The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology (review)

Ray Jackendoff

Language, Volume 78, Number 1, March 2002, pp. 164-170 (Review)

Published by Linguistic Society of America
DOI: *https://doi.org/10.1353/lan.2002.0024*

malist studies as the brief summary of the chapters has hopefully shown. All articles complement the work of the festschrift's honoree, are well-written, and contain interesting data as well as intriguing analyses, pushing the minimalist spirit further ahead.

REFERENCES

BOŠKOVIĆ, ŽELJKO. 1994. D-structure, theta-criterion, and movement into theta-positions. Linguistic Analysis 24.247–86.

——. 1997. Superiority effects with multiple wh-fronting in Serbo-Croatian. Lingua 102.1–20.

CHOMSKY, NOAM. 1995. The minimalist program. Cambridge, MA: MIT Press.

——. 2001. Derivation by phase. Ken Hale: A life in language, ed. by Michael Kenstowicz, 1–52. Cambridge, MA: MIT Press.

GRIMSHAW, JANE, and ARMIN MESTER. 1988. Light verbs and θ-marking. Linguistic Inquiry 19.205–32.

HORNSTEIN, NORBERT. 1995. Logical form: From GB to minimalism. Oxford: Blackwell.

KAYNE, RICHARD S. 1994. The antisymmetry of syntax. Cambridge, MA: MIT Press.

ZAS
Jägerstr. 10–11
10117 Berlin
Germany
[Kleanthes@punksinscience.org]

The mind doesn't work that way: The scope and limits of computational psychology. By JERRY FODOR. Cambridge, MA: MIT Press, 2000. Pp. 126.

Reviewed by RAY JACKENDOFF, *Brandeis University**

As has been his wont in recent years, Jerry Fodor offers here a statement of deepest pessimism about the possibility of doing cognitive science except in a very limited class of subdomains. F is of course justly celebrated for at least two major ideas in cognitive science: *The language of thought* (Fodor 1975) and the modularity hypothesis (Fodor 1983). However, the form in which these ideas have been taken enthusiastically into the lore of the field differs in some important respects from the form in which F couched them and in which he still believes. As I hope to show, the tension between F's actual views and those generally attributed to him plays a major role in the position he advocates here.

Here is a summary of F's argument, as best as I can reconstruct it. The central issue is the problem of 'abduction': how one determines the truth of a proffered proposition and its consistency with one's beliefs. The chief obstacle to successful abduction is that meaning is holistic: One must potentially check the proffered proposition and inferences from it against one's entire network of belief/knowledge. The resulting combinatorial explosion makes it impossible to reliably fix new beliefs and plan new actions within a traditional Turing-style computation. For F, this casts serious doubt on the computational theory of mind, which presumes Turing-style 'symbolic' computation over the syntactic form of mental representations.

F dismisses a number of proposed solutions to the problem of abduction. Connectionist-style computation, he maintains, is actually a step backward, since it cannot even capture the characteristic free combinatoriality of thought, an essential feature of the language of thought hypothesis. Here I concur; Marcus 2001 offers an extended argument to this effect. F also argues that a system of heuristics is unsatisfactory since one needs to perform an abduction to determine which heuristic to apply. I find this argument less convincing; we'll return to it below.

It is worth mentioning that when F speaks of 'Turing-style computation', it is not clear whether he intends to include massively parallel 'symbolic' computation. Such computation is perhaps

* I am grateful to David Olson, Jerry Samet, and especially Merrill Garrett for considerable discussion concerning this book. This work was supported in part by a Fellowship at the Wissenschaftskolleg zu Berlin and in part by NIH Grant DC 03660 to Brandeis University.

mathematically equivalent to serial Turing-style computation, but it is quite different in practical terms. Certainly the brain's form of computation is massively parallel, whether connectionist or symbolic or some combination thereof. It is interesting therefore to ask if such computation is of any practical help in solving the combinatorial explosion of abduction; F does not address this question.

F's dismissals of connectionism and heuristics, however, are just warmups for his principal line of attack. This is aimed against the thesis of 'massive modularity' proposed by such people as Pinker (1997) and Cosmides and Tooby (1992): the idea that the entire mind (or most of it, anyway) is made up of innate domain-specific modules. Most everybody seems to consider this a natural extension of F's modularity hypothesis. What is frequently overlooked is that the final chapter of Fodor 1983 argues at length that, unlike input systems, thought cannot be modular, because of meaning holism.

