

A swan, a pike, and a crawfish walk into a bar*

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

June 23, 2008

The three commentaries of Van Orden, Spivey and Anderson, and Dietrich (with Markman's as a backdrop) form a tableau that reminds me of a fable by Ivan Andreevich Krylov (1769 - 1844), in which a swan, a pike, and a crawfish undertake jointly to move a cart laden with goods. What transpires then is not unexpected: the swan strives skyward, the pike pulls toward the river, and the crawfish scrambles backward. The call for papers for the present ecumenically minded special issue of JETAI was designed to minimize this kind of discord, by charging the authors to examine the possibility of epistemological pluralism in cognitive science — a field whose very diversity makes fundamental disagreement more likely than in other sciences. No doubt, the road mapped out by the editor had been conceived with good intentions in mind, but where did it lead us? It has been said that no good intention must go unpunished. To celebrate this venerable academic tradition (and also because I have a reputation to maintain), the following remarks will therefore be mostly other than conciliatory; caveat lector.

1 Nonlinearity, synergy, and all that jazz

In his opening paragraphs (as well as in the title of his commentary), Van Orden argues for pluralism by invoking the ideas of one of my favorite American philosophers, Charles Sanders Peirce. Van Orden's reading makes Peirce sound pluralist in a postmodern way: he proposes that "the philosophy itself recommends endless loops of self-questioning and recapitulation" and reiterates this sentiment in the closing remarks: "We naturally will remain significantly ignorant about cognition, as we do of every subject matter, which is the clear lesson of history and abduction." This would have been too bad, both for pluralism and for Peirce: in addition to being intellectually shallow, epistemological relativism is self-undermining, hence useless as a foundation for any science, pluralistic or not (as noted by Markman in this issue).¹

Luckily for Peirce's good standing in the philosophy of science (and by extension, for methodological pluralism), Van Orden gets Peirce wrong. As discussed repeatedly and at length by Umberto Eco, Peirce's stance was much more nuanced and indeed pragmatic than the stark relativism that some of his interpreters derive from a misunderstanding of the concept of abduction (Eco, 1979; Eco and Sebeok, 1988; see especially Eco's chapters in Eco, Rorty, Culler, and Brooke-Rose, 1992, where he exposes the futility of

*Response to commentaries on the target article, *On the Nature of Minds, or: Truth and Consequences*, this issue.

¹A concise and highly accessible discussion of the pitfalls of relativism can be found in the recent book *Walking the Tightrope of Reason* by Fogelin (2003).

“unlimited semiosis”²).

But why settle for second-hand opinions about Peirce, even from such sharp minds as Eco’s, when the original text is just a few mouse clicks away? Here’s Peirce himself, writing in *The Fixation of Belief*:

To satisfy our doubts, therefore, it is necessary that a method should be found by which our beliefs may be determined by nothing human, but by some external permanency – by something upon which our thinking has no effect. [...] Such is the method of science. Its fundamental hypothesis, restated in more familiar language, is this: There are Real things, whose characters are entirely independent of our opinions about them; those Reals affect our senses according to regular laws, and, though our sensations are as different as are our relations to the objects, yet, by taking advantage of the laws of perception, we can ascertain by reasoning how things really and truly are; and any man, if he have sufficient experience and he reason enough about it, will be led to the one True conclusion.

— Peirce (1877)

It is this kind of truth that the title of my target paper alluded to. Practitioners of extreme literary “theory” may still wring their hands about everything in the world being subject to unlimited and unconstrained interpretation, or, depending on one’s temperament, rejoice in this predicament of their own invention, but there is no reason these days for cognitive scientists to succumb to such lapse of reason. In particular, abduction, which Peirce identified as the sword that cuts through “endless loops of self-questioning” and which is now known technically as evidential reasoning, has been placed on a solid computational foundation within the Bayesian probabilistic framework (Johnson and Krems, 2001; Holyoak and Morrison, 2005; Nilsson, 2006).³

Of course, to engage with this level of understanding of cognitive computation (whether in organisms situated in the world or in scientists who study them), one needs to be in possession of the proper conceptual tools, namely, of the notions of levels of understanding and of computation — the two main topics of my target article. A prime characteristic of Van Orden’s stance is a repudiation of the need for these tools. All those who attempt it are, however, doomed to share the fate of the person who wanted to get rid of a boomerang by flinging it as far away as possible. How do you say “evidence” or “reasoning” (let alone “Bayesian”) in the language of dynamical systems? Ignoring all the levels but one, and the lowest one at that, makes your explanation less, not more, informative or insightful, and it certainly does not make the levels go away: they are out there *nolens volens* (as observed by Simon, 1973, whom I quoted in the target article). At the end of the day, you find yourself sitting in a puddle of nonlinear differential equations, wondering whether the buzz in your head has anything to do with the smooth, vaguely L-shaped object next to you.

Van Orden praises me for my clear stance on the issues of levels of understanding and of computation, yet the stance that he attributes to me and proceeds to refute (successfully) has nothing to do with the position that I staked out in my article. He writes, “If cognition is computation, then the algorithmic software description, at the rational level of analysis, can be made separately from the mechanistic description of hardware, consistent with Edelman’s (this issue) computational ontology.” In fact, the arguments that I

²In the study of signs, or semiotics, “semiosis” refers to the process of arriving at the meaning of a sign. In Eco’s parlance, semiosis becomes “unlimited” when this process is pursued beyond its reasonable conclusion, e.g., by forcing deeper and deeper “interpretations.”

³For a discussion of the importance for cognitive science of the other main sense of the term “abduction,” which can mean simply “abstraction,” see Edelman (2003).

employed to demonstrate that cognition is computation (in sections 1 and 2 of my paper) precede, and are independent of, my arguments in favor of a multi-level methodology (in section 3). Moreover, I made a point of noting explicitly that the levels of understanding constrain each other.⁴

Van Orden's notion of computation is so deeply idiosyncratic as to be unrecognizable to a computer scientist. Judging from expressions such as "feedforward and computational thinking" or "any algorithmic theory can be rewritten as staged feedforward chains of inputs (causes) and outputs (effects)," he seems to believe that computation must be somehow "feedforward." The truth of the matter is that this epithet is simply inapplicable to the two actual examples of computation that I offer in section 2.1 (it makes equally no sense to affirm or deny that a mechanical or electric resonator is "feedforward"). Moreover, any number of examples can be given of natural systems in which computation is carried out in a feedforward, feedback, sideways, and generally loopy manner, including feedback from, and loopiness that extends into, the system's environment (Brooks and Steels, 1995; Edelman, 2008).

Embodiment and situatedness in the environment are also precisely those properties of natural computation that erase another of Van Orden's baseless distinctions between computation and what he considers to be the right stuff: intentionality and "meaning."⁵ In drawing that distinction, Van Orden refers to Juarrero (1999), who "explains why computation fails to give a successful account [of] meaningful intentional behavior. The crux of the argument concerns whether a philosophical or scientific approach can successfully distinguish intentional from unintentional behavior — a wink from a blink, in her terms." Arguing against computation from this premise is like banging against a wall next to an open door. On the one hand, there are no grounds for any kind of objective reification of "meaning": intentionality is a stance that an observer may or may not take with respect to an observed system, although taking it sometimes helps explain the system's behavior (Dennett, 1987). On the other hand, an embodied, situated system that is subject to evolutionary pressures will be driven by those pressures to take the intentional stance with respect to another system when and where it "pays" to do so, on the average (Dennett, 1995), a point that I took care to highlight in the target paper (in the last paragraph of section 2).⁶

Thus, natural cognitive systems construct their own meaning as they see fit — computationally, of course, because processing and acting on information fed back to the system from its environment cannot help but involve computation. Without understanding the goals of these computations and the various state spaces — some continuous, others, for reasons that I listed in the target article, discrete — in which the computations unfold, there can be no understanding of cognition. Historically, merely extolling the complexity of the observed phenomena or calling them "nonlinear" or "synergetic" has amounted to taking refuge in empty jargon and has generated no explanatory insights into the nature of minds, nor any predictions except portents of more nonlinearities and synergy at every turn. If we are to take the notion of emergence seriously (as I believe we should), we must embrace the notion that explanatory structure emerges (and apparent complexity fades away) when multiple levels of understanding of computational processes interact with each other in a pluralistic epistemological framework.

⁴Here again is the relevant passage from my target article: "In terms of the levels of analysis framework, the contribution of connectionism was to highlight the multiple possible ways of implementing a given function, as well as the interdependence between levels, as when the characteristics of the available implementational substrate constrain the computations that it can perform, and therefore also the tasks that can be addressed (for an example from the domain of vision, see Edelman, 1999, ch.1)."

