

Dietrich, E. (2000). Cognitive Science and the Mechanistic Forces of Darkness, or Why the Computational Science of Mind Suffers the Slings and Arrows of Outrageous Fortune. *Techné: eJournal of the Society for Philosophy and Technology*, Winter, <http://scholar.lib.vt.edu/ejournals/SPT/v5n2/dietrich.html>

**Cognitive Science and the Mechanistic Forces of Darkness, or
Why the Computational Science of Mind Suffers the Slings and Arrows of Outrageous Fortune**

Eric Dietrich

Philosophy Dept.

Binghamton University, Binghamton NY 13902-6000

dietrich@binghamton.edu

1. Introduction. Whither cognitive science?

A recent issue of *Time* magazine (March 29, 1999) was devoted to the twenty greatest "thinkers" of the twentieth century -- scientists, inventors, and engineers. There is one interesting omission: there are no cognitive psychologists or cognitive scientists. (Cognitive science is an amalgam of cognitive, neuro, and developmental psychology, artificial intelligence, philosophy, linguistics, biology, and anthropology.) Freud is there, to be sure. But, while he was very influential, it is not even clear that he was a scientist, let alone a cognitive scientist. There are those who regard Freud as somewhere between incompetent and a charlatan (see Glymour, 1988). In any case, though Freud's positive proposal for the mind's architecture -- namely, that it contains the *unconscious* -- seems correct as far as it goes, it does appear that all the details are wrong. For example: (1) there is a lot more to the mind than the mere unconscious; (2) it is doubtful that there is an id, ego, and superego; (3) most dreams may very likely be meaningless; and (4) human motivation, even unconscious motivation, is about a lot more than sex. In the end, because he was most interested in certain kinds of human mental malfunctioning, Freud is probably best thought of as a physician, a proponent and early explorer of human mental health; he was not an experimental cognitive psychologist.

Piaget is included, too, and that is good. Chomsky is given a tiny paragraph. But where are William James, Edward Tolman, Konrad Lorenz, George Miller, Jerome Bruner, and Alan Newell and Herbert Simon? For that matter, where's B. F. Skinner? Alan Turing, of course, is discussed, too, but his contribution to AI and cognitive science in the *Time* piece is limited to a brief explanation of the Turing test; Turing did far more than that.

Also, in a fold-out section which lists scientific and technological advances from 1900 to 1998, none of the important achievements of cognitive science are even so much as mentioned.

So why is cognitive science missing from this issue of *Time*?

Consider another puzzle. This has been the decade of the brain. But why hasn't there been the decade of the *mind*? There's been no decade of the mind, ever. But it is minds that are important, not brains, not even *working* brains. When it comes to brains, we want to know how working brains produce minds. But even this won't tell us very much about minds, as such. (I will explain why this is so in section 6.2, when I talk about *virtual machines*). How the mind works, what thinking is, and the nature of thoughts are among our greatest mysteries. The science of the mind could arguably be our deepest science (though it is not yet very deep). Yet, neither the mind, nor the science of the mind is very much in evidence in *Time*, or in naming whole decades. What's going on here?

Probably many things are going on. Perhaps *Time* didn't include cognitive science because we don't know much about the mind. It is hard to find two cognitive scientists who agree on any of the details of a theory of mind. But this can't be the explanation. Worries about lack of agreement couldn't really have been that important to *Time*; after all, they did include a philosopher, Ludwig Wittgenstein, and an economist, John Maynard Keynes. No two economist ever agree on anything, and Wittgenstein couldn't even agree with himself!

I suggest that a deep part of what is going on has to do with the *computational hypothesis* in cognitive science. To the extent that there is any agreement in cognitive science, it is agreed that the computational hypothesis is the discipline's foundational assumption. But this hypothesis is so under siege that it is not seen as much of a scientific advance at all, and hence, the founders of cognitive science are not considered important twentieth century scientists. Why is this perfectly nice hypothesis, that never hurt anybody, that is in fact the foundation of most what we know about the mind, so badly regarded? This is what this paper is about.