F sticks to his guns in the present book. He observes that if massive modularity were correct, it would keep abduction under control; each module would have access only to a small subset of one's knowledge and thus would avoid the combinatorial explosion of dealing with everything at once. He then argues that, on a variety of grounds, massive modularity cannot be correct. Therefore, he concludes, abduction remains a crushing objection to the computational theory of mind.

F then trains his armaments on evolutionary psychology. Pinker and Cosmides and Tooby, among many others, have urged that a proper theory of the mind/brain ought to include a reasonable story of how it got to be the way it is; hence considerations of evolutionary plausibility ought to play a role. In general, natural selection tends to come up with good tricks that solve highly specific problems and to evolve complexity by piling new tricks on top of old. It does not tend to evolve maximally general, nonmodular solutions that apply across the board. On these grounds, they argue, given that the brain evolved under the pressures of natural selection, a massively modular theory is prima facie more plausible than a nonmodular one.

F will have none of this. He's right, of course, that the claims of evolutionary psychology are often overstated. Evolutionary considerations should not be the be-all and end-all, the primary criterion for judging theories of the mind/brain. But he's wrong to dismiss them as summarily as he does: they are at least one more criterion among many. F is also right to point out that a full evolutionary story would depend on knowing how genes determine the structure of the brain and how the structure of the brain determines the structure of the thoughts it can think—issues on which we are at the moment thoroughly ignorant. But in fact ALL claims about innate knowledge, including F's input modules, depend ultimately on the solution of these huge important problems, so it is unfair of F to single out evolutionary psychology on this score.

But he goes much further. The following quotes, I think, reveal his basic attitude.

> . . . it simply isn't true that all the sciences are mutually relevant. . . . If, for example, somebody could show how the theory of lunar geography constrains the theory of cellular mitosis, I'm sure he could get it published in *Nature*. But, in all likelihood, there will be no such publication because there is no such constraint, so don't hold your breath. Why should not these banal methodological truisms apply likewise to the relation between theories of evolution and theories of cognition? (83)
>
> . . . the complexity of our minds, or of our behavior, is simply irrelevant to the question of whether our cognitive architecture evolved under selection pressure. I do think it's remarkable that nobody seems to have noticed this. (87)

What's really remarkable is that F tacitly sets aside the vast literature on animal behavior and animal communication, going all the way back to Darwin (1872/1998). Nowadays it is widely accepted that an account of complex innate behavioral repertoires in nonhumans must include a discussion of their efficacy for survival and reproduction and hence of the role of natural selection in their genesis. F is right that nothing in principle FORCES us to this view, but does he think it's misguided? Or does he think that suddenly when we get to humans all bets about evolution are off? Or is it just that he doesn't want to know?

I am inclined to attribute to F the last of these stances because he takes a parallel stance on the problem of abduction:

> I can think of several research strategies that one might consider in this impasse:
> *(i) Do nothing about abduction; wait till someone has a good idea.*
> This is, no doubt, the rational strategy, quite possibly the only one that will work . . .
> [(ii) and (iii): Hope that the problem will be solved eventually by traditional techniques] . . .
> *(iv) For the time, concentrate one's research efforts in those areas of cognitive processing where the effects of globality are minimal . . .*
> I do urge you to consider the fourth option seriously. . . . For one thing, it is compatible with the simultaneous and rigorous pursuit of option (i). (52–53)

In short, he doesn't seem interested in solving the problem; he'd rather do his best to ignore it and advises us to do likewise.

Just for fun, for a moment let's not ignore it. Consider a situation that all academics face. You never know: It might just be the case that some seventeenth-century manuscript hidden away in a Bolivian monastery has important consequences for my research on binary branching in syntax. Will I look? Of course not. More true to life, I have known for twenty years that I should read Kant before purporting to speak authoritatively on innate ideas of space. But have I read him? Well, ya know, there's not enough time to invest in really studying Kant when there are papers to grade and concerts to go to and the rest of life to enjoy. We make decisions like this all the time, consciously or tacitly. There comes a point where we just say, 'Enough'. ('If this were a work of scholarship, and if I were a scholar, I'd try to make some sort of case for these historical claims; but it's not, and I'm not, so I won't' p. 14.)

