

What's the Problem with the Frame Problem?

Sheldon J. Chow

Published online: 2 March 2013
© Springer Science+Business Media Dordrecht 2013

Abstract The frame problem was originally a problem for Artificial Intelligence, but philosophers have interpreted it as an epistemological problem for human cognition. As a result of this reinterpretation, however, specifying the frame problem has become a difficult task. To get a better idea of what the frame problem is, how it gives rise to more general problems of relevance, and how deep these problems run, I expound six guises of the frame problem. I then assess some proposed heuristic solutions to the frame problem; I show that these proposals misunderstand, and fail to address, an important aspect of the frame problem. Finally, I argue that though human cognition does not *solve* the frame problem in its epistemological guise, human cognition *avoids* some of the epistemological worries.

1 Introduction

Loosely understood, the frame problem is the problem of determining which of a potentially infinite number of beliefs are relevant to the task at hand. Put another way, the frame problem is a special problem for cognitive systems of circumscribing (read: framing) the information to be considered in their processes. However, over the years different authors have emphasized different aspects of the frame problem, and there have been many different interpretations of the problem. These different interpretations have not only made specifying what the frame problem is a difficult task (cf. Dennett 1984), but it has produced some confusion in debates on the matter.

This paper has three aims. The first aim is to gain clarity on the issue. To get a better idea of what the frame problem is, how it gives rise to more general problems of relevance, and how deep these problems run, I show that the frame problem is best

S. J. Chow (✉)
Department of Philosophy, Stevenson Hall, The University of Western Ontario, London,
Ontario N6A 5B8, Canada
e-mail: schow55@uwo.ca

understood as a set of closely related problems (cf. Samuels 2010). To this end, I will expound six guises of the frame problem and explain their relation and philosophical significance.¹

The second aim of this paper is to contribute to the debate concerning whether and the extent to which the frame problem is a problem for human cognition. We will see that much of the controversy has to do with the idea that the frame problem poses a tractability problem for computational cognition. Fodor (1983, 1987, 2000) believes that the frame problem (generally understood) is intractable, and therefore detrimental to computational cognitive science. On the other hand, since heuristics are generally understood to ensure computational tractability, heuristics are often invoked as a possible way to circumvent the frame problem. Philosophers such as Samuels (2005, 2010), Carruthers (2006a, 2006b), and Lormand (1990, 1996) suggest that heuristics can serve as techniques that pick out computationally feasible subsets of information to be brought to bear on cognitive tasks. I contend, however, that merely pointing to heuristics in an effort to circumvent the frame problem reveals a misunderstanding of the problem. For the frame problem is not only a problem of explaining how computationally feasible subsets of information can be brought to bear on cognitive tasks, it is also a problem of explaining the reasonable levels of success that humans enjoy in determining what information is relevant to their cognitive tasks; circumscribing what information is to be considered in a computationally tractable way does not ensure that what gets considered is in fact relevant.

The paper's third aim is to argue that human cognition does not *solve* what I take to be one of the most philosophically significant guise of the frame problem, what I call the Epistemological Relevance Problem. I will explain, though, that human cognition *avoids* the worries posed by this interpretation of the frame problem by relying on a highly organized cognitive architecture to guide relevance determinations, and this solves a different, equally significant, aspect of the frame problem. As we shall see, my analysis goes further than the response that "we simply get along as best we can and deal with the problem of planning in ways that, to use Dennett's phrase, is 'good enough for government work'" (Pylyshyn 1996, p. xi). For my position is that, although we do not solve the frame problem under one interpretation, we do solve it under another interpretation, and this goes beyond "government work."

2 The Many Guises of the Frame Problem

The frame problem was first introduced by John McCarthy and Pat Hayes as a problem faced by early logic-based Artificial Intelligence (AI) (McCarthy and Hayes 1969). This was a problem of how to represent, in a logically succinct way, the ways in which properties of objects do not change when acted upon. For any given action, there will be many properties of many objects that remain constant throughout and after the execution of the action. When moving a ball from one location to another, for instance, the ball's color does not change, nor does its shape or its size; other objects in the room do not

¹ In defending Classical computationalism from the purported threat posed by the frame problem, Samuels (2010) also describes the frame problem as constituting a set of related problems. But I do not think his analysis goes deep enough. The analysis I provide here may be seen as an extension of Samuels'.

change; the fact that $2+2=4$ does not change; and so on. If a cognitive system is to maintain a veridical belief-set² about the ball, and the world in general, determinations of which properties change with which actions will have to be coded into its program. In logic-based AI, this was done by coding what are called “frame axioms”, which are sentence-like representations that specify a frame of reference for the action in question (Dennett 1984; Viger 2006a). But the number of frame axioms can become quite cumbersome very quickly, since there will be in general *number of properties* \times *number of actions* frame axioms that would have to be written into the program (Shanahan 2009). Obviously, the more properties or actions to account for, the more cumbersome this will be. Thus, the problem for logic-based AI was how to most compactly represent the non-effects of actions. As such, it is often referred to as a problem of *representation*. Let us call this the *AI Frame Problem*:

AI Frame Problem The problem of how to appropriately and succinctly represent in a system's program the fact that most properties are unaffected by given actions.

The main obstacle to overcoming the AI Frame Problem is the monotonicity of classical logic. As a monotonic logic, classical logic does not allow for different conclusions to be drawn with the addition of new premises. AI researchers have since developed a number of nonmonotonic logics to remedy the situation. According to Shanahan (2009), the introduction of nonmonotonic logics adequately solves the AI Frame Problem, notwithstanding certain technical problems.³ Whether Shanahan is correct is perhaps debatable (see e.g., Fodor 1987, 2000), but I will not adjudicate his claim here. For the original frame problem for AI and its purported solutions are not of much concern to this paper, or to philosophy more generally. Rather, what is of philosophical interest are the computational and epistemological concerns born out of the AI Frame Problem.

Philosophers have interpreted the AI Frame Problem as a general problem of how a cognitive system updates its beliefs if it is to maintain (suitably) veridical beliefs about the world. As Fodor observes, “Despite its provenance in speculative robotology, the frame problem doesn't really have anything in particular to do with action. After all, one's standing cognitive commitments must rationally accommodate each new state of affairs, whether or not it is a state of affairs that is consequent upon one's own behavior” (Fodor 1987, p. 27). This interpretation of the frame problem has an obvious connection to the AI Frame Problem. Upon the arrival of new information, a cognitive system will

² It is obviously too demanding of any system to maintain a completely veridical belief-set. What is required, instead, is a belief-set that is suitably veridical, or reasonably accurate (cf. Samuels 2010).

³ One technical problem that is often cited is the “Yale Shooting Problem.” In this scenario, a gun is loaded at time t , and shooting the gun at Fred at $t+1$ is expected to kill Fred. However, the formalism of nonmonotonic logics cannot uniquely determine whether Fred is dead. This is because there are two ways the formalism can characterize what happens as time passes from t to $t+1$. On the one hand, the formalism can treat the gun's being loaded as persisting through time, and Fred's being alive as changing as a result of being shot by a loaded gun. This is the intuitive case. However, on the other hand, the formalism can treat Fred's being alive as persisting through time, and the gun's being loaded as changing as a result of Fred not dying in spite of the gun being fired at him. The problem is that the formalism has no way to determine which should persist between t and $t+1$: Fred's being alive or the gun's being loaded. Hence, the formalism cannot predict whether Fred is alive or dead (Viger 2006a). We should note, however, that though the Yale Shooting Problem is a problem, it is not insurmountable, and indeed there are many solutions (Shanahan 2009).

have to update its beliefs. In seeing a ball move from one location to another, for instance, a cognitive system will have to update its belief regarding the position of the ball, as well as the ball's new spatial relations with respect to other objects in the room. But again, there are arbitrarily many other beliefs that it will not have to update, such as the colour of the ball, the sizes of the other objects in the room, the relative spatial relations of the unmoved objects, the ambient temperature in the room, that $2+2=4$, and so on. The problem, then, is to determine which of any of the system's beliefs is to be affected in belief-update. Such a determination cannot be made a priori. For instance, are the beliefs about the relative spatial relations of the other objects in the room affected? That depends: was something knocked over by the ball as it moved across the room? Unlike the original representational problem of logic-based AI, however, this problem is *computational* in nature (McDermott 1987)—it concerns the *processes* a cognitive system employs as it updates its beliefs. Let us call this problem the *Update Problem*:

Update Problem The problem of determining which of any of the system's beliefs is to be affected in belief-update.

