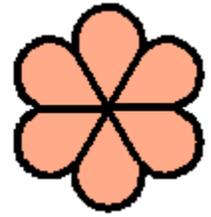


# A Non-Symbolic Theory of Conscious Content: Imagery and activity.



by  
**Nigel J.T. Thomas**

Presented at **Tucson 2000** (Toward a Science of Consciousness) Conference, Tucson AZ, April 14th 2000.

---

Until a few years ago, Cognitive Science was firmly wedded to the notion that cognition must be explained in terms of the computational manipulation of internal representations or symbols. Although many people still believe this, the consensus is no longer solid. Whether it is truly threatened by connectionism is, perhaps, controversial, but there are yet more radical approaches that explicitly reject it. Advocates of "embodied" or "situated" approaches to cognition (e.g., Smith, 1991; Varela *et al*, 1991, Clancey, 1997) argue that thought cannot be understood as entirely internal. Furthermore, it is argued that autonomous robots can be designed to behave more intelligently if representationalist programming techniques are avoided (Brooks, 1991), and that the way our brains control our behavior is better understood in terms of chaos and dynamical systems theory rather than as any sort of computation (e.g., Freeman & Skarda, 1990; Van Gelder & Port, 1995; Van Gelder, 1995; Garson, 1996).

It is controversial whether these approaches to cognition can really be understood coherently without somehow making appeal to the notion of computation over representations, but that is not the question I want to take up here. I am concerned, instead, with how theories of this type can address the nature of our subjective experience of thinking as compared with more traditional (symbolic) cognitive theories. First thoughts may seem to suggest that the newer approaches will measure up poorly. Whereas traditional theorists have always had a lot to say about mental representations and mental processing, non-representational roboticists seem to have little concern with such things, being much more interested in achieving systems capable of autonomous and intelligent *behavior*, regardless of what goes on inside them to achieve this. One criticism of the application of dynamical systems theory to cognition has been that through its rejection of mental representation it effectively abnegates the study of the mind, and heralds a return to the aridities of Behaviorism (Eliasmith, 1996).

However, the relationship between symbolism and the explanation of subjectivity is itself quite complex. On the one hand, one of the main sources of the paradigm was the

Carnegie-Mellon school of Artificial Intelligence research. This work relied, in large part, on the technique of protocol analysis (Newell & Simon, 1972; Ericsson & Simon, 1980). Typically, a subject would be given a puzzle to solve and would talk through the solution process out loud, into a tape recorder. The protocols thus collected would be analyzed to determine the terms and steps of a successful solution strategy, and a computer would then be programmed to solve similar types of problems using a formally similar strategy. The terms consciously used by the human subject became the symbolic tokens manipulated by the program. It was clearly intended that the computational symbols should be taken as modeling the conscious contents in the human solver's mind, and that actual mental contents in humans were to be understood as computational symbols.

Reinforcing this trend was the development and adoption of LISP as the workhorse language of AI. No doubt LISP has many relevant virtues, but amongst them, surely, is the fact that its superficial syntax is excellently adapted to representing natural language sentences and their parsing. A first stab at representing a sentence involves little more than enclosing it in parentheses, and its syntax can plausibly be tackled by complexifying the list structure in just the sorts of ways that LISP provides for. Again, the implicit message, if not necessarily always the explicit claim, was that the English words, phrases and sentences that we consciously hear, or prepare to speak, or say silently to ourselves when we think, are directly equivalent to the LISP atoms and structured lists of an AI program.

In the 1960s and 70s psychological thinking was ripe for a reaction against the Behaviorist paradigm that denied the reality (or, at least, the scientific significance) of subjective thought, and this sort of Artificial Intelligence work seemed to hold out the hope of being able to take the mind seriously once again without abandoning scientific rigor. The case for the psychological and philosophical significance of AI was argued forcefully and influentially on just these grounds by Boden (1977), who presents it as providing a solid scientific foundation for a more 'humanistic' sort of psychology.