How do we know when the appropriate point has arrived? We don't. The perfect reference might be waiting out there at the next foray into the library or the Internet. But eventually the practical need (or the overwhelming desire) to do other things balances the importance of getting the task right—the balance is different for everyone—and we just stop, perhaps with regrets, perhaps not. All the quotes above show F himself choosing to stop.

My suspicion is that inside the mind/brain, deep down in the unconscious processes of thought, things work pretty much the same way. After all, the brain evolved (there's that dirty word again!) to solve problems under constraints of time and energy. That bit of motion I detect over there may be something to eat or something that can eat me or someone to mate with. I don't have all day to make up my mind; if I don't make a move soon I may miss out on something important. Or put in evolutionary terms: Those organisms that managed, on average, to make successful moves on the basis of good and rapid guesses edged out those that didn't. As Dennett 1984 observes, there is no way any device that internally represents the world can take everything into account. But to be an organism in the world one must act. Therefore, any organism that internally represents alternative potential courses of action must in principle have some internal mechanism that calls a halt to further cogitation and selects one. Dennett provocatively suggests that the operation of this mechanism presents itself to human experience as the exercise of free will.

Now it's true that we don't yet have any handle on a computational solution to abduction. And we can of course do as F advises: call a halt to asking how we manage to call a halt to asking. Or we can choose to go on thinking about the problem. It's a choice. It may be that we don't yet have the tools to think about the problem productively. But after all, you don't know whether you have the right tools until something works. In the sixteenth century they didn't have the right tools for thinking about the mechanism of human reproduction, but did that stop people from thinking about it? Should it have? What about in the 1870s? The 1940s? It is impossible to know which of these situations parallels where we are now with abduction. In any event, it would seem to me more constructive to take F's observations about abduction not as a wet blanket but rather as a challenge to find better tools.

Turning from abduction to modularity: I have discussed F's modularity thesis and proposed an alternative version elsewhere (Jackendoff 1987, Ch. 12; Jackendoff 2000; Jackendoff 2002, Ch. 7); I have little to add here. However, I cannot resist mentioning three surprising lapses about modularity in the present book. First consider the following quote:

> A module *sans phrase* is an informationally encapsulated cognitive mechanism, and is presumed innate barring explicit notice to the contrary. A ''Chomskian module'' is an innate database. (58)

Now I take it that language perception is still F's best example of a module; it is the only one discussed seriously in anything of his I am familiar with. But the language perception module per se is NOT innate. Its fast and involuntary operation depends on knowing the rules of one's language, which are acquired. What IS innate is the 'Chomskian module' of universal grammar, the procedure for ACQUIRING the rules of grammar and thus being able to fire up the innate parts of the language processor. Thus the language perception module is not a 'module *sans phrase*', in which case F has given us no examples of such modules. One can, but he hasn't.

A more interesting error is F's 'A priori argument against massive modularity: The input problem' (71–78). The argument runs as follows: How does the mind determine which of a number of 'dumb' domain-specific modules is appropriate for processing a particular input? F argues that such a decision requires a mental process that is 'smarter', i.e. less domain-specific, than the relevant modules themselves. Therefore, he argues, the mind cannot be massively modular since one needs nonmodular mechanisms to feed the modules.

This argument arises, I suspect, from thinking a priori and not looking at actual cases. Consider the language perception module. Its input is some stage of auditory processing, where there is at least a frequency analysis of the auditory signal. But the very same signal is used simultaneously as the input to at least two other modules, the ones for voice recognition and affect (tone-of-voice) perception, which are localized in different brain regions and susceptible to independent impairment through brain damage (Etcoff 1986, 1989). The interface for language perception (I am specific about what 'interface' means in Jackendoff 2002) is sensitive primarily to things like formant transitions; it ignores overall fundamental frequency range and timbre. Any auditory signal with reasonably appropriate formant transitions is picked up by the language module; that's why we are fooled by parrots and crudely synthesized speech. By contrast, the interfaces for voice and affect recognition are sensitive primarily to fundamental frequency range and timbre and ignore (or average out) formant transitions. Any auditory signal with appropriately modulated timbre is picked up by the affect module; that's why we are fooled by sighing violins. In short, there need be no 'übermodule' that sorts auditory inputs out in advance. Nor could there be, since all three modules use the same input in different ways, typically yielding simultaneous judgments of who is speaking, what they are saying, and with what emotional tone.