This interpretation of the frame problem has been most recently described by Samuels (2010); Haugeland (1985) characterizes the frame problem in similar terms.

The Update Problem can be generalized as a problem of relevance, namely a problem of determining what is relevant to a given instance of belief-update. More generally, a cognitive system has to determine what bits of information should bear on a given task, i.e., determine what is relevant to the task. Determining what is relevant is a deep problem, especially for human reasoning since central systems—the cognitive systems that are paradigmatically responsible for general reasoning and decision-making—admittedly allow for free exchange of information, any of which can bear on their operations. Thus, central systems are what is often referred to as *holistic*. Fodor (1983) claims that central systems are holistic in part because they are *isotropic*. A system is isotropic just in case, given an appropriate set of background beliefs, any representation *in principle* can be relevant to any other.⁴ Who won tonight's football game is *prima facie* irrelevant to whether there is beer in your friend's fridge. But if you believe that your friend's favourite football team played tonight, and that your friend usually over-indulges in beer consumption whenever his favourite team wins, then your belief about who won tonight's game actually *is* relevant to your belief about beer in your friend's fridge. Indeed, an itch on your left hand can be relevant to your belief about whether you will get a job promotion, if you carry an appropriate set of (superstitious) background beliefs. Again, there is no way a priori to circumscribe the subset of representations that are relevant in a given occasion of reasoning or inference. I will refer to this construal of the frame problem as the *Generalized Relevance Problem*:

Generalized Relevance Problem The problem of determining, from all the information that can in principle bear, what is relevant to the cognitive task at hand.

It perplexes cognitive science that humans have an astonishing ability for determining (with reasonable levels of success) what is relevant in much of their reasoning, and

⁴ I assume a representational theory of mind, but nothing of what I provide in this paper hangs on this. Those who do not subscribe to any kind of representational theory of mind can simply substitute "information," or something similar, for my talk of "representations."

fixating thereon, without having to spend much time deciding what is and is not relevant, or wasting time cogitating on irrelevant details. In terms of the generalized-relevance guise of the frame problem, it seems that we generally know quite easily what pieces of information bear on—or lie within the frame of—our reasoning tasks. The problem for cognitive science is to understand how we do it.

It is important to be clear, in light of these observations about human performance, what the Generalized Relevance Problem is and what it is not. The problem is not how a cognitive system can possibly determine, from all the information that can in principle bear, *everything* that is relevant to the task at hand. Humans typically do not consider everything that is relevant, and it is normally too cognitively demanding (of human or machine) to do so. This is demonstrably the case—witness the frequency of our errors, or of surprise. What makes matters worse, it may simply be impossible to consider things that are relevant to the task at hand. Suppose for instance that the number of hairs on Caesar's head on March 15, 44BCE is (somehow) relevant to the extent to which the current state of the economy will improve over the next 5 years. Although the former piece of information is in principle relevant to the latter issue, it is impossible to know the number of hairs on Caesar's head on March 15, 44BCE; and so it would be impossible to consider a relevant piece of information when estimating the extent to which the current state of the economy will improve over the next 5 years.

The Generalized Relevance Problem is, rather, a problem of how a cognitive system can make determinations of what is relevant to a given task *with reasonable levels of success*. The problem is for real-world cognitive systems—cognitive systems that are less-than-perfect, and that suffer from real constraints in terms of time and cognitive resources. Therefore, the Generalized Relevance Problem suggests a methodological problem for a cognitive system of making reasonably accurate determinations of relevance efficiently and in a timely fashion. We do not determine everything that is relevant to the task at hand, but enough to get the job done. This means, among other things, that there is minimal waste of time and cognitive resources considering what is irrelevant. (Such an observation might appear to be a mundane point, but we shall soon see that it reveals an additional deep epistemological problem.)

The methodological problem of efficiently making reasonably accurate determinations of relevance primarily concerns the computational burden placed on a system in determining the set of representations to be considered for a given cognitive task. Since relevance cannot be determined a priori, and since any representation held by a cognitive system can in principle be relevant to any other (provided an appropriate set of background beliefs), determining which representations to consider can be computationally intractable. If a system has a sufficiently small set of beliefs to begin with, relevance determinations can be performed without running into major computational problems: the system may simply compute each representation to assess its relevance. But once we consider a cognitive system with sufficiently many and/or complex beliefs—as is the case with humans, who have a wealth of beliefs and stores of other complex representations—assessing every representation quickly becomes infeasible. For such a strategy would entail a vast amount of computations, requiring the use of scarce cognitive resources and taking an unreasonable amount of time. Indeed, assessments of relevance of individual representations are computationally taxing enough, but assessments of relevance *between* representations exponentially add to the complexity (Samuels 2005). Thus, what is needed are computationally

tractable methods to pick out, in a timely fashion, which representations are to be brought to bear on a given cognitive task. Let us call this the *Computational Relevance Problem*.

Computational Relevance Problem The problem of how a cognitive system tractably delimits (i.e., frames) what gets considered in a given cognitive task.

It is important to notice that the Computational Relevance Problem would not be solved or circumvented if a cognitive system were merely to employ a method to differentiate between a relevant representation and an irrelevant one, and then subsequently ignore the representations that fall into the latter category. This is a divide-and-conquer strategy in which the dividing part is too computationally taxing. For not only can such a strategy possibly entail considering each representation held by the system, but even if this is not the case, seeing whether a given representation is relevant means considering the representation and drawing implications. In other words, the system employing such a divide-and-conquer strategy would be performing quite a lot of computations to see whether a given representation is relevant and therefore should be brought to bear. Of course, most beliefs and other representations a system holds will be irrelevant to the cognitive task at hand, and this would then translate to wasting quite a lot of resources on considering what is irrelevant, and performing computations by drawing conclusions that do not need to be drawn. This is what Fodor aptly referred to as “Hamlet’s problem: How to tell when to stop thinking” (Fodor 1987, p. 26).

The Computational Relevance Problem is therefore a problem of how a cognitive system can *ignore* most of what it knows (Dennett 1984); *qua* relevance problem, it concerns how a system ignores *what is irrelevant*. This brings us to the other dimension of the Generalized Relevance Problem mentioned above. The problem here has to do with how a cognitive system considers (mostly) *only* what is relevant. This is what I take to be the “new, deep epistemological problem” that Dennett spoke of (Dennett 1984, p. 130). Again, since relevance cannot be determined a priori, and since any representation held by an agent can in principle be relevant to its task (provided an appropriate set of background beliefs), how does a cognitive system *know* what is relevant? I will refer to this problem as the *Epistemological Relevance Problem*.

Epistemological Relevance Problem The problem of how a cognitive system considers (mostly) only what is relevant, or how a cognitive system knows what is relevant.

It is important to note that this epistemological problem⁵ does not have to do so much with the computational costs of delimiting what gets considered in a given cognitive task (although this remains an important aspect) as it has to do with considering the *right* things.

Thus far I have outlined four different problems that have arisen from the original AI Frame Problem, and I have shown how they are related to one another. There is

⁵ What I am calling here the Epistemological Relevance Problem does not coincide with what Shanahan (2009) calls the epistemological frame problem. According to Shanahan, “The epistemological problem is this: How is it possible for holistic, open-ended, context-sensitive relevance to be captured by a set of propositional, language-like representations of the sort used in classical AI?” It appears that what Shanahan has in mind for the epistemological problem is something closer to the AI Frame Problem, but for humans in determining relevance.

one further interpretation that I have yet to expound. Before I do so, however, I will consider a proposed solution to the frame problem *tout court*, namely the heuristic solution. Considering the heuristic solution, and whether it succeeds, is important for providing context for the remaining interpretation, as well as for understanding whether and the extent to which human cognition in fact solves the frame problem.

