Converting Matter into Mind

William A. Dembski

Introduction

In the *Foundations of Cognitive Science* Herbert Simon and Craig Kaplan offer the following definition:

> Cognitive science is the study of intelligence and intelligent systems, with particular reference to intelligent behavior as computation.

Since this definition hinges on the dual notions of intelligence and computation, it remains scientifically unobjectionable so long as one declines to prejudge the relation between computation and intelligence. As long as the cognitive scientist refuses to prejudge this relationship, his scientific programme assumes the following valid form: he considers presumed instances of intelligence in the world and seeks to model them computationally. This programme takes computation as a convenient paradigm for examining intelligence and then pushes the paradigm to as comprehensive an account of intelligence as the scientific data will allow. If a machine can be constructed which captures (or even extends) the full range of human intelligent behaviors, then the paradigm is fully successful. To the degree that machines fall short of this goal, to that extent the paradigm is unsuccessful, or has failed to realize its potential. Together with the foregoing definition, this approach to intelligence via computation puts cognitive science within the bounds of genuine science.

Now it is possible to prejudge the relation between intelligence and computation. Thus one can presuppose that computation comprehends all of intelligence. Alternatively, one can presuppose that intelligence can never be subsumed under computation. These assumptions have been and will continue to be the source for much fruitful discussion. Such a discussion will be interdisciplinary: to this discussion mathematical logic contributes recursion theory, physics prescribes limits on computational speed, philosophy lays out the mind-body problem, theology raises the question of immaterial souls and spirits, etc. But while all these disciplines inform the debate over the respective boundaries of computation and intelligence, it must be realized that such a debate is primarily philosophical and thus independent of cognitive science qua science. If the director of Carnegie-Mellon's Robotics Institute, H. Moravec, is right when he predicts that in the next century robots will supersede the human race, then this discussion will come to an end, being decided in favor of the view that computation subsumes intelligence. But for now Moravec is playing the prophet. Even this would not be reprehensible, if Moravec were wearing the prophet's mantle. Unfortunately he is wearing the scientist's lab coat, thereby conflating cognitive science qua science with a materialist philosophy of mind.

Cognitive science is legitimate science when it takes an unprejudiced view of the relation between computation and intelligence. Nevertheless, since cognitive scientists as a group are notorious for deciding the issue in advance, I shall henceforth refer to cognitive science qua science as the *science of cognition*. Thus I shall use the phrase *cognitive science* pejoratively, implying that science and philosophy have been conflated because intelligence was prejudged as a form of computation. My view is that cognitive science stands to the science of cognition much as alchemy stood to chemistry. Certainly the alchemist's appeal to magic renders him more ridiculous to modern eyes than the cognitive scientist's appeal to a well-established materialist philosophy. But to my mind the cognitive scientist's conflation of philosophy and science is no less damaging to the science of cognition than the alchemist's conflation of magic and science was to chemistry. The fault of the cognitive scientist does not lie in his being simultaneously a philosopher and a scientist, but in not telling us when he is serving in which capacity. My purpose in this article is to untease that tangled web of philosophy and science which constitutes cognitive science.

The Parable of the Cube

In the *Foundations of Cognitive Science* Simon and Kaplan also offer the following account of artificial intelligence (AI):

> Artificial intelligence is concerned with programming computers to perform in ways that, if observed in human beings, would be regarded as intelligent.

This account is scientifically unobjectionable and assigns to artificial intelligence the main practical business of cognitive science-programming computers to perform tasks thought to require intelligence in humans. Nevertheless, for the cognitive scientist who has prejudged the relation between intelligence and computation, the very phrase *artificial intelligence* becomes tendentious, implying that artificial intelligence has subsumed the whole of human intelligence. Thus cognitive scientists see no way of drawing a fundamental distinction between human and artificial intelligence-with strong emphasis is on the word *fundamental*. The old degree-kind distinction is implicit here. Animal, human, computer, and indeed any finite discursive intelligence (to use a Kantian phrase) become from the point of view of cognitive science instantiations of algorithms. Eventually I shall return to these points. But for now I want to focus on two questions: (1) What is so special about computers that they should constitute the exclusive tool of AI? (2) Why should we expect AI to give us any insights about human intelligence? To answer these questions the idea of a sufficient cause for an intelligence becomes important. To appreciate this idea, we consider two stories, the first a yarn about an imaginary cube, which I call the Parable of the Cube; the second Thomas Huxley's bizarre tale of monkeys with typewriters.

> Imagine you are given a box with one transparent side. Inside the box is a small black cube. Both box and cube are made out of plastic. The box is placed on a viewing stand with the transparent side facing you, much like a

> television. Now you watch. The cube moves around inside the box. Sometimes it is in this corner, sometimes in that. At other times it hangs in mid-air. Yet again it hurls itself against a side of the box. The sides are sturdy and do not break. What's more, the box has been soundproofed, so you cannot hear the little cube bouncing around. How exciting, you say. You are not convinced that the cube's entertainment will rival the television networks.
>
> Suppose next that the cube divides its time between the left and right side of the box. Back and forth it moves. For a time you are hypnotized. Your eyes glaze over. What a dull pastime. Gradually, however, you notice a pattern. You time how long the cube spends on the left. It's always one beat or three beats. Suddenly you remember your Morse code. Behold, that little cube is communicating with you. And not just any old communication. The cube is reciting *Hamlet*-in Morse code. But this is just the beginning. News, mysteries, stock predictions, and soap operas are all part of your newfound entertainment package. In the light of this discovery your television has become passé. Cube watching is now the rage in your household.
>
> The story doesn't end here. Your neighbors start wondering why so much scratch paper is strewn around your home. Obviously you have been receiving coded messages and converting them to English. Soon the secret is out-you have an intelligent cube. People are in awe. They line up outside your doorstep to record the pearls of wisdom that are dropping from your cube's lips, so to speak. The cube has become more than entertainment. It has become a religious guru, expounding the mysteries of religious cubism. This is not simply a smart cube, this is a wise cube. Demand is such that you take the cube and its box on a speaking tour (well not quite, you know what I mean). The cube is hailed as the savior of mankind, its wisdom the uncreated light of the ineffable power. In the end all nations bow down and worship the cube.

In line with the Parable of the Cube let us recall Thomas Huxley's simian typists. Thomas Huxley was Charles Darwin's apologist. Darwin's theory of speciation by natural selection sought at all costs to avoid teleology. The appeal of Darwinism was never, That's the way God did it. The appeal was always, That's the way nature did it without God. Thus one looked to chance, not intelligence, to render Darwinism plausible. Huxley's simians were to provide one such plausibility argument. Huxley claimed that some huge number of monkeys typing away on typewriters would eventually (where "eventually" was a very long time) type the works of Shakespeare. If one assumes the monkeys are typing randomly, not favoring any keys, and not letting one key stroke influence another, Huxley's claim is a simple consequence of a fundamental theorem in probability known as the Strong Law of Large Numbers. Indeed, given enough time one can expect the monkeys to type all the great works of literature, though the bulk of their output will be garbage.

Even with trillions of monkeys typing at blinding speeds over a time span comprising many lifetimes of the known universe, the probability of randomly typing *Hamlet* is still vanishingly small. Thus it is arguable whether Huxley's apologetic for Darwinism was in any way cogent on probabilistic grounds. But the question that is too frequently glossed is, What determines whether the monkeys have finally typed *Hamlet*? The monkeys are assumed unintelligent. Hence they cannot stop and deliver a copy of *Hamlet* when after aeons it finally appears. No. Some intelligent being must examine all the monkeys' output, wade through all the garbage, all the false starts of *Hamlet*, until finally this intelligence comes across a finished copy of *Hamlet*. Now it does no good to claim all that is needed is a simple computer program which has a stored copy of *Hamlet* and compares the monkeys' output with the copy. This merely begs more questions-What intelligence wrote the program? What intelligence installed a copy of *Hamlet* in the computer's memory? Where did the intelligence get the copy of *Hamlet* in the first place?

Humans naturally see meaning and purpose in a work of literature like *Hamlet*, just as they see meaning and purpose in the organisms of nature. What Huxley hoped to show was that such meaning and purpose, Aristotle's teleology and final causes, were in fact illusory. Intelligence was not in any way prior to the random processes of nature. Rather, intelligence was itself a product of nature's randomness, constructing meaning and purpose after the fact. Still, the critical question remains, What intelligence decides whether the monkeys have finally typed *Hamlet*? Without an intelligence to interpret the monkey's output and distinguish the intelligible from the inane, the monkeys will type indefinitely, with one output as inconsequential as the next. Let me put it this way. Huxley's example presupposes an intelligence familiar with the works of Shakespeare. At the same time Huxley wants to demonstrate that random processes, the typing of monkeys, can account for the works of Shakespeare. Thus Huxley's example is supposed to show that the *works* of Shakespeare can be accounted for apart from the *person* of Shakespeare. Huxley wants it both ways. An intelligence must be on hand to know when the monkeys have typed *Hamlet*, and yet *Hamlet* is to stand in need of no author. This is known as having your cake and eating it. Polite society frowns on such obvious bad taste.

