



Original Article

Applying global workspace theory to the frame problem

Murray Shanahan^{a,*}, Bernard Baars^b

^a*Department of Electrical and Electronic Engineering, Imperial College, Exhibition Road, London SW7 2BT, UK*

^b*The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, USA*

Received 5 July 2004; revised 20 October 2004; accepted 21 November 2004

Abstract

The subject of this article is the frame problem, as conceived by certain cognitive scientists and philosophers of mind, notably Fodor for whom it stands as a fundamental obstacle to progress in cognitive science. The challenge is to explain the capacity of so-called informationally unencapsulated cognitive processes to deal effectively with information from potentially any cognitive domain without the burden of having to explicitly sift the relevant from the irrelevant. The paper advocates a global workspace architecture, with its ability to manage massively parallel resources in the context of a serial thread of computation, as an answer to this challenge. Analogical reasoning is given particular attention, since it exemplifies informational unencapsulation in its most extreme form. Because global workspace theory also purports to account for the distinction between conscious and unconscious information processing, the paper advances the tentative conclusion that consciousness may go hand-in-hand with a solution to the frame problem in the biological brain. © 2005 Elsevier B.V. All rights reserved.

Keywords: Frame problem; Relevance problem; Global workspace theory; Consciousness; Analogical reasoning

1. Introduction

The frame problem was originally couched as a difficulty within classical Artificial Intelligence: How can we build a program capable of inferring the effects of an action without reasoning explicitly about all its obvious non-effects? But many philosophers saw

* Corresponding author. Tel.: +44 207 594 6307.

E-mail addresses: m.shanahan@imperial.ac.uk (M. Shanahan), baars@nsi.edu (B. Baars).

the frame problem as symptomatic of a wider difficulty, namely how to account for cognitive processes capable of drawing on information from arbitrary domains of knowledge or expertise. So-called “informationally unencapsulated” processes of this sort, exemplified by analogical reasoning, are especially troublesome for theories of mind that rely on some sort of modular organisation to render them computationally feasible.

However, one thing is clear. If the frame problem is a genuine puzzle, the human brain incorporates a solution to it. In global workspace theory, we find clues to how this solution might work. Global workspace theory posits a functional role for consciousness, which is to facilitate information exchange among multiple, special-purpose, unconscious brain processes (Baars 1997, 1998). These compete for access to a global workspace, which allows selected information to be broadcast back to the whole system. Such an architecture accommodates high-speed, domain-specific processes (or “modules”) while facilitating just the sort of crossing of domain boundaries required to address the philosophers’ frame problem.

The paper is organised as follows. In Sections 2 and 3, the philosophers’ conception of the frame problem is presented. Section 4 challenges the premise that informationally unencapsulated cognitive processes are, in principle, computationally infeasible. In Section 5, global workspace theory is outlined. Arguments and evidence in favour of the theory are reviewed, and the global workspace architecture is commended as a model of combined serial and parallel information flow capable of overcoming the frame problem.

Section 6 concerns analogical reasoning, the epitome of informational unencapsulation, and demonstrates that the most successful of contemporary computational models of analogical reasoning are strikingly compatible with global workspace theory. The concluding discussion addresses a variety of topics including modularity, conscious information processing, and the relationship between parallel and serial computation in a generic account of cognitive function.

2. The frame problem

The frame problem, in its original form, was to address the following question (McCarthy & Hayes, 1969). How is it possible to write a collection of axioms in mathematical logic that captures the effects of actions, without being obliged to include an overwhelming number of axioms that describe the trivial non-effects of those actions? In everyday discourse, we can describe the effect of, say, painting an object simply by detailing how its colour changes. There is no need to state explicitly that the object’s shape remains the same, that the neighbour’s cat remains asleep, that the Sun continues to shine, and so on. This is all taken for granted by common sense.

In mathematical logic, by contrast, nothing is taken for granted. Not only is it necessary to make explicit the changes an action brings about, it is also necessary to make explicit the things that do *not* change, and for most actions this will be a great deal. This was particularly troubling for AI researchers in the 1970s and 1980s who wanted to write programs that worked by carrying out inference with logical formulae. The burden of having to supply a set of explicit formulae describing every non-effect of every action seemed overwhelming, both representationally and computationally.

In the intervening decades, the frame problem in this narrow guise has receded. Researchers in logic-based AI—who have become something of a minority anyway—more-or-less solved the problem in the 1990s using a combination of judicious representation techniques and non-monotonic forms of inference (Reiter, 2001; Sandewall, 1994; Shanahan, 1997).¹ Yet the frame problem lives on in the minds of cognitive scientists. This is largely due to the wider interpretation given to it by certain philosophers of mind, especially Fodor (1983, 1987, 2000).

The credit (or blame) for first bringing the frame problem to the attention of philosophers of mind and cognitive scientists goes to Dennett, who characterised it as the challenge of understanding how “a cognitive creature...with many beliefs about the world” can update those beliefs when it performs an action so that they remain “roughly faithful to the world” (Dennett (1978, p. 125). In *The Modularity of Mind*, Fodor invites the reader to consider a hypothetical robot, and poses a related question: “How...does the machine’s program determine which beliefs the robot ought to re-evaluate given that it has embarked upon some or other course of action?” (Fodor, 1983, p. 114).

Dennett (1984) highlights the issue with a memorable example. He asks us to consider the challenge facing the designers of an imaginary robot whose task is to retrieve an object resting on a wagon in a nearby room. But the room also contains a bomb, which is timed to explode soon. The first version of the robot successfully works out that it must pull the wagon out of the room. Unfortunately, the bomb is on the wagon. And although the robot knows the bomb is on the wagon, it fails to notice that pulling the wagon out brings the bomb along too.

So the designers produce a second version of the robot. This model works out all the *consequences* of its actions before doing anything. But the new robot gets blown up too, because it spends too long in the room working out what will happen when it moves the wagon.

It had just finished deducing that pulling the wagon out of the room would not change to color of the room’s walls, and was embarking on a proof of the further implication that pulling the wagon out would cause its wheels to turn more revolutions than there were wheels on the wagon—when the bomb exploded.²

So the robot builders come up with a third design. This robot is programmed to tell the difference between *relevant* and *irrelevant* implications. When working out the consequences of its actions, it considers only the relevant ones. But to the surprise of its designers, this version of the robot fares no better. Like its predecessor, it sits in the room “thinking” rather than acting.

