

SCOTT HENDRICKS

THE FRAME PROBLEM AND THEORIES OF BELIEF

ABSTRACT. The frame problem is the problem of how we selectively apply relevant knowledge to particular situations in order to generate practical solutions. Some philosophers have thought that the frame problem can be used to rule out, or argue in favor of, a particular theory of belief states. But this is a mistake. Sentential theories of belief are no better or worse off with respect to the frame problem than are alternative theories of belief, most notably, the “map” theory of belief.

1. INTRODUCTION

Human survival requires that we quickly and selectively apply what we know about the world to the situation at hand and generate useful conclusions. This is simple enough in practice, but cognitive science is far from providing a satisfactory account of how we do it. Clark Glymour states the problem this way: “Given an enormous amount of stuff, and some task to be done using some of the stuff, what is the *relevant stuff* for the task?” (1987, p. 65). Among philosophers, this is commonly known as the *frame problem*.¹

The frame problem is a research problem. When we figure out how human beings effortlessly recover information relevant to their present conditions, we will have solved it. In the meantime the frame problem is a useful tool for helping philosophers and cognitive scientists explore various models of human cognition. Competition between opposing models will perhaps be decided in part by looking to see how well each solves the frame problem.

It is not surprising then that we see philosophers appealing to the frame problem as a means for deciding between theories of belief and mental representation. Patricia Churchland writes:

It is now deeply puzzling how the robot might be instructed so as not to be a fool, a problem that in AI research is called the *frame problem*. How do humans manage not to be fools? What does our “common sense” or “intelligence” consist in? The more we try to solve the robot’s problem of sensible behavior, the more it becomes clear that *our* behavior is not guided by explicit sentential instructions in our store of knowledge. Specifying the knowledge store in sentences is a losing strategy.²

According to John Haugeland:

The main suggestion is that the frame problem may be an artifact of assuming that mental representation is quasi-linguistic. Since, however, that is one of the basic enabling assumptions of “classical” artificial intelligence..., there seems little room for compromise: Perhaps the frame problem is really a “pseudo-problem” foisted on unsuspecting epistemology by misguided presumptions underlying AI as a discipline.³

These philosophers claim that no machine – the human brain included – whose medium of representation is sentence-like could instantiate a solution to the frame problem. Haugeland goes on to suggest that a map- or model-like system of mental representation might fare better. It is not an implausible suggestion, for it does seem that a machine that represents the world using maps would have a decided advantage in calling up relevant information on the fly over a machine that must pick through a collection of sentences.

Does the existence of the frame problem help us decide what kind of representational medium beliefs employ? I argue that it does not. Neither sententialism nor the map theory of belief has any advantage over the other in solving the frame problem. To see why, I will first discuss reasons against the thesis that the frame problem generates special problems for sententialism. Second, I explain why a map theory of belief does not make solving the frame problem easier.

One terminological caveat. There is an unfortunate confusion surrounding the expression, “the frame problem.” The confusion concerns which problem should properly be called the “frame problem.” The original frame problem, introduced by McCarthy and Hayes (1969), is essentially the problem of how an agent can keep track of persisting facts when reasoning about change (see Pollock and Cruz, 1999 for a nice treatment).

Researchers in artificial intelligence have continued to use the expression this way. But philosophers have complicated matters. Hayes (1987) points out that philosophers have frequently referred to the problem of relevance (also known as the “problem of control”) as the “frame problem.” Glymour, Churchland and Haugeland, in the quotations above, are doing just this.

Trying to correct the mistake here is impractical. Accordingly, I will follow their use – or misuse – of the expression. But I should emphasize that this so-called confusion is a nothing more than a *terminological* confusion. For a nice discussion of the various problems which have traveled under the name, “the frame problem,” see Lormand (1990, 1996).

2. THE FRAME PROBLEM AND SENTENTIALISM

Sententialism is a claim about the structure and richness of our mental representations – in particular, the mental representations involved in belief and the other propositional attitudes.⁴ Sententialism about *representational structure* is most famously associated with Fodor’s language of thought hypothesis.⁵ Belief contents are inner sentences whose internal structure is roughly isomorphic to linguistic structure. Sentential *richness* is a different, but related issue. The contents of causally distinct belief states are roughly the propositions expressed by the embedded that-clause sentences we use to attribute those states.