It's no surprise that the humanities have a hard time with rabid AI propagandists. Beethoven would not have suffered being told his Ninth Symphony was possible without him. Given Beethoven's high opinion of himself, I am confident of this assertion. As for Shakespeare being told *Hamlet* could make do without him, I'm not sure whether his reaction would have been displeasure or amusement. True artists know that their work is not reducible to any other categories, least of all chance.

Let us now return to the Parable of the Cube. The cube signals intelligent messages in Morse code. Is the cube's signaling spontaneous or does an extrinsic intelligence guide it? Since both the cube and the box are plastic, and since plastic has to date indicated a marked absence of intelligent behavior, we are apt to conclude that the intelligence is extrinsic. Again we infer an intelligence. We do not consider the cube a sufficient cause for the signaling of *Hamlet*, just as the typing of monkeys is an insufficient cause for *Hamlet*. In both cases we have physical systems which express intelligence, but which fail to supply an adequate causal account of the intelligence they express. Can we find a physical system which simultaneously expresses intelligence and provides an adequate causal account for this intelligence? The obvious place to look is the human body, and specifically the brain.

Nerves and Brains

Cubes in boxes and Huxley's simians are examples of physical systems which insofar as they express intelligence fail to account for intelligence. In what way then does the physical system constituting the human brain differ? Why do people attribute thought and intelligence to the matter constituting their brains? Certainly there is a causal connection between brain and behavior. Certainly there is a link between brain and intelligence-lobotomy victims have yet to obtain membership in the National Academy of Sciences. Closer to the truth, however, is the philosophical materialism that permeates today's intellectual climate. With it comes a commitment to explain human intelligence strictly in terms of the human physical system. Given the indisputable connection between brain-states and behavior, the materialist has a facile answer to the mind-body problem: mind = brain.

Philosophical materialism has despite the advent of quantum mechanics yet to part with its predilection for mechanistic explanations. Given this preference, it construes causality strictly in terms of physical interactions. Thus it sees only two possible resolutions of the mind-body problem: (1) The substance dualism of Descartes, i.e., the human body is a machine controlled by an immaterial spirit, much as a pilot drives his vehicle; (2) The monism of Spinoza, i.e., the human body is the whole human. Cartesian dualism is problematic not merely because its ontology includes immaterial souls and spirits, but also because it splinters the human person, assigning to the body a negligible role on the question of intelligence. Confronted with this position modern philosophers choose rather to dispense with immaterial souls and spirits altogether. In this vein the French Enlightenment thinker Pierre Cabanis (1757-1808) offered his famous dictum, *Les nerfs-voilà tout l'homme* (nerves, that's all there is to man).

Nevertheless, a third option exists. This is the historic Judeo-Christian position on mind and body: the human being *unites* physical body and immaterial spirit into a living soul for which the separation of body and spirit is unnatural (in times past this separation was called death). We are to think of a union, not of a Cartesian driver operating his vehicle. I won't defend the historic position just yet, but I must emphasize the obvious: *this position demands an expanded ontology*. Matter by itself, notwithstanding how well it is dressed up with talk of holism, emergence, or supervenience, notwithstanding with what complexity it is organized, is still matter and cannot be transmuted into spirit. I stress this point because many theistic scientists in the name of scientific respectability have reinterpreted the historic position in such a way that spirit becomes an emergent property of the complex physical system constituting the human body. While this reinterpretation deserves attention, it is not the historic position, and it is misleading to attribute it to the theologians of past centuries, or naively to think that had these theologians lived today, they would have eliminated immaterial spirits in favor of a complex systems approach. The historic Judeo-Christian position is inconsistent with both Cartesian dualism and Spinozist monism. The mechanism implicit in these latter views leaves no room for matter and spirit to interact coherently within a single reality. I raise these points now to lay my cards on the table. I shall return to them later.

Cabanis' statement merits a second look. Suppose that autopsies of human beings reveal that their crania are packed with nothing but cotton wadding. Let us assume that in all other ways reality remains unchanged. Thus the great works of literature abound, music flourishes, and science advances. In particular, men are conscious of thinking as before, only now they are aware that their brains are hopelessly inadequate to account for their intelligence, just as a cube in a box is inadequate to account for intelligence. Thus we would have to look elsewhere, perhaps to an immaterial spirit, to account for intelligence. In this way Cabanis would be refuted.

But the brain is clearly not composed of cotton wadding, nor of any material exhibiting comparable simplicity. So why is there any reason to hope that the brain can account for intelligence? The answer is found in the following panegyric to the brain, a literary form common in neuro-physiology texts:

The human cerebral cortex contains something like 1010 to 1014 nerve cells. With that astronomical number of basic units, the cerebral cortex is sometimes referred to as the "great analyzer." If there are a minimum of 1010 nerve cells in the cerebral cortex, that number, 10 billion, is about 2.5 times the human population of the earth. Imagine three planets with the same population as the earth, with telegraph and radio links between every group of people on those planets. With that in mind, one begins to envision the type of situation present in the brain of each individual.

That is only a start, however. Each nerve cell makes contact with some 5,000 or so other nerve cells; that is, each nerve cell has up to 5,000 junctions with neighboring nerve cells, some as many as 50,000 junctions. At those *synaptic junctions* or *synapses*, information is passed between the nerve cells. What is significant about that process is that the information may be modified during its transfer. The number of sites at which information may be altered in some way is, therefore, astronomical, since the number of synaptic junctions within just a gram of brain tissue is of the order of 4 x 1011. The brain's cellular organization shows an almost unbelievable profusion of *connections* between nerve cells. Without such intricate connectivity, learning processes would be impossible.

The brain is considered an adequate explanation for mind and intelligence because of its vast complexity and intricate organization. By being complicated enough, by comprising billions of interrelated components, the brain is supposed to render thought possible.

And here we come to a rub. Precisely because of its vast complexity, no one really knows what is going on in the brain. More precisely, the connection between brain-states and intelligence is a matter of ignorance. This is not to say there is no causal relation between brain and behavior. There is if one looks at isolated, discrete behaviors. But as soon as one moves to the level of goals, intentions, and what philosophers more generally call propositional attitudes, cognitive scientists abandon hope of understanding this higher level through the lower neurological level. Hence they take refuge in notions like *supervenience*, *emergence*, and the now passé *epiphenomenon*. Thus cognition supervenes on neural activity, which in turn supervenes on the underlying physics; alternatively, intelligence emerges out of neural activity, which in turn emerges out of the underlying physical configuration; and consciousness is an epiphenomenon of neural activity.

Those who subscribe to the historic Judeo-Christian position on mind and body are often taken to task for believing that humans possess immaterial spirits. By believing this, they are considered disingenuous, taking refuge in ignorance. Spinoza, for instance, castigates those "who will not cease from asking the causes of causes, until at last you fly to the will of God, the refuge for ignorance." Nevertheless, if the historic position is correct, then those who subscribe to it are by no means ignorant. By looking to immaterial spirits and a transcendent God, they are in fact drawing proper causal connections-if they are right. But regardless whether materialists are right in affirming the brain is a sufficient

reason for intelligence, their ignorance of the precise causal connection between brain and intelligence remains. Granted, it is an ignorance they hope to dispel through research. But it is a hope they have largely abandoned, just because the complexities are so overwhelming.

Thus while the commitment to materialism persists, the hope of explaining human intelligence at the neural level, which for the materialist is the logical level, is not a serious consideration. Karl Lashley will for instance say, when addressing a symposium on the brain-mind relationship, that "our common meeting ground is the faith to which we all subscribe, I believe, that the phenomena of behavior and mind are ultimately describable in concepts of the mathematical and physical sciences." Yet towards the end of his career he will remark, "whether the mind-body relation is regarded as a genuine metaphysical issue or a systematized delusion, it remains a problem for the psychologist (and for the neurologist when he deals with human problems) as it is not for the physicist. . . . How can the brain, as a physico-chemical system, perceive or know anything; or develop the delusion that it does so?" And even though R. W. Gerard's observation is over forty years old, current brain research has yet to remove its sting: "it remains sadly true that most of our present understanding of mind would remain as valid and useful if, for all we know, the cranium were stuffed with cotton wadding."

Brain complexity is not the only problem facing the neurologist, who with Lashley's materialist convictions seeks to connect brain with intelligence. Brains are not uniform. One brain is not isomorphic to the next. While general morphology and structures coincide, brains from one individual to the next differ so much at the neurological qua synaptic level, that a search for common higher-level cognitive correlates holding across brains becomes a task so daunting as to seem hopeless. Even when dealing with a lone brain, it is clear that the same higher-level cognitive behavior has incalculably many distinct neurological antecedents. For example, a multitude of brain-states will induce the same cognitive act (e.g., dialing 911 in case of an emergency). Bioethics enters the picture as well, since brain research entails messing with people's brains in a very real sense. Barring a Nazi regime, unrestricted brain research on humans is not practicable. Finally, much like in quantum mechanics the observer tends to disturb the object being observed, so too brain research is invasive and cannot avoid confounding.

Clean Brains

Enter the clean world of computers. For the way out of this impasse cognitive scientists look to computer science and artificial intelligence. Computers are neat and precise. Unlike brains for which identical copies cannot be mass-produced, computers and their programs can be copied at will. Inasmuch as science thrives on replicability and control, AI offers tremendous practical advantages over neurological research.