“Do something!” they yelled at it. “I am,” it retorted. “I’m busily ignoring some thousands of implications I have determined to be irrelevant. Just as soon as I find an irrelevant implication, I put it on the list of those I must ignore, and...” the bomb went off.³

¹ For an overview see Shanahan (2003).

² Dennett (1984), p. 129.

³ Dennett (1984), p. 130.

In Fodor's words, this robot suffers from "Hamlet's problem: when to stop thinking", and the frame problem is "Hamlet's problem viewed from an engineer's perspective" (Fodor, 1987, p. 140). Significantly, the problem of "when to stop thinking" arises not only when anticipating the consequences of an action. It is a difficulty that besets any *informationally unencapsulated* inference process, that is to say any process for which no *a priori* limit exists to the information that might be pertinent to it. Exemplary among such processes is *analogical reasoning*, "which depends precisely upon the transfer of information among cognitive domains previously assumed to be irrelevant" (Fodor, 1983, p. 105).

3. The computational theory of mind

The concern of this paper is the frame problem in the wide sense intended by Fodor.⁴ Is the frame problem in this sense a real problem, and if so how is it solved in the human brain? If we grant that some mental processes, such as analogical reasoning, are indeed informationally unencapsulated, the first question to be addressed is why cognitive scientists like Fodor consider this property so threatening.

For Fodor and his sympathisers, the frame problem inherits its significance from the Computational Theory of Mind (CTM). According to CTM, cognitive processes are computations. Specifically, they are truth-preserving computations over representations that have a combinatorial syntax. The Computational Theory of Mind is stronger than the claim that the input-output function of every human brain is equivalent to some Turing machine. As Fodor puts it, CTM entails that "the mind is interestingly like a Turing machine" (Fodor, 2000, p. 30).

The frame problem is viewed as a problem for CTM because informationally unencapsulated processes are thought to be computationally infeasible.⁵ Here's Fodor again.

The totality of one's epistemic commitments is *vastly* too large a space to have to search if all one's trying to do is figure out whether, since there are clouds, it would be wise to carry an umbrella. Indeed, the totality of one's epistemic commitments is vastly too large a space to have to search *whatever* it is that one is trying to figure out.⁶

This sort of view is pervasive, and by no means confined to Fodor. Here is Carruthers on the same theme.

⁴ Controversy over the use of the term "frame problem" has led to some acrimonious debate between AI researchers and philosophers (Pylyshyn, 1987). The present paper takes the view that the philosophers have pinpointed an interesting and well-defined question relating to informationally unencapsulated cognitive processes, even if this issue bears only a slight resemblance to the original AI researchers' frame problem.

⁵ The term computationally "infeasible" rather than "intractable" is used here because intractability has a mathematically precise meaning in the context of computational complexity, which we will come to shortly. Philosophers writing on the frame problem are typically informal in their use of such terms.

⁶ Fodor (2000), p. 31.

The computational processes that realize human cognition will need to be tractable ones, of course; for they need to operate in real time with limited computational resources ...And any processor that had to access the full set of the agent's background beliefs (or even a significant sub-set thereof) would be faced with an unmanageable combinatorial explosion.⁷

Famously, such considerations lead Fodor to propose a distinction between the mind's *peripheral processes*, which are supposedly informationally encapsulated and therefore computationally feasible, and its *central processes*, which are informationally unencapsulated and therefore computationally infeasible. The mind's peripheral processes—which do only “moderately” interesting things, like parsing and early vision—can be understood as a system of *modules*, and are amenable to study by cognitive scientists. On the other hand, the mind's central processes—which do all the really interesting and important things like belief revision and analogical reasoning—because they are computationally infeasible are altogether beyond understanding for contemporary cognitive science (Fodor, 1983, 2000).

...it probably isn't true, at least in the general case, that cognitive processes are computations. ...[so] it's a mystery, not just a problem, what model of the mind cognitive science ought to try next.⁸

Fodor's prognosis is dire. But his argument rests on several questionable assumptions. The first is that the Computational Theory of Mind is the only viable research methodology in cognitive science today, “the only game in town” as he put it.⁹ This assumption has been vigorously challenged, and just as robustly defended. But the target of the present paper is the assumption that informationally unencapsulated processes are indeed computationally infeasible, a claim that can be detached from CTM. As Haselager and van Rappard (1988) point out, if we accept the need to account for the “systematicity” of thought (Fodor & Pylyshyn, 1988), then the frame problem is equally an issue for a connectionist theory of mind. The following section deals specifically with Fodor's argument, so its notion of a “cognitive process” is inherited from CTM. But the solution ultimately offered here, based on global workspace architecture, admits a wider variety of forms of computation.

4. Complexity and informational encapsulation

There is no doubt, of course, that some tasks *are* computationally intractable, in a sense that has been made mathematically precise (Garey & Johnson, 1979). To sharpen the discussion, it is worth reviewing the basic computer science. Consider a function F . Suppose it can be proved that an algorithm exists that, for any input string x of length n ,

⁷ Carruthers (2003), p. 505. Although agreeing with Fodor's diagnosis, Carruthers is not sympathetic to his prognosis, and puts forward a mechanism for tackling the problem not dissimilar to the one on offer in the present paper.

⁸ Fodor (2000), p. 23.

⁹ Fodor (1975), p. 406.

can compute $F(x)$ in less than or equal to $T(n)$ steps. So T sets an upper bound on how long the computation will take, in the general case. The rate of growth of T can succinctly be characterised by the dominant term in T , dropping other terms and coefficients. For example, if $T(n) = 5n^2 + 3n + 7$, then the growth of T is characterised by n^2 , and we say that the task of computing F is $O(n^2)$, pronounced “order n^2 ”. If a computational task is $O(n)$, then it is easy, from a computational point-of-view. If a task is $O(n^2)$, then it is not so easy. Even so, any task that is $O(n^k)$ for some constant k is said to be *tractable*. The class of all such problems is denoted P.

Now, there exists an important class of problems for which the best known algorithms are $O(2^n)$, even though polynomial-time algorithms exist for verifying the correctness of a given solution to the problem. This class of problems is denoted NP, and any task that falls into this class is said to be computationally *intractable*. The best known of these is the so-called SAT problem, which is the task of showing that a given expression of propositional logic (of a certain form) is satisfiable. A typical strategy for proving that a task belongs to NP is to show that it can be reduced to the SAT problem in polynomial time. Any such problem is said to be *NP-complete*.

NP-complete problems, though hard from a computational point-of-view, are not the hardest. Beyond NP, we have a class of problems for which it can be proved that no algorithm exists at all that is guaranteed to produce a solution. Such problems are said to be *undecidable*. Examples include the Halting Problem for Turing machines and theorem proving in first-order predicate calculus.

It is generally taken to be a bad thing if a computer programmer becomes entangled with an NP-complete problem, and worse still if they end up wrestling with an undecidable one. Despite this, computer scientists (especially AI researchers) routinely confront both NP-complete and undecidable problems. As far as intractability is concerned, if n is guaranteed to be small, an exponential algorithm is not so worrying. Moreover, it should be remembered that these complexity results are worst-case analyses. Individual instances of an NP-complete problem are frequently soluble in reasonable time, even for large n , especially if clever heuristics are used. Additionally, various practical techniques, including *resource bounded* computation (Russell & Wefald, 1991) and *anytime algorithms* (Zilberstein & Russell, 1993), can be used to mitigate the effects of an unfavourable complexity result in a working system.¹⁰

Now let's return to the supposition that informationally unencapsulated cognitive processes are computationally infeasible. Does this claim stand up to proper scrutiny? In particular, are such processes intractable in the theoretically precise sense? First, let's reconsider the fact that, according to the computational theory of mind, cognitive processes are “truth-preserving computations over representations that have a combinatorial syntax”. In other words, cognitive processes prove theorems in some formal logic. Moreover, for CTM to be capable of explaining the systematicity of thought (Fodor & Pylyshyn, 1988), the logic in question must be at least first order.¹¹ Since it is known that

¹⁰ The inspiration behind these techniques is the concept of bounded rationality introduced in (Simon, 1957).

¹¹ According to (Fodor & Pylyshyn, 1988), systematicity is a property of relational representations. Therefore propositional calculus, though decidable, would be inadequate.

theorem proving is undecidable for first-order predicate calculus, it is natural to ask whether this is the reason informationally unencapsulated cognitive processes are alleged to be infeasible. But this cannot be the basis of the allegation, because the same observation applies to supposedly untroublesome informationally encapsulated cognitive processes. If these also carry out first-order logical inference, they too are undecidable, regardless of the problem size.

Therefore the real concern over computational feasibility is not an accompaniment to the presumption that first-order logical inference is taking place.¹² The real worry, in the context of the frame problem, seems to be that the set of sentences having a potential bearing on an informationally unencapsulated cognitive process is simply “too big”. This suggests that the issue is not really tractability, in the proper sense of the term, for tractability is to do with the rate at which computation time grows with the size of the problem. Rather, the difficulty seems to be the upper bound on n , not the upper bound on $T(n)$, where the n in question is the number of potentially relevant background sentences.

But why should a large such n be a problem? Recall Fodor’s claim that “the totality of one’s epistemic commitments is vastly too large a space to have to search”. Perhaps the perceived difficulty with a large n is the supposed cost of carrying out an exhaustive search of n items of data. Yet plenty of algorithms exist for search problems that have very favourable computational properties. For example, the task of searching a balanced ordered binary tree for a given item of data is $O(\log_2 n)$. Other data structures for search have even better complexity results. This is why Internet search engines can find every website that mentions a given combination of keywords in a fraction of a second, even though they have records for several billion web pages.

Similarly, take Carruthers’ assertion that a process that has to “access the full set of the agent’s background beliefs ... would be faced with an unmanageable combinatorial explosion”. It should be clear by now that the truth of this claim depends entirely on the nature of the process in question. As long as this is not spelled out, the purported infeasibility of informationally unencapsulated cognitive processes remains in question. If isolating relevant information from a mass of background beliefs is a matter of searching a very large data structure, then there is no combinatorial problem.

To see how it might not be merely a matter of searching a large data structure, let’s introduce a distinction between *explicit* belief that is “near the surface” and easily accessible to cognition and the mass of *implicit* potential belief buried deep in the logical consequences of the set of explicit beliefs. Perhaps explicit beliefs can indeed be thought of as stored in a large data structure, and as therefore amenable to efficient search. But it seems reasonable to suppose there could be an implicit belief relevant to an ongoing problem while none of the explicit beliefs entailing it had anything about them to indicate this. How could an informationally unencapsulated process find this implicit belief without having to visit all the logical consequences of every explicit belief?

¹² Exactly why supporters of CTM are not worried by this issue is hard to fathom. Perhaps they are impressed by the armoury of techniques developed by AI researchers to contain the problem. If so, they never seem to mention the fact.

Of course, it would be ridiculous to suppose that just because someone believed *P* they also believed all the logical consequences of *P*, and it's equally ridiculous to suppose human cognition will always successfully exploit all the relevant implications of explicit belief (Cherniak, 1986).¹³ Human beings are fallible, and this is one among many reasons why. But this observation is a red herring. The frame problem is an issue even allowing for the limits of human cognition (Fodor, 2000, Ch. 2). The question is how human beings *ever* manage to select relevant information from a large mass of candidate beliefs in arbitrary domains and bring it to bear on a novel problem. The issue of how far into the consequences of explicit belief a cognitive process can feasibly penetrate is orthogonal. So it cannot be a cornerstone of the computational infeasibility thesis.

Perhaps Fodor and other like-minded cognitive scientists have been led astray by Dennett's caricature of a ploddingly stupid robot trying not to get blown up. Recall that Dennett's imaginary robot explicitly considers each irrelevant implication of its actions before it decides to ignore it. The robot's supposed difficulty is that it doesn't know how to stop thinking. But the design of Dennett's robot is absurd. It carries out a serial computation that exhaustively works through a long list of alternatives one-by-one before it terminates. In the following section, a rather different computational model is brought to bear on the frame problem.

5. Global workspace theory

The discussion of the previous section suggests that a convincing case for the computational infeasibility of informationally unencapsulated cognitive processes has not been made. Proponents of the infeasibility thesis are insufficiently rigorous in their treatment of algorithmic complexity and are unsuccessful in demonstrating that computational problems follow from the nature of the cognitive processes in question. So it is legitimate to regard the existence of such processes as a problem rather than a mystery. Yet it is an injustice to these authors to deny the intuitive force of their challenge. The task remains to explain how an informationally unencapsulated process might be realised in a biological brain. Since the human brain plainly does not use data structures like ordered binary trees, how does a cognitive process like analogical reasoning manage to select among all the information available to it?

However, as it stands this question is somewhat ill-posed. It betrays an assumption that it is the responsibility of the cognitive process itself to make the selection of relevant information. According to the proposal of the present section, a better question would be the following. How is it that the required relevant information is *made available to* a cognitive process like analogical reasoning? The answer offered here relies on the idea of distributing the responsibility for selecting relevant information among multiple parallel processes. The proposed means for achieving this is a *global workspace architecture* (Baars, 1988). The remainder of this section outlines global workspace theory, reviews

¹³ See Gabaix and Laibson (2003) for a discussion of this issue in the context of economics.




Fig. 1. The global workspace architecture.

the relevant empirical evidence, and draws attention to some open questions before returning to the computational issue at hand.

The essence of the global workspace architecture is a model of combined serial and parallel information flow. Multiple parallel *specialist processes* compete and co-operate for access to a *global workspace* (Fig. 1, left). A specialist process can be responsible for some aspect of perception, long-term planning, problem solving, language understanding, language production, action selection, or indeed any posited cognitive process. If granted access to the global workspace, the information a process has to offer is “broadcast” back to the entire set of specialists (Fig. 1, right). The means by which access is granted to the global workspace can be likened to an attention mechanism.

The contents of the global workspace unfolds in a serial manner. But it is the product of massively parallel processing. A sequence of moment-to-moment snapshots of the global workspace would reveal a meaningful progression, and each state would typically be related to its predecessor in a coherent way. Yet the state-to-state relation itself is highly complex, and could not be obtained by a single computational step on a conventional serial machine. Rather, it is the outcome of a selection procedure that has at its disposal the results of numerous separate computations, each of which might have something of value to contribute to the ongoing procession of thoughts.

According to the global workspace hypothesis, the mammalian brain is organised in this way, and using this architectural blueprint it is possible to distinguish between conscious and unconscious information processing. Unconscious information processing is carried out by the parallel specialist processes. Only information that is broadcast via the global workspace is consciously processed.

A considerable body of empirical evidence in favour of this thesis has accumulated in recent years (Baars, 2002; Dehaene & Naccache, 2001). In particular, the hypothesis that consciousness enables global activation has now been supported in dozens of brain imaging studies. The typical experimental paradigm is based on *contrastive analysis* (Baars, 1988), wherein closely matched conscious and unconscious conditions are

compared in waking subjects, either by stimulus manipulation (binocular rivalry, masking, attentional tasks, etc.) or by overpractice of automatic habits.

In all cases tested so far, the conscious condition recruits very widespread cortical resources while the matched unconscious event typically activates local regions only (Baars, 2002, 2003). For example, in an fMRI study by Dehaene et al. (2001) using visual backward masking, unconscious words activated visual word areas only, while the identical conscious words evoked widespread activation of parietal and prefrontal cortex as well. Kreiman et al. (2002) have shown that medial-temporal areas are also activated by pictures, but only if they were conscious, using implanted electrodes in epileptic patients. These regions are specialized for memory and emotion.

In a complementary experimental paradigm, brain response to stimulation has been compared in conscious versus unconscious states. Unconscious states studied include sleep, general anesthesia, epileptic loss of consciousness and vegetative states. Sensory stimulation in all four unconscious states evokes only local cortical responses, but not the global recruitment characteristic of sensory input in conscious subjects (Baars et al., 2003). This general pattern of results has now been shown for vision, hearing, pain perception, touch, and sensorimotor tasks. It appears that conscious events recruit global activity in the cerebral cortex, as predicted by the theory.

Possible neural mechanisms capable of supporting a global workspace are the subject of ongoing research (Dehaene & Naccache, 2001; Dehaene, Sergent, & Changeux, 2003). A naïve interpretation of the architectural proposal in Fig. 1 suggests an analogy with the computer scientist's notion of data transmitted along a bus and stored in a buffer. But this makes little sense at a neurological level. Rather, what is sought is a mechanism by which the activity of small selected portions of cortex—those that have gained access to the global workspace—can systematically influence the activity of the rest of cortex. The activity of a sending neuronal population will not typically be *duplicated* in a receiving population. A better way to understand the relationship between sender and receiver is in terms of *mutual information* (Tononi, Edelman, & Sporns, 1998).

The format of information in the global workspace is another open question. One possibility is that it reflects the underlying structure of the sensorimotor system. Under this assumption, the global workspace would include topographically-structured visual information and audio information with linguistic content. Barsalou (1999) has argued that perceptually-structured information is adequate for conveying conceptual content, while Carruthers (2002) has argued for natural language as a medium for integrating the output of heterogeneous modules.¹⁴ So a perceptual encoding may be sufficient to support the interaction between disparate computational elements required by an informationally unencapsulated process like analogical reasoning.

However, the central thesis of this paper is insensitive to the precise encoding of information in the global workspace. All the present proposal requires is a set of neurons (the global workspace) whose state is capable of simultaneously influencing widely distributed, functionally distinct neural sites, and for each of those sites to be capable in turn of contributing to that state. According to the hypothesis, this is enough to support

¹⁴ Dennett (1991, Ch. 7) assigns natural language a similar role in his notion of a “Joycean machine”.

a blend of serial and parallel computation with the capacity to sift the relevant from the irrelevant without incurring an impossible complexity burden. The interpretation of the content of the state in question is a separate matter, albeit an interesting and important one.

From an engineering standpoint, the global workspace architecture has a pedigree dating back to the earliest days of artificial intelligence. Its origins lie in Newell's *blackboard* metaphor for problem solving (Newell, 1962), which inspired the development of so-called blackboard architectures for AI systems in the 1980s (Hayes-Roth, 1985; Nii, 1986). Nowadays, the blackboard architecture is a standard piece of AI technology, and with the recent design of AI software agents based on global workspace theory, the influence of the blackboard architecture has come full circle (Franklin & Graesser, 1999).

The engineering advantages of a global workspace architecture gel with empirical considerations. First, it is easily deployed on parallel hardware, which is an obvious requirement for a neurologically plausible model of cognition. Second, it ensures that the overall system degrades gracefully when individual processes break down, which is compatible with the brain's robustness. Third, it facilitates the addition of new processes without modification to the existing system, which renders it consistent with evolutionary constraints. Fourth, it accommodates heterogeneous forms of information processing and integrates their results.


The final point is critical here, since it is precisely the ability to manage a high bandwidth of information generated by a large number of distinctive, special-purpose parallel processes that enables a global workspace architecture to realise an informationally unencapsulated cognitive process without falling foul of complexity problems. Notably, the concurrent activation of multiple disparate brain processes is also the signature of *conscious* as opposed to unconscious information processing according to the global workspace hypothesis, and this raises the tantalising possibility that consciousness goes hand-in-hand with a solution to the frame problem in the biological brain.

6. Analogical reasoning

Fodor says little about the computational model behind his claim that informationally unencapsulated cognitive processes are computationally infeasible. Yet there are strong hints of a commitment to a centralised, serial process that somehow has all the requisite information at its disposal, and then has the responsibility of choosing what information to access and when to access it. Although parallel peripheral processes are part of the picture, they are *passive* sources of information that wait to be called upon before delivering their goods (Fig. 2).¹⁵ This centralised, serial model of computation is evoked by Dennett's caricature robot, by the idea of "Hamlet's problem", and by such phrases as "the totality of one's epistemic commitments is vastly too large a space to have to search" (Fodor)

¹⁵ Fodor's peripheral processes perform both input and output. The emphasis of the present discussion is on perception.

Peripheral Processes (Modules)



Central Processes

Fig. 2. A naïve model of information flow.

and “access to the full set of an agent’s background beliefs gives rise to an unmanageable combinatorial explosion” (Carruthers).

By contrast, global workspace theory posits multiple, parallel processes that all contribute *actively* to cognition (Fig. 3). Consider the computational processes that might underlie the likening of a Rorschach inkblot to, say, an elephant. Fodor’s argument hints at a centralised process that poses a series of questions one-at-a-time—is it a face? is it a butterfly? is it a vulva? and so on—until it finally arrives at the idea of an elephant. Instead, the global workspace model posits a specialist, parallel process that is always on the lookout for elephantine shapes.¹⁶ This process is aroused by the presence of an inkblot that actually resembles an elephant, and it responds by announcing its findings. The urgency with which this process commends itself means that the information it has to offer makes its way into the global workspace, and is thereby broadcast back to all the other specialist processes.¹⁷

The Rorschach inkblot example is especially pertinent, since it is closely related to analogical reasoning. As Fodor emphasises, reasoning by analogy is the epitome of an informationally unencapsulated cognitive process because it “depends precisely upon the transfer of information among cognitive domains previously assumed to be irrelevant” (Fodor, 1983, p. 105). Moreover, analogical reasoning is arguably central to human cognitive prowess (Gentner, 2003). So we should be able to offer a *prima facie* case that the global workspace architecture relieves the computational burden allegedly brought on by the informational unencapsulation of analogical reasoning. The Rorschach blot example is a start.¹⁸ But a more convincing case can be made in the context of a more fully

¹⁶ This might seem to be taking specialisation too far. But with 100 billion neurons and many trillions of synapses, the human brain has plenty of scope for truly massive parallelism, and without recourse to “grandmother cell” coding.

¹⁷ The use of words like “aroused”, “announced”, and “urgency” here is purely stylistic, and doesn’t mask the lack of an implementable concept. Competitive access to a global workspace can be realised in a working computational model in various ways. For example, each parallel process can be made responsible for assigning a value to its own relevance in a given situation, and a central winner-takes-all strategy can be used to select the process to be granted access. For a working implementation of the global workspace model see Franklin and Graesser (1999).

¹⁸ The Rorschach blot is an example of literal similarity rather than true analogy.

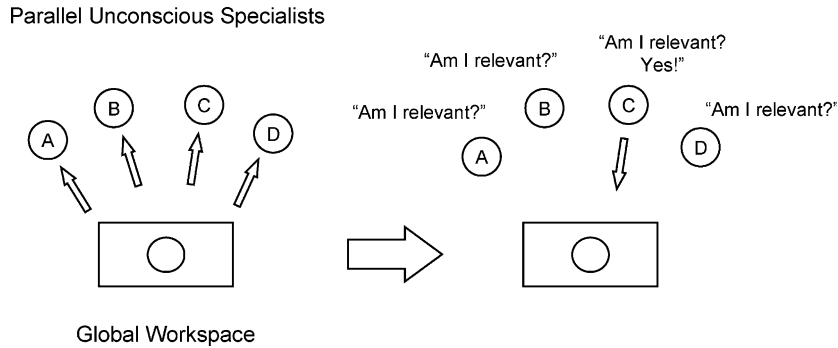


Fig. 3. The global workspace model of information flow.

realised theory of analogical thinking. So what follows is a short review of recent research on the computational aspects of analogical reasoning.¹⁹

6.1. Computational models of analogical reasoning

Most recent computer models of analogical reasoning take their cue from the *structure mapping* theory of analogy (Gentner, 1983). According to Gentner's theory, a central task in analogical reasoning is to find "maximally structurally consistent" mappings between two representations—a *target* located in working memory and a *base* located in long-term memory.²⁰ A *structurally consistent* mapping is one that conforms to certain constraints. In particular, no element in the target/base can be mapped on to more than one element in the base/target, and if two elements are put in correspondence then their sub-elements must also be put in correspondence. Finally, the theory includes a heuristic of *systematicity*, whereby deep and rich mappings are preferred to shallow and sparse ones.²¹

The Structure Mapping Engine (SME) is a computational realisation of Gentner's structure mapping theory (Falkheimer, Forbus, & Gentner, 1989). Its core is an algorithm for finding mappings between a pair of propositional representations that conform to the principles of the theory. It can find a mapping between two given representations in polynomial time, and it assimilates new information incrementally. ACME (Analogical Constraint Mapping Engine) is another computer implementation of the mapping process that adheres to the chief tenets of Gentner's structure mapping theory (Holyoak & Thagard, 1989). But unlike SME, which uses graph matching, ACME views

¹⁹ A related topic in AI is *case-based reasoning* (Kolodner, 1993). However, case-based reasoning tends to operate within domain boundaries, while the nub of the present discussion is the ability to cross domain boundaries.