3 The Heuristic Solution

Many philosophers and cognitive scientists have posed the question: How do we solve the frame problem? However, this question has never been posed within the context of demarcating the various aspects of the frame problem outlined above. The question “How do we solve the frame problem?” might not even be meaningful without specifying which guise of the frame problem is being referred to. It certainly appears as if humans somehow manage to solve all aspects of the frame problem discussed so far, and some philosophers believe that heuristics are the operative solution in human cognition. But we might want to ask whether heuristics really offer a tenable solution to the frame problem, and indeed which, if any, guises of the frame problem humans actually solve despite appearances.

To address these concerns, I will now introduce the heuristic solution to the frame problem as a response to worries emphatically advanced by Fodor concerning a plausible theory of central cognition. My intention will then be to show that the heuristic solution applies only to the Computational Relevance Problem, ignoring the Epistemological Relevance Problem. For ease of exposition, I will use “relevance problems” to refer to all guises of the frame problem here expounded, save for what I have called the AI Frame Problem.

3.1 Fodor's Worries

Fodor (1983, 1987, 2000) argues that only a suitably *informationally encapsulated* system can avoid relevance problems. Roughly stated, informational encapsulation refers to a property of a cognitive system that can draw on only a highly restricted database of information in the course of its processing. There is some debate on how to interpret informational encapsulation and what is entailed by it.⁶ However, for our purposes let us understand it thus: A cognitive system is encapsulated to the extent that, of all the information held by other systems that is relevant to its processing, only a strictly delimited portion is available to it. The peripheral (as opposed to central) cognitive systems dedicated to linguistic and perceptual processing are paradigm cases of encapsulated systems. A common example is the Müller-Lyer illusion: the belief that the two lines are actually equal in length is relevant to determining their relative lengths, but the visual system is encapsulated with respect to this information; it has access to only a highly restricted or proprietary database (of low-level visual information) in the course of its processing.

Relevance problems do not arise for sufficiently informationally encapsulated systems, since there would be a sufficiently small amount of information over which

⁶ For further discussion see Carruthers (2006a, 2006b); Samuels (2005, 2006); Sperber and Wilson (1996).

to compute. To be more precise, encapsulated systems avoid relevance problems in two subtly distinct ways: Not only does the small amount of information contained in the system's database constitute *all* the information that the system *can* consider, thus considerably reducing the number of computations needed for information search, but that small amount of information constitutes the one and only set of background information against which relevance is determined. The more encapsulated a system is, the more tractable its computations will be, and the less relevance problems will be problems.⁷

Nevertheless, Fodor stresses that central cognition cannot be encapsulated given its holistic character and the wealth of complex beliefs and other representations it has access to. Central processing thus appears squarely faced with relevance problems. And this is bad news for the computational theory of mind (CTM), according to Fodor. As he explains, computations are, by definition, operations over local, context-insensitive properties, i.e., the formal syntax of the items being computed. In other words, it is only local properties that are causal determinants in computations. Fodor points out, however, that there are properties that central cognition assesses, such as relevance, simplicity, and conservatism. Unlike the local properties of syntax, these are extra-syntactic, *global* properties of representations in relation to specific contexts. According to Fodor, such contexts can be very large, requiring in the limit that entire belief systems be assessed.⁸ Hence, assessments of global properties quickly produce intractable problems, so the argument goes. At the same time, however, humans somehow perform central cognitive tasks swiftly, with both relative ease and reasonable accuracy. This leads Fodor to conclude that CTM is not a viable theory for central cognition, since CTM fails precisely on what central cognition does best. And since he claims that CTM is the only plausible theory of mind, Fodor is therefore pessimistic about the prospects of a cognitive science for central cognition. Hence:

‘Fodor’s First Law of the Nonexistence of Cognitive Science’ . . . : the more global (e.g., the more isotropic) a cognitive process is, the less anybody understands it. *Very* global processes, like analogical reasoning, aren’t understood at all. (Fodor 1983, p. 107)

As Fodor (2000) later makes clear, we can add assessments of relevance to the “*very* global processes” he has in mind.

In response, however, certain philosophers argue that Fodor overstates what tractable computation requires. Samuels (2005, 2010) and Carruthers (2006a, 2006b), for instance, both claim that, although encapsulation will ensure computational tractability,

⁷ Of course, this does not mean that the system's database will necessarily contain everything that is relevant to its tasks. And so it is possible that an encapsulated system can still have a problem with respect to considering what is relevant. But this poses a relevance problem of a different sort, which I will not be discussing here. In addition, even if an encapsulated system had access to all relevant information (but not only what is relevant), the system might still face the Epistemological Relevance Problem if its database was not organized in such a way that would facilitate determinations of relevance, or that would enable the system to consider (mostly) only what is relevant. (See below the discussion on the remaining guise if the frame problem.) Nevertheless, at the very least, a suitably encapsulated system will avoid the computational problem of tractably delimiting what gets considered in its tasks.

⁸ Fodor (1983, 2000) argues that this implies that entire theories are the units of cognition. In *Modularity of Mind*, he refers to this as a *Quinean* property of central systems. See Samuels (2010) for arguments against this view.

tractability does not necessarily require encapsulation (see also Lormand 1990, 1996; Sperber and Wilson 1996). Rather, they argue that all that tractable computation requires is that computations be suitably *frugal*, in the sense of simply limiting on the amount of information used or processed.

Frugal computational processes can, of course, be realized by “techniques that often, though not invariably, identify a substantial subset of those representations that are relevant to the task at hand” (Samuels 2010, p. 285), namely *heuristics*. The idea, then, is that by employing heuristics a system can avoid computationally taxing assessments of relevance, and they do this by instantiating suitable search and stopping rules that effectively limit the amount of information surveyed in its computations, but that still bring to bear an appropriate subset of information on the system’s tasks. This escapes the need to invoke encapsulation in explaining computational tractability. For a system can very well be unencapsulated, where *any* information *can* be brought to bear on a system’s computations, but the use of heuristics will ensure that only a small (and hopefully relevant) amount of this information *will in fact* be brought to bear in its processing.⁹ Hence, these theorists believe heuristics ensure that, though central cognition may very well be computational and unencapsulated, it operates in a tractable way.

Fodor had recently exclaimed that resorting to heuristic procedures to circumvent relevance problems “is among the most characteristic ways that cognitive scientists have of not recognizing how serious the issues are that the isotropy of relevance raises for a theory of the cognitive mind” (Fodor 2008, p. 116). By referring to what has come to be known as the “sleeping dog” strategy (McDermott 1987),¹⁰ he goes on to argue that the appeal to heuristics begs the question it is supposed to answer. Specifically, Fodor claims that the sleeping dog strategy takes for granted some principle that individuates situations; but the task of individuating situations presupposes that determinations of relevance have already been made (i.e., regarding what is relevant to circumscribing a situation). However, Fodor not only ignores possible heuristics that may not beg the relevance problem, but he fails to consider non-question-begging possibilities to individuate situations. It is possible, for example, that a heuristic can identify a situation as a certain type given a small set of cues, and this identification procedure can be completely based on local (as opposed to global) considerations (cf. Carruthers 2003). There is of course an accompanying risk of the heuristic being wrong in identifying a given situation, but heuristics provide no guarantees.

In an earlier work, Fodor (2000) gives a different, though equally wanting argument against heuristic solutions to relevance problems. In short, his argument is this: In deciding which heuristic to use for a given cognitive task, there is a meta-

⁹ Carruthers believes that we may still speak of systems as informationally encapsulated if we construe encapsulation as a property that simply restricts informational access in processing. In this way, he believes that encapsulation can mean just that most information simply will not be considered in a system’s computations—a notion he calls “wide encapsulation” (Carruthers 2006a, 2006b). However, Samuels criticizes Carruthers’ view, claiming that “not only is it [wide encapsulation] different from what most theorists mean by ‘encapsulation,’ but it’s simply what you get by denying exhaustive search; and since virtually no one thinks exhaustive search is characteristic of human cognition, the present kind of ‘encapsulation’ is neither distinctive nor interesting” (Samuels 2006, p. 45).