Now the obvious question is, How well can computers model the brain? While this is the obvious question, it is not the question that really interests cognitive scientists. The reason is clear. As good materialists we believe that cognition is grounded in neural states. But it is cognition that interests us, not neural states. Moreover, we don't have the slightest idea how neural states correlate with cognition. Thus to simulate with computer programs brain-states of which we have no idea how these relate to cognition is simply to raise more problems than are solved. Simulating brain-states will not throw any light on cognition. This is largely a theoretical consideration. Practically speaking, to model a human brain at the synaptic level is beyond the memory/size capabilities of present machines.

What are cognitive scientists to do? How can they justify the claim that computation provides a sufficient cause for intelligence? Rather than simulate brains, cognitive scientists write computer programs which simulate behaviors typically regarded as requiring intelligence. Thus they bypass the neural level and move directly to the highest cognitive levels: perception, language, problem solving, concept formation, and intentions. Instead of modeling the brain, cognitive scientists model the intelligent behaviors exhibited through those brains. Thus many man-years of programming have been spent developing language translators (unsuccessful), chess playing programs (successful), expert systems (successful to varying degrees), etc. On balance it is fair to say that from the technological side AI has been and will continue to be successful. Nevertheless, as a comprehensive approach to human intelligence, its results have been less impressive. This is not for any lack of ingenuity on the part of computer programmers-some are very clever indeed. But intelligence involves much more than clever programs which are adept at isolated tasks. What goes by the name of AI has only delivered programs with very narrow competence.

Confident that this will change, cognitive scientists adopt the following rationale. If through concrete computer programs (algorithms) they can simulate all important aspects of human intelligence within a complete information-processing package, then they will have proved their case that human intelligence is a species of artificial intelligence. To realize that this view is not all that extreme among cognitive scientists, consider the following comments by Zenon Pylyshyn, professor of psychology and computer science, and director of the Centre for Cognitive Science at the University of Western Ontario. He is regarded as a thoughtful, sober figure in the cognitive science community (as compared to his more propagandistic colleagues):

> I want to maintain that computation is a literal model [*nota bene*] of mental activity, not a *simulation* of behavior, as was sometimes claimed in the early years of cognitive modeling. Unlike the case of simulating, say, a chemical process or a traffic flow, I do not claim merely that the model generates a sequence of predictions of behavior, but rather that it does so in essentially the same way or by virtue of the same functional mechanisms (not, of course, the same biological mechanisms) and *in virtue of* having something that corresponds to the same thoughts or cognitive states as those which govern the behavior of the organism being modeled. Being the same thought entails having the same semantic content (that is, identical thoughts have identical semantic contents).
>
> As dyed-in-the-wool realists, we propose . . . exactly what solid-state physicists do when they find that postulating certain unobservables provides a coherent account of a set of phenomena: we conclude that the [programs] are "psychologically real," that the brain is the kind of system that processes such [programs] and that the [programs] do in fact have a semantic content.

Several comments are in order. Pylyshyn clearly accepts that computation encompasses thought and intelligence. His characterization of cognitive science is, at least in its enunciation, bolder than mine. For he claims that computation is a "literal model" of mental activity, and in

effect repudiates mere "simulation." I consider this distinction spurious since cognitive science has progressed nowhere near the place where it can legitimately make such distinctions. Still, his comments reveal the climate of opinion. His reference in both passages to semantic content is significant, because meaning is the weak underbelly of AI. As we saw with Huxley's simians, the meaning of *Hamlet* was extrinsic to the monkeys' typing. Yet Pylyshyn claims that meaning (semantic content) will be intrinsic to the computer's computation.

Unlike Pylyshyn who claims that computation is a literal model of mental activity, I shall be content to admit that cognitive scientists have proved their case if they offer convincing arguments that machines can simulate the totality of intelligent human behavior in a comprehensive package (not merely a vast assortment of behaviors in isolation). By simulation I mean nothing less than an exhaustive imitation of behaviors requisite for intelligence. I therefore reject all arguments that extrapolate from good chess playing programs or good medical diagnostic programs to the claim that computers can think, have intelligence, display cognitive abilities, evince mentality, etc. Such talk is an abuse of language. I want to see a machine that puts it all together, integrating all those isolated tasks that require intelligence into a comprehensive whole.

Finite Man

I am urging cognitive scientists to fabricate a machine which grasps the whole that is human intelligence. Having made this challenge, I must add a restriction: in maintaining that machine intelligence subsumes human intelligence, cognitive scientists must be limited to machines that are physically possible. There is a vast difference between machines that can be physically realized and machines that exist only in the never-never land of abstraction. This never-never land of abstraction is known to mathematicians as the set of *partial recursive functions*. These functions constitute the maximal collection of computable objects. The branch of mathematics known as *recursion theory* studies these partial recursive functions and provides the theoretical underpinnings for computer science. Now any real computer running a real program has a limited amount of time and memory with which to complete its computations. Real computers are constrained by limited resources. Abstract computers, the partial recursive functions, suffer no such constraint.

Since the partial recursive functions contain everything that is computable, it follows that any real computer is just an abstract computer in disguise. The converse, however, does not hold. For instance, a computation that requires 101000 additions and multiplications is beyond the capability of any machine which could be fit into the known universe. Given the size of the universe (under 1080 elementary particles), a duration of many billions of years, the maximum speed of information-flow (the speed of light), and the smallest level at which information can be reliably stored (certainly no smaller than the atomic level), no such computation can be realized. On the other hand, such a computation is readily accomplished by some partial recursive function. Implicit here is the question of computational complexity, a facet of computer science which today is playing an increasingly dominant role.

Now this distinction between physically realizable and abstract machines becomes important when we consider the intrinsic finiteness of human behavior. It is common to claim that humans are finite beings. This can be argued. Scripture, for example, indicates that humans are made in the image of an infinite God. Pascal writes, "by space the universe encompasses and swallows me up like an atom; [but] by thought I comprehend the world." Yet regardless what we believe about man's finiteness generally, man's behaviors are finite. And this is the point of departure for the sciences of man. Science cannot deal with, to use Kant's terminology, noumenal man; it can only deal with phenomenal man. Desiring a monopoly on human intelligence, cognitive scientists are quick to presuppose that phenomenal man is the whole man. Phenomenal man is the man we can observe, the man known through his behaviors. Granted, this is the only man scientists can deal with. But is phenomenal man the whole man?

Let us be clear that that human behavior and sensory experience are intrinsically finite. One can understand the finiteness of human behavior on many levels. At the atomic level man is a finite bundle of atoms: reconstruct an individual atom by atom, giving the atoms proper relative positions and momenta, and you have a perfect clone. This construction is of course utterly infeasible; moreover, quantum effects may render it theoretically impossible. An equally infeasible finite reconstruction of the human organism which, however, has a better chance of avoiding quantum indeterminacy can be made at the chemico-molecular level (cf. molecular biology and biochemistry).

At the other extreme one can argue that since language can fully describe human behavior, and since language is intrinsically finite (there are only so many words to choose from, any sentence is of finite length, only so many sentences can be uttered in any lifetime), it therefore follows that human behavior is intrinsically finite. Another argument for the finiteness of human behavior can be made from the way human sensibilia can be encoded. Compact discs can for instance store audio (e.g., music) and visual (e.g., photographs) experience suitably encoded as a finite, discrete string of information, which when properly decoded can be played back with an arbitrary degree of resolution.

The level at which I prefer to understand the finiteness of human behavior is neurological. This approach is in line with the earlier quote by Pierre Cabanis, *Les nerfs-voilà tout l'homme*. At this level behavior and experience result from the firing of a finite number of nerve cells which can fire only so many times a second. Continuity of experience is therefore a myth. Experience is fundamentally discrete. It is because the number neurons and their rate of firing is finite that we experience the digitally encoded sound off compact discs as music rather than a shower of staccatos. For the same reason we experience a movie as continuous action rather than a discrete set of frames.

Let us for the moment play along with Cabanis, reducing man to his neurology. At this level of analysis not only do human behavior and sensory experience become finite, but also the total number of possible human beings becomes finite. The following loose combinatorial analysis argues the point. Let n be an upper bound on the number of neurons in any human, f an upper bound on their firing rate (i.e., number of firings a neuron is capable of per second), and l the maximum life span of any human (in seconds). Then during any firing interval there are 2n possible ways the n neurons can fire, and over a maximal life span there are (2n)fl = 2nfl possible ways the n neurons can fire in succession (let us call such successions of neural firings *behavioral sequences*). If one assumes that man equals phenomenal man, then 2nfl possible behavioral sequences include all conceivable human lives. A fortiori, there are at most that many human beings, for human beings with exactly the same behaviors and experiences are identical (we assume materialism of course).

