

Volume: 12 ([next](#), [prev](#)) **Issue:** 040 ([next](#), [prev](#)) **Article:** 7 ([next](#) [prev](#) [first](#)) **Alternate versions:** [ASCII Summary](#)

Topic:

Article:

PSYCOLOQUY (ISSN 1055-0143) is sponsored by the American Psychological Association (APA).

THE UBIQUITY OF COMPUTATION

Commentary on Harnad on Symbolism-Connectionism

Eric Dietrich
Program in Philosophy, Computers, and
Cognitive Science
Binghamton Univ.

dietrich@binghamton.edu

Abstract

I argue that from a natural, explanatory perspective, computation is ubiquitous. This is due to our finite ability to measure states and processes in the world. While it is true that in many of our sciences we frequently use continuous descriptions, such descriptions are not necessary, and in cases where information is being studied, continuous descriptions can hamper research. Then, I discuss briefly the empirical nature of the computational hypothesis, and defend it against Harnad's (2001) criticisms. I also point out that computation in the physical world is a semantical process, and that criticisms of it based on its alleged purely syntactic nature are misguided.

REPRINT OF: Dietrich, E. (1993). The ubiquity of computation.
Think 2: 12-78 (Special Issue on "Connectionism versus Symbolism"
D.M.W. Powers & P.A. Flach, eds.).
<http://cwis.kub.nl/~fdl/research/ti/docs/think/2-1/index.stm>

1. For many years now, Harnad (2001) has argued that transduction is special among cognitive capacities -- special enough to block Searle's (1980) Chinese Room Argument. His arguments (as well as Searle's) have been important and useful, but not correct, it seems to me. Their arguments have provided the modern impetus for getting clear about computationalism and the nature of computing. This task has proven to be quite difficult. Which is simply to say that dealing with Harnad's arguments (as well as Searle's) has been difficult. Turing (1964/1950), it turns out, only got us started. But Harnad's (and Searle's) arguments ultimately fail. Turing, it turns out, was on the right track.

2. To begin, I will state for the record what computationalism is (or ought to be). It is the hypothesis that all cognition, all thinking, is the execution of some partial recursive functions. Our jobs as psychologists, AI researchers, cognitive scientists, cognitive ethologists, neuroscientists, etc. is to discover which functions from this vast set of functions best describe the various cognitive capacities we observe. (I don't for a minute think that we should approach our task explicitly thinking of functions; we shouldn't try to be mathematicians, necessarily.)

3. Computationalism has to be true because every process is a computation: computation is ubiquitous. At least that's the hypothesis. Here's my argument (this argument is derived from Fields, 1989). It's impossible to describe any process (transduction included) to an arbitrary degree of precision. Past a certain degree, noise will render all future descriptions meaningless. This point is basically Heisenberg's Principle, and hence can be put another way: all measurements of any process are ultimately nonclassical. Measurement is nonclassical when it affects the internal states or dynamics of the system being measured. The claim is that all measurement becomes nonclassical for every physical system if one attempts to specify, or measure, states to an arbitrary degree of precision. Fields puts it this way: "In the limit of an infinitely precise measurement of [a] state, all information about the next state ... is lost" (p. 176). There is, therefore, an upper bound on the precision of the measurements we can make and hence on the descriptions we can deploy for a given system or process. In short, the number of states we can describe is at most countably infinite. Hence, for any process or system, there will always be some virtual machine that completely describes that process's behavior. We can, in theory, implement that virtual machine on a computer (but not always in practice) because of:

1. Church's thesis, and
2. Turing's proof of the universality of the Turing machine. QED.

(Please note: computationalism is a hypothesis about cognition. It is not a metaphysical thesis about the nature of the world. The argument in the preceding paragraph derived from Fields is a metaphysical thesis about the world (actually, Fields thinks it is a physical thesis). Hence, Fields' argument is not an argument for computationalism, strictly speaking. The argument for computationalism is a straight-forward instance of universal instantiation, based on the assumption that Fields' argument is correct.)

4. Now, many have regarded the ubiquity of computation as a testament to its vacuity, Harnad and Searle among them. They say that if everything is a computation, then the thesis is vacuous, and also that computationalism is vacuous. But neither of these claims is vacuous. The first claim is a general, universal hypothesis made in the time-honored tradition of science everywhere. The claims that everything is made of atoms, or that all objects tend to continue their motion unless otherwise disturbed, or that all species evolved are well-known scientific hypotheses. They are not vacuous at all, but are instead profound and deep claims about the world. Furthermore, the inference from, e.g., 'everything is made of atoms' to 'this keyboard is made of atoms' is a case of universal instantiation and constitutes a test of the hypothesis: if it should turn out that my keyboard is not made of atoms, then the hypothesis is false, and our physics is in deep trouble.

5. So it is with computationalism and Fields' thesis. The claim that everything is a computation is not vacuous. And the inference from it to the claim that thinking is computation is likewise not vacuous. Finally, the inference to my thinking is computing is a test, because if it should turn out that my thinking is not computing then both computationalism and Fields' thesis are false. Testing for individual counterexamples is probably the only way we have of proceeding, by the way.