But there were already other trends within AI, and other roots of cognitivism. Turing, after all, had suggested that intelligent behavior rather than human like inner processing would be the criterion of machine intelligence (Turing, 1950). In this vein, although still committed to symbolic computation, Minsky advocated AI as engineering rather than as psychology. Any means that could be devised to get machines to behave intelligently would be worth pursuing, regardless of whether they were the means used by humans (McCorduck, 1979). But although Minsky and his heirs did not want to depend on psychology for their inspiration, cognitive psychological theorists often did find inspiration in the ingenious programs that such engineers developed. If the relevant internal symbols and processing structures did not show up in introspection, that might mean no more that they reflected deeper, unconscious, (and probably more fundamental) mental processes than those being modeled at Carnegie-Mellon.

Chomsky's linguistic theories were an also an important influence on the formation of

cognitive science, and reinforced the trend toward thinking of cognitive theory as being concerned with non-conscious symbol processing. Chomsky, famously, distinguished depth grammar, which is universal and innate, from the surface grammar of the languages that we actually speak, and consciously think in. Cognitive theories inspired by this picture would naturally focus on the computational representation of depth grammar structures, and on the transformation processes between depth and surface. We are, of course, conscious of neither of these things, but the natural language thoughts of which we *are* conscious appear, from this point of view, as little more than epiphenomena of the non-conscious representations and processes deeper down. The philosophical commitments of this view were spelled out in Fodor's classic *The Language of Thought* (1975). It is clear that the innate "mentalese" language that he proposes (i.e., the symbolic computational knowledge representation system) is not to be understood as consciously experienced. If it were, the book's subtle arguments would be quite superfluous!

It appears to me that this history has led to a continuing tension within traditional 'symbolic' cognitive science between seeing the computational symbol structures as equivalent to conscious thoughts, and seeing them as modeling processes that go on below the conscious level. The former tradition allows researchers to see and present themselves as working on explaining the mind in the sense in which it is pre-theoretically understood by most of us (i.e., as conscious mind). However, the latter view gives theorists and programmers much more freedom to apply their ingenuity, and it is the only possible view when one is tackling processes whose underlying mechanisms clearly are not conscious, such as perceptual processes. Thus, in practice, most symbolic programmers today probably give little attention to whether the symbols they work with are reasonable candidates for models of conscious contents of human minds. I want to suggest that they are wise to ignore this issue. It was a mistake from the first to think that such computational symbols, embodied, at the physical level, by the movements of electrons in silicon or ions in the brain, would somehow quicken and glow with refulgent consciousness<sup>(1)</sup> (or even intentionality - Cummins, 1989, 1996), once an intelligently behaving system was achieved.

This is not to say (I do not say) that such an artificial computational system might not actually be conscious, just that it is a mistake to expect the computational symbols it manipulates to be the very things of which it is conscious. It should that this mistake has been made by critics of AI at least as much as by its proponents. It is, surely, Searle's mistake in his notorious *Chinese Room* argument (1980). Searle correctly realizes that however intelligently a program behaves, and however intimate he may become with the symbols it manipulates, and their vicissitudes, he will never see them glow with the light of intrinsic intentionality<sup>(2)</sup>. He concludes that such a program can never constitute a mind. Others, similarly disappointed by the failure to find the glow even in the brain, conclude that understanding consciousness is a problem too "hard" for cognitive science (or any other sort of natural science)<sup>(3)</sup>. Yet others bang the table and beg to differ. They are all looking for intentionality and consciousness in the wrong place.