Before we go any further, though, I want to say for the record what the computational hypothesis is.

2. The computational hypothesis.

The computational hypothesis (also known as *computationalism*) is a version of *functionalism* where all the functions are computable. It claims that cognition is the execution of Turing-computable functions defined over various kinds of representational entities. Period. There is a long and rather complicated story about how computationalism works, what "Turing-computable" means, and how it figures in the definition. I will spare you these here (but see Dietrich, 1990, and Dietrich and Markman, 2000). All I need for present purposes so to say what computationalism is *not* :

Computationalism is only a foundational hypothesis. Computationalism does not get specific about which particular functions cognition is. Indeed we aren't sure which functions cognition is. Therefore, computationalism does not tell us what models to build, nor which experiments to run. All computationalism gives us is a framework within which to work.

Computationalism, as with computation on garden variety computers, is not committed to mental representations (internal encodings of information) of any particular variety. Rather, computationalism is compatible with many of different kinds of representations from numerical quantities to propositional nodes in a semantic network (for more on this see Markman and Dietrich, 2000).

In sum, assuming computationalism leaves all the hard work left to do. Which means it is not really a theory. Computationalism is a *theory schema*. We still need to roll up our sleeves and get down to the difficult business of developing a theory of mind. Computationalism does tell us what this theory will look like -- but only broadly.

3. The real problem with computationalism

Computationalism is attacked from without and from within cognitive science. The vigor of the attacks, the large number of researchers and scholars involved, and the weakness of the arguments used in the attacks together with the weakness of proposed alternatives suggest to me that more is going here than meets the eye. We do not have here simply a case of a hypothesis that seems to reasonable men and women to be false, or for which there is little evidence, or which is so radical that agnosticism seems the prudent stance. Instead, we have a hypothesis which is, I think, regarded as deeply *anti-human*, and hence repugnant. I suspect that the real reason everyone is out gunning for computationalism is because it violates, not our common sense nor some well-developed scientific intuition, but rather our conception of what it means to be human. I believe, in short, that the real problem with the general acceptance of computationalism is its perceived association with what I call *The Mechanistic Forces of Darkness*. Here's a quote that expresses well this felt repugnance:

[AI]'s real significance or worth [lies] solely in what it may contribute to the advancement of technology, to our ability to manipulate reality (including human reality), [but] it is not for all that an innocuous intellectual endeavor and is not without posing a serious danger to a properly human mode of existence. Because the human being is a self-interpreting or self-defining being and because, in addition, human understanding has a natural tendency to misunderstand itself (by interpreting itself to itself in terms of the objectified by-products of its own idealizing imagination; e.g., in terms of computers or logic machines) -- because of this there is a strong possibility that, fascinated with their own technological prowess, moderns may very well attempt to understand themselves on the model of a computational machine and, *in so doing*, actually make themselves over into a kind of machine and fabricate for themselves a machinelike society structured solely in accordance with the dictates of calculative, instrumental rationality (Madison, 1991).

Who can resist an enemy so appealing that you want to emulate it?

Here are two other quotes from a recent book:

"... human cognition is too rich to be simulated by computer programs" (Horgan and Tienson, 1996, p. 1).

and

"...human (and other natural) cognition is too subtle and sophisticated to conform to programmable representational rules" (Horgan and Tienson, 1996, p. 145).

These two quotes are just bald assertions. The authors make no attempt to justify or argue for them. Apparently, the authors think the truth of their statements is obvious -- and from a human-centered perspective, it is.

In sum, I think that computationalism's troubles are due to its perceived anti-humanism. We fear the mechanistic forces of darkness which AI and cognitive science represent. Our fear of such forces goes hand in hand with our refusal to see ourselves as part of the natural order. I don't mean to belittle this fear; I mean to take it seriously -- but I do think it is uncalled for ... and dangerous.

