ELSEVIER
Artificial Intelligence 106 (1998) 353–358
Artificial Intelligence

Book Review

B.C. Smith's, *On the Origin of Objects* *

R.P. Loui [1]
Box 1045, Washington University, St. Louis, MO 63130, USA

Most who read this review will want to know what Brian Cantwell Smith's book has to say. But that would miss Smith's point. Much of this book's contribution is its style. Furthermore, his book is relevant to the AI audience as much for its style as for its ideas: it shows how AI researchers are peculiarly capable of defining a philosophical discussion. Thus, this review is not just about what *On the Origin of Objects* says, but also about how it says it.

Brian Smith apparently realizes three important things:

(1) there has never been a philosophy of computing,
(2) now is the time for such a thing, and
(3) the philosophers of computing will likely be the broadly reflective scholars who have been running from artificial intelligence (ever since the field has been contracting intellectually to meet the technical standards imposed from other fields).

Brian Smith is not alone in his three-fold realization about the impending birth of computing philosophy and the parturience of AI in this respect. Brian Smith is, however, the first to the punch, and *On the Origin of Objects* is his challenge for all future philosophical prose.

On the Origin of Objects is thus an important book, even a beautiful book. It reasserts its author as one of the deepest and erudite thinkers of computing. It is also, to this reviewer, an intellectually uninteresting book and thoroughly frustrating to read. These are two separate points: First, the book is a meditation on some metaphysical questions posed by symbol systems, and the author admits repeatedly that this is a purely metaphysical exercise (I would have demanded an apology rather than an admission). The problem is that metaphysics is a love-it or hate-it area of philosophy that has no practical implications (unlike, for example, epistemology which is crucially involved in explicating formal criteria for knowledge and belief). Second, the book frustrates this reviewer because the author, as is his reputation from prior work, is incapable of getting to the point. It is not just that Brian Smith seems to lean toward the semi-literary style of, for example, Friedrich Nietzsche. It is probably due to a mismatch of interests between writer and reader

* MIT Press, Cambridge, MA, 1996. $37.50 (cloth). $17.50 (paper). 440 pages. ISBN 0-262-69209-0.

[1] E-mail: loui@cs.wustl.edu.

PII: S0004-3702(98)00102-7

(a mismatch I suspect Brian Smith will have with most readers in AI). Both of these points are developed below.

What is the idea? Basically, Brian Smith believes that there is one "real world out there" (his "metaphysical monism") while there is an apparent arbitrariness in how we symbolize it (our "ontological pluralism", p. 375). He feels that this mismatch of monism and pluralism is an intolerable situation which must be remedied. First he rejects situation semantics (which might be associated with recent Stanford philosophers of language who puzzled in the last decade over how pronoun reference depends on context). Then he proposes something intermediate between the world and any theory-of-the-world.

> ... material objects ... occupy what we might call a middle ground, halfway between the physical world, and the intentional world. The [result is a] median nature of materiality ... (p. 361)
> ... objects must be held by subjects in a middle distance, not too close and not too far away ... (p. 366)

At this new level, there is "registration". It is a level at which reference to the world can be made, but not through linguistic commitment or through symbols that have intentionality (reference). It is the level that philosophers of science might call pre-theoretic: there is an observer, but there are not yet data, since data are not theory-neutral. This level is made possible by the embedding of an observer in the "real world", and of course, Brian Smith aims to permit that observer to be an AI program as well as a biological system.

> ... we are expanding our registrational capacities—building instruments and other devices that mediate our full participation in the world. (p. 360)
> ... it is not just the ontology of computation that is at stake; it is the nature of ontology itself. (p. 42)

Making sense of this level "of registration" is the central problem that occupies the author for several hundred pages (roughly the second half of the book). Along the way, gems are dropped that reflect an incredibly rich understanding of computing. I consider these gems to be more valuable than the details of the main line of Smith's investigation (and I inventory some of them below). If Brian Smith were not paralyzed by metaphysical problems, he could deliver the seminal work on the philosophy of computing just by collecting and displaying these gems of his insight.

As it stands, however, the work is more about representation than it is about computing. The author says as much, somewhat improperly attributing this emphasis to his background in AI (p. 14n., improper because some people in AI work not only on knowledge representation, but also on reasoning). Really, the work is a long rejoinder to John Searle, Fred Dretske, Jerry Fodor, Hilary Putnam, and the other philosophers who bankrupted the philosophy of mind with their dubious work on representation in the past two decades. I am reminded of a response by Putnam to Fodor in the *Notre Dame Journal of Symbolic Logic* in which an exasperated Putnam finally admits that Fodor's solipsist linguistic universe is nonsense even to Putnam. Be assured that every word of *On the Origin of Objects* makes

sense, though I would not argue that the effort at making sense of it is always generously rewarded. Although *On the Origin of Objects* is an improvement over this overexposed mainstream philosophical work on representation, it is nevertheless mired in the same questions.