In fact, the practice of referring to computational symbols in cognitive systems as *mental representations* has been severely misleading. Once this point is taken on board, it should be apparent that non-representational cognitive theories are no worse off than are (properly interpreted) symbolic theories when it comes to explaining consciousness. None of them do so directly. That does not mean that they are not relevant. I believe that Turing's underlying insight was correct: if you can get the behavior right then, to all intents and purposes you have created a conscious mind. The proper task of Artificial Intelligence research is getting the behavior right. But what constitutes getting it right? Certainly not just fooling people in a teletyped conversation. In what follows I will attempt to sketch a theory of how conscious content, specifically mental imagery, might be explained in terms of certain *behaviors*. Imagery (which I take to include verbal imagery, inner speech) is the quintessential conscious thought content. What follows will focus mainly on vision and visual imagery, but with the understanding that equivalent considerations apply to all modalities. The theory draws upon work in situated robotics to some extent, but it is ultimately neutral about the engineering problem of how to get a machine to behave in the relevant ways, and about the biological problem of how to best explain how our brains get us to behave in the relevant ways. Symbolicists, dynamicists, connectionists, and the rest can continue to duke that out.

\*\*\*\*\*

I should probably say here that, given the available length for this presentation, the theory of imagery I advocate, what I call *Perceptual Activity Theory*, can only be sketched in the barest outline. A more detailed account can be found in my forthcoming article in *Cognitive Science* (Thomas, in press). There I also argue for the theory on rather different grounds, and give a fairly extensive empirical and conceptual critique of the better-known rival theories.

But any workable theory of mental imagery is going to be parasitic upon a theory of perception. Almost certainly, they share mechanisms to a considerable degree. Not only is there a good deal of empirical evidence to support this (Farah, 1988; Kosslyn, 1994)<sup>(4)</sup>; a phenomenological resemblance to perceptual experience is, surely, criterial for imagery (Finke, 1989; Thomas, 1997). Since the time of Alhazen visual perception has been seen as a matter of getting representations into the head (Lindberg, 1976). Originally, of course, these representations were optical images, and the problem was to understand how these got into eyes. However, once that process ceased to be mysterious, it became apparent that only the first step had been taken toward a full scientific understanding of vision. The almost universal response has been to try to understand how the meaningful content in the optical image gets itself further into the head, and is transformed into a more cognitively useful representational form. Thus we find the following given as a textbook *definition* of computer vision research:

**Computer vision** is the construction of explicit, meaningful descriptions of

physical objects from images. (Ballard & Brown, 1982)<sup>(5)</sup>

From this point of view, percepts are inner representations, and we should thus expect mental images to be the same. The principal cognitive science debate about imagery, the so-called analog-propositional dispute, has been about the proper *format* for such representations. Pylyshyn (1978, 1981) argues that a discursive symbolic description is sufficient to account for the phenomenology and the experimental findings concerning imagery; Kosslyn (1980) argues that it is *not* sufficient, and that the symbolic description must be translated into a "quasi-pictorial" format when imagery is accessed.

Clearly if we are looking for a non-representationalist account of the mind (as the first part of this paper suggests we should) we will also need a different view of perception. Fortunately, there is an alternative at hand, already closely allied with the non-representational robotics movement (Scassellati, 1998). The so called **Active Vision** approach to machine perception (see, e.g., Bajcsy, 1988; Ballard, 1991; Aloimonos, 1992; Swain & Stricker, 1993; Landy, Maloney, & Pavel, 1996)<sup>(6)</sup> takes the primary task of perception to be to provide for the intelligent control of behavior in the environment, and it rejects the requirement of the building of inner descriptions in favor of a view wherein:

Visual sensory data is analyzed purposefully in order to answer specific queries posed by the [robot] observer. The observer constantly adjusts its vantage point in order to uncover the piece of information that is immediately most pressing. (Blake & Yuille, 1992).

In order to behave intelligently *in its current situation* a robot may need to determine certain specific, behaviorally relevant facts about its environment - Is there a clear path *that way*? How much further will I need to extend my arm in order to grasp *this* object? Etc. - and it is equipped with perceptual transducers in order to give it the capability of answering such questions. The transducers are actively deployed to find the answers. The machine does not first build up an inner representation of the environment and then query that; rather it queries the environment itself, *when it needs specific information*. The world is actively explored rather than passively registered. Furthermore, the answers obtained do not amount to representations in the traditional sense, but may rather be just a simple YES or NO ("There is/isn't a clear path") or perhaps some parametric value ("Extend the arm one third of its reach"): answers that are not meaningful in a context free sense, but only in relation to the question that elicited them.