I think that what is going on with computationalism is like what happened to Darwin's evolutionary hypothesis. Darwin came along and said we were fancy chimpanzees. Now along come the cognitive scientists saying that we are fancy calculators. The attitude toward such mechanistic hypotheses is not that they seem false given the data, but rather they *must be false, regardless of the data*. People have a deep dislike of such hypotheses because they violate our sense that humans are special, and more than mere animals, more than mere mechanisms.

Digging deeper, I think our steadfast refusal to see ourselves as part of the natural order works with two other ingredients in a positive feedback loop to *generate* cognitive science's perceived anti-humanism. The other two ingredients are: 1) AI's and cognitive science's tendency to oversimplify cognition, and 2) a confusion about the nature of computers. In the next section, I discuss this feedback loop in detail.

By the way, and for the record, AI and cognitive science really do have some robust failures. Here are the main five. I think these failures also contribute to cognitive science's troubles, but these failures are not part of the feedback loop mentioned above. Instead, these failures are merely the failures of a very young science faced with problems of staggering difficulty.

1) We have failed to explain the plasticity of human intelligence. We only have vague ideas about how the mind works on novel problems and in novel situations, and how it works so well with degraded input. We cannot say where new representations come from or how new concepts are acquired. We have done a better job explaining generative processes like syntax than content bound processes like semantics.

2) We have failed to tell an integrative story about cognitive and sensorimotor behavior. For most of the history of cognitive science, cognition got the lion's share of research attention. Sensorimotor behavior and robotics were considered as an afterthought. Many researchers, however, have come to the conclusion that this list of priorities is exactly backward. These researchers suggested that cognitive science must concentrate first on the sensorimotor aspects of organisms and the development of systems that interact with their environment. Only when these processes are understood fully, can cognitive science graduate to the study of higher level processes.

3) We have failed to tell an integrative story about brains and cognition. Again, throughout the history of cognitive science, the mind got all the attention. Understanding how the brain carried out cognitive processes was assumed to be an implementational detail. Again, researchers have suggested that the priorities of the field should be reversed. The behavior and structure of brains seem crucial to the flow and structure of cognition. The slogan is that mind emerges from the brain. Cognitive science, it is argued, should be the science of this emergence.

4) Our explanations of human development and maturation typically do not characterize the trajectory of development. Instead, developmental theories often capture snapshots at different stages of development, and then posit mechanisms to jump from one stage the next.

5) We have failed to make an intelligent machine. If cognition is computation over representations, then why is it so proving so hard to make an intelligent computer.

Note the way failures 1) and 5) work together. To many, the computer just doesn't seem like the right sort of thing for grounding research into the nature of plasticity and representational content. The sentiment is that to the extent that plasticity is crucial to cognition, then computation must be the wrong way to think about cognition. Here is an analogy: if the Wright brothers had failed at building a flying machine, then it would have been reasonable to question their theory of flight (namely that it requires lift). On the other hand, it would not have been reasonable to abandon the theory of lift: the theory might have been correct (as is indeed the case), but it might have turned out that building a flying machine was harder than it looked at first.

To sum up: many are dismayed at the robustness and complexity of human cognition, and they do not think that computation is up to the task of explaining it. But we have only been at it for forty-five years or so. That is nowhere near long enough to judge computational cognitive science a failure. Still, these failures loom large. It is against this background that a feedback loop of misunderstandings, poor methodology, and unwavering belief in human specialness works to generate undeserved animus against artificial intelligence and cognitive science.

4. The three ingredients of the feedback loop and how they work together.

I claim that three ingredients together form a feedback loop that is responsible for most of the attacks on cognitive science, and, hence, on cognitive science's low status among the sciences. I discuss the three in this order: 1) AI's and cognitive science's tendency to oversimplify cognition, 2) the belief that computers are basically logic machines, and 3) our refusal to see ourselves as part of the natural order.