⁵The last of the cognitive sciences that clings to the imaginary distinction between syntax and semantics is formalist (Chomskyan) linguistics, and even there it is being rolled back under the onslaught of the new Empiricism (Goldsmith, 2007).

⁶As it happens, a lucid illustration of Dennett's idea of intentional stance can be found in this issue, in Spivey and Anderson's discussion of the Game of Life, which, as they point out, gives rise to a "powerful *illusion of meaningfulness*." Thus, Van Orden's hopes notwithstanding, no blink is intrinsically a wink.

2 Playing Bach his way

Spivey and Anderson (S&A) in their commentary argue in favor of an exclusively bottom-up, emergentist “complex dynamical systems approach” to cognitive science. Despite their claim of neutrality (S&A describe their agenda as “neither pluralist nor anti-pluralist”), their approach must be labeled as anti-pluralist because they explicitly reject the need for multiple levels of explanation. Indeed, S&A attempt to resuscitate the standard trope of extreme “connectionism”⁷: that what really matters in cognition — both epistemologically and ontologically — is supposedly *only* what happens on some kind of base level of cognitive computation.⁸

What is that base level? By S&A’s own logic, it should be what physicists tell us it is: the particles of the Standard Model, strings, twistors, loop quantum gravity, or whatever. S&A clearly do hear this call of the wild — they describe themselves as “encouraged (like experimental physicists) to get our laboratory measurements ‘closer’ to the original continuous trajectory, closer to the ‘ontic stream’” — but wisely stop short of heeding it (and thereby abandoning brain science in favor of physics): instead they identify the base level with “neural patterns.”

This choice (which I applaud; cf. Edelman, 2008, ch.3) can only be justified by observing that the level of neural firing affords a maximum in explanatory power relative to the levels above and below it. According to the pluralist position that I outlined in the target article, this explanatory sweet spot is real, but is only one among several along the hierarchy of levels. Crucially, none of the levels (certainly not the base level) is capable of explaining all by itself everything that needs to be explained.

By holding that the base-level explanations suffice, S&A assume an extreme reductionist stance, as when they claim that “human-to-human conversation can also be generally described solely at that lower level of organization, with the invocation of higher-level structures serving as nothing more than a descriptive shortcut, not as a necessary link in the causal chain.” Moreover, they explicitly object to explanations that invoke downward causation: “The risk of granting the symbolic level too much causal power to influence its own underlying subsymbolic level is that one can find oneself sliding down this slippery slope towards claiming that . . . being Jewish actually causes the DNA to replicate itself only through the mother’s family line!”

In response, I must point out that the possibility of abusing a theoretical move is not a valid reason for banning it altogether. To pick up S&A’s example, while being Jewish — which is a self-construal state brought about by a variety of personal and societal factors — does not make a person retroactively parthenogenetic or matrilineal, it does indeed *cause* things to happen on the level of DNA expression. For instance, neuron-level calcium current triggered by hearing a tune can cause immediate early gene (IEG) expression in neurons (Deisseroth, Mermelstein, Xia, and Tsien, 2003) — or not, depending for example on whether or not you *believe* you are Jewish and *perceive* the tune as that of *Hava Nagila* (say). Indeed, the importance of downward causation is not limited to cognition but rather is a general rule in biology, as noted by Noble (2008): “We need the insights obtained from higher-level analysis in order to succeed even at the lower levels. The reason is that higher levels in biological systems impose boundary conditions on the lower levels. Without understanding those conditions and their effects, we will be seriously restricted in understanding the logic of living systems.”

The existence of multiple, distinct levels on which cognitive computation must be taking place is dictated by the interpretability constraint in communication situations, which arise in any computational system

⁷S&A manage to keep a straight face while invoking this term, which I myself cannot imagine being used outside scare quotes, for reasons detailed in my target paper.

⁸In the way of clarification, here and in the target paper the term “epistemology” refers to ways of *knowing about* some things, while “ontology” refers to the way those things actually *are*.

that is distributed or partially compartmentalized, be it a human brain or a society of humans. Due to this constraint, which I identified in the target article, and which S&A disappointingly choose not to discuss, communicating entities must resort to discrete representations, on the pain of becoming bogged down in computational intractability. As I noted, these discrete representations may take the form of symbolic dynamics, but they cannot be understood solely in terms of the underlying continuous dynamics. In this sense, emergent symbols are real (just as the Jewishness construct is), that is, ontic and not just epistemic.⁹