F does not address this case. The one case he does address is Cosmides and Tooby's proposed 'cheater detection module'. My own sense is that cheater detection is too small to be a module in the same sense as language perception; it is more like a good trick within a larger module concerned with social cognition—it is on the order of the default principle in language that agents are expressed as subjects. But F's argument applies to a larger social module as well (I omit long stretches of repetitive snide remarks):

> We've been assuming that there's something in the input to the [cheater detection module] that turns it on; some property of its input representations to which it is selectively sensitive and that ''carries the information'' that the current distal array constitutes a social exchange. Question: What feature could this be? How does a module decide whether what it's looking at *is* a social exchange? . . . Nobody has *any idea* what kind of cerebration is required for figuring out which distal stimulations are social exchanges. . . . (75–76)

Actually, I think we DO have some idea: it's a social exchange whenever PERSONS are involved. And there are a lot of good tricks for detecting persons, among them face recognition (which is a good trick since we are fooled by 'happy face' buttons), use of language, and perceptual mechanisms for picking up 'animate' motion (which can be fooled by something as simple as abstract cartoon shapes moving in the right manner [Heider & Simmel 1944, Bloom & Veres 1999]). F is correct in seeing an a priori problem, but his solution in terms of a nonmodular input system is hardly a priori, and in fact incorrect. The right solution is that the interfaces that provide modules with their inputs are dumb too.

F's argument against heuristic solutions to abduction is similar: He asserts that one needs some 'smart' process to decide when it's appropriate to initiate a 'dumb' heuristic. I suspect the same solution applies: What makes a heuristic 'dumb' is in part that it is triggered by a good trick that works pretty well most of the time. This is the message I get from such work as

Gigerenzer 2000 and Kahneman, Slovic, and Tversky 1982, for instance; whatever the differences between their positions, they do seem to agree on this particular point. So F is again far off the mark.

A final point about modularity concerns its connection with evolution. F observes that natural selection could serve to 'instruct' modules that concern the structure of the world, such as vision:

> Suppose, as the experimental evidence rather suggests, that human infants are born believing that unsupported objects generally fall, and that the auditory location of a sound source generally predicts its visual location; and [3 more examples] . . . and so forth. . . . [I]f the beliefs in question are innate, the only instructional mechanism that's on offer as a candidate is natural selection. . . . Innate modules thus require a detailed epistemic fit between what's in the mind and what's in the world. Only a correspondingly detailed instruction of the mind by the world could possibly produce that, since we may assume with confidence that the world is prior to the mind. (93–94)

But then he goes on to one of the most astonishing arguments in the book: Language is exempt from this procedure of 'instruction' because

> . . . the facts that make a speaker/hearer's innate beliefs about the universals of language true (or false) *aren't* facts about the world; they're facts about the minds of the creature's conspecifics. . . . According to the usual Chomskian story, the conspecificity of speaker and hearer is what guarantees that what they innately believe about one another's language is true, and hence that their offspring (who are generally [sic!] also conspecifics) will be able to learn the language that they share. If that is so, then there is no particular need for what the language organ believes to have been shaped by natural selection. That's why Chomsky can (and, if I read him right, in fact does) hold both that human language is innate and modular, and that it is not an adaptation. My guess is that all these claims are true. (95–96)

He applies the same argument to the putative 'theory of mind' module: 'The same endowment that determines my innate theory of how your mind works also determines that your mind works the way that my innate theory of your mind says that it does' (97).

I don't know where to start unraveling the confusions in this argument. First of all, since when are the minds of conspecifics not part of the world? Second, couching the discourse in terms of a speaker's (and even the language organ's!) 'beliefs about the universals of language' has a strange flavor to which we return below. But the major problem has to do with a serious misunderstanding about how aspects of social interaction could come to be innate.

Consider first theory of mind. F takes it that our (and his Granny's) innate intuitions about how the mind works are basically a correct account of how the mind works; he has repeated this assertion endlessly over the past twenty years. But then, why do so many people have the intuition that language isn't innate, why do so many cultures believe in a soul, why is cognitive science so hard, why do we over-attribute mental lives to animals and computers, and above all, why do we make so many mistakes in assessing the thoughts, opinions, and motives of others on the basis of their behavior? More plausible to me is that the 'theory of mind' is another set of good tricks rather than a set of true beliefs. Like visual good tricks, these evolved to provide a more flexible and nuanced behavioral repertoire, this time in a social context.