4.1. AI's tendency to oversimplify.

The quickest way to state AI's tendency to oversimplify is to say that logical positivism, while dead in philosophy (and for good reasons), has deeply infected AI. And oversimplification is one of the hallmarks of logical positivism.

Specifically, AI simply spends too much time and energy developing logical models of cognitive processes. Of course, logic is only applicable to a very small number of cognitive processes and so AI tends to focus on these. As evidence for this claim, fully one-half of the papers received by the *Journal of Experimental and Theoretical AI* (which I edit) report some logical result or other in the form of a logical theorem. There are entire conferences and journals devoted to using logic to explain cognitive processes such as temporal reasoning and what is called common sense reasoning -- allegedly, the sort of reasoning we all use everyday.

Humans do naturally reason logically, from time to time, so exploring various logics is a reasonable thing to do. Still, the amount of research on logic is excessive. One is reminded of the old joke about the drunk looking for his keys under a street lamp. When a passer-by offered to help him find his keys, the drunk said that he had lost his keys over there in the dark alley. Puzzled, the passer-by said "Why don't we look for them over there, then?" To which the drunk replied "The light's better here." Logic research on the cognition seems to be looking where the light is good.

4.2. A confusion about computers.

Anti-computationalists, like everyone else, know that deep down in the guts of every computer, exists some Boolean logic, in the form of logic gates and various circuits, governing its behavior. They then draw the false, but prima facie plausible, conclusion that computers are

essentially logic machines, hence it is *no accident* that AI spends a lot of time using logic to characterize human cognition, hence AI *essentially* misdescribes human thinking (because human thinking is much more than logical inferencing), hence AI *essentially misdescribes what it means to be human*.

Furthermore, virtually every single computer in existence is in fact a tool: a word-processor, a game, or an emailer. This makes sense because these things can be done via logic. Essentially, modern computers are *logic-based tools* .

4.3. The natural order and human specialness.

Most anti-computationalists do not want a theory of the human mind that in their eyes does not do justice to the marvelousness, the uniqueness, the specialness of human beings. They instead want a theory that justifies their belief in our specialness. Human cognition is obviously powerful -- considerably more powerful than that of even our closest chimpanzee cousins. A normal five year old human child communicates in a far more intricate fashion than even the most well-trained, adult chimp. And this is just for starters. So, yes, humans are different. And, yes, we are special, as all life is special. But we are, for all that, animals. Indeed, we are computational animals, furry robots. And no amount of wishing otherwise will change this fact.

We all know that the belief that we are special has been damnably hard to hang on to. First, Copernicus and Galileo kicked us out of the center of the universe. Then Kepler squashed the perfect circles of Earth's and the other planets' orbits around the sun into ugly ellipses. Then Darwin said that we were a kind of ape. Now along come the cognitive scientists claiming that we share important similarities with fancy calculators. No one wants to hear this. We want a theory of the mind that enshrines us as the pinnacle of creation, that explains why humans are special, rather than why we aren't.

The feedback loop now kicks in. It works like this. AI researchers (at least many of them) spend a lot of Procrustean effort trying to force as much cognition as possible into some logic mold or other, and then ignoring the rest of cognition. This is no accident, say the anti-computationalists, because as is "well-known," computers are essentially logic machines in the first place, which we deploy as fancy tools. So, to many, it does look like AI is up to its eyeballs in "logocentrism". But it is obvious to the most casual observer that there is more to human thinking -- there is more to *being* human -- than being a logic-based tool, or even more than being a rational machine (animal). Hence, AI and indeed all of computational cognitive science is completely misdescribing what it means to be human. Since humans are not merely word processors, Gameboys, nor emailers -- since humans aren't just logic-based tools of any sort - - it follows that humans aren't computers. Hence cognitive science must be wrong.

