FOREWORD

Contemporary research in artificial intelligence and cognitive science has been dominated by the conception that minds either are computers or at least operate on the basis of the same principles that govern digital machines. This view has been elaborated and defended by many familar figures within these fields, including Allen Newell and Herbert Simon, John Haugeland, Jerry Fodor and Zenon Pylyshyn, Margaret Boden, Daniel Dennett, David J. Chalmers, Philip Johnson-Laird, Steven Pinker and scores unnamed, who have committed themselves to its defense. If general agreement were the measure of truth, it would be difficult to deny that computationalism, thus understood, appears to be true.

Indeed, even most of those who disagree with the computational conception tend to accept some of its elements. The alternative known as connectionism, for example, which has been endorsed by David Rumelhart, Paul Churchland, and Paul Smolensky, among others, retains the conception of cognition as *computation across representations,* even while it rejects the contention that computationalism has properly understood the nature of representations themselves. The difference between them turns out to be that connectionists hold that representations are distributed as patterns of activation across sets of neural nodes, where cognition is now taken as *computation across distributed representations.*

Even granting the difference between nodular and distributed forms of representations, computationalists and connectionists can agree on many aspects of cognition. The classic formulation of the computational conception, for example, has been advanced by John Haugeland, who has suggested *that thinking is reasoning, that reasoning is reckoning, that reckoning is computation, that computation is cognition, and that the boundaries ofcomputability define the boundaries of thought.* As long as the differences between nodular and distributed representations are acknowledged, there are no reasons—no evident reasons, at least—why a connectionist should not adopt these computationalist contentions.

In order to establish whether minds are or operate on the basis of the same principles that govern computing machines, however, it is necessary to accomplish three tasks. First, discover the principles that govern computing machines. Second, discover the principles that govern human minds. And, third, compare them to ascertain whether they are similar or the same. That much should be obvious. But while leading computationalists have shown considerable ingenuity in elaborating and defending the conception of minds as computers, they have not always been attentive to the study of thought processes themselves. Their underlying attitude has been that no theoretical alternative is possible.

Part I: *Semiotic Systems.*

The essays collected here are intended to demonstrate that this attitude is no longer justified. The conception of *minds as semiotic (or "sign-using") systems* should lay this myth to rest, once and for all. According to this approach, which builds on the theory of signs advanced by Charles S. Peirce (1839-1914), *minds* are the kinds of things that are able to use signs, where *signs* are things that stand for other things (in some respect or other). Since there are at least three different ways in which things can stand for other things, there turn out to be at least three different kinds of minds, which use different kinds of signs.

"Semiotic systems" as *systems that can use signs* ("minds"), of course, differ from "semiotic systems" as *systems of signs of kinds that semiotic systems can use* ("signs"). The meaning of a sign for a system is understood on the basis of its effects on (actual or potential) behavior, which requires a dispositional conception of meaning that is non-behavioristic, non-extensional, and non-reductionistic. The semiotic conception thus satisfies Fodor's condition that a theory of cognition ought to connect the intensional properties of mental states with their causal influence on behavior. It would be a blunder to dismiss the semiotic account merely on the basis of mistaken preconceptions about dispositions.

Thus, the three essays that appear in Part I provide an introduction to the theory of semiotic systems that includes a discussion of the theory of signs and of the kinds of minds that are distinguishable on that basis. Since there are at least three *kinds of signs*—namely, iconic, indexical, and symbolic—which have the ability to use things as signs on the basis of their relations of resemblance, of cause-or-effect, or of habitual association to other things, there also appear to be at least three *kinds of minds*—namely, iconic, indexical, and symbolic—which have the ability to use signs of those kinds. And there appear to be two higher kinds of mentality, known as transformational mentality and as metamentality.

The conception of cognition as computation across representations, of course, would be severely undermined if what computers can compute does not count as "representations", at least for those machines, where a distinction has to be drawn between "signs" that are significant for *the users of* machines and "signs" that are significant *for use by* those machines. There is nothing surprising about the idea that inputs and outputs are significant for the users of machines, which are designed to fulfill the expectations we impose upon them. The question that matters is whether those inputs and outputs are meaningful for those machines.

