Book Reviews and Response

Z.W. Pylyshyn, *Computation and Cognition: Toward a Foundation for Cognitive Science* (MIT Press, Cambridge, MA, 1986); 292 pages, $33.75 (hardcover), $9.95 (paperback).

Reviewed by: Alan K. Mackworth
Department of Computer Science, University of British Columbia, Vancouver, BC, Canada V6T 1W5

The prototypical working scientist may often safely ignore foundational treatises and debates. Smugly ensconced in the enfolding security of a received framework, he (typically, he is a "he") twiddles symbols or knobs, enjoying the simple pleasures of puzzling out nature's secrets. Tempting as that idyll may be, the prototypical researcher in AI or cognitive science should renounce it. For us, it's premature—both illusive and elusive.

Pylyshyn's goal is to show that cognition, taken broadly to include perception, problem solving, planning, language use and action, is a natural scientific domain. He argues that an autonomous science of cognitive systems must have a unified vocabulary and a coherent set of principles with a range of established theoretical and experimental methods. The principles must be as parsimonious, exact and constraining as possible to allow explanation and prediction of a wide range of cognitive phenomena as well as falsification and revision of the theory. Although this is commonly accepted, most cognitive experimenting and theorizing, by psychologists and computer scientists, is radically underconstrained and thereby flawed. Pylyshyn puts forward a strong testable hypothesis: cognition is computation. This is risky. The popular tactic of seeing computation as a *metaphor* for cognition is much safer. Unfortunately, metaphors are not scientific hypotheses: they cannot be wrong.

Pylyshyn goes further. Even if the computation hypothesis is adopted, he does not allow the claim that any program that simulates some behavior stands as an explanatory theory. A theory must explain and predict behavior by exhibiting a computational embodiment that is "strongly equivalent" to the cognitive process. "Strong equivalence" means isomorphism at the descriptive level of representation-based algorithms and "functional architecture". The operational test of "cognitive penetrability" is one way to determine independently whether a capability is to be modelled by rule-governed computation or embedded directly in the functional architecture. In summary, the hypothesis is that cognition is representation-based computation on a specific functional architecture. The process-architecture interface is system-dependent but task-independent; various methods are proposed for determining it. Computational cognitive theories that presuppose an architectural primitive which is demonstrably cognitively penetrable are, under this view, simply inadequate.

When Pylyshyn's thesis is baldly summarized in this impoverished way, the reader may come to hold one of two incompatible negative views of it. The first is, put crudely, "So what? It's obvious. We all know that." The second is, "No. It's all wrong. It doesn't explain emotion, learning, connectionism, or my grandmother." A reader holding the first view should realize that most of it may feel obvious in hindsight and it is not, of course, entirely original but by summing it up and spelling it out so clearly, Pylyshyn has moved us closer to a science of cognition. A reader holding the second view should carefully consider the importance of focussing a fledgling, sometimes floundering if not foundering, science on the proper questions. Pylyshyn provides a coherent foundational framework for mainstream cognitive science. Sceptics and critics of the mainstream orthodoxy have here a set of clear targets. Other foundational frameworks must address the same issues although the answers could be quite different.

Having built the framework, Pylyshyn then uses the tools it provides to study a range of philosophical, theoretical and experimental approaches to the design and analysis of cognitive systems. For practitioners of AI the key conclusion is that implemented code is only a necessary component of a useful contribution to cognitive science. He outlines a sufficient set of such components.

Some sections of the book are hard to read. They are simply more obscure and discursive than they need be, even granting the intrinsic difficulty of these matters. For the reader with the necessary time, patience, energy and motivation, I recommend reading the book in its entirety. In fairness, I warn the gentle reader that those resources are required in good measure. The faint of heart should start with the excellent epilogue and summary in Chapter 9 and then read the Preface and Chapters 1–5 until any one of the requisite resources is exhausted. As with any complex detective story, with a plot entangled in clues and cues, it can be more rewarding to read the *denouement* first. Much of the material in Chapters 6, 7 and 8 could be skimmed lightly by AI researchers not much interested in private psychological turf wars over behaviorism, direct perception, analogues and mental imagery, although the material there follows coherently from, and supports, the foundational thesis.