The number 2nfl is huge, even for modest n, f, and l. Thus with billions of people, even billions of universes, we should not expect to see two human lives approximating each other, much less repeated. Often the vast complexity of human behavior exhibited in such huge numbers is taken to justify the reduction of humans to neural firing sequences, as though complexity and organization in themselves provide a sufficient reason for intelligence. We have noted that this reduction gets sidestepped by introducing terms like *supervenience* and *emergence*, which are supposed to distinguish higher level "intelligent" behavior from its physiological underpinnings. But the conclusion remains that all human behavior finds its immediate, efficient cause at the neurological level. And at this level behavior is finite.

The procedure for specifying 2nfl as an upper bound for behavioral sequences could stand some refinements. Instead of choosing n to enumerate an individual's neurons, one might have chosen n to enumerate the synaptic interconnections at which neurotransmitter is released, thus increasing n by a few orders of magnitude. I have implicitly assumed that neurons neither are born nor die, and that interconnections between neurons are stable. Again this is an artificial assumption. But anyone who challenges it need only increase n to include all those neurons which will be born or die as well as all potential interconnections, restricting at any given time attention to those neurons and interconnections that are currently active. I have also implicitly assumed that neural firing is limited to discrete intervals: a behavioral sequence proceeds in discrete time intervals wherein each neuron either fires or fails to fire. Lags between firings of distinct neurons are therefore ignored-this becomes not unreasonable if the firing intervals are made sufficiently brief. Thus to justify the assumption of synchronous firing among neurons the firing rate f may also need to be increased. Suffice it to say, there is an upper bound (however crude) on all the behavioral sequences that can conceivably constitute phenomenal man.

Consider now an individual named Frank who comprises n neurons, and let F be the collection of Frank's n neurons. Define a *behavioral instant* in Frank's life as an n-tuple (Ba : a OE F) where for each neuron a, Ba indicates whether neuron a fired during that interval (more precisely, Ba is a boolean variable taking the value 0 or 1 depending on whether neuron a fails to fire or fires, respectively). Frank's life (behavioral sequence) then consists of the *behavioral process* B(a,t) where t is a discrete time variable (f = 103 is an upper bound on the firing rate of neurons; thus we can take t as multiples of 10-3). Certainly this approach to Frank is solipsistic. Frank *is* his neural firings, and it doesn't matter a bit what the world is doing. Of course, we assume that the world is impinging on Frank and therefore affecting his Ba's over time. But this is irrelevant to our analysis.

Finally, if we assume that Frank's life is bounded by l seconds, then Frank's actual life is at most one of 2nfl possible lives he might have lived. Moreover, it follows from elementary combinatorics that Frank's whole life can be encoded in a string of 0's and 1's of length nfl, e.g.,

$$S_{Frank} = 1000101101 \ldots 10,$$

where the ellipsis represents nfl minus 12 digits (12 being the number of digits actually displayed). *This is your life, Frank*. If we choose n, f, and l big enough, then such sequences of length nfl can encode all human lives. In particular, there is a base 2 number of length nfl that encodes me-even my yet unlived future, even the writing of this essay. This number is of course just SBill.

Our analysis of Frank is a classic example of a brain in a vat. The brain receives stimuli and emits responses. These stimuli and responses occur over time and can be arranged in sequence. The string SFrank captures that sequence. It is important to note that strings like SFrank, SBill, SJane, and SSusan are assigned according to a rationale. By encoding experience and behavior, these strings capture the life of an individual. If one accepts that God's final judgment of humanity is according to the deeds done in the body, then such sequences are sufficient evidence for God to reach a verdict. These numbers are not assigned in the way telephone numbers or social security numbers are assigned. There the only requirement is that the number have so many digits and be unique to the individual. The rationale for SFrank goes much deeper. It captures Frank's life.

The Dilemma of Humanism

Phenomenal man is computational man. Computational man, however, has yet to be computed on an IBM or Cray computer. Currently, computational man exists solely in some abstract machine from the realm of partial recursive functions. Leaving the point which concerns the cognitive scientists for the moment aside, namely, how human intelligence is circumscribed by physically realizable machines, let us consider how the reduction of phenomenal man even to abstract machines threatens the humanist, who on the one hand thinks man is wonderful, and on the other staunchly retains a philosophical materialism. I find the humanist's assumptions inconsistent. If his philosophical materialism is correct, then there is nothing about man to transcend the constitution and dynamics of his physical system, the human body. Thus humanist man is in the end just phenomenal man. And this man, as we have demonstrated, is just computational man. Now the inconsistency lies in the fact that computational man is not all that wonderful, as humanists readily admit.

Thus when humanists like Hubert Dreyfus, Joseph Weizenbaum, and Theodore Roszak declaim against the dehumanization fostered by too high a view of machines and too low a view of human mentality, they inevitably sidestep the question whether some big enough abstract machine can capture the human being. They refuse to admit that unless man in some way transcends matter, the reduction of man to machine is indeed valid. Humanists attribute to man dignity and worth. Humanists look at man as the end of all man's longings. Man is ultimate. Thus any talk of transcendence is deemed a projection of impulses already present in man. But when humanists limit their attention to man as a product of the material universe, and refuse to acknowledge transcendence in man, or better yet, a transcendent creator who has made man in his image, they bare their necks to their cognitive-scientist opponents. For despite the rhapsodic flights and poetic rapture wherewith humanists celebrate the grandeur of man, man the product of nature, man the physical system is mechanical man.

The humanist wants to believe that humans possess a certain something which computers do not, that the computer cannot vitiate his exalted view of humanity. Nonsense. If his materialism is correct, then humans can be trivially realized as abstract machines, i.e., partial recursive functions in some programming framework. Such a reduction to even wildly complex abstract machines renders the things he holds dear-dignity, freedom, and value-null and void. I don't think I've overstated the case. An atomistic view of intelligence is ruinous to any exalted view

of man. This is evident from our solipsistic analysis of Frank. On materialist assumptions all of Frank is encompassed in SFrank. Behavioral sequences can even accommodate contingency. Thus SFrank¢ is Frank's life if he had gotten that promotion, SFrank" is Frank's life if he had not been dropped on his head as a baby, etc. When we consider all possible behavioral sequences for Frank-sequences constrained only by his genetic makeup and possible experiences-we arrive at a set {SFrank , SFrank¢ , SFrank" , … }, where the ellipsis is finite. Frank's intelligence is entirely encompassed in this finite set.

Now a trivial consequence of recursion theory is that all relations on finite sets are computable. Thus under materialist assumptions, whatever one may mean by Frank's intelligence can be encompassed within the framework of computer science. Whatever Hubert Dreyfus meant by his title, *What Computers Can't Do*, unless he is willing to look beyond the matter which constitutes the human body, he cannot legitimately mean that humans can display intelligence inaccessible to machines. In fact, one of his primary points is that computers must fall short of humans because they cannot possess a human body. But this is really a minor point, since it is not the body that is at issue, but the experience of that body, and this experience can be adequately realized in some coding scheme, like the one we indicated for Frank.

Finiteness really shatters the humanist dream. My aim here has been to force the humanist to own up to the unpleasant implications of a philosophical materialism to which he so often subscribes. Most people, I am afraid, do not realize the full import of the revolution that is mathematical recursion theory, or in its applied form, computer science. *Computers are the ultimate machine*. Church's Thesis, a guidepost in computer science, guarantees that computers are the *ne plus ultra* of machines. Every machine, much like Frank's behavior and experience, can be discretized. Once discretized, it can be simulated on a finite state computer. This is a point that can be justified at length, but let me instead direct the reader's attention to an example which should make the claim plausible. Aircraft companies routinely use supercomputers to simulate the flow of air over wing designs. This method for evaluating designs is reliable, successful, and tidy. In the process, reality is encoded computationally and simulated. Given sufficient resolution of the computational representation, the simulation is so fine that if reality could be reconstructed from the simulation, it would be indistinguishable from the original reality.

Church's thesis, which is unanimously confirmed by over 50 years of theoretical and practical experience in mathematics and computer science, indicates that what any machine can do, a computer can do. Thus it does no good to hope that the brain may turn out to be a better machine-a better "something" if "machine" is unpalatable-than a computer. Anyone who offers such alternatives simply does not know his computer science, or his neuro-physiology, or both. The brain only has so many neurons, each of which has only so many synaptic interconnections; these neurons have a maximal rate of firing and are subject to threshold effects-there are no unlimited degrees of firing. Once one has a finite state object, its dynamics are fully representable on a computer (maybe not on a computer we can realize with the usual integrated circuits, but certainly on an abstract machine). The brain is such an object. On materialist assumptions it is illegitimate to reject a computational model of mentality. Once, however, one admits computation as encompassing intelligence, it becomes illegitimate to ascribe to intelligence and humanity honors it can deserve only by not being a machine, honors like dignity and freedom.

The Problem of Supervenience and Personal Identity

I have been assuming that the humanist is dissatisfied with the idea of man being a machine. Let us now suppose he accepts my account of phenomenal man as an abstract machine. Let us say that on his materialist assumptions he is driven to the conclusion that man is a computational machine, albeit a very complex, highly organized computational machine. He will want to retain the things he holds dear, like dignity and freedom, but he will now have to redefine them to fit a computational paradigm. How shall he do it? The method of choice currently is to appeal to supervenience. Supervenience encompasses a multiplicity of notions like emergence, hierarchy, systems theory, holism, etc. For the purposes of this discussion I shall limit myself to supervenience.