¹⁰ Roughly, the sleeping dog strategy is a rule to the effect of: Do not alter (update) a representation unless there is explicit indication of change, according to the dictum “let sleeping dogs lie.”

problem of deciding how to decide which heuristic to use. According to Fodor, not only is this the beginning of an infinite regress, but the meta-problem of deciding how to decide is faced with the relevance problems associated with holism and global inference just as the first problem is. However, here again Fodor fails to consider the possibility of heuristics being cued by a small set of abstract features of the problem (cf. Gigerenzer and Goldstein 1999). This is not a failsafe plan and it will sometimes misfire, but it would work nonetheless without invoking any regress. Or we might understand heuristics as cognitive procedures that can be expressed as rules that need not be represented in cognition, but are primitive processes, perhaps built into the architecture of cognition. And if this is the case, then there may not be any deciding which heuristic to use (one, as they say, *just does it*), and *ipso facto* no deciding how to decide which heuristic to use.

Therefore, despite Fodor's worries, heuristics may very well play a crucial role in achieving frugality in cognition, as is certainly the case for our AI counterparts. At the very least, heuristics appear to circumvent the problem of delimiting the amount of information a system considers in its processing—heuristics can simply pick out an appropriately limited set of representations from the total set possessed by an agent according to certain parameters, and the system can make use of them accordingly. There will be no guarantees that all and only relevant representations will be picked out (Samuels 2005, 2010), but again, such is the nature of heuristics. To be sure, humans are fallible, so we should not expect perfection in their cognition. Thus, it appears as if heuristics do indeed avoid relevance problems.

However, appearances can deceive, and we must therefore look at the details of these claims. Although Fodor's arguments against heuristic solutions to relevance problems were seen to be poor, this does not mean that the purported heuristic solution is in fact a solution. To see if heuristics do indeed circumvent relevance problems, let us take a closer look at what is being proposed in response to Fodor's worries about a computational theory of central cognition.

3.2 What Problem Do Heuristics Solve?

In specifying the heuristics that make cognition tractable, Carruthers (2006a, 2006b) draws on the work of Gigerenzer and his collaborators (e.g., Gigerenzer et al. 1999). Carruthers appears to take Gigerenzer's analysis of "fast and frugal" heuristics as suggestive of the general processes that we might expect to find in human cognition to enable suitably frugal computations.¹¹ However, the domains of inference for which (most of) Gigerenzer's heuristics perform best are not general enough to support the claim that the said heuristics are commonly employed for most human cognitive tasks. The tasks for which Gigerenzer's heuristics are best suited are decision tasks of choice between two alternatives, or of some value to be estimated among alternatives, and wherein the cognizer has limited familiarity of the alternatives. These tasks constitute only a small subset of the cognitive tasks humans

¹¹ Gigerenzer likewise believes that the majority of human inference and decision-making involves heuristics (Gigerenzer and Todd 1999; see also Gigerenzer 2007). But Gigerenzer's analysis extends to reasoning and inference *qua* central cognition, whereas Carruthers, in arguing for the massive modularity thesis (Sperber 1994), suggests that Gigerenzer-style heuristics can be deployed by the various modules that purportedly compose central cognition.

commonly face (cf. Sterelny 2003). Even for those instances when we have to make choices or estimate values, we are not always forced to choose among pairs of alternatives—we are often faced with many options. Moreover, when we are faced with tasks of choice or value-estimation, we are usually confronted with things we not only recognize but that we are very familiar with. Indeed, many of our choices demand an intimate acquaintance with the alternatives (cf. Sterelny 2003, 2006). It therefore seems as if Gigerenzer's proposed heuristics have only a limited role in our reasoning.¹² So we might set Carruthers' proposed solution to one side.

Samuels (2005, 2010), on the other hand, does not specify any heuristics that he thinks can ensure computational frugality in determining what representations get considered in a given cognitive task. He rests his arguments instead on an intuitive, though vague, appeal to web-search-engine-like heuristic or approximation techniques (Samuels 2010; cf. Shanahan 2009). The prospect of web-search-engine-like heuristics initially looks promising to deliver the desired frugality since actual web search engines are able to search through billions of websites in fractions of a second. Yet it is not entirely clear from Samuels' account exactly what role such heuristics would play in *relevance determinations* in human cognition. To be sure, web-search-engine-like heuristics can certainly search and retrieve information quickly, efficiently, and frugally, but these heuristics per se do not determine the relevance of the information they retrieve. This is made apparent by Samuels' (correct) assertion that the representations that a heuristic will bring to bear on a given cognitive task will not be invariably relevant. To put the point more directly, web-search-engine-like heuristics may very well circumvent the Computational Relevance Problem, but this solution does not address the Epistemological Relevance Problem (cf. Shanahan 2009).

Certainly, one can point to the advances in AI, computer programming, and other real-world computational applications, and proclaim that they are real instances where a program successfully deploys heuristics without having to go through computationally taxing processing to determine relevance beforehand (much like how web search engines are real-world examples of success). Samuels makes this point, as does Lormand (1996). However, citing concrete success stories for heuristics fails to acknowledge the fact that the AI researchers and computer programmers have already determined the relevance of certain information to certain tasks, and that they design heuristics to operate according to *those* determinations. This is what I take Fodor's point to be: In describing the frame problem as a problem of formalizing computationally relevant facts among irrelevant ones—and hence delimiting the candidates for belief-updating when certain events occur—he asserts,

The programmer decides, case by case, which properties get specified in the data base; but the decision is unsystematic and unprincipled. . . . [T]here appears to be no disciplined way to justify the exclusion [of properties that cause a system to update indefinitely many facts about the world]. (Fodor 1987, p. 33)

¹² This is in stark contrast to Gigerenzer's assertion that his research on heuristics is the study of "the way that real people make the *majority* of their inferences and decisions" (Gigerenzer and Todd 1999, p. 15, my emphasis). This also undermines Carruthers' suggestion that Gigerenzer's fast and frugal heuristics can provide tractable processes for cognition *tout court*.

The question is not: How do AI and computer programs determine what information is relevant? We already know the answer to this, namely certain methods are *programmed into* them. The question is rather: How did the AI *researchers* and computer programmers determine what information is relevant to what tasks in the first place? In the absence of an answer to this question, the epistemological problem remains.

Both Samuels (2005) and Carruthers (2006a) also appeal to the notion of *satisficing* to explain how human cognition can be computationally tractable. However, such an appeal is empty as it stands. For satisficing is really a substitution for a stopping rule for search among alternatives; this rule instructs search to stop once a certain goal is reached.¹³ But relevance determinations involve more than just search. Even if we generalize the notion of satisficing (which seems to be the case in many discussions in cognitive science) to a procedure that arrives at solutions that are not optimized but are in some sense “good enough”—processing (e.g., evaluation) continues until some specified aspiration level is reached—the very process *that* satisfices remains unaccounted for. That is, asserting that a system satisfices says nothing of how the goals or aspiration levels are specified, nor of how it is determined when these goals or levels are reached; and these are the very problems of relevance under consideration. Hence, heuristics can certainly satisfice, but it is a further question how they make relevance determinations in their processing. On this issue Samuels and Carruthers are silent.

Fodor rightly points out that (what I am calling here) the Epistemological Relevance Problem “goes as deep as the analysis of rationality” (Fodor 1987, p. 27). We can see why: the demands to be satisfied in determining what is relevant just are the demands of rationality. To repeat, humans appear to be largely rational—in Fodor’s words, “witness our not all being dead” (Fodor 2008, p. 116)—and determining what is relevant is part and parcel of rationality. Samuels (2005, 2010) demurs at Fodor’s suggestion that the totality of one’s beliefs must be consulted in determining relevance, since this is the “only guaranteed way” (Fodor 2000, p. 36) of classically computing a global property like relevance. As Samuels remarks,

guarantees are *beside the point*. Why suppose that we always successfully compute the global properties on which abduction depends? Presumably we do not. And one very plausible suggestion is that we fail to do so when the cognitive demands required are just too great. In particular, for all that is known, we may well fail under precisely those circumstances that the classical view would predict—namely, when too much of a belief system needs to be consulted in order to compute the simplicity or conservatism of a given belief. (Samuels 2005, p. 119)

Samuels is right that guarantees are beside the point (cf. Shanahan 2009; Sperber and Wilson 1996), and that we often fail in determining relevance. Nevertheless, we are still left to explain our successes in determining what is relevant to our cognitive

¹³ More precisely, a satisficing procedure sets an aspiration level appropriate to the task and goals of the agent; once this aspiration level is met, processing stops. As conceived by Simon (e.g., 1979), satisficing is a method to delimit the amount of time and resources devoted to a search task by avoiding exhaustive search and foregoing optimizing or maximizing aspiration levels.

tasks. And, again, our successes are not few. To paraphrase Fodor (1987),¹⁴ the moral is not that the heuristic solution to the frame problem—or to relevance problems more generally—is wrong; it is that the heuristic solution is empty unless we have, together with the heuristic strategy (or strategies), some idea of what is to count as relevant. Thus, despite the fact that heuristics can probably avoid the Computational Relevance Problem, we are still faced with the Epistemological Relevance Problem.