In short, *Computation and Cognition* is recommended as an effective, if occasionally indigestible, antidote for methodological smugness.

Z.W. Pylyshyn, *Computation and Cognition*: *Toward a Foundation of Cognitive Science* (MIT Press, Cambridge, MA, 1986); 292 pages, $33.75 (hardcover), $9.95 (paperback).

Reviewed by: Mark Stefik
System Sciences Laboratory, Xerox Palo Alto Research Center, Palo Alto, CA 94304, U.S.A.

Foundations for Understanding Cognizers

Pylyshyn's book is about foundations. His ambitions for this at first appear modest. As he says,

> A more realistic way to put it might be to say that the material contained here is to the foundations of cognitive science what a hole in the ground is to the foundation of a house. It may not seem like what you eventually hope to have, but you do have to start some place.

Pylyshyn is interested in phenomena of cognition. He proposes to study entities that exhibit intelligent behavior.

> Just as the domain of biology includes something like all living things so the domain of cognitive science may be something like "knowing things," or as George Miller (1984) colorfully dubbed it, the "informavores."

Pylyshyn himself re-dubs these things "cognizers" which gives us cognitive science as the study of cognizers. This sounds somewhat circular; it shifts the matter from a question about the central properties of cognitive science to the central properties of cognizers. Speaking like a philosopher of science, Pylyshyn notices that although cognizers have mass, gravity is already well-covered by physics; similarly phenomena involving sound, or physical hardness, or chemical reactions are already explained by established and separate disciplines. So which properties are central? Pylyshyn doesn't address that question right away. But he sets a direction, arguing that the natural category "cognizer" encompasses a populatiuon of objects such as people and computers.

In part Pylyshyn defines his position in opposition to others. For example, he observes that the philosophical discussions of cognitivism are

> concerned with certain issues, such as the proper understanding of "meaning" or the role of consciousness . . . While these questions may be fascinating . . . they may well be precisely the kind of questions that are irrelevant to making scientific progress. Indeed,

> they may be like the question, how it is possible to have action at a distance . . . debated with much acrimony in the eighteenth century. The question remains unresolved today; people may have decided that the question is not worth pursuing . . . .

Cognizers act on the basis of representations. Pylyshyn identifies this as their central property. He identifies these representations as rich structures of significance; not only are they "state" in a computational sense, they are about the environment and they correspond to things we call beliefs and goals. His central thesis is that cognizers.

> instantiate such representations physically as cognitive codes and . . . their behaviour is a causal consequence of operations carried out on these codes.

To this reviewer that sounds very much like the physical symbols system hypothesis, rephrased in computerese. One can ask whether this idea explains enough. Pylyshyn sets out to show a range of implications of the idea. Pylyshyn identifies his approach as a "representational theory of mind."

From Philosophy to Computation

Pylyshyn illustrates his points with brilliantly simple, concrete examples. One such example that is extended several times concerns a pedestrian who sees a car skid and crash into a pole. He runs across the street, looks into the wreckage, races off to a telephone, and dials 911. Pylyshyn wants to know what kind of theory would be appropriate for explaining this sequence of events. He first notes that the activities of the Cognizer are not usefully explained by the laws of energy and momentum from physics. Those principles describe the skidding of the car, but the racing of the pedestrian is not connected to the accident by transfer of energy or momentum.

Pylyshyn raises several variations of this story, where the pedestrian dials 91 instead of 911, or in which the pedestrian uses a pushbutton phone instead of a dial phone. Pylyshyn recognizes that it is a weak staw man, indeed, who argues that cognition might be a simple extrapolation of physics. The core of his concern is with explanations that can account for computational quandaries of mind, more than the physical.

What is needed, he argues, is that a computer program should "actually perform at a sufficiently fine and theoretically motivated level of resolution." The early GPS papers (Newell and Simon [1]) are a model of this approach.

The Parts of the Case

Pylyshyn's argument goes through several stages. He argues that one needs to decide which kinds of regularities need explaining. He recommends starting with phenomena similar to those explained by what he calls folk psychology or "grandmother's psychology," that is, stories that account for human behavior in terms of what a person is expected to know or feel. Here, he connects action with stored representation. People do the things they do because of what they know and believe. If they knew or believed different things, they would act differently. A theory that doesn't do some of this starts off on the wrong foot.

Representations are encoded. Pylyshyn describes the nature of encoding and action by analogy with the programming of computers. For many AIers, this set of assumptions at this level of generality passes with neither question nor examination. How could it be wrong? Pylyshyn is interested in understanding cases where it could be wrong, and his sequence of points builds to some problematic examples.

Pylyshyn aruges that mind is understood in levels. He discusses (1) a biological or physical level, (2) a symbolic or functional level, and (3) an intentional or semantic level. The biological level corresponds to the physical substrate of memory. The functional level reflects the situation that many different physical configurations can mean the same thing. Within this level, Pylyshyn wants to define classes of representations according to their equivalent roles in explaining behavior. He wants to be able to assign "symbolic codes" to these classes so as to explain how actions connect with behavior. Pylyshyn equates his semantic level with Newell's knowledge level.

By describing these levels he is trying to establish what he calls a functional architecture. It is worth noting that a description at this level of generality encompasses many possible models of mind. For example, it includes GPS-like models from [1]. It also includes, for example, models of scripts and plans, or say, any of John Anderson's models. The computational framework doesn't rule much out.

Pylyshyn then goes on to develop equivalence criteria, for deciding whether a computational model explains a process. He demands not only a kind of similarity in possible behavior, but also a close complexity equivalence. By this he means more than measures of time and space as in the complexity of computer algorithms. He wants a close correspondence in the actual steps and processes. Two factors make this prescription challenging in practice: limited theory and difficulties in observation. Theories of complexity for parallel algorithms are at a very early stage of development. The combinations of parallel, serial, and distributed processes and systems that are studied systematically in computer science are much simpler than those that arise in biological systems. Secondly, these processes are difficult to observe in biological systems. This combination makes the prescription difficult and limiting. For this reviewer it was not made sufficiently clear just how much leverage Pylyshyn expects from these concepts and techniques because he did not acknowledge the practical impediments.

Pylyshyn is particularly interested in understanding how to tell which behaviors are alterable by changes in knowledge or belief. He characterizes the ones that can be changed as being "cognitively penetrable."

Pylyshyn then applies this development to approach the question of how and whether the processes of perception and mental imagery are cognitively penetrable. In this he wants to contrast his theory with other notions of "imaginal thinking," because as he says, other investigators have felt that a computational view might have to give way to something different. He argues that experiments that illustrate cognitive penetrability support the relevance of a representation theory of mind for partially explaining these phenomena.

Whose Foundations? A Closing Parable

Laying foundations is a challenging task. There is recurring debate in cognitive science about foundations. Frameworks for understanding the mind are as different as the proverbial accounts from a group of blind men, each describing an elephant from where he stands. "It's a hose." "It's a rope." "It's a tree." "It's a hill."

A recent example of this occurred in the opening session of the Symposium "How can slow components think so fast?" organized by the American Association for Artificial Intelligence at Stanford in March 1988. Allen Newell gave the opening presentation describing some general constraints on architectures for cognition.

Beginning with the biological puzzle suggested by the title of the session, Newell developed a set of levels of organization building up from "organelles" which perform operations in 10^{-4} seconds through neurons, neural networks, on up to brain-sized organizations and beyond. At each successive level of organization the larger operations require approximately ten times as long to carry out.

Newell then observed that there is almost no time available for a neural system to produce cognitive behavior. He defined cognitive behavior in terms of "deliberate actions," that is, actions that require bringing knowledge to bear, such as making an appropriate move in chess. In many examples, an interval of about one second (say 150 milleseconds to 10 seconds) is long enough for humans to take deliberate actions. After accounting for time necessary for nondeliberative processes being observed, Newell concludes that deliberation itself requires as little as 100 msecs. From this he surmises that cognitive brain architectures must usually employ learning mechanisms that enable most deliberation to be carried out by recognition.

Following Newell's presentation three speakers were invited to comment on it. This was where the "blind men" phenomenon became apparent. Jack Byrne did not comment on Newell's presentation directly, but rather reported on detailed mechanisms for learning and memory in the nerve "circuitry" for one portion of the simple nervous system of a sea mollusk. Bernardo Huberman gave a whirlwind presentation of results on dynamical systems and how dynamical models can make useful predictions about chaotic behavior, learning, and interactions of agents in an environment.

Only the last speaker, Steve Kosslyn, addressed Newell's presentation directly and this started a debate. Kosslyn argued that while Newell's level analysis provides limits and sanity checks on theories of mind, he doubted its utility for understanding the architecture of the mind. He went on to discuss experimental techniques for mapping human and animal behavior to active areas of brain structure.

To this observer and reviewer the most striking point in the discussion was not any contradiction between Newell's predictions and the mapping of brain activity. Rather, there was a large gap between the phenomena of interest to the different parties. Newell's arguments said nothing about the organization of learning across large areas of a brain; Kosslyn's arguments did not give a fine enough grain to localize behavior and provided essentially no insights into dynamics. Neither gave predictions that would bear on the nerve mechanisms of Bryne's sea mollusk nerve networks; Huberman's theories and equations could apply as well to large networks of many kinds, whether they be computers, nerves, or astronomical bodies.

So where in all of this is cognitive science, and what has this to do with Pylyshyn's book?

To take the second question first, we note that Pylyshyn acknowledges an intellectual debt to Newell, and spent time with him during some of the writing of the book. If Pylyshyn had given an address on this at the Symposium, this reviewer predicts that his claims of foundations would provoke a similar debate.

Pylyshyn's book brings order to the debate by proposing to carve cognitive science "at its natural joints." Thus, he argues that a human mind operates at different levels and can be understood best by recognizing those levels: the physical, the symbolic, and the intentional.

A framework like Pylyshyn's can be challenged by raising different kinds of questions. Are there natural joints for understanding the mind? If yes, are Pylyshyn's and Newell's joints the right ones? Do we know enough about nerves, brain organization, and cognition to adequately evaluate the proposal?

Pylyshyn himself acknowledges this issue in the introduction by his analogy about house foundations and a hole in the ground. Carving nature at joints helps to sort out the kinds of questions we ask. However, we know very little about either mental phenomena or brain structure and the theories and data we have explain a small subset of the range of phenomena.

To be too satisfied with the explanations can preclude attention to other questions. Furthermore, it's not enough that joints separate phenomena into levels. It is better if they also suggest how we can understand the connections between levels.

In the context of this young and active field, any synthesis is incomplete. Cognitive scientists must still contend with the issues of determining which kinds of questions are crucial, and which have answers within their grasp. Controversy arises when claims about foundations are seen as implicitly determining the main questions to be investigated, the methods to be used, and the scientific disciplines said to be central. If the debaters in the AAAI symposium were to offer recommendations for how to understand the organized complexity of the mind, we could imagine very different recommendations. "Study nerve networks." "Study learning and chaotic networks." "Study brain areas and function."

Recommendations

Even given that there are other approaches and foundations for understanding the mind, *Computation and Cognition* is a worthwhile book.

I have two main complaints. The first is that the prose is often dense and wordy. (If I said that it was not cognitively penetrable, I would be misusing that phrase.) The book needs a crisp summary of the definitions of levels and a crisp summary of the cases for perception and mental imagery. So many of Pylyshyn's words were about the stance of his theory relative to others (in terms of the philosophy of science) that it was hard to find out what his position was.

My other complaint is that I wanted a more detailed theory of some aspects of mind. For that, one must turn to other authors. Pylyshyn's book stays at the abstract level. For example, his concern with perception is mainly with whether a representational theory of mind is relevant to explaining it. There are no diagrams of the eye or statements of any detail about how human vision might work. Similarly, although Pylyshyn is interested in drawing from folk psychology, there is a very large gap between grandma's prediction of what cousin Fred might do and a predictive model for any aspect of Fred's mind. Pylyshyn eschews details.

My main praise of Pylyshyn's book is that he provides compelling, short, and clear examples for the points he chooses to make. For students of AI the book shows how phenomena that might seem quite ordinary can be quite remarkable when seen through a psychological perspective informed by computing.

In reviews [1] of an earlier version fo Pylyshyn's work, reviewers of both theoretical and experimental persuasions picked at various aspects of Pylyshyn's approach. They discussed ambiguities in the application of the cognitive-penetrability criterion. They gave examples of computational activities in various kinds of nerve and brain cells that challenged the neat separation Pylyshyn's levels seem to propose between neuroscience and cognitive psychology. The examples show systems affected by information infusion that are not thought of as being cognitive (e.g., the digestive system) and specialized systems of parallel and distributed computation that make crucial use of representations other than beliefs and goals. In responding to these earlier criticisms, Pylyshyn has tightened up his definitions a bit.

If Pylyshyn decides to give us an updated edition of this book, this reviewer would like to see a presentation that reaches into more experimental examples. At this stage of our understanding, it seems more fruitful and illuminating to attend to observations than arguments about philosophy. It is not that the computational perspective is wrong. Rather, the programming and system examples from nature are richer than the ones we have invented ourselves. Thus, it's not that we need to export computer science to psychology and neuroscience; rather, some more traffic in both directions is needed. Making the examples accessible to an AI audience would energize that field at just the right time, as its computational options are becoming more interesting. In the absence of a new edition, this reviewer suggests that AI readers augment the book with its earlier reviews [2, 3].

REFERENCES

1. Newell, A. and Simon, H.A., *Human Problem Solving* (Prentice-Hall, Englewood Cliffs, NJ, 1972).
2. Open Peer Commentary on Cognition and Computation, *Behavioral and Brain Sci.* **3** (1) (1980) 1133–1153.
3. Haugeland, J., Review of Computation and Cognition, *Philosophy of Science* **54** (2) (1987) 309–311.

On *Computation and Cognition*: *Toward a Foundation of Cognitive Science*, a response to the reviews by A.K. Mackworth and M.J. Stefik.

Response by: Zenon Pylyshyn
Centre for Cognitive Science, University of Western Ontario, London, Ontario, Canada N6A 5C2

Why Foundations, and Where Next?

The two reviews of my book that appear in this issue raise a number of questions that prompt this brief commentary. As to the particular evaluations that are contained in the reviews, I have scarcely any disagreements. I agree with Stefik that there are many perspectives one can take on the major foundational issues, and that there is much room for maneuvering at this stage in the development of the science. I also agree with Mackworth's comments—including his recommendations for how the book might be read. I take the opportunity provided by the editor to comment on these reviews as an invitation to suggest how the book might have been written, now that I have had the advantage of hindsight, a large number of published reviews, and the intervening approximately five years during which a great deal of discussion has taken place concerning the foundations of cognitive science.

