Available online at www.sciencedirect.com
ScienceDirect
Artificial Intelligence 171 (2007) 1110–1113
Artificial Intelligence
www.elsevier.com/locate/artint

Book review

Douglas Hofstadter, I am a Strange Loop, Basic Books, New York, 2007.

Marvin Minsky, The Emotion Machine, Simon & Schuster, New York, 2006.

Instead of a review

Daniel Dennett

Center for Cognitive Studies, Tufts University, Medford, MA 02155-7059, USA

Available online 10 October 2007

I am much too close, personally and intellectually, to Doug Hofstadter and Marvin Minsky to write a proper academic review of these wonderful books, so this will be an appreciation, an avowedly partisan paean to both the projects and the methods of these two original and bold thinkers. Such a celebration might seem unnecessary, just one more bouquet tossed onto the stage by an ardent fan, if it weren't for the fact that in spite of—or perhaps because of—their fame, many researchers who ought to know better ignore or dismiss them, and this is a golden opportunity for me to try to change their minds. I travel in several quite different academic circles and I find that each gang has its particular way of not taking these thinkers seriously. The neuroscientists deplore the absence of rigorous experiments and the refusal of both Hofstadter and Minsky to canvass the relevant experimental literature thoroughly and explicitly. *Where are the data?* The philosophers of mind, at the other extreme, find few formal arguments and a frustratingly cavalier refusal by the authors to define their terms at the outset. *Where are the proofs?* The cognitive psychologists, in the middle, are offended by the fact that neither Hofstadter nor Minsky sees the need to adjudicate between all the competing models and theories that have been painstakingly developed and defended, and instead offer their own impressionistic and oversimplified sketches. *Where are the models that make testable predictions?* The artificial intelligence crowd wants to see a running demo program. *Where is the code?* It's all just *speculation!* And then there is the style, too clever and playful, apparently written for bright high school students, not professors and graduate students. How can any self-respecting academic researcher admit to having learned anything from these crowd-pleasers? This isn't research, it's *edutainment!*

The books amply support these charges, but they aren't bugs, they're features. Hofstadter and Minsky have both developed models and theories that should satisfy the most hard-bitten empirical scientist, but now, enlightened by those projects, they are operating in a domain that many researchers find unbearably anarchic; nobody knows the rules. But that doesn't mean we should shun this lawless territory. And their super-accessible styles are integral to the task; these authors aren't out to impress anybody; they're out to explain, using whatever might work. (And finally, isn't Dennett just biased because of the heady praise both authors give his own work? That's a distinct possibility, but I will try to overcome that suspicion by explaining just what I think is valuable, indeed necessary, in the contributions both Hofstadter and Minsky are making to the daunting task of explaining how the human brain makes the human mind.)

E-mail address: ddennett@tufts.edu.

doi:10.1016/j.artint.2007.10.001

The difficulty with human minds is that they are much too familiar to us. We already "know all about them" and are effortlessly fluent in describing and explaining how they operate *using the terms of folk psychology that we learned as children*. We "see" and "hear" and "think"; we "make up our minds" and "decide" and "remember" and "intend" and "imagine" and "dream" and "wonder" and "hope"; we are beset by "pains" and "urges" and "anxiety" and "boredom." But when we ask ourselves how we do all these things, we usually go around in circles. For instance, as Minsky notes:

> . . . what is a goal, and how can we *have* one? If you try to answer such questions in everyday words like "*a goal is a thing that one wants to achieve,"* you will find yourself in circles because, then, you must ask what *wanting* is—and then you find that you're trying to describe this in terms of other words like *motive*, *desire, purpose, aim, hope, aspire, yearn,* and *crave.*
>
> More generally, you get caught in this trap whenever you try to describe a state of mind in terms of other psychology words because these never lead to talking about the underlying machinery (p. 187).

Now it is possible to spend one's whole academic career exploring the epicycles of these circles. Many philosophers of mind have done just that, filling hundreds of books with carefully analyzed accounts of the implications and assumptions of these everyday "psychology words" and never even trying to ask what kind of underlying machinery there might be and how we manage to catalogue its products and actions so deftly without having a clue what we're talking about. Some of them haven't asked this question because they firmly believe that there *is no* machinery; it's all done by some kind of utterly non-mechanical mind-stuff. Or at least: the mental realm and the physical or causal realm are "incommensurable, " or there is some sort of "radical emergence" that makes it utterly impossible to go back and forth between mind talk and neuron talk. Hofstadter has a typically clever (and tongue-in-cheek) way of exposing and resisting this defeatism. He coins the terms *thinkodynamics* and *statistical mentalics*. The idea is that like thermodynamics, thinkodynamics is conducted in terms of large-scale structures and patterns, while like statistical mechanics, statistical mentalics shows how this all "reduces" to the kazillion little steps at the lower level (p. 34). There *is* a sort of "emergence" here, but it is not the sort that mysterians like to imagine; it is like the emergence of traffic jams and Conway Life patterns and, well, the temperature of a gas: don't try to define the macro-level phenomena as *specific* complexes of micro-phenomena (the way we can define an atom of some element as a satisfyingly specific organization of sub-atomic particles). The macro-level phenomena are statistically robust patterns that permit vast variation in their micro-elements. And, one has to do a little squinting to see these patterns. They are only approximately tracked by the micro-machinery. All this is music to my ears and it is, really, philosophy.[1] As it stands it isn't an empirical theory of anything, and it doesn't prove anything, but it does *open up the idea of possibilities that otherwise might be ignored or underestimated.* It is a bit of conceptual calisthenics, exercising our imaginations to rethink the issues and cast off some of the assumptions that have been hobbling us.

Some may think that they don't need any such calisthenics. They already have a firm, clear grasp on the nature of the phenomena they are trying to explain. Perhaps they are right, but then where are their theories or models? We're drowning in data, but nobody seems to know how to refine it into evidence for any comprehensive theory. So people work on truncated issues, leaving till later the daunting task of merging their local visions into a single non-contradictory model. It's a defensible research strategy (but see Allen Newell's classic article "You can't play Twenty Questions with nature and win" [5]). People shouldn't try to do what they are no good at doing, and it is also true that this genre of free-wheeling "informed speculation" can be an utter waste of time in the hands of somebody who isn't well-informed and hasn't got a disciplined imagination. Very few people are equipped to do this kind of work well, but both Hofstadter and Minsky have had years of practice using the best available thinking tools. They haven't just marinated their minds in Lisp and other computer languages; they have built and dismantled computer models of a wide variety of phenomena, and sympathetically explored the efforts of others. Among the treasures of Minsky's book are the morals he draws from the pioneering work in GOFAI (Good Old Fashioned AI, Haugeland [4]). Although those avenues are now widely branded as having been disproven and discredited, Minsky extracts the *conceptual* fruits they yielded in spite of their now almost comical simplicity. Yes, they were naive, and sometimes offensively hubristic, but they also illuminated otherwise murky corners of the mind, and Minsky does a fine job of saving these insights for another generation of theory-mongers. Minsky's evolving model-sketch of the mind now is cast in term of a gaggle

[1] See my "Real patterns," [3], for some different paths to the same happy territory.