S&A worry that “in creating models that use symbolic dynamics, setting the exact thresholds of the categorical regions . . . is a tricky business, with dire consequences if done incorrectly.” This concern is, however, misplaced, because it conflates epistemology with ontology. The confusion between the methodological choice of the scientist who studies a system and the unavoidable necessity (which I explained in the target paper) of discrete representations in the system itself is regrettable. Science aside, in successful natural cognitive systems, the categorical regions will have been set right *because* those systems are successful and because communication (and hence success) requires discrete (therefore readily interpretable) representations. S&A are of course free to choose on which level to describe certain things. However, in opting for an exclusive lowest-level description, they (along with Van Orden) necessarily ignore the way those things are. You can guess what my own choice is. As Hofstadter (1985) wrote in the context of a similar debate, quoting the famous harpsichordist Wanda Landowska, “You play Bach your way, and I’ll play him his way.”

3 No first philosophy

Dietrich in his commentary argues for pluralism¹⁰ by exposing the apparent insanity of its antithesis: “It is madness — collective madness, it turns out — to suppose that one theory can explain all that needs to be explained about the brain and its emergent mind.” I am uneasy with this sweeping statement, because it does not distinguish between theories and explanatory frameworks. Computation is not a theory of the mind — it is a framework within which theories can (or rather, as I argued in the target article, must) be developed.

On the level of theories, Dietrich is, I believe, exactly right. The possibility that no single theory may turn out to be capable of covering all of vision (let alone all of the mind) had been voiced by Marr (1981), who distinguished between Type 1 theories (general and concisely statable) and Type 2 explanations (bags of tricks, or one damn thing after another, an expression one sometimes hears in relation to artificial intelligence). Marr himself much preferred Type 1 theories, but was open to settling for Type 2 in cases where there seemed to be no alternative. A hog’s breath is arguably better than no breath at all,¹¹ and a non-minimalistic, less than elegant yet effective Type 2 explanation is better than ignorance (it is also, as Markman points out in his commentary, the likely kind of explanation of evolved entities).

With this distinction in mind, I consider Dietrich’s statements such as “we have no idea how the neurochemistry of the brain produces thoughts” or “we are clueless about creativity” as uncalled-for by the state of the art in cognitive sciences, a book-length survey of which I offer elsewhere (Edelman, 2008). I suspect that he is just being playful here, trying (and succeeding) to draw out dour folks such as myself. Moreover,

⁹Such emergence has been demonstrated in game-theoretic studies such as that of Puglisi, Baronchelli, and Loreto (2008), who write: “The linguistic level emerges as totally self-organized and is the product of the (cultural) negotiation process among the individuals.”

¹⁰By defending radical pluralism rather than the more, how should I put it, pluralistic, garden variety that I profess here, Dietrich, himself an inveterate pluralist (Markman and Dietrich, 2000; Dietrich and Markman, 2003), makes me seem what in Romanian is described as *mai catolic decât Papa* (“more catholic than the Pope”).

¹¹This phrase, *sans* the qualification, is the motto of the Hog’s Breath Saloon, 400 Front Street, Key West, FL.

H. L. Mencken’s remark on the intelligence of the American public applies, *mutatis mutandis*, to the present case: no philosopher ever lost readership by underestimating what science knows about the brain/mind.¹²

More seriously, my impression that philosophy is afoot here is confirmed by the following quote from Dietrich: “Understanding consciousness in any deep but still standard scientific way (i.e., reducing it to neural processes) is not in the cards, basically because dualism will always *seem* true.” I might add that it will always *seem* to us that the sun revolves around the earth; so what? In arguing for pluralism from the premise that “philosophy shows conclusively that there exist incommensurable points of view on matters of great significance,” Dietrich seems to be trying to crown philosophy (or maybe just common sense) as the ultimate arbiter of scientific truth, an act that I must resist. For whatever it is worth, in that I am in good company:

My position is a naturalistic one; I see philosophy not as a priori [. . .] groundwork for science, but as continuous with science. I see philosophy and science as in the same boat — a boat which, to revert to Neurath’s figure as I so often do, we can rebuild only at sea while staying afloat in it.¹³ There is no external vantage point, no first philosophy.

— W. V. O. Quine (1969, pp.126-127)

4 Conclusion

The closing line of Krylov’s fable — “The cart’s still there today” — has become a proverb in the Russian language, commonly used to bemoan the lack of progress in a collaborative undertaking. I submit that in the present case, however, the situation is not all that gloomy. The swan, the pike, and the crawfish should just go off and have a drink together, because the cart of cognitive science already is pretty much where it needs to be: right underneath the tree of knowledge, from whose trunk spring tier after tier of leafy branches and whose roots reach deep into the loam of computation. When we return from the pub, we can help each other onto the cart and pick some fruit.