It will come as no surprise that I believe the main proposals contained in *Computation and Cognition* still hold. In fact, I am distressed to find them being referred to as the *orthodoxy* [1]. However, if the book were being written now, the arguments in support of these proposals would have to be expanded in several directions in response to recent advances. Despite the fact that the basic assumptions laid out in *C & C* appear to many (e.g., apparently to both Mackworth and Stefik) to be nearly truisms, vigorous opposition to the tri-level view of computation and cognition continues. This is the view, familiar to computer scientists, that both processes have a stable level of organization at the semantic, syntactic, and physical levels; or as Allen Newell puts it, the knowledge level, the symbol level, and the physical level. Opposition to it comes from the same sources discussed in *C & C*: from philosophers who accept intentional (or knowledge level) explanations but do not accept the symbolic level as a genuine level of explanation; from those (including both philosophers and biologists) who eschew both intentional and symbolic explanations, from psychologists who believe that most cognitive processing (e.g. that involving mental imagery) will have to be modeled in terms of processes which do not manipulates discrete symbols—perhaps ones that use an "analogue" medium; and even from some who accept the symbol level but object to the use of knowledge ascriptions (all of these are referenced in *C &*

C). However, by far the most influential recent source of objection has come from the connectionist or PDP movement [2]—a movement whose acceptance, particularly by psychologists and by the popular science press, threatened for a time to undermine the current computational approach in cognitive science.

This is not the proper forum for a discussion of the general issues raised by these objectors, especially since Jerry Fodor and I have just published a rather long analysis of the shortcomings of the connectionist "alternative" [3], and I have also summarized some of my recent views on these issues [4, 5]. But throughout such debates one should not lose sight of the point of trying to illuminate the foundational assumptions and roots of our present "orthodoxy". The reason, as has been said in connection with the study of history more generally, is that otherwise we may be doomed to repeat our mistakes. Indeed, there is every reason to believe that we are in the midst of repeating many of the major errors of the first "cybernetic revolution" as well as the first 50 years of behaviorist psychology.

Two failures have combined to make this new attack on the orthodoxy popular, both of which illustrate the relevance of understanding the foundational assumptions of one's science. The first is the failure to understand what the traditional view, which has come down to us from Turing and other early theorists, is committed to and, even more importantly *why*. The second is the failure to understand the reasons why the first round of cybernetic models were inadequate, and consequently the loss of an opportunity to learn from their failures.

While the differences between the connectionist models and those of the first series of "neural-like" models (e.g., the first *perceptron*, or Uttley's early *conditional probability computer*) may appear to be technically significant, the two approaches are in fact very similar in spirit and in their underlying epistemology. Both are motivated by the same intuitively appealing idea, namely the idea of building a cognizing system from elementary ("neural-like") elements that become organized by virtue of the statistical properties of events occurring at their inputs. In the meantime the lessons learned from the deep discoveries of Turing, mathematicians studying the foundations of mathematics and computing, logicians, and other theorists, are misunderstood and dismissed as "old fashioned". These discoveries include the need for a sufficiently rich symbolic representational system, in order to encode the semantic domains over which the systems are computing, and the need for a system of transformations that preserve the semantic interpretations of the codes (e.g., numerals and operations which mirror the functions of number theory, in the case of Turing's original machine; logical forms for encoding propositional contents and transformations that preserve semantic coherence, in the case of the various systems designed to carry out both logical and heuristic inferences).

The fundamental set of assumptions embodied in the classical view of computing has been discussed by many people, probably most effectively in the

AI tradittion by Allen Newell, who refers to the *Physical Symbol System Hypothesis*, and in the philosophy of mind tradition by Jerry Fodor, who refers to it as the *Language of Thought Hypothesis*. However, instead of focusing on these hypotheses and the elaborate reasons that have been given for adopting them, the connectionism movement has focused on *current implementations* of classical architectures and has dismissed the entire set of ideas on the basis of such irrelevancies as the speed, damage resistance, determinism, lack of parallelism, etc of many of the current models, as well as the apparent differences between present-day computers and the nervous system. These are all properties that, however true they may be of many contemporary models, are irrelevant to the architectural ideas discussed in *C & C*, and certainly do not provide arguments that favor connectionism over the classical ideas about computation, such as those that derive from the insights of von Neumann and Turing.