Let us briefly pause to see where we have gotten to. We have been discussing what heuristics are supposed to do for us in cognition. The part of cognition that is of concern is central cognition, since this is where problems of relevance arise. We saw that philosophers such as Samuels, Carruthers, and Lormand believe that heuristics offer a solution to problems of relevance.¹⁵ I argued, however, that heuristics do not offer the kind of solution that Samuels and Carruthers think; specifically, given the distinctions of the kinds of relevance problems I described above, heuristics do not offer a complete solution to relevance problems, as they fail to solve or circumvent the Epistemological Relevance Problem.

We also saw that Fodor believes that heuristics fail as a solution to relevance problems in central cognition. Yet Fodor's rejection of the heuristic solution is not motivated by the concerns I gave. Rather, Fodor believes that typical tasks in central cognition require global assessments of information, requiring in the limit consideration of one's entire set of beliefs, and this is something that heuristics cannot do. If Fodor is right, then it is hopeless to look to heuristics as a solution to relevance problems in central cognition. Nevertheless, I believe that the demands of global assessment that Fodor believes is placed on central cognition are too lofty (I am not sure if anyone except for Fodor thinks that central cognition is global in his sense; cf. Samuels 2010). Moreover, despite my arguments against Samuels', Carruthers', and Lormand's suppositions about the work heuristics do for us, nothing of what I have argued shows, nor was meant to show, that heuristics have no role in circumventing the epistemological worries associated with relevance problems. So there are two issues that remain to be addressed. The first is whether human cognition solves the Epistemological Relevance Problem. The second is what role heuristics play, if any, in solving relevance problems.

4 How Do We Solve Relevance Problems?

4.1 The Epistemological Relevance Problem Again

Let us begin by asking whether heuristics help in solving the Epistemological Relevance Problem in ways that we have yet to consider. The answer to this question is “No,” unfortunately. In fact, I believe that heuristic solutions fail generally for two reasons, which I will discuss in turn.

¹⁴ “The moral is not that the sleeping-dog strategy is wrong; it is that the sleeping-dog strategy is empty unless we have, together with the strategy, some idea of what is to count as a fact for the purposes at hand” (Fodor 1987, p. 31).

¹⁵ Notwithstanding that Carruthers believes that central cognition is massively modular.

First, of all our epistemic commitments, we in fact almost never *know* the bearing that a given commitment may have on a given cognitive task, at least not a priori. This is just to repeat Fodor's concerns, and this is precisely what makes the Epistemological Relevance Problem a problem. It is possible that something that presently appears to be irrelevant to a given task will be recognized as relevant only after the appropriate research is conducted, and once we have sufficiently built up our background knowledge. The times when we do or can know all that is relevant to a task will likely be confined to those instances when the problem is well-defined and specially circumscribed. This, it seems, is why science—the paradigmatic example of a holistic enterprise—requires teams of people doing lots of work over long periods of time; and even then, science does not always get everything right. In this sense, I agree with Fodor (2000) that solving the frame problem (or at least the epistemological aspect of concern here) is tantamount to considering the whole background of one's epistemic commitments, for this is the only way to guarantee that one has not missed anything relevant. Yet I disagree with him that our getting on in the world is proof that we actually solve the Epistemological Relevance Problem, i.e., that we actually somehow (noncomputationally) bring to bear the whole background of our epistemic commitments.

At this point, one may wonder whether the Epistemological Relevance Problem is really a problem. Indeed, this interpretation of the frame problem assumes a notion of relevance that is objective, and it claims that we are adept at knowing (at least some of) what is relevant in this sense to a given cognitive task. For lack of a better term, I will call this kind of relevance *objective relevance* (cf. Gabbay and Woods 2003). Objective relevance, as I am proposing to understand it here, refers to a kind of relevance that exists independently of cognizers. When x is objectively relevant to y , x bears on y in ways that support certain counterfactual propositions, such as “if x were different, then y would have been different”. Again, paradigmatic examples of objective relevance come from science. For example, the motions of terrestrial objects are relevant to the motions of the planets. The way in which the motions of terrestrial objects are relevant to the motions of the planets is an objective fact, though we did not know it until Newton came along. This relevance relation supports the counterfactual proposition, “if the motions of terrestrial objects were different, then the motions of the planets would have been different”.¹⁶ Now, when we understand the Epistemological Relevance Problem as claiming that we are adept at knowing (at least some of) what is objectively relevant in the sense described here, this might seem unrealistic, and consequently one may doubt the notion of relevance assumed by the Epistemological Relevance Problem. For indeed, not only will there undoubtedly be a great many things that are objectively relevant to a given phenomenon which we will not or cannot ever know, but it often takes a lot of work to discover what is objectively relevant to what, and it may assume too much of human cognitive abilities to assert that we are adept at knowing (at least some of) what is objectively relevant. Furthermore, one may argue that it is unrealistic to demand that humans pay

¹⁶ Notice that objective relevance also cannot be determined a priori; but rather than against a set of beliefs, relevance is determined against a background of facts and phenomena. The motions of terrestrial objects, for instance, are not objectively relevant per se; it is objectively relevant with respect certain phenomena, such as planetary motion.

attention to objective relevance, if this is tantamount to considering the whole background of one's epistemic commitments.¹⁷

However, I do not believe that these concerns are problematic for the Epistemological Relevance Problem. First, it is clear that science is concerned with objective relevance, and this is partly what leads me to agree with Fodor that the Epistemological Relevance Problem goes as deep as the analysis of rationality, and that solving it is tantamount to considering the whole background of one's epistemic commitments. Second, it may not be too much to demand of humans to pursue objective relevance, regardless of whether it is practically unattainable. The reason for this involves a long story having to do with an appropriate theory of rationality, but the point can be put briefly by invoking the notion of a *regulatory ideal*. As Cliff Hooker explains, a regulatory ideal is a Kantian notion "of a goal (destination or process) to be striven for, whether or not it can be actually achieved, because it is judged valuable, and attempts to move in its direction are feasible and have at least some beneficial consequences" (Hooker 1994, p. 218; cf. Hooker 2011, p. 162). If pursuing objective relevance is thus understood as a regulatory ideal, then it does not matter whether it is practically attainable. Certainly, we must have practically satisfiable *performance standards* that place normative demands in pursuit of satisfying ideals (Hooker 1994, 2011), and it is doubtful that pursuing objective relevance should make its way into these standards. But we ought not confuse performance standards with regulatory ideals. None of this is to say that an appropriate theory of rationality *must* have the pursuit of objective relevance as a regulatory ideal, but rather that, insofar as it is a *possible* regulatory ideal, we may very well demand of humans to pursue objective relevance even so it lies outside the bounds of what is practically attainable.¹⁸

The second reason why I think heuristic solutions to the Epistemological Relevance Problem fail is that there is nothing to suggest that heuristics, or cognition generally, is in the business of ensuring that nothing relevant has been missed (cf. Samuels 2005, 2010). Thus, I reject Fodor's claim that we in fact solve the problem, and that cognitive science struggles futilely in investigating out how we do it. To be sure, there is nothing to suggest that cognition attempts to solve the Epistemological Relevance Problem in the first place; in which case, it is not much of a problem. And indeed, we already know how to solve it, namely the hard way, notwithstanding that this is not practical in almost any case. It seems, rather, that most of human cognition is content with, and seems to get by on, *satisficing* judgments of relevance. And judgments that are in some sense good enough should not be confused with solving the Epistemological Relevance Problem. Ensuring certainty in our relevance determinations is severely cognitively taxing, requiring time and resources that we simply do not have in managing all of our day-to-day cognitive tasks. Consequently, complete or holistic relevance is an issue deserving serious attention in our cognitive lives only when we are doing certain high-stakes reasoning, such as in science or philosophy.