Acknowledgments

I thank Barb Finlay for a much-needed reality check on my expository style, Rick Dale for editorial comments on a draft of this paper, and Eric Dietrich and JETAI for hosting this invigorating exchange of views.

References

Brooks, R. and L. Steels (1995). The artificial life route to artificial intelligence: Building embodied, situated agents.

¹²Here is the original quote from Mencken: “No one in this world, so far as I know — and I have searched the records for years, and employed agents to help me — has ever lost money by underestimating the intelligence of the great masses of the plain people” (*Notes on Journalism*, Chicago Tribune, 19 September 1926).

¹³Philosopher Otto Neurath’s boat metaphor, which originally appeared in *Protokollsaetze*, *Erkenntnis* 3:204-214 (1932), is best explained by another quote from Quine (1960, p.3): “We are like sailors who on the open sea must reconstruct their ship but are never able to start afresh from the bottom. Where a beam is taken away a new one must at once be put there, and for this the rest of the ship is used as support. In this way, by using the old beams and driftwood the ship can be shaped entirely anew, but only by gradual reconstruction.”

- Deisseroth, K., P. G. Mermelstein, H. Xia, and R. W. Tsien (2003). Signaling from synapse to nucleus: the logic behind the mechanisms. *Current Opinion in Neurobiology* 13, 354–365.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster.
- Dietrich, E. and A. B. Markman (2003). Discrete thoughts: Why cognition must use discrete representations. *Mind and Language* 18, 95–119.
- Eco, U. (1979). *The role of the reader: Explorations in the semiotics of texts*. Bloomington, IN: University of Indiana Press.
- Eco, U., R. Rorty, J. Culler, and C. Brooke-Rose (1992). *Interpretation and overinterpretation*. Cambridge: Cambridge University Press.
- Eco, U. and T. A. Sebeok (1988). *The Sign of Three: Dupin, Holmes, Peirce*. Indiana University Press.
- Edelman, S. (1999). *Representation and recognition in vision*. Cambridge, MA: MIT Press.
- Edelman, S. (2003). But will it scale up? Not without representations. *Adaptive Behavior* 11, 273–275. A commentary on *The dynamics of active categorical perception in an evolved model agent* by R. Beer.
- Edelman, S. (2008). *Computing the mind: how the mind really works*. New York: Oxford University Press.
- Fogelin, R. (2003). *Walking the Tightrope of Reason: The Precarious Life of a Rational Animal*. Oxford: Oxford University Press.
- Goldsmith, J. A. (2007). Towards a new empiricism. In J. B. de Carvalho (Ed.), *Recherches linguistiques à Vincennes*, Volume 36.
- Hofstadter, D. R. (1985). Waking up from the Boolean dream, or, subcognition as computation. Chapter 26, pp. 631–665. Harmondsworth, England: Viking. Metamagical Themas.
- Holyoak, K. J. and R. G. Morrison (Eds.) (2005). *The Cambridge handbook of thinking and reasoning*. New York: Cambridge University Press.
- Johnson, T. R. and J. F. Krems (2001). Use of current explanations in multicausal abductive reasoning. *Cognitive Science* 25, 903–939.
- Markman, A. B. and E. Dietrich (2000). In defense of representation. *Cognitive Psychology* 40, 138–171.
- Marr, D. (1981). Artificial intelligence: a personal view. In J. Haugeland (Ed.), *Mind Design*, Chapter 4, pp. 129–142. Cambridge, MA: MIT Press.
- Nilsson, N. J. (2006). How are we to know? Available online at <http://ai.stanford.edu/~nilsson/hawtk/hawtk-webpage.htm>.
- Noble, D. (2008). Claude Bernard, the first systems biologist, and the future of physiology. *Experimental Physiology* 93, 16–26.
- Peirce, C. S. (1877). The fixation of belief. *Popular Science Monthly* 12, 1–15.

Puglisi, A., A. Baronchelli, and V. Loreto (2008). Cultural route to the emergence of linguistic categories. *Proceedings of the National Academy of Science* 105, 7936–7940.

Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.

Quine, W. V. O. (1969). Natural kinds. In *Ontological relativity and other essays*, pp. 114–138. New York, NY: Columbia University Press.