Mackworth correctly points out that scientists working within an accepted paradigm feel that they can safely ignore foundational questions. As if to illustrate this premise, Stefik asserts that "At this stage of our understanding, it seems more fruitful and illuminating to attend to observations than arguments about philosophy." These are among the most common criticisms I have heard of the enterprise attempted in *C & C*. To the extent that this is a plea for further research, good ideas, and more solid results, then one can scarcely disagree. Yet in cognitive science, unlike, say, in solid state physics, every result comes wrapped in a great number of contentious methodological and philosophical commitments, typically unacknowledged. That's partly because the field is young and diffuse and in the process of defining itself. It's also partly because the questions that are being asked typically have had a long history in philosophical inquiry. Because of this, one cannot begin the process of *scientific* inquiry in cognitive science without taking a stand on a dozen different classical philosophical issues, most of which represent positions that are deeply ingrained in our culture. One of these is a belief in the doctrine that behavioral capacities of organisms can be traced directly to the record of causal events in their immediate environment, particularly to statistical properties of this record. Another is a deep skepticism of mentalism and of related notions such as "meaning", and a consequent belief in the primacy of biological facts over other sorts of facts and considerations for understanding intelligent behavior (e.g., facts about the systematicity of both actual and potential behaviors, the systematicity of cognitive capacities, the nature of the information available in an organism's past experience in relation to the organism's current capacities, as well as considerations of the descriptive adequacy, generality, and sufficiency of models).

There is no question that we badly need results, and that students need detailed expositions of the best current work. I can understand Stefik's wish that there might be more of this sort of exposition in *C & C*. But *C & C* is not intended to be that kind of book (however, see [6] for some attempts at a more expository book). It is a book which, to be sure, purports to be thoroughly informed by results of current work in several fields, but which does not present an exposition of the details of that work. Rather it attempts to map out the territory where work is progressing, abstracted from particular models—except where the latter are needed to illustrate a point (the book does, after all, make most of its points through a very large number of references to empirical results and to particular models). It attempts to lay bare the assumptions that the most sophisticated practitioners have implicitly adopted.

While it would be counterproductive for a large part of the field of cognitive science (or AI) to spend its time engaged in foundational analysis, it is also an illusion that foundational assumptions can be avoided by rolling up one's sleeves and getting down to work—say by proving statistical convergence theorems for some randomly connected network of entities, whose ultimate explanatory capacity has known shortcomings. We need results, but we also need to be nudged now and then into an awareness of the implicit assumptions within which we are working, and with their entailments. And that was the purpose of *C & C*, a purpose which I believe it can fill in a modest way with a certain amount of patient squeezing and shoving, along the lines recommended by Alan Mackworth.

REFERENCES

1. Cummins, R., Critical notice: Review of Compution and Cognition, *Can. J. Phil.* **18** (1988) 147–162.
2. Rumelhart, D. and McClelland, J. (Eds.) *Parallel Distributed Processing* (MIT Press, (Bradford Books), Cambridge, MA, 1986).
3. Fodor, J.A. and Pylyshyn, Z.W., Connectionism and cognitive architecture: A critical analysis, *Cognition* **28** (1988) 3-71; also in: S. Pinker and J. Mehler (Eds.), *Connections and Symbols* (MIT Press (Bradford Books), Cambridge, MA, 1988).
4. Pylyshyn, Z.W., Computing in cognitive science, in: M. Posner (Ed.), *Foundations of Cognitive Science* (MIT Press, Cambridge, MA, 1989).
5. Pylyshyn, Z.W., The role of cognitive architecture in theories of cognition, in: K. VanLehn (Ed.), *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition* (Erlbaum, Hillsdale, NJ, 1989).
6. Posner, M. (Ed.), *Foundations of Cognitive Science* (MIT Press (Bradford Books), Cambridge, MA, 1989).