Epistemology, the philosophy of logic, and the philosophy of science are the main branches of philosophy that are useful to artificial intelligence (soon, too, the philosophy of social organization and the philosophies of economics and law). The philosophy of mind and the philosophy of language have meanwhile drifted toward metaphysics, where one finds topics as diverse as "reality", "God", "causation", "identity", "truth" and "time". Most philosophically facile AI researchers are part of the extended diaspora of Viennese logical positivists and logical empiricists and have little patience for metaphysics. They understand the arbitrariness of models; they interpret scientific theory instrumentally; they appreciate Daniel Dennett's pluralism of stances (if the intentional stance is worth taking, then take it; otherwise, take some other stance); and they can appreciate that their work is rampant with construction and convention whenever they take the time to give constructivism and conventionalism fair hearing.

Brian Smith's fetish for metaphysics is something he shares more with observers of AI than practitioners of AI. He makes no effort to convince the reader that metaphysical puzzles are worth the time. This is an especially hopeful act in an age when scholars are constantly suspicious of whether a question is properly posed, and when scholars have had thick experiences with the muddles that result from badly posed questions. One has to admit that an effort to divine our "true" relation to the "real world" smacks of an anachronistic or provincial conservative foundationalism in an age of liberally constructed and virtual realities.

> ... in fact there may not be any compelling reason to believe there is even a metaphysical fact of the matter ... (p. 55)
>
> No wonder the religious right are on the rise. It can sometimes seem as if they, alone, recognize the power and depth of the unsatisfied longing [for anchoring one's metaphysics]. (p. 92)

The principal beneficiaries of this essay in AI will be those who situate robots, then make some kind of strong AI claim. These are the "designers [who] despair of assigning content" (p. 64). Thus, Rodney Brooks is the only AI researcher to figure significantly in the book. Thus, too, an MIT AI Lab colleague marvels over the book, suggesting that Brian Smith has glimpsed such magnitude that it has rendered him mute.

This brings me to the point about style. Most who love the book will be enamored because of Brian Smith's writing. It shows that a computer scientist can not only think to the highest standards held by the humanities, but can also stand prose-to-prose with them. On the one hand, Brian Smith can turn a phrase as well as celebrated Chilean author Isabel Allende, and is equally confident deploying the well-worn phrase judiciously: "a path between the Scylla of naive realism and the Charybdis of pure constructivism", (p. 3); "to combine decisive theoretical bite with exquisite resolution", (p. 11); "logic, truth, and mathematics ... [are] not illicit steroids taken on the sly", (p. 93); "this profusion of practice has a certain plucky integrity" (p. 361); "... the limit case of pure encounter is

utterly ineffable" (p. 367). On the other hand, too many things are "elided", "exquisite", "ineffable", "inexorable", "in a fog", or "plucky" to Brian Smith. The text is infected with the poetic nonsense of the sixties:

(1) overreaching nonparallelism for its own sake "to do justice to the tundra, to gardening, to politics, to rock" (p. ix), "the moral, the erotic, the political, the artistic, and the sheerly obstreperous", (p. 3 and again on p. 361, similarly. Recall that Pablo Neruda, in his Nobel-poet address, castigated that decade's poets for disingenuous profundity of this kind); and

(2) attraction to monosyllabically cool phrases like "flex" and "slop", "zest" and "spunk" (though these are infinitely preferable to Putnam's neologism, "incommensurability").

One wishes Brian Smith could just chain substantive claims. Even an aphorism reveals its content better than this author's grand italicized claims. These claims are breathtaking but unconnected and rare. They are like random footholds in the rock he asks us to climb with him:

> Computation is not a subject matter. (p. 73)
> The idea that metaphysics must be univocal, perspective-independent, from-nowhere [nodding to Tom Nagel], value-free, and the like, is itself a relic of the old (predecessor) metaphysics; it is not an inherent truth. (p. 89)
> The laws of physics are fundamentally deictic [context-dependent in their referring]. (p. 168)
> Physicists will have to look to a theory of intentionality for an account of the notion of individual, not the other way around. (p. 180)
> There are no physical objects (individuals), only material ones. (p. 183)
> Intentionality is a way of exploiting local freedom or slop in order to establish coordination with what is beyond effective reach. (p. 208)
> Ontology is the projection of registration onto the world. Representation is the projection of registration onto the subject or vehicle. (p. 349)