¹⁷ Thanks to Richard Samuels for bringing these points to my attention.

¹⁸ I do not claim that Cliff Hooker would endorse my arguments regarding objective relevance as a possible regulatory ideal. Whether he would, I suppose, depends on whether pursuing objective relevance is, in Hooker's terminology, a *degenerate idealization*, which is a non-perspicuously represented deviation of real behavior by real agents. I believe that pursuing objective relevance is not a degenerate idealization, but I am not prepared to argue the point here.

I therefore believe that any heuristic solution—or any solution generally—to the Epistemological Relevance Problem is empty in virtue of the problem being, by all practical considerations, unsolvable¹⁹ given our limited cognitive wherewithal (at least for the majority of interesting tasks), as well as being a problem that we are typically not in the business of solving in the first place. Therefore, we can disagree with Fodor (2000) that our inability to understand how we solve (the epistemological aspect of) the frame problem is detrimental to computational cognitive science; for indeed, we do not solve the problem, and so Fodor’s worries are moot.

But now the question arises: If we do not solve the Epistemological Relevance Problem, then how do we manage our day-to-day tasks and make reasonable decisions? How do we tend to bring to bear what is relevant if we do not know what is in fact relevant? This is just the issue that I had accused Carruthers (2006a, 2006b) and Samuels (2005, 2010) of skirting in their appeal to heuristics to circumscribe what gets considered in our cognitive tasks, but I will not likewise skirt the issue here. For indeed, we humans enjoy reasonable levels of success in our reasoning, and this must be explained.

4.2 A Further Guise of the Frame Problem

The explanation I will offer proceeds from Dennett’s account of the frame problem. He gives an example of making a midnight snack (complete with beverage) to illustrate that such mundane tasks require copious amounts of knowledge:

Now of course I couldn’t do this without knowing a good deal—about bread, spreading mayonnaise, opening the fridge, the friction and inertia that will keep the turkey between the bread slices and the bread on the plate as I carry the plate over to the table beside my easy chair. I also need to know about how to get the beer out of the bottle into the glass. Thanks to my previous accumulation of experience in the world, fortunately, I am equipped with all this worldly knowledge. . . . Such utterly banal facts escape our notice as we act and plan. (Dennett 1984, p. 134)

Dennett observes that it is implausible to think that such knowledge is stored in a human as separate declarative statements. We must therefore possess a system of representing this knowledge in such a way that it is accessible on demand. Dennett therefore believes that the frame problem for AI researchers is to design a system that is organized in such a way that achieves the efficient representation and access we observe in humans:

[The frame] problem concerns how to represent (so it can be used) all that hard-won empirical information Even if you have excellent knowledge (and not mere belief) about the changing world, how can this knowledge be represented so that it can be efficaciously brought to bear? . . . A walking encyclopedia will walk over a cliff, for all its knowledge of cliffs and the effects of gravity, unless it is designed in such a fashion that it can find the right bits of knowledge at the

¹⁹ By this I mean that we cannot solve the problem, given our limited cognitive wherewithal. I do not mean “unsolvable” in a sense that implies that the problem is not recursive.

right times, so it can plan its engagements with the real world. (Dennett 1984, pp. 140–1)

In this light, the frame problem can be understood as the problem of how cognition achieves the informational organization, and enables access to the relevant information at the right times; this seems to be required for human-like cognitive performance.²⁰ In light of these considerations, we have yet a further aspect of the frame problem, what I will call the *Representational Relevance Problem*:

Representational Relevance Problem The problem of how a cognitive system embodies the informational organization, and enables access to the right information for its cognitive tasks, that seems to be required for human-like cognitive performance.

Unlike the Epistemological Relevance Problem, the Representational Relevance Problem is indeed a real, nontrivial problem, somehow solved by human cognition. This aspect of the frame problem is what makes the frame problem (generally understood) worthwhile for cognitive science to investigate it and to understand how we manage to solve it. Perhaps we failed to recognize the importance of this representational guise of the frame problem because we were distracted by the computational framework of early AI research, and of the computational turn in philosophy. This framework requires that explicit sentence-like representations be stored in memory, and the appropriate (relevant) ones be summoned according to the cognitive task at hand. Certainly, scarcely anyone believes that this is how humans actually store beliefs or knowledge. Yet this is the framework within which the frame problem was first realized—the problem I identified at the beginning of this paper as the AI Frame Problem. We might note that the Representational Relevance Problem is a lot more like the original AI Frame Problem than the other guises presented above insofar as the representational problem and the AI problem both concern the representation of knowledge. At the same time, we must understand that, among other distinctions between these two interpretations of the frame problem, the Representational Relevance Problem is not committed to sentence-like representations, and moreover, the Representational Relevance Problem bears an intimate relation to the computational problem concerning the facilitation of processes to determine what information to bring to bear.

Recall that the Computational Relevance Problem is the problem of how a cognitive system tractably delimits what gets considered in a given cognitive task. We saw above that various heuristic solutions to this computational problem have been proposed, most notably by Samuels (2005, 2010) and Carruthers (2006a, 2006b). I had said that such heuristic solutions may very well solve the computational problem, but they do not solve the epistemological problem. But now that we have set the epistemological problem aside, we might see more clearly why heuristics *can* solve the computational problem,

²⁰ As a point of interest, this idea was intimated by Fodor in *The Language of Thought*, though his purposes for broaching the issue were slightly different. He there commented that “a fundamental and pervasive feature of higher cognitive processes [is] *the intelligent management of internal representations*” (Fodor 1975, p. 164, emphasis in original).

and indeed can solve the problem in an interesting way.²¹ Of course, the computational problem can be solved in an uninteresting way; example: for a given cognitive task, limit what gets considered by referring only to the top three items on one's grocery list. Certainly, this technique (I balk at calling it a heuristic) will avoid computational worries of tractability, but it likely will not be of any help for the task at hand (unless the task concerns groceries, but even then it is questionable whether the technique will be of any help). On the other hand, the computational problem can be solved in an interesting way, insofar as what actually gets considered—what information a suitably tractable technique picks out—bears in helpful ways. Heuristics solve the Computational Relevance Problem in this interesting way. Within the current discussion this is no longer tantamount to solving the Epistemological Relevance Problem. I claim instead that the computational problem is interestingly solved in virtue of the representational problem being solved—it is because human cognition (somehow) manages to solve the Representational Relevance Problem that effective computational methods (heuristics) can be deployed, to not only delimit what is considered in a cognitive task, but to bring to bear *helpful* information.

Imagine a very large bookcase filled with hundreds of books, and suppose you want to learn about the mystery-solving abilities of Sherlock Holmes. You know that within the many books in the bookcase there is lots of information about Sherlock Holmes. Perhaps you want to retrieve *The Hound of Baskervilles*, which is contained in the collection. Where to begin? Well, it depends on how the books are organized on the shelves. Finding out how the bookcase is organized is easy enough by empirical means and a bit of induction. If the books are organized according to author, you might start by looking among the shelves near the top of the bookcase for Doyle's works. However, if the books are organized according to genre, you would have to do a little more work to find where the mystery novels are located. This may require you to glean information about the subject matter of other books: You see *Pride and Prejudice*, and so you go on to some other section of the bookcase; you see *Murder on the Orient Express*, and you know you are on the right track. But now what? Do you start to scan left or right, up or down? Again, it depends. Are the books within genres organized according to author or title? The point here is that the efficiency and success of whatever method(s) you employ in your search will depend on how and the extent to which the books are organized. The method that you employ in one case may not be as efficient, or may not be successful at all, when applied in another case. If the books are placed randomly within the bookcase, any method you employ will not match the efficiency and success of methods applied to an organized bookcase, except perhaps by pure luck. (Notice also that consulting *The Hound of Baskervilles*, or indeed all of the Sherlock Holmes stories, may not be sufficient for the task of learning about the mystery-solving abilities of Sherlock Holmes. One may want to additionally consult biographies on Doyle, commentaries on the Holmes stories, the

²¹ I note in passing that, although I speak here of heuristic solutions generally, I take Samuels' gesture toward web-search-engine-like techniques to be more informative than Carruthers' gesture toward Gigerenzer's fast and frugal heuristics. As discussed above, Gigerenzer's fast and frugal heuristics do not seem to apply to the representative tasks of human cognition.

methods of real-life detectives, and maybe more. How might one begin to search for all of this? How might one be sure not to miss anything relevant? We can see here relevance problems recapitulated.)