It is useful to think of this querying or exploring of the environment as the performance of measurements or tests, and to think of the sensory organs, the transducers, as instruments used to perform them. Or rather, we should say, the transducers, the sense organs, comprise parts of instruments. What they actually test for in any given circumstance depends not only on what sort of energies they are capable of transducing,

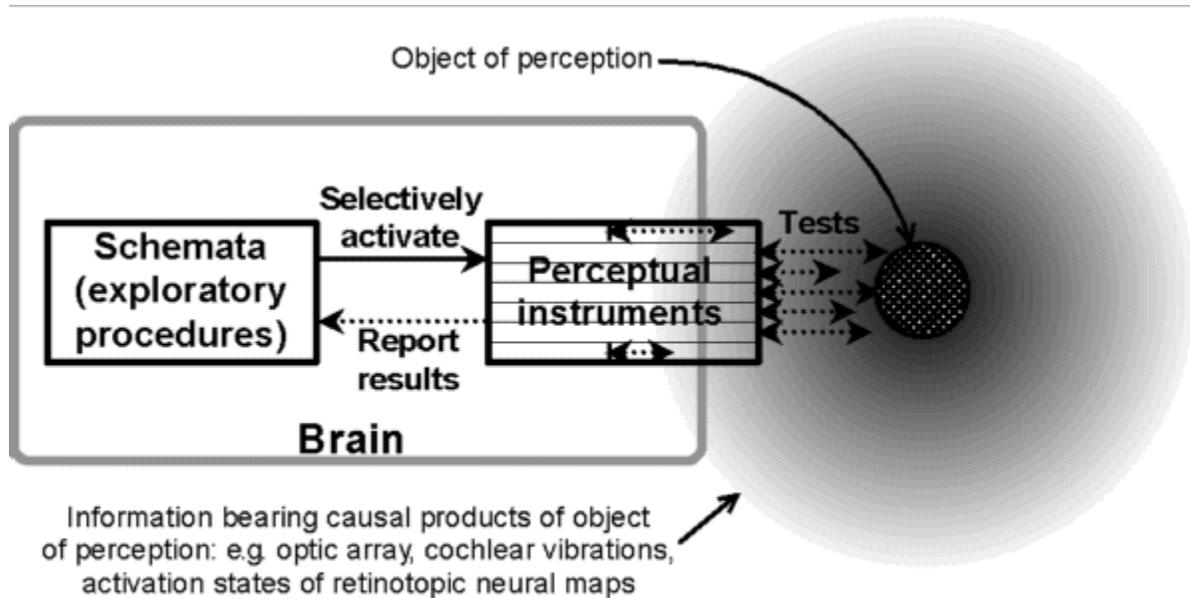
but also on how they are deployed: how they are oriented and moved relative to the environmental point of interest, how their sensitivity is dynamically modulated and calibrated, and how their output is analyzed. Thus an instrument comprises not just a sensor, but also the motor system (or musculature) that moves it, outward signal paths (or efferent nerves) that control these systems, and the central algorithm that controls this deployment and analyses the incoming signal. Thus a single sensor, according to the algorithm that is currently in control of it, may subserve *many* different perceptual instruments directed at ascertaining quite different sorts of facts about the environment. (Ballard (1991) calls this "sensor fission.")

Sometimes, as in haptic perception, a sensor may interact directly with the environmental object of interest. More often, perhaps, it will interact with some reliably correlated causal product of it: thus, visual instruments interact directly with structural features of the "optic array" of light (Gibson, 1979), but the usefulness of this is that such features are reliable indicators of the layout of the more tangible environment. In some cases the correlated causal product might be within the perceiver's body (heat in the flesh as correlated with a nearby fire, for example, or sound induced vibrations in the cochlea) or even within its brain (the map of the retinal image in V1 may be an example), but it would be a serious mistake to think of them, merely in virtue of such location, as percepts, or meaningful, potentially conscious representations, or even as computationally manipulable symbols (Thomas, in press).