6. One can see that the claim that everything is a computation is not vacuous by noticing what it would mean if it were true and what it would take to make it false. With respect to computationalism, this last point is very important. Many researchers misunderstand (or refuse to accept) the empirical nature of computationalism. Computationalism would be false if thinking (e.g., my thinking) involves executing functions that are equivalent to the halting problem. And we have every reason to believe we could demonstrate this if it were

true. More importantly, it seems clear, I think, that computationalism's status as a viable hypothesis would be seriously jeopardized if it should turn out that humans routinely compute functions that are equivalent to arbitrarily large instances of the travelling salesman problem. Strictly speaking, such a scenario is compatible with computationalism, but assuming P does not equal NP , such a scenario means that we seriously misapprehend computational architectures. In such a case, I for one would be prepared to re-evaluate my commitment to computationalism.

7. In sum, computationalism is a non-trivial, falsifiable, thesis about cognition, whose truth follows from a profound fact about measurement and the physical world, and for which there is ample evidence.

8. Where does this leave Harnad's arguments and his hybrid project? I'll take these in turn. All of Harnad's arguments against computationalism and for what he calls the Total Turing Test (TTT) seem to me to be question begging or just bald assertions. For example, defining computation as the "manipulation of physical 'symbol tokens' on the basis of syntactic rules that operate only on the shapes of the symbols..." (para. 2) is clearly question-begging. I, for one, have no idea what 'syntactic rules operating only on shapes' even means with respect to real computation, and such a phrase clearly packs into the notion of computation the very aspect that allows Harnad (and Searle) to reject it. And anyway, it is by now almost a commonplace that computation is fully semantical (e.g., see Dietrich, 1990; 1989; Dyer, 1990; and Smith, 1988; in press).

9. Another question-begging feature of Harnad's argument is his reliance on the distinction between simulation and the real thing: if computationalism is true, this distinction is vacuous for cognition. And nowhere does he give us a compelling argument that computers merely simulate thinking. He just asserts that they do.

10. Furthermore, Harnad gives us no reason to believe either that "certain forms of performance may be unattainable from 'ungrounded' symbol systems" (para. 5), or that parallelism is essential to realizing certain nontrivial cognitive abilities (para. 6). And for this last claim, there is considerable evidence that it is false. In their paper 'Multi-layer feedforward networks are universal approximators' (1989), Hornik, Stinchcombe, and White give a slow, rigorous proof that standard feedforward networks with as little as one hidden layer are capable of approximating any (Borel) measurable function, i.e., that they are as strong as Turing machines. And Selmer Bringsjord (1991) has offered the following quick, but convincing argument that they are equivalent: neural nets are computationally equivalent to cellular automata, which in turn are computationally equivalent to K-tape Turing machines, from which it follows that the Church-Turing thesis remains true even for neural nets; hence, no claims about different strengths for different types of machines need be entertained.

11. Finally, what about Harnad's hybrid project? I think it sounds interesting and even worthy of funding. His hybrid system isn't required for the reasons he thinks it is, i.e., his system isn't required because a more standard system would be incapable of thinking. But his hybrid might be methodologically necessary because it is possible that only with such a system might we humans figure out how to implement intelligent thought. In any case, I think there is room for everyone. Figuring out how the brain produces a mind is a problem that dwarfs even such chestnuts as the search for a unified field theory or predicting the weather. With such an immense task on our hands, it won't do to lock any researchers out, especially those with some positive project they're pursuing, like Harnad.

REFERENCES

Bringsjord, S. (1991) 'Is the Connectionist-Logicist Clash One of AI's Wonderful Red Herrings?' In: *Journal of Experimental and Theoretical Artificial Intelligence* 3.4: 319--349.

Dietrich, E. (1989) 'Semantics and the Computational Paradigm in Cognitive Psychology'. In: *Synthese* 79:

Dietrich, E. (1990) 'Computationalism'. In: Social Epistemology 4: 135--154.

Dyer, M. G. (1990) 'Intentionality and Computationalism: Minds, Machines, Searle and Harnad,' In: Journal of Experimental and Theoretical Artificial Intelligence 2.4. pp. 303--319.

Fields, C. (1989) 'Consequences of nonclassical measurement for the algorithmic description of continuous dynamical systems'. In: Jour. of Experimental and Theoretical AI, vol. 1, n. 3: 171--178.

Harnad, S. (2001) Grounding symbols in the analog world with neural nets -- A hybrid model. PSYCOLOQUY 12(034) <http://www.cogsci.soton.ac.uk/cgi/psyc/newpsy?12.034>

Hornik, K., Stinchcombe, M., and White, H. (1989) 'Multi-layer Feedforward Networks are Universal Approximators'. In: Neuronetworks 2.

Searle, J. R. (1980) "Minds, brains and programs." Behavioral and Brain Sciences 3: 417-424.
<http://www.cogsci.soton.ac.uk/bbs/Archive/bbs.searle2.html>
<http://www.bbsonline.org/documents/a/00/00/04/84/index.html>

Smith, B. C. (1988) 'The Semantics of Clocks'. In J. Fetzer (ed.) Aspects of Artificial Intelligence, Dordrecht, Holland: Kluwer, pp. 3--31.

Smith, B. C. "Is Computation Formal?" (unpublished but widely-circulated manuscript, the contents of which will appear in volume II of "The Age of Significance", forthcoming from MIT Press).

Turing, A. M. (1964) "Computing machinery and intelligence". In: Minds and machines, A . Anderson (ed.), Engelwood Cliffs NJ: Prentice Hall. <http://cogprints.soton.ac.uk/documents/disk0/00/00/04/99/index.html>

Volume: 12 ([next](#), [prev](#)) **Issue:** 040 ([next](#), [prev](#)) **Article:** 7 ([next](#) [prev](#) [first](#)) **Alternate versions:** [ASCII Summary](#)

Topic:

Article: