

# Mostly Harmless\*

— a review of *Action in Perception* (A. Noë, MIT Press, 2004) —

Shimon Edelman  
Department of Psychology  
232 Uris Hall, Cornell University  
Ithaca, NY 14853-7601, USA  
<http://kybele.psych.cornell.edu/~edelman>

June 14, 2005

Alva Noë's book [14] is, in equal proportions, a philosophical argument against the established construal of visual perception, and a promulgation of an alternative view, according to which

“[...] we ought to reject the idea — widespread both in philosophy and science — that perception is a process *in the brain* whereby the perceptual system constructs an *internal representation* of the world. [...] What perception is, however, is not a process in the brain, but a kind of skillful activity on the part of the animal as a whole” ([p.2]; all the page numbers refer to [14]).

Radical as this claim may seem to a casual reader, the actual alternative that emerges from integrating Noë's arguments with the state of the art in vision research is in fact rather tame, if not entirely mainstream (hence the title of my review). The main reason for this happy, if somewhat anticlimactic, conclusion is that the first part of Noë's thesis — that perception cannot be fully internal — is perfectly in tune with quite a few existing theories of cognition [3, 13, 6, 5, 7]. The unifying framework behind these theories is this: they all take heed of the injunction that the external world should bear the main burden of being its own representation [17, 15, 16].

Some of the many existing arguments in favor of this approach are computational: reconstructing the world in an internal memory space is infeasible, ultimately futile, and actually unnecessary [17, 6]. Others are psychological: people's performance in perceptual tasks suggests that no reconstruction is attempted [2, 6]. Excellent philosophical works compatible with the idea that the world is its own best representation are also available [13, 5]. Much of the present book is a useful contribution to this latter literature, made all the more valuable by the fine balance it strikes between the roles of internal and external components of representations (rather than rejecting the former entirely; cf. [6]).

What about the second part of Noë's thesis — that “to perceive you must be in possession of *sensorimotor bodily skill*” [p.11]? In my understanding, this statement (and many others similar to it in the book) unpacks into the following postulate: to perceive the world as a human does, you must reside in a human body, and, moreover, be in constant, full sensorimotor control of that body. We thus discover that there are actually two independent claims here: (i) the body must be human, because it is human perception we are interested in, and, (ii) active control of that body is required to support perception. Let us examine these in turn.

---

\*The final version of this article will appear in *Artificial Life*, published by The MIT Press.

Embodiment in a *human* form is clearly only an issue if it is specifically human-like perception that is at stake. This view hardly needs special justification: denying that animal brains, alien brains, or computers, are capable, respectively, of animal-, alien-, or a machine-like perception would amount to an untenable anthropocentrism [19], or worse (cf. *What do you mean... human?* [4]).<sup>1</sup> Indeed, Noë acknowledges that “to perceive *like us* [...] you must have a body like ours” [p.25; my emphasis].

Throughout the book, Noë works very hard to convince the reader that disembodied perception is inconceivable. Moreover, he argues, a perceiver must be not merely embodied, but actively controlling the body at all times; only such “enactive” perception counts as the real thing. Often, the argumentation involves gentle, yet firm and consistent name-calling, as in the many pronouncements that distinguish between real, or genuine, and, presumably, faux perception; here is one example: “genuine perceptual experience depends ... on our exercise of sensorimotor knowledge” [see pp.10-13]. In some cases, the semantics game seems to get ontologically risqué, as in the following passages, which, I confess, I fail to comprehend: “They [Kanizsa lines] are *perceptually* present without being actually perceived” [p.61]; “An ecological approach to color (as well as other appearances) treats it as *natural*, but as nonphysical” [p.154]. All this, and the occasional heavy dosage of philosophical talk (as in Chapter 3, which is all about “P-properties” of objects), will probably do little to persuade a skeptical reader to adopt the enactive perception doctrine.

My own skepticism is fed by the conviction that the dependence of the functioning of the mind on being in the world can be qualified in an important sense: once my mind, perceptual systems and all, is fully formed — say, so as to become functionally equivalent to that of an average adult — the world can be safely detached (at least temporarily), without destroying the mind.<sup>2</sup> If during such disconnection from the real world my optic nerve is artificially stimulated in a manner isomorphic to the spatiotemporal pattern of activity that would have been induced in it by a great white shark swimming in a deep blue sea, then by Zarquon I’ll see a shark. Moreover, my perception of the shark will be none the less “genuine” for being artificially induced than if I were to see the shark on my HDTV set, vegging out in front of which is a convenient homegrown substitute for a brain-in-a-vat experience.