From this point of view, any robot or organism capable of a rich behavioral repertoire in a complex environment will need to have a large number of perceptual instruments at its disposal, even though it may not need many types of actual sensor (five should be about enough!). But with large numbers of instruments, coordinating their use becomes an important issue. A sensor employed in making one sort of test may not always be simultaneously available for making another, and each instrument also draws on a necessarily limited pool of computational resources for control and analysis functions. It becomes important to deploy them strategically, using each instrument when and where its deployment is most likely to bring in currently useful information. Thus some sort of overall control program for the perceptual apparatus as a whole will be required. In a robot or organism capable of learning to improve its functioning in the environment this control system will need to be modifiable in the light of experience, revising its strategic plans for the deployment of perceptual instruments according to whether or not particular strategies have proved efficient or have led to desirable outcomes in the past. In a sense, then, the control system can be said to embody expectations about the world, but expectations that are constantly modified in the light of actual findings. Following Neisser (1976), I call such a self-modifying control structure a *schema*.

Note, however, that the schema is *not* to be identified with the conscious percept. Although it may be being constantly updated from perceptual input, its job is not to represent the world to *us* (to our selves, our consciousness, "higher level systems," our Cartesian minds, whatever), but to control the effective deployment of our perceptual

instruments. This is the behavior that needs to be got right for perceptual consciousness to arise. For present purposes it hardly matters how this control function is realized or explained: whether in terms of dynamical systems theory, neural networks, "traditional" symbolic programming, or whatever. That is pragmatic issue for engineers or biologists. The only caveat is that it is important not to think of any computational symbols involved in the workings of the schema as being the mental representations of thought itself. The suggestion is that perceptual consciousness consists not in the internal presence of any sort of model of the outside world, but in the complex process of interaction between the schema, the array of perceptual instruments, and the environment itself.



**Figure 1**  
(Dotted arrows represent processes inhibited during imagery)

The general form of such a [schema](#) controlled perceptual system is shown in figure 1. The two-headed dotted arrows toward the right represent the perceptual tests that can be made upon the perceived object itself, or its correlated causal products, querying the environment and returning results. The ongoing cyclical interaction between the schema, the array of instruments, and the environment is intended as equivalent to the "perceptual cycle" of Neisser (1976), and is hypothesized to be the material basis of conscious experience. The schema is responsible for sending out the orders, as it were, for the making of particular perceptual tests. However, it seems plausible that it might be possible for this process sometimes to become decoupled from actual perceptual testing. If the operation of the perceptual instruments were inhibited, or if their actual output were ignored (or partially ignored) in favor of the expectations embodied in the schema, then the cyclical process could become decoupled from the current environment, and the experience it produced would reflect memories previously laid down in the schema rather than the structure of the current environment. This, I suggest, is the material basis of imagery experience. If the schema is busily sending out orders to make the tests that would confirm the presence of a cat, no confirmatory results come back to it, yet it

continues sending out cat-relevant orders anyway, then we are imagining a cat. Such activity has intentionality because it is directed at an object type (cats), and explaining such intrinsic intentionality is at least the first step to explaining consciousness. There will, of course, be a considerable (although never complete) degree of overlap between the physiological processes involved in this imagery and those involved in actually perceiving a cat, so it should be no surprise that the experiences are also significantly similar.

Provided this decoupled perceptual activity is not confused with actual perception, it can serve a very useful cognitive function. It makes it possible to be conscious of, to think about, things that are not actually present to sense. It makes it possible to imagine, to daydream, to plan: it is the substrate of conscious thought (Ellis, 1995). In language using creatures, imagined speech provides an especially powerful medium for complex and wide ranging thinking. But all this imagining is ultimately a form of behavior, albeit highly complex and largely covert behavior. If this paper is on the right track then filling in the details of this story and explaining how the relevant behavior is controlled and coordinated remains a large task for cognitive science. However, it is a task that seems likely to be amenable to the sorts of intellectual and experimental tools that cognitive scientists have already developed. Despite what some people may think, cognitive science may yet have the resources to explain the conscious mind.