The author refers to his prose as an argument (pp. 22, 82). It is not, at least not in the sense of a derivation or line of reasoning. It is closer to a coredump of Brian Smith's psychology (complete with signposts of his fancies and depressions, p. 17). It is clearly a remembrance of a personal journey, and the author is an able (if ever-present) guide. Smith needn't hide so many of the important vistas along self-centered path-following. This important claim appears in a side-bar, for example:

> the metaphysical view to be advanced in this book (is) that reductive explanations are, in point of fact, false. ... They are metaphysically precluded. (p. 82)

What the prose really is, is "playfully gradualist dialectic" (p. 375) (though dialectic must be understood in the broadest sense here: more Karl Marx than Nicholas Rescher). And what the book really is, is best depicted on p. 315 where appears the reproduction of artwork owned by Adrian Cussins, a philosopher of mind. Like the book, the painting is non-representational. In fact, it is tortured in its relation to representation. Like *On the Origin of Objects*, it is energetic and honest in its attempt at artistry, but a bit of a scatter.

> One way to understand this book is as the metaphysics of a northerner—and as such much more aimed at rock than at rocks. (p. 101n.)

To the author's credit, many of the claims are precise enough to be opposed: Surely object-oriented programming succeeds operating systems in its theorizing about the ontology of computing, and does not do so de novo (p. 44); What if the explicit representation $\langle 0, 0 \rangle$ were compiled away (p. 51)? Is Smith sure that formal manipulation is constitutively defined in terms of semantics (p. 72n.)? Is the world of falling interest rates part of metaphysical reality, or socially constructed reality (p. 97)? And so on. His more importantly unpalatable claims are undebatable, resting entirely on his aesthetics: that "something must settle the matter" of "intentional contents" (p. 62); that "there is no way to proceed ... developing a comprehensive theory of computation ..., except by taking on metaphysics and ontology directly" (p. 69); that "anti-metaphysical stances ... [unify] at unimaginable cost" (p. 92).

The gems that forebode a philosophy of computing are too many to enumerate. I agree with all of the following (except the one marked):

> ... process [is] arguably the most potent notion in all of computer science. (p. 32)
> Programs ... are treated as much as prescriptions as descriptions. (p. 36)
> To decide whether a given system is or is not an implementation of a standard program ... requires more sophisticated individuation criteria. (p. 47)
> An inscription error [is] a tendency for a theorist ... to read [ontological] assumptions or their consequences back off the system, as if that constituted an independent empirical discovery or theoretical result. (p. 50)
> ... I want to emphasize ... that computer science and AI are, by their own admission, largely engineering disciplines. (p. 66)
> Computers turn out in the end to be rather like cars: objects of inestimable social and political and economic and personal importance, but not the focus of enduring scientific or intellectual inquiry (p. 74). (This reviewer disagrees, having made the transportation science disanalogy in order to show that computing alters intellectual foundations genuinely.)
> ... the formal conception of "game" [underlies] so much AI, economic modeling, and structuralism. (p. 78)
> ... the constructive aspect of computer science ... is fully consonant with ... adopting a partially (dis)connected stance towards a partially (dis)connected subject matter. (p. 359)
> ... designers, playwrights, artists, ... are drawn into the act ... . Few fields, if any, are being left behind; ... it would be a mistake to think that these people are just users of computation. On the contrary, they are participating in its invention. ... The line between specifically computational expertise and general computational literacy is fading ... (p. 360)
> ... notions of mathematical proof [are] being revised ... . Other distinctions are collapsing, such as those between and among theories, models, simulations, implementations ... (p. 360)
> ... we are post-Newtonian, in the sense of being inappropriately wedded to a particular reductionism of scientism, inapplicable to so rich an intentional phenomenon.

> Another generation of scientists may be the last thing we need. Maybe, instead, we need a new generation of magicians. (p. 362)

In the end, Brian Cantwell Smith has written a book that is beautiful more for its art (and he has done a fabulous job with his own graphic design) than for its intellectual contribution. It does raise the standard of what will be considered the philosophy of computing. It has succeeded in titillating those philosophers who seek a more literary feeling. And it has catapulted Brian Smith, for a time (as long as the philosophy of computing is mistaken to be the modern philosophy of mind), ahead of all who would today claim to be computing's principal philosopher.