Why might the frame problem be a *special* problem for the sentential approach to belief content? Churchland puts forward the following consideration.⁶ Suppose you build a robot that stores its knowledge about the world in the way suggested by the sentential approach – its knowledge is a collection of sentences. For the robot to use this knowledge for practical, real-time problem-solving, it must be able to distinguish relevant sentences from irrelevant sentences, and relevant inferences from irrelevant inferences. Churchland claims it is implausible to suppose we could ever find mechanisms that would allow the machine to do this. Although she does not say much more on the issue, we need not look far for a plausible reconstruction of her concern.

The research question posed by the frame problem concerns how a cognizer – human or artificial – succeeds in using the information it stores exactly when it should use that information. To answer the question, cognitive scientists have to tell us both how cognizers store information about the world, and what mechanisms are in place for using the information once it is stored. If you answer the first question using sententialism, then our prospects for answering the second look gloomy. In the first place, a review of the AI literature reveals that every attempt to construct mechanisms for using information stored in sentences such that a machine could perform practical, real-time, problem-solving, has been a miserable failure. This is a pessimistic induction: So far, attempts to construct the necessary mechanisms have failed; therefore the mechanisms probably do not exist.

Second, there are no *general* principles for deciding what counts as a good inference in a particular situation. This is revealed by examining the role of *context* in intelligent “real-life” problem-solving.⁷ What counts as the correct use of the relevant knowledge is highly sensitive to context. But every context is highly peculiar and idiosyncratic. There are, in other words, no general principles for characterizing what will *matter* in one context rather than another. And thus, there can be no general principles for determining what counts as the right use of right knowledge at any given time – in short, there are no general principles of inference which generate solutions to all the “real-life” practical problems normal humans face.

Unfortunately, if there are no such principles, then likewise there cannot be any mechanisms for using information stored in sentential form to solve “real-life” practical problems. After all, if we are going to avoid the charge that the mechanisms for useful sentence retrieval are really just very smart homunculi, the mechanisms have to be stupid procedures that follow useful general principles.

Although this line of argument could stand some further refinement, that is not necessary, for an important difficulty with the argument is already easy to see. Suppose we grant that context makes impossible the existence of general principles of

inference for generating conclusions to practical problems. Does the non-existence of sentence-using mechanisms follow from the non-existence of general principles of inference? Clearly it does not. Let me explain.

Suppose you think that successful knowledge control should be implicit in the manner of knowledge storage. It is like storing boxes in the trunk of your car. If you have to deliver the boxes in a certain order, you should start by packing the box to be delivered last. Thus, when it comes time to deliver the first box, it is the most readily accessible box in the trunk. In doing this, you have ensured that the organization of boxes in the trunk facilitates retrieving them conveniently at the right time. We want to do something similar in the case of sentence storage. The advantage of knowledge organization is that so long as we organize the sentences correctly, the mechanisms for exploiting this organization can be very basic. All the important work is done in the manner of storage.

Here is the reason this is helpful: There does not have to be any set of principles of inference for always generating the desired practical conclusions. There does not have to be, in other words, any principle for how to organize the robot's sentences. Suppose I am designing the robot in question, and you ask me, "Why did you put *that* sentence *there*?" I will answer, "That sentence has to be there so it will be available in such-and-such circumstances." If you push the matter and complain, "But that won't make it available in these *other* circumstances," then I will gerrymander the organization a bit until it comes out right, and so on until I get the robot to perform well. If this is how nature designed us, then things look pretty bad for a cognitive science that aims to articulate a set of laws governing all practical and theoretical reasoning.⁸ But the issue of whether such laws exist is relatively independent from whether human beings could have been designed to solve practical problems using mental sentences. The only connection is that if such laws do not exist, then humans get the job done on the basis of a very untidy bit of machinery. And this untidy machinery may or may not use sentences.