What then is supervenience? Supervenience begins with a simple motto: "No difference without a physical difference." Supervenience, however, is not a crass form of physicalism. Philosopher Paul Teller cashes out this motto nicely:

> Imagine that in some given case or situation you get to play God and decide what's true. To organize your work you divide truths into two (not necessarily exhaustive) kinds. The first you call truths of kind P (for a mnemonic think of these as the *P*hysical truths . . .); and the second you call truths of kind S (for a mnemonic think of the truths of some *S*pecial science or discipline, such as psychology, sociology, ethics, aesthetics, etc.). You begin your work by choosing all the truths of kind P which will hold for the case. Then you turn to the truths of kind S. But lo! Having chosen truths of kind P, the truths of kind S have already been fixed. There remains nothing more for you to do. . . .
>
> [Consider] all cases of actual watches turned out by the same assembly line and set identically. The truths of kind P will in this case be the physical truths about the watches' structure, and the truths of kind S will be truths describing the watches' time-keeping properties. Of course, with identical physical structure and setting, the watches will keep the same time. I will say that for collections of cases of this kind truths of kind S *supervene* on truths of kind P.

To say the truths of type S supervene on truths of type P has the following succinct logical formulation:

> Here P ranges over physical predicates, S over nonphysical predicates, and u and v over objects in the real world.

Supervenience is to be understood hierarchically: what happens at a lower level (cf. P) constrains what happens at a higher level (cf. S). Thus the cognitive scientist might say that human intelligence (the higher level stuff) supervenes on human neuro-physiology (the lower level stuff).

Since supervenience is a transitive relation, and since human neuro-physiology in turn supervenes on human molecular biology which in turn supervenes on human atomic physics which in turn supervenes on human elementary particle physics, it follows that human intelligence supervenes on the underlying fundamental physics. At this point it is usually granted that we have bottomed out, having reached the lowest level of explanation.

Supervenience is not without philosophical difficulties. First and foremost among these is that supervenience is *not* a reductive analysis. For this reason certain philosophers (including this author) regard supervenience as mysticism in scientific dress. Philosopher of language Stephen Schiffer is unrelenting in this charge:

> "Supervenience" is a primitive metaphysical relation between properties that is distinct from causation and more like some primitive form of entailment. . . . I therefore find it more than a little ironic, and puzzling, that supervenience is nowadays being heralded as a way of making non-pleonastic, irreducibly non-natural mental properties cohere with an acceptably naturalistic solution to the mind-body problem. . . . The appeal to a special primitive relation of "supervenience" . . . is obscurantist. Supervenience is just epiphenomenalism without causation.

How can supervenience be so wicked, especially since it is touted by so many naturalistic-minded philosophers? Supervenience makes no pretence at reductive analysis. It simply says that the lower level conditions the higher level-how it does it, we don't know. Supervenience offers no causal account of how lower levels constrain higher levels. If such an account were on hand, we should have a reductive analysis and be able to dispense with talk of supervenience-the idea of reduction has after all been around for some time, certainly preceding supervenience. Admitting ignorance of how lower levels affect upper levels and being willing to forego reductive analysis, Schiffer regards as philosophical treason. Those who employ supervenience as a philosophical research strategy Schiffer charges with dualism, obscurantism, metaphysics, and epiphenomenalism.

Given our overriding interest in the relation between intellect and brain, we ought to ponder whether supervenience is in any legitimate sense applicable to the mind-body problem. Mind is supposed to supervene on body just as time-keeping properties supervene on the physical structure of a watch. But is this a fair analogy? Examples like those of physical watches conjoined with nonphysical time-keeping properties are supposed to capture the idea of supervenience. Now there is a fundamental difference in the way time-keeping supervenes on the physical object which constitutes a watch and the way intellect can be said to supervene on the physical object which constitutes a human body. Time-keeping supervenes on a watch because, *and only because*, our intellect contributes *temporal concepts* to the physical object which constitutes that watch. The hierarchy of levels basic to supervenience are levels we construct through our intellect. There seems therefore a self-referential paradox in saying of this intellect which constructs so many instances of supervenience that it itself supervenes on a physical system. The intellect plays a distinguished role in any supervenience account, and it is not clear that it is legitimate to turn it on itself and thereby proclaim that the very instrument we need to establish supervenience itself supervenes.

The question of falsifiability also comes up. Let us say the intellect supervenes on the brain. How can we know this? What evidence would count to disprove this assertion? At the very least we would need an exhaustive account of the correspondence between brain states and mental states; for without an exhaustive account there would always remain the nagging uncertainty whether lower level properties are fully determinative of higher level properties-full determination of the higher by means of the lower is the definition of supervenience. Such an account would decide the mind-body problem one way or the other (cf. our cotton wadding example). But once we have such an exhaustive account, we can dispense with the notion of supervenience-such an exhaustive account will be a reductive analysis. What more then is the claim that intellect supervenes on the brain than bald assertion? Any scientific justification of supervenience will demonstrate far more than mere supervenience-it will tell a causal story. What more is supervenience than a materialist faith which makes lower levels determinative of higher levels?

No treatment of supervenience would be complete without at least touching on the ever popular Doppelgänger examples. These examples are philosophical thought experiments, science fiction stories if you will, that address the following question: What relation does an exact physical duplicate (the Doppelgänger) of a human being bear to the original human being? The human body is after all just an organized hunk of matter. What if we construct an atom for atom copy of this hunk, imparting to each atom the right relative momentum and energy state? In this construction, have we duplicated the original human's mental states? Does the Doppelgänger have a soul? Does the Doppelgänger experience pain? Would it be right to construct a Doppelgänger, freeze him, and later use his bodily organs for transplants in the original? Let us say we have the technology to construct Doppelgängers at will. Is it morally acceptable to build a teleportation device which sends an individual, say, to Mars by transmitting a complete specification of his body to Mars, constructing the Doppelgänger on Mars, and then destroying the original on earth?-after all, we don't want more than one of you in the universe at a given time. What is lost by destroying the original and letting your Doppelgänger run free? Stories about Doppelgängers can be multiplied almost endlessly.

Doppelgänger examples address the philosophical problem of personal identity: What does mean for you to be you? Depending on one's point of view these examples can be entertaining or disturbing. Certainly a teleportation device like the one described would count decisively for supervenience and against immaterial souls and spirits. But we do well to remember that thought experiments are *thought* experiments precisely because they impracticable. Thought experiments are not scientific experiments, and therefore cannot decide scientific questions. They are useful for raising interesting questions and may inspire concrete scientific experiments. But they are hypothetical in the extreme. Willard Quine has some sobering words on the matter:

> The method of science fiction has its uses in philosophy, but . . . I wonder whether the limits of the method are properly heeded. To seek what is "logically required" for sameness of person under *unprecedented circumstances* is to suggest that words have some logical force beyond what our past needs have invested them with.

Another reason for not being unduly swayed by Doppelgänger examples is quantum mechanics. Quantum mechanics, with the limitation it places on measurement at the micro-level, makes it highly doubtful whether human technology is capable of building the scanning and reconstituting devices necessary for the construction of Doppelgängers. Commenting on the teleportation device in *The Emperor's New Mind*, physicist Roger Penrose writes,

> Is there anything in the laws of physics which could render teleportation *in principle* impossible? Perhaps . . . there is nothing in principle against transmitting a person, and a person's consciousness, by such means, but that the "copying" process involved would inevitably destroy the original? Might it then be that the preserving of *two* viable copies is what is impossible in principle? . . . I believe that [these considerations] provide one pointer, indicating a certain essential role for *quantum mechanics* in the understanding of mental phenomena.

Penrose, the physicist, conflates fundamental physics with consciousness and mental phenomena so as to give physics almost a mystical role. Still, his observations should be considered before lending too much credence to the Doppelgänger examples.

The question still remains, Is your physical replica you? I am unwilling to answer this question without qualification. For me quantum mechanics, nonlinear dynamics (chaos), human physiology, and probability theory all conspire to make the premise this question requires me to grant-namely, the existence of my Doppelgänger-about as plausible as Greek mythology. Nevertheless, I take anything to be possible and admit that all my beliefs are falsifiable given the right circumstances. If I should be confronted with my Doppelgänger, and if this double were constructed by purely mechanical means at the hands of human technicians, I should decide in favor of supervenience. But for me the important question is how the Doppelgänger came into existence. No doubt I'm biased, but without a causal story of the Doppelgänger's origin, I would attribute his existence to God. But once God is back in the picture, I have no problem attributing to my Doppelgänger an immaterial spirit. So we're back where we started.

The Dilemma of Semi-Materialism

Earlier I described three approaches to the mind-body problem: the substance dualism of Descartes, the monism of Spinoza, and the historic Judeo-Christian position. I want now to focus on a fourth option which has of late been gaining currency in theistic circles. I shall refer to this view as semi-materialism. By semi-materialism I mean a philosophical position which on the one hand acknowledges the God of Scripture, but on the other denies that man's soul and spirit have an ontology distinct from (i.e., not derivative from) the body. Semi-materialism is a melding of traditional theology and supervenience. God is still creator, sovereign, and transcendent, but man is now fully realized in his human body.