The lessons for understanding of how human cognition manages to solve the Computational Relevance Problem should be obvious. The methods used to solve the Computational Relevance Problem (e.g., heuristics) will be efficient and successful (in the interesting sense mentioned above) only insofar as the stockpile of information to which the methods are being applied is organized in specific ways. Any sort of heuristic can operate over generous amounts of information, and thereby enable fast and frugal cognition in the uninteresting sense. But structured or organized information plays a crucial role in enabling heuristics to perform well, i.e., (reasonably) robustly and successfully. Contrariwise, information that is unstructured or unorganized in specific ways will retard the efficiency and success of search and processing. This is why I claim that it is in virtue of the Representational Relevance Problem being solved that the Computational Relevance Problem is solved. For if computational methods work to bring to bear *helpful* information, a cognitive system must requisitely embody the informational organization to enable access to the right sorts of information for its cognitive tasks.

We might now begin to see how we enjoy our characteristic levels of success in our reasoning. We may not be in the business of solving the Epistemological Relevance Problem, but, given the present account, we have more constraints on our reasoning than we know. The organization of our cognitive architecture appears to be enough for humans to get by on. There will certainly be times when we fail to process objectively relevant information (in the sense described above), or when we process information that is not very objectively relevant at all. In some cases we may end up processing some objectively *irrelevant* information. Moreover, cognitive limitations, fatigue, stress, and other extraneous factors will sometimes prevent us from processing relevant information. Nevertheless, our cognitive architecture provides a network that informs us of what to consider, and under satisfactory conditions its processes are constrained and guided accordingly (sometimes via heuristics). The situation may not be ideal, but it is good enough for us to get by on—indeed, such is to be expected from satisficing organisms. On the other hand, when we enter into certain high-stakes arenas, such as science or philosophy, we alter our standards, and whatever is picked out by the organization of our cognition may no longer be good enough. In such circumstances, objective relevance is sought, and again, this is why progress and getting things right are much more difficult to achieve in these endeavors.

All this should not be taken to imply that solving the Representational Relevance Problem only requires some sort of know-how. For instance, Dennett's midnight snack example, and all my talk of successful action, may give the impression that solving the Representational Relevance Problem just consists in knowing how to successfully navigate the world (thus accessing the right information at the right times). Though this is definitely part of the story I am attempting to convey here, it is not the complete story. For the Representational Relevance Problem is not simply a problem of know-how; importantly it is a problem involving propositional knowledge as well. When Dennett fixes himself a midnight snack, he knows *that* mayonnaise does not dissolve knives on contact, and *that* opening the refrigerator does not

cause a nuclear holocaust in the kitchen (Dennett 1984, p. 135); and Dennett's knowledge of how to make a midnight snack entails he knows *that* these facts are not relevant to making a midnight snack. To turn away from examples concerning practical action, we might consider cognitive tasks that do not result in action. Predicting which team will win tonight's football game is like this. You know *that* the home team's quarterback is injured, and *that* this fact bears on your prediction. Solving the Representational Relevance Problem does not always involve, and is not just about, knowledge-how.

At this point, one may be reminded of the philosophical debate concerning whether knowledge-how is a species of knowledge-that (see e.g., Ryle 1949; Stanley and Williamson 2001; Snowdon 2004; Noë 2005). It is beyond the scope of this paper to wade into the muddy waters of this debate. I am inclined to agree with Snowdon (2004), who argues that a great deal of know-how consists in possessing propositional knowledge. Knowing how to bake bread, for example, may consist *inter alia* in knowing that flour is the main ingredient. If something like this is true for much of our knowledge-how, then perhaps the know-how part of solving the Representational Relevance Problem partly consists in possessing propositional knowledge. I believe this is in fact the case. At the same time, however, the propositional knowledge involved in the know-how part of solving the Representational Relevance Problem does not have to be declarative; such knowledge can be *tacit*. We may not always be in a position to offer up propositions concerning how we determine what is relevant to our tasks, but we still know how to do it. As Fodor (1968) warns, we should not "confuse knowing that with being able to explain how" (p. 634). Thus, knowing what to bring to bear on our tasks may be partly constituted by an epistemic relation to a proposition without that relation being of the knowledge-that sort desired by intellectualists (where we are able to report these propositions).

These matters deserve more attention than I can devote here. The crucial point is that solving the Representational Relevance Problem requires us to bear an epistemic relation to *highly organized knowledge*. As Dennett observed, it is implausible to think that our knowledge is stored as separate declarative statements. But how else might our knowledge be represented and organized? This, of course, is just the problem with the Representational Relevance Problem—human cognition manages to solve it, yet we struggle to understand how.

4.3 A Gesture Toward a Solution

Before concluding this paper, I should like to suggest a possible way in which human cognition manages to solve the Representational Relevance Problem. Since I lack the space to fully develop and defend the account, what I give here is nothing more than a gesture toward a promising possibility.

I believe that something along the lines of the *file model* of cognition provides an appropriate account of mental representation (e.g., Evans 1982; Fodor 2008; Lawlor 2001; Récanati 1993; see also Kahneman et al. 1992; Pylyshyn 2003; Treisman 1982). This model can support a picture of cognition wherein vast connections exist between representations. Every symbol would be part of an interconnected network, where activation can spread through the system depending on the strengths and the

organization of the connections between symbols in the network. However, in principle, the structure allows for any content to activate any other content, which is a hallmark of central cognition (cf. Viger 2006c). This type of picture appears to be just what is needed to solve the Representational Relevance Problem.

Now, what our cognitive architecture picks out as the right things to consider—as what is relevant to the task at hand—must often turn out to be objectively relevant. This is evident from how humans get on in the world, and the success rate of many human inferences. The correlation is certainly not a result of chance. Again, providing an account of why this is so is beyond the scope of the present paper. However, insofar as our representational systems track things in the world, they will reflect the organized structure of information in the world, including objective relevance relations. Of course, we must acknowledge that some of our representational systems do not always track the truth (see e.g., Plantinga 1993; Viger 2006b). And moreover, though we have a reasonable success rate at determining what is relevant, this should not overshadow our many mistakes and errors. All this is consistent with the manifest fact that we often get things wrong. However, all we need to explain our reasonable levels of success at determining what is relevant to our tasks is that we get things right enough to satisfy our goals (cognitive and practical); and I believe our getting things right enough can be accounted for by our representational systems that track truth, *to whatever extent those systems do so*.

5 Concluding Remarks

Given the explication of the Representational Relevance Problem presented above, we might see that the problem has shifted from understanding how humans manage to determine relevance (the epistemological problem) to understanding how humans manage their knowledge to bring to bear the right bits of information at the right times. These are different issues—we do not, and need not, solve the first (as I have argued), but we do solve the second. And moreover, it is in virtue of solving the representational problem that we are able to solve the computational problems associated with the frame problem.

Understanding how humans solve the Representational Relevance Problem should be a central interest of philosophers and cognitive scientists. Understanding the sort of architecture that allows us to consider the right information at the right times most certainly would be valuable. This would require turning our focus from trying to discover what information is objectively relevant to trying to find out what information is *useful* for a cognitive system. Useful information does not have to be relevant in the sense that the epistemological problem is concerned about, but useful information will be what is needed to guide successful action and everyday thought. The current position of Jupiter may very well be relevant to how I should position my hand when I pick up a glass of beer (see Fodor 2008, p. 117; cf. Kyburg 1996), but I do not need to know that to pick up a glass of beer. Likewise, the number of titles housed in the Library of Congress may be relevant to the outcome of tonight's football game, but I do not need to know that to make a prediction on the outcome.