5. The unfortunate consequences of the Fear of the Mechanistic Forces of Darkness.

There are two consequences to the fear of the mechanistic forces of darkness and the feedback loop between it and the other two ingredients. The first consequence is that the *real* computational hypothesis is *not* the actual focus of the attacks, but rather the focus of the attacks is "Computerism". Computerism is the view that humans are a variant of the kinds of computers we have on our desks. Here is computerism defined:

You are a computer. Your mother was a computer. And computers, as we all know, are just rigid, rule-following logic machines we use as tools, exactly like the thing on your desk.

Lest you think I am (merely) joking. Here is another quote from Horgan and Tienson (1996), two well-known anti-computationalists:

"[According to classical, computational cognitive science,] [c]ognitive processing conforms to precise, exceptionless rules, stable over ... representations and articulable in the format of a computer program." (1996, p. 24).

But Horgan and Tienson are wrong. They haven't even approximately described classical, computational cognitive science. In real cognitive science, the rules are not exceptionless, the representations are not necessarily propositional (which is what they mean when they say "articulable in the format of a computer program"), and the words "computer program" have the connotation that human cognitive processing could be stated cleanly in C++, Java, or Lisp. This connotation is very deceptive, and it is profoundly in error. Computerism is obviously false and hence easy to attack. Attacking it while seeming to attack computationalism is how many anti-computationalist make their living.

It is troubling that many confuse computerism with computationalism. No cognitive scientist thinks that humans are much like modern computers of any variety. The claim, to repeat, is that thinking is the computing of recursive functions (Turing-computable functions) of the right sort.

The second, and deeper, consequence of the fear of the mechanistic forces of darkness is the undeserved popularity of the three new contenders in cognitive science and their love affair with "emergence." The three new contenders are dynamic systems, embodied cognition, and connectionism (artificial neural nets). Dynamic systems is an approach to cognitive science that says that what matters in theorizing about cognition is the underlying physical processes of the neurons which can be described best using differential equations. Embodied cognition says that what matters in theorizing about cognition is that minds are housed in bodies which must interact with the world, and this interaction forms the basis of all thought. And,

connectionism says that what matters in theorizing about cognition is the informational processes of the neurons which can be described best using vector calculus.

All three of these new approaches to understanding the mind have the following consequence:

What is really interesting about human cognition *emerges* from some underlying substrate, and we need only study the substrate.

I call this *the emergent cognition principle*. Its allure derives from the fact that by assuming it, cognition turns out to be basically a free lunch, hence basically mysterious, hence non-mechanizable. Humans thus remain special -- and mysterious.

It's the emergent cognition principle that does all the work here. Most anti-computationalists see the emergent mind as something that is *not* a natural kind and hence not really subject to scientific investigation. Emergence is seen as producing something that in some sense reduces to the working brain, but which really is not a proper object of investigation. Instead it is best to investigate and theorize about the substrate -- the brain and its neurons -- out of which the mind emerges. The dark result is that the scientific investigation of cognition and the mind is held back; our understanding of our true computational selves is thus delayed. And this has consequences ranging from incorrect treatments of mental illness, through the heavy costs of misunderstanding human decision-making, rationality, and creativity, to misplaced energies in dealing with our social problems such as war, crime, over-population, and environmental damage.

6. There are no free lunches: A quick defense of computational cognitive science.

Certainly one solid defense is to object to the straw-man ploy of many anti-computationalists. As I said, they frequently misdescribed computationalism as computerism and attack the latter instead of the former. But only the latter is easily assailable. AI and cognitive science have two other important defenses, too. First, anti-computationalists are ignoring a vast portion of AI and cognitive science that is not logic-based, and second, the alleged fact that computers are basically and essentially logic-based tools is incorrect, and deeply so.