It is important to understand that semi-materialism is not solely a question of methodology. Treating the human person as a physical system is not merely a scientific research strategy for the semi-materialist. The semi-materialist accepts supervenience-no difference without a physical difference-and therefore holds that talk of souls and spirits by the ancients is a prescientific way of describing consciousness as it emerges from the human physical system. Thus apart from man's moral responsibility to God, the semi-materialist has no great quarrel with the cognitive scientist. Both are content to view man as strictly a physical system. On the question of God they do of course differ. But semi-materialism compartmentalizes anthropology and theology so that whenever traditional theology conflicts with a supervenient anthropology, the former gets reinterpreted to jibe with the latter.

The late Donald MacKay was an outstanding example of a semi-materialist. MacKay was a professor of communications at Keele University in England who specialized in brain physiology. Eleven years ago he wrote a book entitled *Human Science & Human Dignity*. Throughout the book he emphasized the need to examine the data of science and theology dispassionately. His goal was to develop an integrated view of man. How was this integration to be accomplished? The following passage is revealing:

> [With] a hierarchy of levels there is no question of keeping the different explanations in "watertight compartments": what someone has called "conceptual apartheid." Although their categories are different and they are not making the same statements, by calling them hierarchic we commit ourselves to the view that there is a definite *correspondence* between them. In particular, no change can take place in the conscious experience reported in a higher-level story without some corresponding change in the stories to be told at the lower level (though again, not conversely). On this view, the way to an integrated understanding of man is not to hunt for gaps in the scientific picture into which entities like "the soul" might fit, but rather to discover, if we can, how the stories at different levels correlate.

Although MacKay speaks of *correspondence* between levels, he really means something much stronger, namely *determination*. To see this, I call the reader's attention to the sentence, "no change can take place in the conscious experience reported in a higher-level story without some corresponding change in the stories to be told at the lower level (though again, not conversely)." The phrase "not conversely" is decisive; it demonstrates that he takes the lower levels as fixing the upper levels. This is supervenience. In fact, it is the same brand of supervenience we described in the last section. It follows that MacKay's view faces the same objections we raised in the last section against supervenience. These objections he fails to address since in his integration of theology and anthropology he takes supervenience as a given, reinterpreting theology in its light.

Now theology is not so malleable an instrument as to yield to scientific pressure. The problems of trying to reconcile a supervenient anthropology with a traditional theology invade the whole of theology. Thus much of what MacKay calls the "traditional imagery" associated with death has to be discarded or reinterpreted. What are we to make of the incarnation of Christ? Do Jesus' soul and spirit fit into the semi-materialist's hierarchy of levels? What about miracles? If we accept that God can interact causally with the material universe, why should it be inconceivable that a human spirit can interact causally with a human body? MacKay accepts a general resurrection mankind. Yet within the semi-materialist framework it is not clear how humans are anything more than lumps of matter in motion, which at the resurrection are simply reconstituted and again set in motion.

For me the chief difficulty with semi-materialism is that from God's perspective it trivializes man. Because of its supervenient anthropology, semi-materialism gives us a man whose soul and spirit are not only inseparable from the body, but actually derived from the body. Now why should this man be trivial from God's perspective? To answer this question let us return to the Parable of the Cube. Suppose I am watching the cube move around inside the box. Its motion can be explained in several ways: (1) The cube just sits there at rest. (2) The cube traces some predictable trajectory. (3) The cube moves randomly. (4) I control the cube's motion, say with a joystick. (5) Some other intelligent agent controls the cube's motion. These cases are exhaustive, though as we shall see immediately, they are not exclusive. Moreover, depending on one's view of causality, (3) may be vacuous.

Cases (1) and (2) are really superfluous since they can be subsumed under case (4). For case (1) this is obvious. If case (4) holds, then I have complete control over the cube's motion. I now decide to keep the cube at rest, say by leaving the handle of my joystick alone. This is just case (1). In this way (4) subsumes (1). What about (2)? To say that the cube's motion is predictable is to say that there is some description that prescribes the cube's motion. Moreover, since the motion is predictable, I actually have that description. Thus I can take that description, follow its instructions, and cause the cube to move in the prescribed manner. But this is just case (2). Thus (4) subsumes (2) as well.

Now my contention is that of the three remaining cases only case (5) is interesting. In (3) the cube's motion is so unpredictable and erratic that I never expect to receive a coherent message. This is like the monkeys' random typing. I may look for patterns in the cube's motion, but as soon as I think I've got the hang of what the cube is doing, it disappoints me and does something else. I see no rhyme nor reason to what the cube is doing. This case is thoroughly unsatisfying to my intellect. Case (4) is also uninteresting. All I'm doing is moving the cube around. I feel like moving it here, so I move it here. Next I feel like moving it there, so I move it there. If I've got a copy of *Hamlet*, I can move the cube in such a way that its motion encodes *Hamlet*. But this is no fun-I might just as well read *Hamlet* directly. If I had more cubes, I might make some interesting designs. If I had multi-colored cubes I could let my imagination run wild and pretend I'm an artist. But with only one cube the situation is dull indeed. Only when another intelligence is moving the cube and communicating with me through its motion does cube watching become interesting. It was case (5) that towards the beginning of this essay resulted in the frenzy we called religious cubism.

Now consider God's relation to the material universe. God created the universe. The universe is finite. How does God view the universe? He knows it in every detail. He sees the end from the beginning; it holds no surprises for him. Now instead of a cube in a box, let us imagine a universe in a box. In this case God is the observer outside the box. This is certainly legitimate, since God is transcendent, in no way conditioned by his creation. Whereas we were looking at the motion of but one cube, God is looking at the simultaneous motion of all the various bits and pieces of matter that constitute the universe, seeing them in all their many configurations. God is particularly interested in humans, so he pays special attention to those bits of matter that constitute us.

When God is watching us, which of the three remaining cases holds? Case (3) is clearly out of the question. Randomness exists only where there is ignorance. No surprises await God. He sees history at a glance. Thus for theism case (3) is vacuous. What about case (4)? In this case the universe is a giant toy which God controls with a sophisticated joystick. The only problem is that this toy cannot amuse God. Just as a cube in a box makes for a dull toy, and cannot amuse us unless another intelligence influences it, so a universe subject only to God's intelligence is a dull toy for God, only in the extreme. In fact, for our limited intellects, moving a cube inside a box presents a greater challenge than for God to run the universe, whether by natural law or by direct intervention. There is no novelty, no thrill, no satisfaction for God in simply controlling the universe as a giant toy. To him it is like a cube in a box, only simpler.

This leaves us with case (5). To us a cube in a box is only interesting when an intelligence other than ourselves uses it to communicate with us. The same holds for the material universe and God. The only reason the universe is interesting to God is because there are intelligent beings, namely us, who express themselves through the universe, namely the matter that constitutes our bodies. If these intelligences are not external to the universe, then we land in case (4): a toy universe populated by toy people subject to a bored God who cannot be amused. Only case (5) entails a non-trivial creation in the light of its creator. There are no other possibilities.

This analysis gives the lie to a sentiment common in scientific circles. Accordingly, as we learn more and more about the immensity of the universe, we should think less and less of ourselves. Alternatively, the universe is such a big place that we must be insignificant. The foregoing analysis shows that without humans, intelligent creatures created in the image of God, the universe itself is insignificant, at least in the sight of God. Size has simply no bearing on significance, least of all to the mind of God. It is because we are here that the universe is significant. However small we become in relation to the universe is simply of no consequence. It is worth repeating Pascal's famous dictum: "By space the universe encompasses and swallows me up like an atom; by thought I comprehend the world."

The historic Judeo-Christian position on mind and body entails that God's view of the universe corresponds to case (5). To what case does the semi-materialist's view correspond? For the semi-materialist how does God view the universe, and more particularly man? By assuming supervenience, the semi-materialist has made it clear that he will not hunt for gaps in the scientific picture of man; he will not look for places into which he can fit soul and spirit. His refusal to look beyond the physical aspect of man is not, as we have already noted, simply a question of scientific methodology. Semi-materialism is an attempt by theists to unite science and theology in a consistent framework.

By a process of elimination we find that the semi-materialist's universe is a case (4) universe, the toy universe populated by toy people. Logic yields him no alternatives. The semi-materialist must forego case (3)-a random universe which God cannot predict-unless he wants to question God's omniscience. He must also forego case (5), since he derives human intelligence solely from its expression through matter. This leaves case (4). But since matter is finite and the dynamics of matter taken by itself is trivial to God's intellect, a case (4) universe makes for insipid theology. I would go so far as to say that a case (4) view of man ruins both anthropology and theology. The consequences of a case (4) universe are far reaching.

This was brought home to me after a recent conference at Trinity Evangelical Divinity School. At this conference James I. Packer delivered a talk entitled "Evangelicals and the Way of Salvation: New Challenges to the Gospel, Universalism, and Justification by Faith." The talk created something of a stir since in it Packer called his compatriot John R. W. Stott to account for the latter's recent expressed views on conditionalism. Apparently Stott has subscribed to a form of conditionalism for some time, but has only recently gone public with these views. According to

conditionalism, at the end of the age the righteous are raised to immortality and eternal life (eternal life being something they do not now possess), while the unrighteous are annihilated, their existence being erased from the warp and woof of reality. This took many of us by surprise since Stott is a leading and respected Christian thinker.