So the only special problem for the sentential approach posed by the frame problem is the pessimistic induction: So far, cognitive science – AI in particular – has failed to articulate a set of sentence-manipulating mechanisms for solving the frame problem. Thus we should conclude that no such mechanisms are available. Unfortunately, I do not think that many will find the induction very compelling. It may provide a motivation for entertaining alternatives to the sentential approach, but it is a far cry from showing that there is any in principle obstacle to the possibility of a sentential cognizer. In sum, Churchland's appeal to the frame problem has not given us a strong motive for abandoning sententialism. There may be other reasons for doing so; reasons independent of how sententialism fares in solving the frame problem. But that is a different matter altogether.

On the other hand, a clear reason for abandoning the sentential theory would be the existence of an alternative account of belief content which can solve the frame problem, or at least offer a stronger hope for solving the problem. It is to this possibility I now turn.

3. MODELS AND THE FRAME PROBLEM

Haugeland has suggested that rejecting sententialism about belief content in favor of a map- or model-like representational medium might make solving the frame problem easier.⁹ Sentential representations are distinctive in being “independent” of each other, but it is precisely the “independence” of mental sentences that makes the frame problem seem so difficult to solve.

The notion of independence can be brought out in two ways. First, when a cognizer harbors independent representations, it is frequently appropriate to distinguish between explicit and implicit knowledge. Suppose that Lucy harbors the mental sentence that Tucson is south of Phoenix. Presumably, with little effort, she could infer that Phoenix is north of Tucson. This fact – the fact that she *could* draw the inference with ease – is sometimes captured by saying that Lucy knows *implicitly* that

Phoenix is north of Tucson. Implicit content is the information implied by a set of explicit contents. So if our mental contents are sentential, and thus independent of each other, then we can draw a distinction between our implicit and our explicit knowledge.

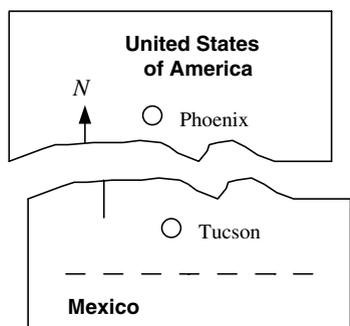
The notion of independent representations can be drawn out in a second way. If we represent the world using independent representations (e.g. mental sentences), then updating one representation does not of necessity change any other. For the acquisition of new information to have the effect it should, every alteration must be made separately. Suppose Winnie believes two things: she believes that New York is east of Boston, and that Boston is west of New York. Later, she gives up the first belief when a friend tells her, “Winnie, New York is *west* of Boston.” Of course in order to behave rationally she should give up the second belief as well. But on the sentential model this change must be made separately, for the information that Boston is west of New York, since it is the content of a separate belief, is independent from the information that New York is east (or west) of Boston.

Independence, then, has the consequence that in addition to storing our information about the world in sentence form, we need to figure out how cognition *uses* (retrieves and manipulates) those sentences. This consequence is not a surprising one. It corresponds exactly with the two questions cognitive science needs to address in order to answer the research question posed by the frame problem: (i) What medium of representation? and (ii) How do we use those representations?

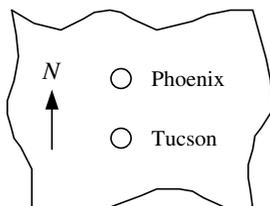
Now notice that if instead our information about the world is borne by inner *models* – inner structures which are isomorphic to what is represented¹⁰ – problems concerning use seem to disappear. According to Haugeland, the advantage of a model-like system of representation is that there is no distinction between explicitly represented properties and implicitly represented properties – model-like representation is “complicit.”¹¹ This means that when one feature of the model is updated, other features are updated automatically – the parts

of the representation are of necessity not independent of each other in the way that sentences are.