In maintaining that these social capacities did not need to be shaped by natural selection, F offers no alternative account. It is as if they came into existence by miraculous virgin birth. As mentioned above, a long tradition tracing back to Darwin takes it that social capacities in other animals ARE products of natural selection. The difficulty, of course, is understanding how a relation based on mutuality could evolve: there are no preexisting 'facts about the world' that could 'instruct' such a capacity.

In part the difficulty lies in the metaphor. It is really misleading to think of 'the world instructing the organism', as if the world has the information and the organism takes it in. Rather, a more appropriate metaphor for the evolutionary process is an enduring community of conspecifics tinkering with its way of dealing with the world. Those tinkerings that happen to lead to relative reproductive success are perpetuated proportionately. Among the tinkerings are ways of dealing with conspecifics, including communicating with them. A glimmer of how communicative systems could have emerged through such tinkering is offered by the beginnings of research in computer modeling of the evolutionary emergence of communicative systems (e.g. Steels 1998;

Nowak et al. 2000). On another front, I have proposed (Jackendoff 1999; 2002, Ch. 8) an incremental scenario for evolution of the language faculty which permits the complexity of language to emerge gradually as the outcome of continued tinkering rather than in one or two giant miraculous leaps.

None of these considerations about the evolution of social capacities, of course, play any role in F's discussion. He really doesn't seem to show much interest in the details. But I have saved his most basic error for last, one that pervades the texture of this book and everything F has written over the past 25 years. Unlike most working cognitive scientists (and like most philosophers), F believes that thoughts are somehow connected directly to the outside world, even though the neurons that instantiate those thoughts clearly have no direct access to the world. That is, thoughts manage to be ABOUT the world because of the mystical relation of INTENTIONALITY. Furthermore, F thinks of everything in the mind as a propositional attitude; for example we have seen him speaking of a child having 'innate beliefs about the universals of language'. These two aspects of the language of thought hypothesis have generally been overlooked in the lore, but they are central to F's thinking.

I have discussed in detail elsewhere why I think this view is a profound mistake (Jackendoff 1983, Ch. 2; 1987, Ch. 7; 1992, Ch. 8; 2002, Ch. 10) and why I think much of F's despair about cognitive science arises from his insistence on this view. Here I will confine myself to a couple of remarks to suggest how far his approach is out of touch with actual research practice. Consider a rule of grammar, for instance the rule that unstressed vowels are reduced in English. As a linguist, I am acquainted with this rule as a proposition. But for a nonlinguist speaker of English, the rule hardly has that status. Saying that an English speaker knows (or believes) that unstressed vowels are reduced is like saying that someone who can sing *'Happy Birthday'* knows (or believes) that the fifth note is a perfect fourth above the first note: it's a referentially opaque claim about regularities in the individual's mental organization, as evidenced by his or her behavior. That is, knowledge of language is NOT a collection of propositional attitudes—which is why Chomsky occasionally suggests substituting another term like 'cognize' for 'know'.

Moreover, linguists in the Chomskian tradition do not conceive of language universals as something a child has knowledge ABOUT, as though they are something out there in the world (or other people's minds). Rather, observed language universals are an external reflection of the possibilities provided by UG in the human brain. This completely inverts F's logic about the relation between the mind and the world.

I believe it fair to say that people working in, say, vision, take a similar attitude: they are concerned not with how we get a TRUE picture of the world, but with how our brains use cues on the retina to construct our visual experience, a representation of the world that evolution has found to work well enough in our ordinary environment. Such an approach to vision has a pedigree dating back at least to the gestalt psychologists of the early twentieth century (not to mention Kant), but F heaps scorn on it:

> . . . it's a widely advertised piece of neo-Darwinist anti-intellectualism (see, e.g. Patricia Churchland 1987) that ''looked at from an evolutionary point of view, the principal function of nervous systems is to get the body parts where they should be in order that the organism may survive. . . . Truth, whatever that is, definitely takes the hindmost.'' . . . To repeat: there is nothing in the ''evolutionary,'' or the ''biological,'' or the ''scientific'' worldview that shows, or *even suggests*, that the proper function of cognition is other than the fixation of true beliefs. (68)

But there's plenty in all these worldviews that contradicts F, and it's far from anti-intellectual. It's just anti-traditional Anglo-American philosophy of mind. F, steeped (or mired) in that tradition, is simply unable to see beyond it.

Now it is indeed a serious challenge to extend this 'engineering' stance from perception (including language perception) to thought. But I think it's not insuperable. Dennett 1991 offers one approach; Jackendoff 1983; Jackendoff 1987; and Jackendoff 2002, Ch. 10 offer another, not incompatible one.