Conditionalism is of course a recurrent heresy in the Church, and to see it associated with so distinguished a name was a source of puzzlement. Actually, it ought not to have been puzzling. Ten years before the conference Stott had openly subscribed to a supervenience view of mind and body. Granted, he did not use the word "supervenience." But in writing the foreword to MacKay's book *Human Science & Human Dignity*, Stott left no doubt about accepting MacKay's "hierarchy of levels." MacKay's book arose out of his 1977 *London Lectures in Contemporary Christianity*. Commenting on the book and the lectures, Stott wrote,

> I listened to Professor MacKay's lectures with absorbed interest. His keen mind penetrates the heart of every argument, and coolly, dispassionately, he exposes logical fallacies wherever he detects them, in Christians and non-Christians alike. . . . He is determined to hold fast to the truth in its wholeness. His well-known rejection of reductionism . . . is matched by his resolve to face and to integrate all the available data. Above all, while readily acknowledging that from one point of view a human being is an animal, and from another a mechanism, he refuses to stop there. In order to do full justice to human beings, he introduces us to the concept of a "hierarchy of levels" at which human life is to be understood and experienced.

Let me stress it again-this is supervenience. I have argued that supervenience plus God entails a case (4) universe. I shall now go so far as to charge that conditionalism and annihilationism are not merely consistent with a case (4) universe, but logically necessary. My justification for this claim is simply this: for a just God to make a strictly finite material human being-the only human being a case (4) universe has to offer-suffer the torments of hell for eternity is to render infinite punishment for finite fault. The logic of a case (4) universe requires an untraditional view of hell.

The problem of evil is always a problem, but in a case (4) universe it becomes a catastrophe. Since we assume God is all-powerful and all-knowing, he can constitute and reconstitute matter any way he likes. Since there are no immaterial souls or spirits, our intelligence, behaviors, diseases, social structures, national boundaries, successes, wars, sins, etc., etc. all derive from the way matter is constituted. In each of these instances much is to be desired. With all the pain and suffering, why doesn't God do something? This is a valid question. In a case (5) universe we take comfort in knowing that events in the material universe are the unfolding of a drama that is grounded in eternity, of which we are a part. But in a case (4) universe all such comfort is a sham. All the bereaved mother can say is, God could have kept my child from dying, but he didn't. All the criminal can say is, God could have altered my sociopathic brain-states, removed me from the crime-infested ghetto in which I grew up, and thereby given me the opportunity to be an upstanding member of society-but he didn't. Of salvation it can only be said, It is God's choice whether to constitute your brain to be favorably of ill disposed towards him. So much for the doctrine of predestination.

In a case (4) universe, how does God respond to prayer? Suppose I am suffering from an addiction and pray to God that he remove it. Since I have no immaterial soul or spirit, the only help God can offer is material. Well, what does God do in response to my prayer? Does he miraculously reconstitute my brain in some way or alter my body chemistry so that my addiction is removed, though otherwise keeping me the same person?-here lurks the problem of personal identity. Or does he do nothing? Is it simply that praying is physiologically good for me and that the prayer accomplishes its end simply in the praying? Is it that God has built into my body a predisposition that makes prayer good for me? Are God's responses to prayer simply secondary causes he built into the universe at the point of creation?

The Historic Position

Earlier I described the historic Judeo-Christian position on mind and body as holding that the human being unites physical body and immaterial spirit into a living soul for which separation of body and spirit is unnatural and entails death. I also emphasized that this position demands an expanded ontology: unlike semi-materialism with its commitment to supervenience, the historic position does not see spirit as a derivative of the complex physical system that makes up the human body. My purpose here is not to expound this historic anthropology, but to trace a bit of its history and examine how it has gone from the prevailing position in the West to the status of a quaint relic.

The position as I have stated it is but a straightforward restatement of the *Genesis* account of man's creation:

> The LORD God formed the man from the dust of the ground [body] and breathed into his nostrils the breath of life [spirit], and the man became a living being [soul].

In the New Testament we find both Paul and James echoing this passage. Thus Paul writes,

> If Christ is in you, your body is dead because of sin, yet your spirit is alive because of righteousness,

whereas James writes,

> As the body without the spirit is dead, so faith without works is dead.

Now the Bible is not a book of systematic theology, nor does it explicitly block all modern moves at reinterpretation. Thus it does not call supervenience by name, nor does it explicitly reject supervenience as a plausible account of spirit and soul. But if we trace the course of western theology, we see that theologians before the age of modern science held views which cannot be reconciled with semi-materialism and

its concomitant supervenience. Thus Augustine does more than echo the Genesis account of man when he states his own position on death, a position which becomes increasingly difficult to reconcile with supervenience:

> As regards bodily death, that is, the separation of the soul from the body, it is good unto none while it is being endured. . . . For the very violence with which body and soul are wrenched asunder, which in the living had been conjoined and closely intertwined, brings with it a harsh experience, jarring horridly on nature so long as it continues, till there comes a total loss of sensation, which arose from the very interpenetration of spirit and flesh.

The idea of spirit and flesh interpenetrating has a distinctly different feel from supervenience.

With Aquinas we find the historic position coming into full bloom. Thus he writes,

> It must necessarily be allowed that the principle of intellectual operation which we call the soul is a principle both incorporeal and subsistent.

Even this statement might be reconciled with semi-materialism if one conceives of the soul as an abstract algorithm (say ensconced in a Platonic heaven) which the body qua machine instantiates. But this reinterpretation becomes implausible in the light of his following comment:

> The intellectual soul, because it can comprehend universals, has a power extending to the infinite; therefore it cannot be limited by nature either to certain fixed natural judgments, or to certain fixed means. . . .

Computer algorithms are finitary and therefore clearly "limited by nature either to certain fixed natural judgments, or to certain fixed means." But for Aquinas the intellectual soul transcends the finite, having "a power extending to the infinite." At this point in the evolution of theology I see no way of reconciling the historic position with semi-materialism.

Even Descartes appreciated the importance of the historic position. In the *Discourse on Method* he writes,

> Although machines can perform certain things as well as or perhaps better than any of us can do, they infallibly fall short in others, by which means we may discover that they did not act from knowledge [cf. the intellect], but only from the disposition of their organs [cf. programming, algorithms, state of the machine]. For while *reason is a universal instrument which can serve for all contingencies*, these organs have need of some special adaptation for every particular action. From this it follows that it is morally impossible that there should be sufficient diversity in any machine to allow it to act in all events of life in the same way as our reason causes us to act.

Descartes' point of departure from the historic position was his commitment to mechanism and its consequent substance dualism which compartmentalized reality into physical and spiritual parts so that the two could no longer interact coherently. Because of his mechanism Descartes wanted only the most tenuous connection between physical and spiritual reality, looking for an immaterial soul to interact with a physical body solely at the now infamous pineal gland. But mechanism is opposed to all gaps in physical causality. Had he held to the older view that causality cannot be fundamentally understood in purely physical terms (a view, by the way, not inconsistent with modern quantum mechanics) he would not have propounded his substance dualism, which to philosophers unsympathetic to theology is easily truncated by removing the spiritual component completely.

Finally in Kant we have a decisive break with the historic position. Subscribing completely to the mechanism inherent in Newtonian mechanics, Kant refused to consider causality outside of space and time. Since spirits do not reside in space and time, it follows that they can have no influence on what occurs in the physical world-at least no influence of which we can have any knowledge. With Kant the only knowable reality is physical reality embedded in space and time. A reality in which matter and spirit can freely interact, where they can interpenetrate, to use Augustine's idea, is disallowed. Though Kant's critical philosophy has been to some extent discredited because of his total and absolute acceptance of Euclidean geometry and Newtonian mechanics as determinative of reality, the sense that what is knowable is solely the physical, and that it is known solely through the physical persists. Its modern day outworkings include materialism generally, physicalism, and of course supervenience.

In this day is the historic position still tenable? In holding it, does one subscribe to a god-of-the-gaps solution to the mind-body problem? The modern Zeitgeist holds that since the 17th century science has been closing in on theology, constantly shrinking its legitimate domain of competence. Are the scientific gaps in our knowledge of the relation between mind and body at the vanishing point so as effectively to banish spirit from the human person? In response, I must say that not only do I take the historic position as still tenable, but I take it as the only tenable position for the theist who claims to adhere to anything like a traditional theology.

God-of-the-gaps solutions have since the Church's incompetent handling of Copernicus and Galileo left a bad taste in the mouth of sincere intellectuals. God-of-the-gaps solutions are particularly embarrassing when scientists are told what they can't do, and then go ahead and do it. Nevertheless, there are gaps, and then there are gaps-not all gaps are created equal. Thus there are gaps that science has decisively filled, e.g., heliocentrism has with finality displaced geocentrism. Then again, there are gaps which science claims to have filled which on closer inspection in fact failed to be filled. Thus Newtonian mechanics claimed to give an exhaustive account of the dynamics of matter. For two centuries scientists claimed this gap was filled. But in the 20th century along came general relativity and quantum mechanics. The sense that science is closing in on all bankable truth is therefore misleading. Science is much more like the stock market than a bank, exhibiting erratic fluctuations and sharp dips-Newtonian stock, for instance, has fallen sharply this century.

