



Response

The scope of the cognitive sciences:  
Reply to 6 reviews of *The MIT Encyclopedia of the  
Cognitive Sciences*

F.C. Keil

*Department of Psychology, Yale University, 2 Hillhouse Avenue, P.O. Box 208205,  
New Haven, CT 06520-8205, USA*

Although there have been several reviews of the *The MIT Encyclopedia of the Cognitive Sciences*, the six reviews in this issue of *Artificial Intelligence* represent an unusual opportunity to see in one collection how scholars from a wide range of perspectives evaluate MITECS. I found it very useful to consider the reviews side by side and am grateful to the reviewers for providing a number of new insights into the nature of the cognitive sciences. It is also gratifying to see such generally positive assessments from five of the six reviewers (Carr, Dorr, Husbands, Okamoto and Peterson) and it is intriguing to consider the more negative comments by Lakoff. In this essay, rather than consider in detail the many points raised by the reviews, I examine more globally how a project like MITECS might be evaluated and how it seemed to fare in light of these reviews.

There are four primary ways in which an attempt to provide a comprehensive portrait of the cognitive sciences might be evaluated:

- (1) Does it give a correct picture of the phenomena of interest to cognitive scientists?
- (2) Even if it covers the right phenomena does it also cover the most important approaches to those phenomena?
- (3) Even if covers the right phenomena and the right approaches does it do so in a way that clearly accurately represents the relevant literatures?
- (4) Even if each of the articles succeeds on all accounts individually, do they as a collection link together in ways that give readers a sense of the cognitive sciences overall?

The six reviews of MITECS touch on each of these four kinds of potential problems; but clearly with different emphases. Moreover, in at least some cases, it is not clear which of these criticisms is really intended. Rob Wilson's comment provides a very useful description of how MITECS emerged as a project and how that process might be expected to relate to some of the concerns raised by reviewers. It shows how the evolution of

---

*E-mail address:* frank.keil@yale.edu (F.C. Keil).

MITECS was not driven by some agenda that we had in mind to promote a particular part line of Cognitive Science and exclude others. Instead, our goal was to get the best sense we could of what the cognitive sciences are, a little bit of where they are going, and of the extent to which the whole is greater than the sum of the parts. We also wanted to provide a way for those who want to learn to get a sense of what some area of the Cognitive Sciences is about and, for the more ambitious, what the enterprise as a whole looks like. Our methods of selecting associate editors and contributors were all designed to get a truly non-partisan sense of the cognitive sciences. Given the nature of the various fields that make up the cognitive sciences and the sometimes passionate disagreements that occur within some of those fields, I agree with Rob that we covered most of the major themes and did so without some single overriding bias. That said, there are more subtle ways in which we might have missed the mark that concern each of the four points listed above. Let me turn to each of those:

### **Does MITECS give a correct picture of the phenomena that are of interest to cognitive scientists?**

Not all areas of academic inquiry agree on correct entities of study or even on the existence of roughly the same set of entities. Many in the humanities disagree violently on whether a field such as literary theory is really about intentions of authors, Marxist economic principles, or oppression. In the natural sciences, it is much less common to see such debates. Physicists and biologists do not tend to disagree so much about the metaphysical reality of what they study. Where do the Cognitive Sciences sit in this respect? Several of the Cognitive Sciences seem to have little disagreement. The neurosciences and most of psychology agree on the phenomena of interest, as do most of the philosophers. We found strikingly little discord here on what should be included. Perhaps more surprisingly, we also found pretty good consensus with culture, cognition and evolution despite what appears to be much more potential for polemic in the area. Our reviewers also seem to agree that the articles in those areas were on target, especially Carr, Okamoto, and Dorr. With respect to computational intelligence, there is a concern raised by both Peterson and Husbands that the computational approach has been overemphasized at the expense of other approaches to artificial minds, such as those that look at entire active systems in context. Here, I am less clear on whether the real concern is about a neglect of the right phenomena or about whether the wrong approaches are being taken to right phenomena. It is clear, as argued by Rob Wilson, that an enormous range of senses of computation are part of MITECS, including ones that do not at all endorse symbolic forms of computation. What phenomena are missing? Does looking at intelligent devices in the context of actions in real world settings bring into relief phenomena that are otherwise obscure? Possibly. That is, indeed, the promise of such an alternative. As of yet, however, I cannot see any clear consensus on what those phenomena are.

Finally, we turn to linguistics. Here, as noted, we had some of the most vigorous debates among contributors and with us as the editors. But it is my sense that the debates were not usually about the phenomena. People tended to agree that there are universal patterns of language structure that need explanation, an explanation that will include a formal

description of those structural properties. They also tended to agree on the importance of linking those formal descriptions to language acquisition and learnability issues more broadly. I do think we all owe a great deal to Chomsky for enabling us to see more clearly these phenomena and why they are of interest. There is more agreement on these phenomena than there was 40 years ago, and that agreement represents a major advance. One might interpret Lakoff's remarks about the wrong paradigm as partly about studying the wrong phenomena, but it is hard to see a real case being made along those lines. It is a well known argument in the philosophy of science that new paradigms can highlight new methods and new phenomena in ways that mutually enhance the other; but in the case of cognitive linguistics, it is not clear what phenomena it embraces that other forms of linguistics and cognitive science do not. Metaphor, for example, while a very popular topic in some cognitive linguistics circles, has certainly been a phenomenon of interest to many others as well.