At the same time, solving the Representational Relevance Problem allows us to *avoid* the epistemological worries of the Epistemological Relevance Problem. The

epistemological problem is concerned with how a cognitive system considers (mostly) only what is relevant, or how a cognitive system knows what is relevant. But when we acknowledge at once that we do not solve this problem and that we have the sort of cognitive architecture that enables access to the right sorts of information at the right times, then the epistemological “worries” cease to be worries. Solving the representational problem is all that is required for human-like performance in thought and action, and I do not think that we should expect, or even desire, more than that.

Acknowledgments Many thanks to Chris Viger for all his comments and suggestions. Thanks also to Richard Samuels and Rob Stainton for comments and input on earlier drafts. And thanks to an anonymous referee of this journal whose comments helped to improve the latter part of this paper.

References

- Carruthers, P. 2003. On Fodor’s problem. *Mind Lang* 18: 502–523.
- Carruthers, P. 2006a. *The architecture of the mind: Massive modularity and the flexibility of thought*. Oxford: Oxford University Press.
- Carruthers, P. 2006b. Simple heuristics meet massive modularity. In *The innate mind: Culture and cognition*, ed. P. Carruthers, S. Laurence, and S. Stich, 181–198. Oxford: Oxford University Press.
- Dennett, D.C. 1984. Cognitive wheels: The frame problem of AI. In *Minds, machines, and evolution*, ed. C. Hookway, 129–152. Cambridge, MA: Cambridge University Press.
- Evans, G. 1982. *The varieties of reference*. Oxford: Oxford University Press.
- Fodor, J.A. 1968. The appeal to tacit knowledge in psychological explanation. *The Journal of Philosophy* 65: 627–640.
- Fodor, J.A. 1975. *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J.A. 1983. *The modularity of mind*. Cambridge, MA: The MIT Press.
- Fodor, J.A. 1987. Modules, frames, fridgeons, sleeping dogs and the music of the spheres. In *Modularity in knowledge, representation and natural-language understanding*, ed. J.L. Garfield, 25–36. Cambridge, MA: The MIT Press. Reprinted from *The robot’s dilemma: The frame problem in artificial intelligence*, Z. Pylyshyn, ed., 1987, Norwood, NJ: Ablex.
- Fodor, J.A. 2000. *The mind doesn’t work that way: The scope and limits of computational psychology*. Cambridge, MA: The MIT Press.
- Fodor, J.A. 2008. *LOT 2: The language of thought revisited*. Oxford: Clarendon.
- Gabbay, D., and J. Woods. 2003. *A practical logic of cognitive systems, volume 1. Agenda relevance: A study in formal pragmatics*. Amsterdam: Elsevier.
- Gigerenzer, G. 2007. *Gut feelings: The intelligence of the unconscious*. New York: Viking.
- Gigerenzer, G., and D.G. Goldstein. 1999. Betting on one good reason: The Take the Best heuristic. In *Simple heuristics that make us smart*, ed. G. Gigerenzer, P.M. Todd, and the ABC Research Group, 75–95. New York: Oxford University Press.
- Gigerenzer, G., and P.M. Todd. 1999. Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*, ed. G. Gigerenzer, P.M. Todd, and the ABC Research Group, 3–34. New York: Oxford University Press.
- Gigerenzer, G., P.M. Todd, and the ABC Research Group (eds.). 1999. *Simple heuristics that make us smart*. New York: Oxford University Press.
- Haugeland, J. 1985. *Artificial intelligence: The very idea*. Cambridge, MA: The MIT Press.
- Hooker, C.A. 1994. Idealisation, naturalism, and rationality: Some lessons from minimal rationality. *Synthese* 99: 181–231.
- Hooker, C. 2011. Rationality as effective organisation of interaction and its naturalist framework. *Axiomathes* 21: 99–172.
- Kahneman, D., A. Treisman, and B.J. Gibbs. 1992. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 24: 175–219.
- Kyburg Jr., H.E. 1996. Dennett’s beer. In *The robot’s dilemma revisited: The frame problem in artificial intelligence*, ed. K.M. Ford and Z.W. Pylyshyn, 49–60. Norwood, NJ: Ablex.

- Lawlor, K. 2001. *New thoughts about old things: Cognitive policies as the ground of singular concepts*. New York: Garland Publishing.
- Lormand, E. 1990. Framing the frame problem. *Synthese* 82: 353–374.
- Lormand, E. 1996. The holobrophobe's dilemma. In *The robot's dilemma revisited: The frame problem in artificial intelligence*, ed. K.M. Ford and Z.W. Pylyshyn, 61–88. Norwood, NJ: Ablex.
- McCarthy, J., and P. Hayes. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine intelligence*, ed. B. Meltzer and D. Michie, 463–502. Edinburgh: Edinburgh University Press.
- McDermott, D. 1987. We've been framed: Or, why AI is innocent of the frame problem. In *The robot's dilemma: The frame problem in artificial intelligence*, ed. Z.W. Pylyshyn, 113–122. Norwood, NJ: Ablex.
- Noë, A. 2005. Against intellectualism. *Analysis* 65: 278–290.
- Plantinga, A. 1993. *Warrant and proper function*. New York: Oxford University Press.
- Pylyshyn, Z.W. 1996. The frame problem blues. Once more, with feeling. In *The robot's dilemma revisited: The robot's dilemma revisited: The frame problem in artificial intelligence*, ed. K.M. Ford and Z.W. Pylyshyn, xi–xviii. Norwood, NJ: Ablex.
- Pylyshyn, Z.W. 2003. *Seeing and visualizing: It's not what you think*. Cambridge: The MIT Press.
- Récanati, F. 1993. *Direct reference: From language to thought*. Oxford: Blackwell Publishing.
- Ryle, G. 1949. *The concept of mind*. London: Hutchinson.
- Samuels, R. 2005. The complexity of cognition: Tractability arguments for massive modularity. In *The innate mind: Structure and contents*, ed. P. Carruthers, S. Laurence, and S. Stich, 107–121. Oxford: Oxford University Press.
- Samuels, R. 2006. Is the human mind massively modular? In *Contemporary debates in cognitive science*, ed. R.J. Stainton, 37–56. Malden, MA: Blackwell.
- Samuels, R. 2010. Classical computationalism and the many problems of cognitive relevance. *Studies in History and Philosophy of Science* 41: 280–293.
- Shanahan, M. 2009. The frame problem. In E.N. Zalta (Ed.) *The Stanford encyclopedia of philosophy* (Winter 2009 Ed.). <<http://plato.stanford.edu/archives/win2009/entries/frame-problem/>>.
- Simon, H.A. 1979. *Models of thought*. New Haven: Yale University Press.
- Snowdon, P. 2004. Knowing how and knowing that: A distinction reconsidered. *Proceedings of the Aristotelian Society* 104: 1–29.
- Sperber, D. 1994. The modularity of thought and the epidemiology of representations. In *Mapping the mind: Domain specificity in cognition and culture*, ed. L.A. Hirschfeld and S.A. Gelman, 39–67. Cambridge, MA: Cambridge University Press.
- Sperber, D., and D. Wilson. 1996. Fodor's frame problem and relevance theory (reply to Chiappe & Kukla). *The Behavioral and Brain Sciences* 19: 530–532.
- Stanley, J., and T. Williamson. 2001. Knowing how. *The Journal of Philosophy* 97: 411–444.
- Sterelny, K. 2003. *Thought in a hostile world*. Malden, MA: Blackwell.
- Sterelny, K. 2006. Cognitive load and human decision, or, three ways of rolling the rock uphill. In *The innate mind: Culture and cognition*, ed. P. Carruthers, S. Laurence, and S. Stich, 218–233. Oxford: Oxford University Press.
- Treisman, A. 1982. Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance* 8: 194–214.
- Viger, C. 2006a. Frame problem. In *Encyclopedia of language and linguistics*, ed. K. Brown. Oxford: Elsevier.
- Viger, C. 2006b. Is the aim of perception to provide accurate representations? A case for the “no” side. In *Contemporary debates in cognitive science*, ed. R.J. Stainton, 275–288. Malden, MA: Blackwell Publishing.
- Viger, C. 2006c. Symbols: What cognition requires of representationalism. *Protosociology Int J Interdiscipl Res* 22: 40–59.