²⁰ The ensuing discussion takes place against the backdrop of a language-like, viewpoint-free representational medium (even when realised in a distributed connectionist network). But global workspace theory is equally applicable to an analogical, viewer-centred representational medium.

²¹ The "systematicity" of an analogical mapping in Gentner's terminology has no direct connection with the "systematicity" of thought (Fodor & Pylyshyn, 1988).

the mapping process as a constraint satisfaction problem and employs connectionist techniques to solve it.

Both the SME and ACME models find mappings between two *given* representations, and neither addresses the question of *retrieval* (or *access*), that is to say how to find candidate base representations from a large pool of possibilities. Retrieval is the aspect of analogical reasoning responsible for seeding “the transfer of information among cognitive domains previously assumed to be irrelevant”. So the retrieval process is inherently informationally unencapsulated, and is therefore the locus of the frame problem for analogical reasoning. To quote Thagard, Holyoak, Nelson, and Gochfeld (1990), “the problem of analogical retrieval is how to get what you want from memory without getting more than you can use”.²²

The retrieval problem has been tackled in the context of ACME by a computational model called ARCS (Thagard et al., 1990), and in the context of SME by a computational model called MAC/FAC (Forbus et al., 1994). The ARCS model (Analogical Retrieval by Constraint Satisfaction) employs the same combination of constraint satisfaction and connectionism as the ACME system from which it is descended. The ARCS retrieval process is executed in two stages. The first stage finds candidate base representations using pre-coded notions of semantic similarity, while the second stage evaluates how well these candidates satisfy a number of generic constraints on analogical mapping related to those advocated by Gentner’s theory. The MAC/FAC model is a similar two-stage retrieval process. In the MAC stage (Many Are Called), a crude but computationally cheap filter is used to isolate a manageably small number of potential base representations for mapping onto a given target. In the ensuing FAC stage (Few Are Chosen), SME is used to select the best among the candidate mappings recommended by the MAC stage.

Both the SME-MAC/FAC model and the ACME-ARCS model can account for a good deal of well-established psychological data about analogical reasoning, as enumerated in the key papers already cited. However, they also have limitations. Of especial interest here, as we’ll see later, are certain criticisms levelled at these models by Hummel and Holyoak (1997) and Keane, Ledgeway, and Duff (1994). Both sets of authors draw attention to the unrealistic demands these models place on working memory capacity. Here is a summary from (Hummel & Holyoak, 1997).

Both ACME and SME form explicit representations of very large numbers of possible local matches between elements of source and target analogs, most of which are then discarded as mapping proceeds. Such processes do not seem compatible with the generally accepted limits of working memory.

This, alongside other concerns, motivated the development of IAM (Keane et al., 1994) and LISA (Hummel & Holyoak, 1997). Both IAM (Incremental Analogy Machine) and LISA (Learning and Inference with Schemas and Analogies) carry out mapping serially and incrementally, processing small propositional structures one at a time. For this reason, they are able to emulate human analogical reasoning performance more closely than their

²² Thagard et al. (1999), p. 261. As they go on to say, “this is the core problem of memory retrieval in general”. It’s also a good redescription of what, following Fodor, this article means by the frame problem itself.

predecessors, duplicating the difficulties humans have in forming “unnatural analogues” and human sensitivity to the order in which the parts of the target and base representations are processed.

The IAM model tackles mapping only. However, the LISA model tackles retrieval and mapping together, reflecting another of Hummel and Holyoak’s desiderata, namely the need for a graceful integration of these two aspects of analogical reasoning. The result is a system in which retrieval is naturally cast as a parallel process while mapping is inherently serial. For this reason, the LISA model—which is arguably the most accurate and psychologically plausible of the current computational models of analogical reasoning²³—is also the one that maps most cleanly and elegantly onto a global workspace architecture.

6.2. *The interplay of parallel and serial computation*

The thrust of the ongoing argument is that the global workspace architecture eases the computational burden on an informationally unencapsulated process through its inherent parallelism. Accordingly, this section asks whether it is possible to map a computational model of analogical reasoning to the global workspace architecture in order to realise the potential for parallel execution. The main focus of attention is the retrieval process. As their respective authors emphasize, the ARCS, MAC/FAC, and LISA models of retrieval are each amenable to parallel implementation. But in the context of the LISA model, we shall also look at the interplay of parallel and serial computation.

We begin with ARCS. Following Thagard et al. (1990), the time complexity analysis for a single run of ARCS must consider two components—the time required to form a network of mapping constraints across two potential representations and the time taken for the corresponding network to settle.²⁴ The former is purportedly $O(n^2m^4)$ for a serial machine, where n is the total number of candidate base representations and m is the number of propositions in the largest base or target representation. This is not a favourable statistic, but it is the worst-case analysis and the authors’ experimental data is more encouraging. Moreover, it seems likely that a parallel implementation would reduce the worst-case statistic to better proportions, although this possibility is not discussed in (Thagard et al., 1990).

The explicitly claimed advantage of parallelism for ARCS hinges on the second component of the complexity analysis, namely the time required for the network to settle (to find a solution). The authors do not offer a worst-case analysis of this statistic, but they appeal to empirical results to justify the claim that parallel implementation would reduce it from $O(n)$ to constant time. Although parallelism is clearly helpful for this component of the computation, the dominant term in the overall analysis is still n^2m^4 , and this is worrisome if n is very large. This remains true in spite of the authors’ empirical evidence, which was gathered with a comparatively small set of potential base representations.

²³ See Keane and Eysenck (2000, p. 436), for example.

²⁴ Thagard et al. (1990), p. 282. The authors also offer a space complexity analysis, which isn’t relevant here.

Therefore, since a very large n is precisely what the frame problem is all about, there is only modest support for the ongoing argument from the ARCS direction.