---

## References

- Akins, K. (1996). Of Sensory Systems and the "Aboutness" of Mental States. *Journal of Philosophy* 91: 337-372.
- Aloimonos, Y. (Ed.). (1993). *Active perception*. Hillsdale, NJ: Erlbaum.
- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, 76, 996-1005.
- Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48, 57-86.
- Ballard, D. H. & Brown C. M. (1982). *Computer Vision*. Englewood Cliffs, NJ: Prentice-Hall.
- Blake, A., & Yuille, A. (Eds.). (1992). *Active vision*. Cambridge, MA: MIT Press.
- Boden, M. A. (1977). *Artificial Intelligence and Natural Man*. Hassocks, U.K.: Harvester.
- Brooks, R. A. (1991). Intelligence Without Representation. *Artificial*

*Intelligence* (47) 139-159.

Chalmers, D.J. (1996). *The conscious mind*. Oxford: Oxford University Press.

Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1994). A critique of pure vision. In C. Koch & J. Davis (Eds.), *Large scale neuronal theories of the brain*. Cambridge, MA: MIT Press.

Clancey, W. J. (1997). *Situated cognition*. Cambridge: Cambridge University Press.

Cotterill, R. M. J. (1997). On the mechanism of consciousness. *Journal of Consciousness Studies*, 4, 231-247.

Cummins, R. (1989). *Meaning and mental representation*. Cambridge, MA: MIT Press.

Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.

Eliasmith, C. (1996). The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition. *Philosophical Psychology* 9: 441-463.

Ellis, R.D. (1995). *Questioning Consciousness: the interplay of imagery, cognition, and emotion in the human brain* Amsterdam: John Benjamins.

Ericsson, K. A. & Simon, H. A. (1980). Verbal Reports as Data. *Psychological Review* (87) 215-251.

Farah, M. J., (1988). Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review*, 95, 307-317.

Finke, R. A. (1989). *Principles of mental imagery*. Cambridge, MA: MIT Press.

Fodor, J. A. (1975). *The Language of Thought*. New York: Thomas Crowell.

Freeman, W.J. & Skarda, C.A. (1990). Representations: Who Needs Them? In: J.L. McGaugh, N.M. Weinberger & G. Lynch (eds.). *Brain Organization and Memory*. New York: Oxford University Press. pp. 375-380.

Garson, J.W. (1996). Cognition Poised at the Edge of Chaos: A Complex Alternative to a Symbolic Mind. *Philosophical Psychology* 9: 301-322.

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston,

MA: Houghton Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Kosslyn S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.

Kosslyn S. M. (1994). *Image and brain*. Cambridge, MA: MIT Press.

Landy, M. S., Maloney, L. T., & Pavel, M. (Eds.). (1996). *Exploratory vision: The active eye*. New York: Springer-Verlag.

Lindberg, D. C. (1976). *Theories of Vision from Al-Kindi to Kepler*. Chicago: University of Chicago Press.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

McCorduck, P. (1979). *Machines Who Think*. San Francisco: Freeman.

Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Pylyshyn, Z. W. (1978). Imagery and artificial intelligence. *Minnesota Studies in the Philosophy of Science*, 9, 19-55.

Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, 88, 16-45.

Ryle, G. (1949). *The concept of mind*. London: Hutchinson.

Scassellati, B. (1998). A Binocular, Foveated Active Vision System. *MIT-AI Memo 1628*. [WWW document] URL <http://www.ai.mit.edu/projects/humanoid-robotics-group/cog/cog-publications/scaz-3heads.pdf>

Searle, J. R. (1980). Minds, Brains and Programs. *Behavioral & Brain Sciences* (3) 417-424.

Searle, J. R. (1992). *The Rediscovery of the Mind*. Cambridge MA: MIT Press.

Smith, B. C. (1991). The Owl and the Electric Encyclopedia. *Artificial Intelligence* (47) 252-288.

Swain, M. J., & Stricker, M. A. (1993). Promising directions in active vision. *International Journal of Computer Vision*, 11, 109-126.