Consistently with his original enactive thesis, Noë resists the notion that brains in vats can be true perceivers. This is supported by an interesting discussion of the role of causation in perception (based on a thought experiment from a 1961 paper by Grice). Alas, this discussion proceeds to appeal all too often to intuitions as to what constitutes “genuine” perception [p.170]. What about empirical evidence, over and above intuition, that can be brought to bear on this issue? Noë seems to believe that the only available empirical results are the phosphenes that Penfield in 1954 induced in his patients by direct brain stimulation: “At the present time, however, we are not able directly to generate more complicated experiences.” [p.211]. As a matter of fact, Newsome et al. [18] showed in 1990 that a minute electrical current injected in the proper place in a monkey’s brain induces it to perceive coherent visual motion where there is none. Since then, targeted microstimulation has been shown capable of inducing a variety of perceptual states, both in monkeys and, recently, in humans (e.g., [9]).

An important part of Noë’s arguments against the established conception of vision and in favor of the enactive framework rests on computational considerations. Sadly, here too the text is behind the times. On the one hand, Noë criticizes theories that consider vision as inverse optics [p.20] — an approach whose decline started nigh on twenty years ago; it was effectively abandoned by the computational vision community around the turn of the century [6]. On the other hand, in his eagerness to secure a place for vision outside the

---

<sup>1</sup>The author of that edifying article, who for many years edited the *Astounding* science fiction magazine, used to ask his authors to write him a story that would describe an alien that thinks *as well* as a human, but not *like* a human.

<sup>2</sup>Consider: if that were not the case, I would be able to commit suicide merely by going down into my basement and switching off the light.

brain, he wields a theoretical tool that is, from an explanatory standpoint, an empty slogan — J. J. Gibson’s idea of direct perception: “In actively encountering the way in which how things look varies with movement, we *directly* encounter how things are” [p.85]. As pointed out long ago [21, 10], the claim that perception is “direct” amounts to shirking the responsibility of actually explaining how it is that the relevant information afforded by the environment is isolated from irrelevant variation and acted upon.

Not all is lost, however. Gibson’s deep insights into the environmental cues that can support behavior can be salvaged if supplemented by a detailed computational analysis of the problems at hand, of the possible algorithmic solutions, and of their implementations [11]; Noë, on the very next page, intuitively just the right kind of computational approach — or at least that’s how I interpret his remark on misperception: “. . . one misperceives because one draws on the wrong sensorimotor skills and expectations” [p.86]. Veridical perception, thus, is what happens routinely, when one draws on the *right* skills and expectations. When translated into the parlance of computational vision, this amounts to espousing a statistical approach [1, 8], which is readily relatable to the framework outlined by Kevin O’Regan and by Noë himself in a recent paper [16]. When Noë returns to Gibson a few pages later [p.105], it is to echo the main message of that paper: “To perceive ... is to perceive structure in sensorimotor contingencies.” This is very true, but: the computational challenges — undertaken by the brain, not circumvented by any kind of “direct” magic — should be appreciated and understood. A classical example of one such challenge, the problem of cyclopean vision (how to avoid seeing double when looking through two eyes), is discussed on p.45. Noë quite mistakenly likens it to retinal image inversion. Unlike that non-problem, cyclopean vision has a real and inescapable computational component: finding correspondences between the projections in the two images of the same distal features; without some solution to the correspondence problem, the depth information cannot be retrieved [12].