So let's move on to the MAC/FAC model. Both the MAC and FAC stages employ a set of parallel *matchers* that compare the target to each potential base in long-term memory and a single *selector* that chooses among the outputs of the matchers. Since the matchers are completely independent, the time complexity of a maximally parallel implementation reduces to that of finding a single match hypothesis between a given target and base. According to Forbus, Gentner, and Law (1994), this is $O(m)$ in the worst case on a parallel machine and should typically be better than $O(\log m)$, where m is the number of propositions in the larger of the base or target representation.²⁵

This is a far more pleasing complexity result than was obtained for ARCS. It suggests that analogical reasoning, if implemented on massively parallel hardware, is computationally feasible, even for a very large pool of potential base representations. Moreover, in global workspace terms, the array of parallel matchers in MAC/FAC corresponds directly to specialist, unconscious processes, the selector process corresponds to the attention mechanism that gates access to the global workspace, and the working memory in which the final mapping is carried out corresponds to the global workspace itself.

The fact that analogical reasoning can be rendered computationally feasible through implementation on a form of parallel hardware that is compatible with global workspace theory is very encouraging. However, other models of parallel computation would fit MAC/FAC just as well as the global workspace architecture. The aim here is to demonstrate a more intimate link between the demands of informationally unencapsulated cognitive processes in general and specific features of the global workspace architecture. In this respect, the LISA model turns out to be more pertinent than MAC/FAC.

Although they don't offer a formal complexity analysis, Hummel and Holyoak (1997) state that: "During mapping, driver propositions are activated serially ... At the same time, recipient propositions respond in parallel to the semantic patterns generated by the role-filler bindings of an active driver proposition." (In Hummel and Holyoak's terminology, "driver propositions" belong to the representation resident in working memory, while "recipient propositions" belong to representations held in long-term memory.) So LISA effectively carries out retrieval in parallel and mapping serially. Consequently, as with MAC/FAC, retrieval time will not grow with the size of the pool of potential base representations, given sufficient parallelism in the underlying hardware.

Additionally, the combination of parallel and serial processing makes LISA more compatible with global workspace theory than any of its rivals. The currently active proposition in LISA's working memory corresponds to the current contents of the global workspace.²⁶ The parallel activation of propositions in LISA's long-term memory corresponds to the broadcast of the contents of the global workspace. And the distributed propositions in LISA's long-term memory correspond to parallel, unconscious processes in global workspace theory. But most importantly, the necessarily *serial* presentation of

²⁵ Forbus et al. (1994), p. 160.

²⁶ Assuming rehearsal in working memory is a form of conscious information processing (Baars & Franklin, 2003), this suggests an explanation for the fact that, while some aspects of analogical reasoning are undoubtedly carried out unconsciously, others plainly are not.

propositions in LISA's limited capacity working memory matches the necessarily *serial* presentation of coherent material in the limited capacity global workspace. It is surely no accident that the features of the LISA model that make it more psychologically plausible than its competitors are the very same features that also align it so closely with global workspace theory.

7. Discussion

Let's review the argument so far. We set out by undermining the in-principle claim that informationally unencapsulated cognitive processes are computationally infeasible. It turned out that the case put forward by Fodor and others is too weak to sustain such a conclusion. The way the biological brain handles such processes is thereby demoted from an out-and-out mystery to a scientific challenge. The global workspace architecture, with its blend of parallel and serial computation, was then proposed as an answer to this challenge.

In place of the naïve serial, exhaustive search that seems to be behind the intuitions of those authors who see the frame problem a serious obstacle, the global workspace architecture supports the idea of a single, serial thread of computation that can also draw on the resources of a massively parallel set of distributed processes. A review of current computational models of analogical reasoning—often taken to be the epitome of informational unencapsulation—demonstrated a close fit between global workspace theory and the most psychologically plausible of these models, namely that of Hummel and Holyoak (1997).

Of course, analogical reasoning is only one among many informationally unencapsulated cognitive processes. One advantage of bringing a computational model of a specific cognitive process within the wider compass of global workspace theory is that it embeds that model in a generic account of cognitive function. Instead of seeing analogical reasoning as an isolated process, it comes to be seen as interleaved with many other cognitive activities—such as perception, action selection, planning, belief revision, language production, language comprehension, and so on—each of which can exploit the brain's diversity of massively parallel specialist processes.

This brings us to the question of the implications of the present proposal for the so-called *modular* theories of the human mind to which many contemporary cognitive scientists subscribe in one form or another (Carruthers, 2002; Fodor, 1983, 2000; Karmiloff-Smith, 1992; Mithen, 1996; Pinker, 1997; Sperber, 1994; Tooby & Cosmides, 1992). Although modular theories come in various guises, nearly all share two central tenets. First, that some proportion of human mental structure comprises a set of special-purpose modules with a narrow remit. Second, that the flexibility and creativity of human thought demands an additional mechanism capable of transcending rigid modular boundaries.

The second of these tenets is the flash-point for the present debate, and authors differ widely in the mechanisms they propose. Fodor (1983, 2000) assigns cross-modular capabilities to the mind's "central processes", which he believes to be beyond the reach of present-day cognitive science, laying much of the blame for this on the frame problem.

Karmiloff-Smith (1992), writing from a developmental perspective, hypothesises a stage of “representational redescription” whereby knowledge that was previously represented implicitly for a child becomes explicitly represented *to* the child. Sperber (1994) posits a “metarepresentation module”, which can manipulate representations about representations and concepts of concepts, allowing knowledge from multiple domains to be blended. Mithen (1996), drawing on archaeological evidence, introduces the idea of human “cognitive fluidity”, which permits “thoughts and knowledge generated by specialized intelligences [to] flow freely around the mind”.²⁷ Carruthers (2003) suggests that “natural language is the medium for non-domain-specific thinking, serving to integrate the outputs of a variety of domain-specific conceptual faculties”.²⁸

Global workspace theory, with its commitment to multiple, parallel, specialist processes, is consistent with a modular view of the human mind. What the global workspace architecture has to offer in addition is a model of information flow that explains how an informationally unencapsulated process can draw on just the information that is relevant to the ongoing situation without being swamped by irrelevant rubbish. This is achieved by distributing the responsibility for deciding relevance to the parallel specialists themselves. The resulting massive parallelism confers great computational advantage without compromising the serial flow of conscious thought, which corresponds to the sequential contents of the limited capacity global workspace.

In the context of a global workspace architecture, there are no distinguished central processes with special privileges or powers. Rather, every high-level cognitive process can be analysed in terms of the interleaved operation of multiple, specialist, parallel processes, all of which enjoy equal access to the global workspace. Analogical reasoning, for example, can recruit the sort of capacity to integrate spatial relations allegedly located in prefrontal cortex (Waltz et al., 1999), alongside whatever other brain resources it may require. The significance of informational encapsulation is thereby much diminished, and the spectre of the frame problem is dissolved.