I do think F is right that there's a problem about abduction. And I agree that some of the claims of evolutionary psychology have been overstated. And I do think he's right that there are

aspects of mental function that are not best captured by local computation. One candidate might be attention, which seems to involve a global, cross-modality apportioning of resources (and which F never, to my knowledge, has discussed). I can also understand F's feelings of gloom and doom; I have my own with respect to large swaths of linguistics (see Jackendoff 2002). But F's dialectic here is so sloppy, so empty of actual examples, and so loaded with polemic that on balance it's hard to consider this book a useful contribution to ongoing discourse.

REFERENCES

BLOOM, PAUL, and CSABA VERES. 1999. The perceived intentionality of groups. Cognition 71. B1–B9.

CHURCHLAND, PATRICIA. 1987. Epistemology in the age of neuroscience. Journal of Philosophy 84. 544–55.

COSMIDES, LEDA, and JOHN TOOBY. 1992. Cognitive adaptations for social exchange. The adapted mind, ed. by John Barkow, Leda Cosmides, and John Tooby, 163–228. New York: Oxford University Press.

DARWIN, CHARLES. 1872/1998. The expression of the emotions in man and animals, 3rd edn. New York: Oxford University Press.

DENNETT, DANIEL C. 1984. Elbow room: The varieties of free will worth wanting. Cambridge, MA: MIT Press.

——. 1991. Consciousness explained. New York: Little, Brown.

ETCOFF, NANCY. 1986. The neuropsychology of emotional expression. Advances in clinical neuropsychology, vol. 3, ed. by G. Goldstein and R. E. Tartar. New York: Plenum.

——. 1989. Asymmetries in recognition of emotion. Handbook of neuropsychology, vol. 3, ed. by F. Boller and J. Grafman, 363–82. New York: Elsevier Science Publishers.

FODOR, JERRY. 1975. The language of thought. Cambridge, MA: Harvard University Press.

——. 1983. The modularity of mind. Cambridge, MA: MIT Press.

GIGERENZER, GERD. 2000. Adaptive thinking: Rationality in the real world. New York: Oxford University Press.

HEIDER, FRITZ, and MARIANNE SIMMEL. 1944. An experimental study of apparent behavior. American Journal of Psychology 57. 243–49.

JACKENDOFF, RAY. 1983. Semantics and cognition. Cambridge, MA: MIT Press.

——. 1987. Consciousness and the computational mind. Cambridge, MA: MIT Press.

——. 1992. Languages of the mind. Cambridge, MA: MIT Press.

——. 1999. Possible stages in the evolution of the language capacity. Trends in Cognitive Sciences 3. 272–9.

——. 2000. Fodorian modularity and representational modularity. Language and the brain, ed. by Y. Grodzinsky, L. Shapiro, and D. Swinney, 3–30. San Diego: Academic Press.

——. 2002. Foundations of language. Oxford: Oxford University Press.

KAHNEMAN, DANIEL; PAUL SLOVIC; and AMOS TVERSKY (eds.) 1982. Judgment under uncertainty: Heuristics and biases. Cambridge: Cambridge University Press.

MARCUS, GARY. 2001. The algebraic mind. Cambridge, MA: MIT Press.

NOWAK, MARTIN A.; JOSHUA B. PLOTKIN; and VINCENT A. A. JANSEN. 2000. The evolution of syntactic communication. Nature 404.495–98.

PINKER, STEVEN. 1997. How the mind works. New York: Norton.

STEELS, LUC. 1998. Synthesizing the origins of language and meaning using coevolution, self-organization, and level formation. Approaches to the evolution of language, ed. by J. Hurford, M. Studdert-Kennedy, and C. Knight, 384–404. Cambridge: Cambridge University Press.

Volen Center for Complex Systems
Brandeis University
Waltham, MA 02454
[jackendoff@brandeis.edu]

Where mathematics comes from: How the embodied mind brings mathematics into being. By GEORGE LAKOFF and RAFAEL E. NÚÑEZ. New York: Basic Books, 2000. Pp. xvii, 492.

Reviewed by D. TERENCE LANGENDOEN, *University of Arizona*

This is an extremely ambitious book. Its goal is to launch the discipline of 'mathematical idea analysis', a branch of cognitive science devoted to the understanding of the concepts of mathemat-