### **Does MITECS cover the right approaches to various phenomena?**

Here, there is certainly more potential for disagreement. After all, even in the hardest of the "hard sciences", such as particle physics, there are major disagreements about methods. Our goal was not to showcase any one method but to give a feel for the range of substantive approaches that were perceived by most practitioners as making real advances. In some areas, there seems to be little disagreement about the range of approaches described. In neuroscience, the contrasting approaches are pretty well understood by everyone. In psychology, there may be more debate about how even to describe some approaches (such as the links between dynamical system theories and connectionist theories), but the terrain of the field is apparent to most of its practitioners. In the case of cultural psychology, the reverse case seems to be true. Scholars agree to a large extent on many of the most interesting phenomena, but are still struggling to find a consensus on what are the most promising approaches. MITECS doesn't fail to cover a key approach in this area as much as it illustrates some of the struggles still going on in attempts to find more fruitful approaches. In the area of artificial intelligence, there may have been a neglect of "non-computational" approaches to intelligence, i.e. those that reject the computational perspective. My own sense here is that these alternative approaches still have the flavor of showing fascinating "demos" rather than articulating clearly in positive terms what those approaches do have at their theoretical core; but I take to heart the reviewers' observations and would certainly want to explore this approach in a second edition and ask potential contributors to try to lay out in specific terms the distinctive properties of those alternative approaches and the constituencies that support them. Although I do not really grasp what criteria make various artificial "systems-in-action" truly non-computational as opposed to simply another computational approach within an already highly diverse set, I am open to that possibility and look forward to learning more about it.

Linguistics is a different case altogether. There are dramatic differences of opinion on how best to study the nature of language and our entries reflect many major controversies in the field. Certainly many of our contributors would be astonished and/or greatly amused to be seen as part of the "formalist nativist" paradigm, as Lakoff suggests they are. As noted

by Rob Wilson, we were constantly struck by how much more often the linguistics entries resulted in animated reviews and disagreements. Such disagreements may well be sign of a healthy evolutionary process in a relatively new field, but they hardly point towards one dominant view. But perhaps the embodied approach or “paradigm” that Lakoff offers was slighted. How could we tell? Presumably because the approach has had a major influence on many in the field. In that case, even if they disagree, scholars should be citing this alternative as a worthy adversary. We did not see such a pattern in the literature. Nor did we see clear successful “demos” of the sort that we do see in AI that make me inclined to wonder about whether another approach is needed to explain those successful results. I respect greatly many of the papers that Lakoff cites but I do not see them neatly coalescing into an alternative approach.

### **Does MITECS cover approaches or phenomena accurately and clearly?**

Here I am pleased there is a consensus among all reviewers that MITECS does a very good job of having clear readable entries that accurately and faithfully describe the relevant literatures. As editors we find this particularly gratifying since we devoted enormous energies to those goals, on several occasions writing reviews as long as the entries themselves. Our contributors were almost all wonderfully responsive to our reviews and to those of “insider” experts and “outsider” members from other cognitive science areas. If nothing else, readers of MITECS are able to get unusually clear, accurate, concise interpretations of most of the major literatures in the cognitive sciences.

### **Does MITECS give an appropriate sense of how the whole of the Cognitive Sciences is greater than the sum of its parts?**

Even if MITECS were perfect on the first three dimensions I have discussed, it might fail to adequately capture how the cognitive sciences intersect to create a truly interdisciplinary program of research and scholarship. Our reviewers wondered the same thing, but were cautious in attributing the lack of connections to failings in MITECS or to failings in the Cognitive Sciences. I think a bit of both is at work. There is no doubt that more connections will emerge in the future and so the field itself is surely incomplete. It is clear from the entries in MITECS that many phenomena (e.g., pattern recognition, coordinated action, learning, communication, numerical computation, tracking of individuals) are relevant to all of the cognitive sciences and that scholars need to be converging more in their discussions. But it is also true, that as editors we are sure to have missed some connections or to have not prodded contributors enough to talk more about possible links. There was plenty of such prodding and some extremely helpful responses; but in other cases, contributors felt that they had done enough to describe their own cherished area in such unforgivingly tight amounts of space. It is one of the hopes of the MITECS project that, by having all of these entries so concisely in one volume and by containing at least some of the links, MITECS itself can be a force towards more meaningful cross disciplinary advances. That said, there is also the danger of succumbing to a cognitive science version

of meaning holism, or the frame problem, in which everything ultimately connects to everything else. We quite deliberately excluded showing links that were technically correct but didn't seem to be truly meaningful. A major challenge for the future is to determine the truly important and significant links that constitute the real highways of interaction in the cognitive sciences.

Overall, given the four ways that I have set up for evaluating MITECS, the reviews seem to be most critical about whether we have covered a fully representative set of approaches. Even here there is not a clear consensus among the reviewers, but perhaps there is an emerging theme that “non-computational” AI merits closer attention. With respect to the fourth issue of cross disciplinary coverage, while there were only modest criticisms, it may be here that the most work needs to be done both by the field and in the next edition of MITECS.