6.1. Non-Logic AI.

For one thing, there is a lot of AI that is not based on logic. If one half of Jetai's papers are based on logic, what about the other half? Consider research on analogy, i.e., seeing that one thing is like another. Analogy research is really one of the success stories of AI and cognitive science. Very briefly, analogy is when one concept reminds a cognizer of another concept. One famous example is Rutherford being reminded of comets by studying the paths of alpha-particles. It is now known with a high degree of certainty that two concepts are analogous when their *structural descriptions* map onto one another (Gentner, 1983). A structural description is a tree-like knowledge representation made up of multi-place predicates. There is an enormous amount of data that analogy is mapping of structural descriptions. Also, there are several robust computer models of such analogical structure mapping (e.g., Falkenhainer, et al., 1989). And none of this has anything to do with logic or logical reasoning or non-monotonic logic or anything else logical. And there are several more such success stories about other cognitive capacities.

6.2. Virtual Machines.

Everyone seems to ignore one of the very deepest points about computers: A computer, any computer, comprises a hierarchy of *virtual machines*, each different from the one below and above it, and each supervening on the one below it. Note that "virtual" does not mean "not real." Computers are *not* just logic machines, and they are *not* just number crunchers. Like I mentioned earlier, this is an extremely deep point: several dissertations could be, and should be, written on the topic of virtual machines. Though I cannot do the details justice here, I will say a few things.

A word processor, like *Microsoft Word 98*, is a virtual machine that exists on top of some other virtual machines, such as an operating system, which in turn exist on top of some hardware machine. The hardware machine is not more real than the word processor. Each virtual machine has a methodology and mode of explanation unique to it (for example, the explanation and debugging of your word-processor is very different from the explanation and debugging of your operating system, or your disk drive). These methodologies and modes of explanation cannot be reduced cleanly in any epistemological sense to those of the machines below. To say that everything reduces to Boolean logic in a computer is exactly like saying psychology and biology reduce to physics. The claim is no doubt true in a technical, ontological sense, but it is epistemologically wrong. Trying to reduce the behavior of, say, an analogy system (or a word processor) to the Boolean algebra of the gates in the supporting chip would be to completely lose what was important about the analogy system in the first place. If you try to reduce without remainder a virtual machine to the machines below it, you are going to wind up with an incomprehensibly complex mess, and you will be methodologically stymied.

I now urgently call your attention to an important distinction, ignoring which breeds monsters: virtual machines are *not* like the emergence of minds the anti-computationalists hope

for in the emergent cognition principle. A full explanation of this would take a separate book, but fortunately, I have to be quicker.

To say that minds emerge according to the emergent cognition principle is to say that minds are basically *epiphenomenal*. Saying that minds are epiphenomenal can be interpreted in two different ways. The first way is metaphysical. To say that minds are epiphenomenal is to say that minds do not logically supervene on the physical, which is to say that they aren't part of the physical world, though they may in fact be associated somehow, some way, with brains (c.f., Chalmers, pp. 150 - 155). This is *not* the sense meant here. I don't doubt that many anti-computationalist believe that minds do not logically supervene on the physical, but that is not a requirement for belonging to their club. Instead, and this is the second way, epiphenomenalism is to be interpreted epistemologically: minds logically supervene on the physical (i.e., on brains) but are nevertheless *theoretically inert* -- all the interesting theorizing is done at the level *below* the mind, the level of neural processing and the like. Here's a quote:

...emergent structures share properties of universality which are to a large extent *independent* of the specific physical properties of the underlying substrate (Petitot, 1995, emphasis in the original).

Petitot means "epistemologically" independent (though, this is less than clear; he might mean metaphysically independent, too, which would be stronger and render him a dualist).