For the students of perception, the reward for solving the computational problems of vision will be not early retirement, but a bigger — possibly much bigger — problem of the phenomenology of vision (in the example just mentioned, this is the problem of explaining the unitary nature of binocular vision). As Noë points out, “The obstacle to phenomenology is that the transparency of experience makes it seem puzzling how we can ever make experience itself the object of our inquiry” [p.72]. The desire to overcome this obstacle permeates the entire book, surfacing in a key formulation on p.164: “Looks are genuine, relational properties of things. But looks are not relations between things and your mind; they are relations between objects and the environment in which you find yourself a perceiver.” I cannot really relate to this statement of the nature of phenomenal perception, so instead of arguing with it, I would like to suggest an alternative, which is rooted in the old identity theory of the mind, and yet is compatible both with the latest neuroscience and the latest computational theories of vision: looks are points in multidimensional spaces, spanned by activities of neurons, which resonate to patterned sequences of world events spread over the space-time neighborhood of the perceiver’s brain, and are affected by remembered traces of such patterns;<sup>3</sup> cf. [6, 20, 7].

Having thus hedged the enactive perception thesis, let me propose a modification to it that would make it acceptable, I dare say, to a majority of my colleagues<sup>4</sup> in the cognitive sciences: it is the *development* of perception (specifically, of *human-like* perception) that requires embodiment in, and control over, a whole human body. A phrase found in several places in the book gets it almost right: “To perceive you must be in possession of *sensorimotor bodily skill*.” I would put it thus: to perceive, you must be in possession of sufficient prior experience of exercising such skill. Like many other important things in life, perception is, to a large extent, not so much about being as about having been.

---

<sup>3</sup>For perceivers devoid of any neurons, such as aliens or machines, substitute the appropriate elementary computational units.

<sup>4</sup>Excluding the die-hard innatist fringe, that is.

## References

- [1] H. B. Barlow. Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. *Vision Research*, 30:1561–1571, 1990.
- [2] S. J. Blackmore, G. Brelstaff, K. Nelson, and T. Troscianko. Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, 24:1075–1081, 1995.
- [3] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–160, 1991.
- [4] J. W. Campbell. What do you mean. . . human? *Astounding Science Fiction*, September 1959. Editorial.
- [5] A. Clark. *A theory of sentience*. Oxford University Press, Oxford, 2000.
- [6] S. Edelman. *Representation and recognition in vision*. MIT Press, Cambridge, MA, 1999.
- [7] S. Edelman. Constraining the neural representation of the visual world. *Trends in Cognitive Sciences*, 6:125–131, 2002.
- [8] D. Knill and W. Richards, editors. *Perception as Bayesian Inference*. Cambridge University Press, Cambridge, 1996.
- [9] H. W. Lee, S. B. Hong, D. W. Seo, W. S. Tae, and S. C. Hong. Mapping of functional organization in human visual cortex: Electrical cortical stimulation. *Neurology*, 54:849–854, 2000.
- [10] D. Marr. *Vision*. W. H. Freeman, San Francisco, CA, 1982.
- [11] D. Marr and T. Poggio. From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.*, 15:470–488, 1977.
- [12] D. Marr and T. Poggio. A computational theory of human stereo vision. *Proceedings of the Royal Society of London B*, 204:301–328, 1979.
- [13] R. Millikan. *White Queen Psychology and other essays for Alice*. MIT Press, Cambridge, MA, 1995.
- [14] A. Noë. *Action in Perception*. MIT Press, Cambridge, MA, 2004.
- [15] J. K. O'Regan. Solving the real mysteries of visual perception: The world as an outside memory. *Canadian J. of Psychology*, 46:461–488, 1992.
- [16] J. K. O'Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24:883–917, 2001.
- [17] W. Reitman, R. Nado, and B. Wilcox. Machine perception: what makes it so hard for computers to see? In C. Wade Savage, editor, *Perception and cognition: issues in the foundations of psychology*, volume IX of *Minnesota studies in the philosophy of science*, pages 65–87. University of Minnesota Press, Minneapolis, MN, 1978.
- [18] C. D. Salzman, K. H. Britten, and W. T. Newsome. Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177, 1990.

- [19] M. Scriven. The compleat robot: A prolegomena to androidology. In F. J. Crosson, editor, *Human and artificial intelligence*, pages 117–140. Appleton Century Crofts, New York, 1970.
- [20] J. J. C. Smart. The identity theory of mind. In E. N. Zalta, editor, *Stanford Encyclopedia of Philosophy*. Stanford University, 2001. URL <http://plato.stanford.edu/archives/spr2001/entries/mind-identity/>.
- [21] S. Ullman. Against direct perception. *Behavioral and Brain Sciences*, 3:373–416, 1980.