Now there is still another type of gap, and these are the gaps which theology imposes on human knowledge. The Scriptures refer to these as *mysteries*. On this point Wittgenstein has a pertinent observation:

> It may easily look as if every doubt merely *revealed* an existing gap in the foundations; so that secure understanding is only possible if we first doubt everything that *can* be doubted, and then remove all these doubts.

Science seeks to remove all doubts and fill all gaps, but on its own terms. Now there are gaps which theology says science shall never fill. In prescribing such gaps, theology issues a challenge to science. On the basis of these gaps science can break theology (though it cannot make theology); to break theology it need merely fill these gaps. Let me be so bold as to say that theology is falsifiable by science. Though such a claim runs counter to the ever popular compartmentalization of science and theology, it is true.

The apostle Paul recognized the falsifiability of theology when he noted, "If Christ has not been raised, our preaching is useless and so is your faith." Theology cannot be reconciled with an arbitrary collection of facts. If some purely naturalistic therapy could, for instance, be devised that would rid the world of behavior which in times past was attributed to sin, then the whole moral force of Scripture and the atonement of Christ would be called into question. For the present discussion, if cognitive scientists could devise a computer which captured a sufficiently broad spectrum of human cognitive abilities, I would say the cognitive scientists had proven their point. Certainly the historic position would be discredited. Let me hasten to add, however, that cognitive science is not only far from achieving such a goal, but, as I would argue on theoretical grounds, attempting to solve a problem without a scientific solution. I'll touch on these points in the following closing section.

Concluding Remarks

The philosophy that drives cognitive science is a materialism committed to explaining man via computation. To justify this philosophy the cognitive scientist writes computer programs that attempt to capture intelligent human behavior. The grander and more encompassing such programs, the better. Computers, however, are not cheap. The government agencies that fund research don't distribute money like drunken sailors. Thus to justify hefty research grants, cognitive scientists make promises they can't keep. Just as in the late 50's and early 60's, when language translators were going to make keeping up with the Soviets a breeze, the same euphoria persists today. Thus, as we noted earlier, H. Moravec, the director of Carnegie-Mellon's Robotics Institute, feels no compunction when he predicts that the next century will be populated by robots that displace the human race. And if this be so, is it not incumbent on the research agencies to fund Moravec and hasten his prophecy? Or consider C. G. Langton of the Santa Fe Institute who has recently edited a book entitled *Artificial Life*, the proceedings of a conference by the same name. Artificial life conjures visions of Frankenstein's monster, of man tearing himself out of nature and raising himself by his bootstraps. Even the U.S. Department of Energy felt the need to sponsor this vision. But on closer examination one sees that artificial life is pretty computer pictures and clever computer algorithms.

Cognitive scientists suffer a conflict of interest. Coupled with the need for a quick fix in the form of research monies, cognitive science too often degenerates into propaganda. Claims are inflated and difficult problems get swept under the rug. Even if human intelligence is physically realizable through an electronic computer, a hypothesis I reject, it is by no means obvious that human intelligence is capable of realizing it. Thus even if some super-intelligence could build a computer that, to use Pylyshyn's phrase, is a "literal model of mental activity," it is not at all clear whether cognitive scientists have the brains, if you will, to get the job done. The problems are daunting; in fact so daunting that humility is the order of the day. The tendency is to inflate one's position at the expense of truth. This is bad. I have written on this topic elsewhere:

> Inflation is a problem in science as well as economics. Partial results, admissions of ignorance, and uncertainty in general do not elicit the adulation of confident, bold assertions. The tendency is to inflate the supposed validity of one's cause. Politicians are rewarded for the confidence they evince, however ill founded, and penalized for their lapses into diffidence, however well founded. Art authenticators are expected to deliver a definitive verdict on a work of art and will do so, particularly if they are the acknowledged authorities on the artist in question, even though in certain instances they may be less than justified in handing down such verdicts, instances confirmed by the numerous fakes that have ended up in museums. So too [cognitive] scientists are notorious for overplaying their cards.

We should demand a distinction between philosophy and science; not a separation but a distinction. We should be very clear on what the cognitive scientist has indeed accomplished, and what doctrines the cognitive scientist qua philosopher is advocating. We should also take a high view of intelligence. Because the cognitive scientist's accomplishments are in fact so meager, it is easier for him to vitiate human intelligence than to admit his insignificant progress. This must not be permitted. Consider for example the following remarks by Roger Schank, now at Northwestern University, but formerly director of Yale's Artificial Intelligence Project:

> There is a difference between merely matching or displaying a set of English sentences in response to a specific initial English sentence and understanding the meaning of such a sentence. But what is the dividing line? When does mere pushing around of meaningless symbols inside the computer become understanding? If it were possible to get a computer to respond reasonably to sentences such as the above [sentences about characters in a simple story], could it be said to understand them? We already have created programs that enable a computer to respond to such phrases at fairly deep levels, in different syntactical arrangements, and with different expressions for the same events. It is hard to claim that the computer understands what *love* is or what *sadness* is. It is hard for most people to claim that they understand what *love* and *sadness* are.

I would take issue with Schank about his computer program responding to phrases "at fairly deep levels": the stories to which his scripts program is responding are too contrived and simple to deserve anything like the designation "deep." More objectionable, however, is his conclusion that because humans understand love and sadness imperfectly, humans ought therefore to admit that computers can understand

these concepts as well. Schank's very notion of understanding is defective. This humiliation of the human intellect to bolster the negligible successes of cognitive scientists is all too common.

Those who hold to the historic position on mind and body should find encouragement in what I call the *Law of Priority in Creation*. I would like to see this law elevated to a status comparable with the laws of thermodynamics. The law is not new with me. It is found in Scripture:

> Jesus has been found worthy of greater honor than Moses, just as the builder of a house has greater honor than the house itself.

The creator is always *strictly greater* than the creature. It is not possible for the creature to equal the creator, much less surpass the creator. The Law of Priority in Creation is a conservation law. It states in the clearest possible terms that you can't get something for nothing. There are no free lunches. Bootstrapping has never worked.

With the rise of atheistic evolutionism, the West has en masse repudiated this law. The creature is henceforth greater than the creator, for man has surpassed inanimate nature, whose creation he is. Cognitive scientists also repudiate this law in their work. Their dream is to build a computer that will shame us, that will so surpass us intellectually as we do the apes. The Law of Priority in Creation, however, repudiates their programme. For the computers they build, the programs they write all testify of a creative genius in man which surpasses the objects it creates. For any computer program that is supposed to rival the human intellect, I merely point to the human author who conceived the program. According to the Law of Priority in Creation the cognitive scientist's programme is self-refuting.

Finally a few comments about the subtitle of this essay are in order. Alchemy was the programme of the middle ages for transmuting base into precious metals. It sought to accomplish this by a combination of naturalistic and mantical means. The key to this transmutation was the philosopher's stone, an imaginary substance thought capable of performing the desired transmutation. In connection to cognitive science, the following observation is almost prophetic:

> Alchemists became obsessed with their quest for the secret of transmutation; some adopted deceptive methods of experimentation, many gained a livelihood from hopeful patrons. As a result, alchemy fell into disrepute.

Cognitive science has become today's alchemy. Cognitive scientists are obsessed with transmuting matter into mind. Unsound philosophy has deceived them into believing that the philosopher's stone is found in the computer. They gain a livelihood from hopeful patrons at the Defense Department, the National Science Foundation, and other funding agencies. Cognitive science has, unfortunately, yet to fall into disrepute.

I am describing the spurious philosophical enterprise called cognitive science. Just as alchemy was legitimized when it gave up its grandiose ambitions and turned to chemistry, so too, one may hope, cognitive science will cast off its pretensions and turn to what I have called the science of cognition. In taking information processing as its paradigm for examining human cognition, the science of cognition is a branch of computer science-it is legitimate and cannot be impugned. I encourage scientists to press on in the science of cognition and determine just how much of human cognition can be represented computationally. Such a research programme does not threaten me. I am, however, committed to viewing computers and the programs they run as tools for my intellect, much as hammers are tools for my hands, and not as my peers. Cognitive science degenerates into a spurious philosophical enterprise when computers are no longer viewed as tools, but as potential peers or superiors.

Church's thesis tells me that man qua scientist can do no better than understand the human intellect in terms of an information processing model. But it is the height of presumption for man qua philosopher to claim this model is all encompassing. The cognitive scientist finds this unacceptable because he does not like what he deems as arbitrary limits imposed on the pursuit of knowledge. Actually, no such limit has been imposed. Let him pursue a legitimate scientific research programme as far as he can. But let him remember that the facts point resoundingly to a very imperfect understanding of man in purely scientific categories; that sound philosophy is consistent with this finding, indicating that scientific categories may well be inadequate for a complete understanding of man; and that historic Judeo-Christian theology, by looking to transcendence in both man and God, affirms this state of affairs will continue.