibility, wasn't it? Then there must be a finite number of these nested Chinese boxes, and some day all the progeny of Adam and Eve will exhaust the supply. In retrospect, this is a somewhat comical argument, but at the time it was surely a compelling appeal to intuition. What else could *conceivably* explain the inheritance of features? The majestically complex story of DNA and its surrounding micro-machinery was simply unimaginable at the time. It is still virtually unimaginable, except with the aid of all our technological prostheses and the theories they have permitted.

What moral do I draw from this? That philosophers who ignore the available aids to imagination may relinquish their right to be taken seriously when they proclaim their insights into metaphysical necessity.

Computer simulations, like virtual holes-in-walls, can be very misleading. It is tempting to read more into the observed behavior of the system than one should, and to suppose the phenomenon one is apparently observing is more robust and multi-dimensional than it is. No doubt every prosthesis has its shortcomings, and we have learned to be wary of the distorting powers of lenses, microphones, and other extenders of our senses. But the potential of computers for treacherous misdirection in their role as imagination-extenders has not yet been assayed. The microscopist pays for magnification with a variety of tunnel vision; what are the comparable costs to those who come to rely on computers to extend their thought?

I have a confession to make. I don't yet know the computer language LISP, the *lingua franca* programming language of AI. Since I have often written about AI, that is a little like a Kremlinologist admitting he can't speak Russian. "Learn LISP, Dan," a person in the field urged me the other day. "It will change your life." That's what I'm afraid of. Another adept in the field has spoken to me of the need to "marinate" in LISP for a year or so, before the full effects on one's conceptual powers is felt. This would give my imagination a new flavor, I'm sure, but would it also tenderize some of my favorite mind-fibers in ways I would regret (if I could notice them)? How could one tell?

In *The Defense*, Vladimir Nabokov presents a vision of the world as seen through the eyes of a chess grandmaster, Luzhin. The people and furnishings of this world array themselves in rows and files, protecting one another from capture: a shop is just a knight's move across the street, and perhaps one can get out of debt by casting. So obsessively immersed is Luzhin in the system of chess concepts that his imagination has been crippled. Learning a new system of concepts, like learning a foreign language, changes one's habits of thought—breaking some bad habits, but only by creating others. In the process, one changes one's cognitive style—one's style of thought. That is an ill-understood change, and so far we don't have a vocabulary for describing such phenomena.

But that is just the *sort* of vocabulary that bubbles up in apparently inexhaustible profusion from the daily practice of the artists and artisans of computer science. And once the terms have been given life in practice, one can abandon the context of their birth—in this case the computer. Many of the leaders in the field of AI are no longer writing programs themselves: They don't waste their time debugging miles of code; they just sit around thinking about this and that with the aid of the new concepts. They've become—*philosophers*! The topics they work on are strangely familiar (to a philosopher) but recast in novel terms. There is a good deal of exploration to be done with the aid of these new concepts. I think I want to come along for the trip, and not just as a tourist, so I guess I'll have to learn LISP. I hope I don't come to regret the bargain. ■