Consider an analogy. Think of the patterns in a kaleidoscope or the patterns in the sky made by clouds or passing jets (rainbows would also work). The patterns exist at one level, but the explanation of the patterns exists at a lower level. We explain these patterns by *reducing* them to the behaviors of their constituent parts. For example, we explain kaleidoscope

patterns by explaining how the lens, mirrors, and colored glass (now plastic) work. We explain contrails as condensed water vapor which is in turn explained by detailing the behavior of water molecules in the atmosphere. In each of these cases, the patterns are themselves theoretically *inert*. There is no science of contrails or of cloud patterns in the sky or of kaleidoscope patterns. Anti-computationalists enamored of this view of the mind.

But virtual machines are completely different. Whereas defenders of the emergent cognition principle hold that minds cannot be reduced without remainder to brains and that minds are not focus of scientific explanation, cognitive scientists believe that minds do ontologically reduce without remainder to brains but that explaining minds is a separate enterprise from explaining brains. This is exactly how virtual machines work on your computer. The behavior of your word processor can be reduced without remainder to the behavior of the underlying hardware, but no one in their right mind would do so, for at the hardware level all the interesting aspects of your word processor disappear. The description languages of the two levels is entirely different. As a simple example, suppose your word processor has a bug of some sort. You call the software representative of the company and she tells you how to fix the bug. This will be a software fix. She didn't tell you to get out your soldering gun and do any hardware work. If you had to do that, the bug wouldn't be a bug in your word processor in the first place (and she would have no doubt told you to call the computer manufacturer).

So the mind is a hierarchical suite of virtual machines logically supervening on the brain. Explaining the latter is *not* explaining the former, and vice versa. Certainly the mind is a working brain, and there are going to be in-principle reductive explanations of aspects of the mind in terms of the brain. But these explanations won't be at all useful without accompanying explanations of how the mind works couched at the level of the mind. That is, we will definitely need cognitive, computational explanations.

7. Glorious Machines, Humans, and the Natural Order.

I wish *Time* had included more cognitive science and more cognitive scientists. I wish the anti-computationalist would go home. For my part, I'm all for embodied cognition, dynamic systems, and neural nets. In their place, these three are enormously powerful explanatory methodologies. No one scientific methodology is going to explain something as complicated as the mind. But for all that, the idea that cognition is computation remains the single greatest advance we have ever had when it comes to understanding the mind. And we have yet to fully understand its implications. Indeed, we have yet to even fully understand computation.

The hypothesis that the brain is a computer, that you are, and that your mother is, is almost certainly correct. Where does that leave us? It's the end of the twentieth century; perhaps we could now drop the false dichotomy of "mere machine" or "divine creation" which continues to haunt us. Perhaps we could think of ourselves as we really are: glorious machines whose complexities make us precious.

References.

- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press, Oxford.
- Dietrich, E. (1990). Computationalism, *Social Epistemology*. 4 (2), pp. 135-154. (with commentary)
- Dietrich, E. and A. Markman (2000). Cognitive Dynamics: Computation and representation regained. In Dietrich, E. and Markman, A. (eds.) *Cognitive Dynamics: Conceptual change in humans and machines*. Lawrence Erlbaum.
- Falkenhainer, B., Forbus, K., & Gentner, D., (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41 (1), 1-63.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7, 155-170.
- Glymour, Clark, (1988). How Freud left science. In J. Earman, A. I. Janis, G. Massey, N. Rescher (eds.) *Philosophical Problems of the Internal and External World*, University of Pittsburgh Press.
- Horgan, T. and Tienson, J. (1996). *Connectionism and the Philosophy of Psychology*. MIT Press, Cambridge, MA.
- Madison, G. (1991). Merleau-Ponty's deconstruction of logocentrism. In M. Dillon (ed.) *Merleau-Ponty: Vivant*. SUNY Press, Albany, NY.
- Markman, A. and E. Dietrich (2000). In defense of representations. *Cognitive Psychology*, 40, pp. 138-171.
- Petitot, J. (1995). Morphodynamics and attractor syntax: constituency in visual perception and cognitive grammar. In Port and van Gelder (eds.) *Mind as Motion*. MIT Press, Cambridge, MA.