Thomas, N. J. T. (1997). Mental Imagery. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* [Online serial] URL <http://plato.stanford.edu/entries/mental-imagery/>

Thomas, N. J. T. (in press). [Are Theories of Imagery Theories of Imagination? An Active Perception Approach to Conscious Mental Content.](#) *Cognitive Science*.

Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind* (59) 433-460.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Van Gelder, T. (1995). What Might Cognition Be, If Not Computation? *Journal of Philosophy* 92: 345-381.

Van Gelder, T. & Port, R. (eds.) (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.

---

## Notes

1. This, of course, is an allusion to Ryle's (1949: 6.2) seminal remarks on the absurdity of inner conscious objects.
2. Searle (1992), rightly in my view, regards intrinsic intentionality and consciousness as intimately related concepts.
3. Although what I say here does not adequately represent his arguments, this popular position is most powerfully articulated in the much-admired work of Chalmers (1996).
4. At least in the case of *visual* imagery, which has received by far the most attention.
5. See also Marr (1982). In this context, "images" means optical images: the TV camera equivalent of retinal images. In the fast paced world of AI, 1982 may seem a long time ago, but, in fact, a very similar definition of "machine vision" is offered today in the (as yet still in draft) online version of the *MIT Encyclopedia of Cognitive Science*. Note that the descriptions in question are supposed to be *meaningful*: i.e. the computational symbols that comprise them are supposed to have (or, at least, model) intentionality.

6. This work has an important precursor in the perceptual theory of Gibson (1966, 1979). It also finds significant parallels in recent biologically inspired thinking about perception (Churchland, Ramachandaran, & Sejnowski, 1994; Akins, 1996; Cotterill, 1997). The formative influence on my own thinking about imagery was the Gibson inspired work of Neisser (1976).

Thanks to Steve Lehar for helpful discussion.

---

---

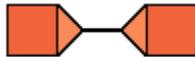
Please post comments & new link suggestions to the  
discussion board:  
(I may post a reply and/or send back email, as  
appropriate.)



or you can [email](#) me.

---

**[Return to Home Page:](#)**  
**[Imagination, Mental Imagery, Consciousness, Cognition:](#)**  
**[Science, Philosophy & History.](#)**



Last revised: May 23, 2002.

---

(See also Non-Conscious, below). Non-Conscious It is widely recognised that we are not conscious of much of what goes on in our brains. Book-length discussions include Eagleman, 2011, and Wilson, 2002.<sup>3</sup> The term non-conscious is used in this thesis to refer to all consciously inaccessible mental processes. According to Dreyfus, the consequence of GOFAI's symbolic representation is a world of innumerable isolated facts, related only as stipulated by the programmer and therefore inflexible with regard to context and purpose. The Computationalist model bears little or no resemblance to the world as we experience it.<sup>16</sup> From the late 1970s onwards, GOFAI research did indeed encounter a number of problems. of conscious visual imagery in light of the fact that its underlying neural substrates and mechanisms are distinct from those of conscious visual perception. and imagery is not reflected in the phenomenology of the states but is reflected in the manner in which perception and imagery represent the scene (perceptually, versus imagistically). symbolic linguistic representations. An alternative to Kosslyn's view of "mental pictures" and Pylyshyn's symbolic, linguistic representations takes mental images to represent the layout of perceived space, in Gibson's sense (see Neisser and Kerr, 1973). Perceptual activity theory presents imagery as non-discursive and relates it closely to seeing as. It is thus well placed to be the basis for a general theory of imagination and its role in creative thought. It rejects the "traditional" symbolic computational view of mental contents, but is compatible with recent situated cognition and active vision approaches in robotics. This theory is developed and elucidated. Three related key aspects of imagination (non-discursiveness, creativity, and seeing as) raise difficulties for the other theories. Perceptual activity theory presents imagery as non-discursive and relates it closely to seeing as. It is thus well placed to be the basis for a general theory of imagination and its role in creative thought. \*