Acknowledgements

Thanks to Igor Aleksander, Ron Chrisley, Stan Franklin, and Mercedes Lahnstein. Thanks also to the paper’s three anonymous reviewers.

References

- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
 Baars, B. J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford University Press.
 Baars, B. J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Science*, 6(1), 47–52.

²⁷ Mithen (1996), p. 71.

²⁸ Carruthers (2003), p. 657.

- Baars, B. J. (2003). How brain reveals mind: Neuroimaging supports the central role of conscious experience. *Journal of Consciousness Studies*, 10(9-10), 100–114.
- Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Science*, 7(4), 166–172.
- Baars, B. J., Ramsøy, T. Z., & Laureys, S. (2003). Brain, consciousness and the observing self. *Trends in Neurosciences*, 26(12), 671–675.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Carruthers, P. (2002). The cognitive functions of language. *Behavioral and Brain Sciences*, 25(6), 657–674.
- Carruthers, P. (2003). On Fodor's problem. *Mind and Language*, 18(5), 502–523.
- Cherniak, C. (1986). *Minimal rationality*. Cambridge, MA: MIT Press.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1–37.
- Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, 4, 752–758.
- Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Science*, 100(14), 8520–8525.
- Dennett, D. (1978). *Brainstorms*. Cambridge, MA: MIT Press.
- Dennett, D. (1984). Cognitive wheels: The frame problem in artificial intelligence. In C. Hookway (Ed.), *Minds, machines and evolution* (pp. 129–151). Cambridge: Cambridge University Press.
- Dennett, D. (1991). *Consciousness explained*. Harmondsworth: Penguin.
- Falkenheimer, B., Forbus, K., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1–63.
- Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1987). Modules, frames, fridgeons, sleeping dogs, and the music of the spheres. In Z. W. Pylyshyn (Ed.), *The robot's dilemma: The frame problem in artificial intelligence* (pp. 139–149). Norwood, NJ: Ablex.
- Fodor, J. A. (2000). *The mind doesn't work that way*. Cambridge, MA: MIT Press.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critique. *Cognition*, 28, 3–71.
- Forbus, K., Gentner, D., & Law, K. (1994). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141–205.
- Franklin, S., & Graesser, A. (1999). A software agent model of consciousness. *Consciousness and Cognition*, 8, 285–301.
- Gabaix, X., & Laibson, D. (2003). A new challenge for economics: 'The Frame Problem'. In I. Brocas, & J. D. Carrillo (Eds.), *The psychology of economic decisions, volume one: rationality and well-being* (pp. 169–183). Oxford University Press.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. San Francisco, CA: Freeman.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155–170.
- Gentner, D. (2003). Why we're so smart. In D. Gentner, & S. Goldin (Eds.), *Language in mind* (pp. 195–235). Cambridge, MA: MIT Press.
- Haselager, W. F. G., & van Rappard, J. F. H. (1998). Connectionism, systematicity, and the frame problem. *Minds and Machines*, 8(2), 161–179.
- Hayes-Roth, B. (1985). A blackboard architecture for control. *Artificial Intelligence*, 26(3), 251–321.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295–355.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427–466.
- Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT Press.
- Keane, M. T., & Eysenck, M. W. (2000). *Cognitive psychology: A student's handbook* (4th ed.). Psychology Press.

- Keane, M. T., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, 18, 387–438.
- Kolodner, J. L. (1993). *Case-based reasoning*. Los Altos, CA: Morgan Kaufmann.
- Kreiman, G., Fried, I., & Koch, C. (2002). Single-neuron correlates of subjective vision in the human medial temporal lobe. *Proceedings of the National Academy of Science*, 99(12), 8378–8383.
- McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In D. Michie, & B. Meltzer, *Machine intelligence* (Vol. 4) (pp. 463–502). Edinburgh University Press, 463–502.
- Mithen, S. (1996). *The prehistory of the mind*. London: Thames and Hudson.
- Newell, A. (1962). Some problems of basic organization in problem-solving systems. In *Proceedings of the second conference on self-organizing systems* (pp. 393–342).
- Nii, H. P. (1986). The blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*, 7(2), 38–53.
- Pinker, S. (1997). *How the mind works*. Harmondsworth: Penguin.
- Pylyshyn, Z. W. (Ed.). (1987). *The robot's dilemma: The frame problem in artificial intelligence*. Norwood, NJ: Ablex.
- Reiter, R. (2001). *Knowledge in action: Logical foundations for specifying and implementing dynamical systems*. Cambridge, MA: MIT Press.
- Russell, S., & Wefald, E. (1991). *Do the right thing: Studies in limited rationality*. Cambridge, MA: MIT Press.
- Sandewall, E. (1994). *Features and fluents: The representation of knowledge about dynamical systems*, Vol. 1. Oxford: Oxford University Press.
- Shanahan, M. P. (1997). *Solving the frame problem: A mathematical investigation of the common sense law of inertia*. Cambridge, MA: MIT Press.
- Shanahan, M. P. (2003). The frame problem. In L. Nadel (Ed.), *The Macmillan encyclopedia of cognitive science* (pp. 144–150). New York: Macmillan, 144–150.
- Simon, H. (1957). *Models of Man*. New York: Wiley.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld, & L. A. Gelman (Eds.), *Mapping the mind: Domain specificity and culture* (pp. 39–67). Cambridge: Cambridge University Press.
- Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, 46, 259–310.
- Tononi, G., Edelman, G. M., & Sporns, O. (1998). Complexity and coherency: Integrating information in the brain. *Trends in Cognitive Science*, 2(12), 474–484.
- Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind* (pp. 19–136). Oxford: Oxford University Press.
- Waltz, J. A., Knowlton, B. J., Holyoak, K. J., Boone, K. B., Mishkin, F. S., de Menezes Santos, M., et al. (1999). A system for relational reasoning in human prefrontal cortex. *Psychological Science*, 10(2), 119–125.
- Zilberstein, S., & Russell, S. J. (1993). Anytime sensing, planning and action: A practical model for robot control. In *Proceedings 1993 international joint conference on artificial intelligence (IJCAI 93)* (pp. 1402–1407).