Thus, the papers in Part I concern one of the most important issues at stake here. The arguments presented there—in relation to definitions and definability, the physical symbol system hypothesis, and syntactical and semantical theories of language—are intended to establish the existence of what I shall call *the static difference* between them, which emanates from the capacity of digital machines to process marks as opposed to the ability of semiotic systems to process signs:

THE STATIC DIFFERENCE

ARGUMENT 1: Computers are mark-manipulating systems, minds are not.

Premise 1 Computers manipulate marks on the basis of their shapes, sizes, and relative locations.

Premise 2: These shapes, sizes, and relative locations exert causal influence upon computers, but do not stand for anything for those systems.

Premise 3: Minds operate by utilizing signs that stand for other things in some respect or other for them as sign-using (or "semiotic") systems.

Conclusion 1: Computers are not semiotic (sign-using) systems.

Conclusion 2: Computers are not the possessors of minds.

Thus, even if minds effected transitions between thoughts as the computational conception commends, the conception of minds as computational systems and as semiotic systems would continue to distinguish between them, as Part I explains.

Part II: *Computers and Cognition.*

If some computationalists have displayed the tendency to take for granted that minds operate on the basis of the same principles that govern computing machines, a matter that requires both empirical investigation and conceptual clarification for its vindication, others have been more cautious. Nevertheless, the dominant paradigm still appears to be vulnerable to detailed explorations of the nature of thought processes, on the one hand, and of the nature of computational procedures, on the other. The studies presented here strongly suggest—in my view, prove conclusively—that the principles that govern thought processes differ from those that control computational procedures. While Part I focuses on representations, Part II focuses on the nature of computation itself.

The core question is whether thinking is computing or whether computing is thinking. Since computing depends upon programs, and programs, in turn, depend upon algorithms, the answer hinges upon the existence and ubiquity of *mental algorithms.* Suppose, for example, that certain kinds of thinking—such as dreams, daydreams, and even ordinary thought processes as well as perception and memory—are *not* controlled by algorithmic procedures. Then even if particular kinds of problem solving (such as the evaluation of proofs in logic and mathematics, for example) *are* controlled by algorithmic procedures, it would remain the case that computing is at best one special kind of thinking.

Thus, the papers in Part II concern another of the most important issues at stake here. The arguments presented there—in relation to exacting comparisons between the properties of algorithms and of programs as causal implementations of algorithms, and of the nature of thought processes and of computational procedures—are intended to establish the existence of what I shall call *the dynamic dijference* between them, which arises because computers process information by means of algorithms, while minds, by contrast, are non-algorithmic systems:

THE DYNAMIC DIFFERENCE

ARGUMENT 2:	Computers are governed by algorithms, but minds are not.
Premise 1:	Computers are governed by programs, which are causal models of algorithms.
Premise 2:	Algorithms are effective decision procedures for arriving at definite solutions to problems in a finite number of steps.
Premise 3:	Most human thought processes, including dreams, daydreams, and ordinary thinking, are not procedures for arriving at solutions to problems in a finite number of steps.
Conclusion 1:	Most human thought processes are not governed by programs as causal models of algorithms.
Conclusion 2:	Minds are not computers.

Indeed, the conception of cognition that emerges from these analyses is that *cognition is a causal process involving the use of signs.* The computational conception appears to be no more than an overgeneralization based upon a special kind of thinking that is not characteristic of thought processes. Thinking is *not* reducible to reasoning, reasoning is *not* reducible to reckoning, reckoning is *not* reducible to computation, computation is *not* cognition, and the boundaries of computability do *not* define the boundaries of thought. Ironically, even computers themselves are best understood as special kinds of signs, where their own significance presupposes the existence of interpretations, interpreters, or minds.

The computational conception, according to which cognition is computation across representation, thus appears to be fatally flawed. First, as the papers in Part I have displayed, the kinds of representations that computers compute are not signs that are suitable for cognition. Second, as the papers in Part II have established, even if the kinds of representations that computers compute were suitable for cognition, computing is at best a special kind of thinking that does not adequately represent the character of cognition. Neither the nature of representations as the objects of cognition (thought) nor the nature of transitions between them (transitions between thoughts) are computational in character.