The point can be brought out using an example offered by a proponent of a “map”¹² theory of belief content, David Lewis.¹³ Compare an ordinary map and a sheet of paper with a list of sentences on it. We will suppose that both aim at describing some geographical state of affairs – say, the relations between the cities of Tucson and Phoenix. Now notice, I can tear off a part of the list so that the only information represented on that piece of paper is the information that Tucson is south of Phoenix, while the information ‘Phoenix is north of Tucson’ is preserved on the other piece of paper. The same is not true for my map. I cannot tear off a part of the map so that only the information that Tucson is south of Phoenix is represented. Suppose I tear the map like so:



Now, neither portion of the original map successfully represents the information that Tucson is south of Phoenix. The best I can do is to tear the map this way:



This piece of the original map successfully represents the information that Tucson is south of Phoenix. Yet it continues to represent the information that Phoenix is north of Tucson

and it cannot help but do so. Both pieces of information are *complicit* in the map, even though they can be distinguished when represented by sentences. Haugeland's suggestion is that much of the information carried by our beliefs is complicit in this way. That is what it means to suggest that the representational medium is model- or map-like. And this is what is supposed to make an internal model, or "isomorph," useful.

The map proposal seems to make the frame problem easier to solve in the following way. Map-like belief represents the world by employing structures that are isomorphic to what is represented and the information possessed by such structures is "complicit." Thus the frame problem does not appear to be an obstacle to the map theory because the hard question posed to the sentential theory – How do you organize sentences? – has no analog on the map theory. It is perhaps easiest to see this point specifically in connection with knowledge recovery. With some initial prompts in place – perception, or the tugging of wants and goals – the isomorph does the rest, just as easily as the world itself would. Information is "recovered" in virtue of the relevant parts of the isomorph being exercised in a kind of simulation. Colin McGinn confirms this picture: "A thinking system, we might say, is a *simulation engine* – a device that mimics, copies, replicates, duplicates, imitates, parallels reality" (1989, p. 176).

4. PROBLEMS FOR THE MAP THEORY

If a system of mental representation makes it easier to see how an agent can solve the frame problem, it has a decided advantage over other theories. In what follows I will argue the map theory does not possess the advantages described in the section above.

We are considering the claim that belief is a self-contained working model of the world. How does a cognitive system use that model? Answering this question is not as easy as it seemed earlier. The problem of how we get at the information contained in the model is as difficult as the problem of how we get information from the world itself.

The map approach was supposed to give us a leg up in solving the frame problem because, according to that approach, accessing the right knowledge in the right circumstances is as natural as an agent's merely being in the world and doing its thing – the mental model parallels reality. This is McGinn's point about *simulation* in the quotation above.

But this gets things the wrong way around. In general, we start with questions about how rational agents – humans in particular – navigate their world intelligently. Cognitive science appears on the scene and offers us the following general answer to these questions: Rational agents are, in part, cognitive systems; and cognitive systems, in virtue of the activity of the parts that possess representational properties, make possible intelligent navigation. In short, cognitive science aims at explaining how agents ensconced in their world do the things they do, and it tries to provide this explanation by appealing to the representational paradigm. The map approach, on the other hand, appears to have left us with the problem of explaining how a cognitive system can navigate its way around a worldly isomorph, and thus, this approach leaves us with a problem that is as difficult as the problem the representational paradigm was originally intended to solve. This is, in general, not the way good science proceeds. And it is certainly no answer to the frame problem.

Notice the objection is not a “homunculus regress” objection. I am not claiming that in order to explain how the inner model gets used, the map theorist must posit a homunculus. McGinn correctly remarks that “we do not need to presuppose beliefs *about* mental models in order for mental models to act as the machinery of belief. Mental models do not need to be *interpreted* – they just need to be used” (1989, p. 200). True enough – nothing about the map approach implies that a homunculus is *necessary* to use a mental model. The present worry is slightly different. At present we have no reason to think that building a machine that can use mental models in a natural and efficient way is going to be any easier than building a machine that uses sentences.

The map theorist, though, may respond that it is unfair to the map theory to assume that the internal model was meant to be a *self-contained* working model. So let us examine a more plausible version of the map theory. Instead of a self-contained working model, the map theorist might posit an internal model that is isomorphic to only certain aspects of the world. The rest of the isomorphism is taken over by the mechanisms for doing things with the model. Rather than packing all the representational properties into the model, we want our representation of the world to arise partly out of the model and partly out of our *manipulations of* the model. McGinn's model theory runs along these lines. He writes:

The procedures that operate on these models themselves model external processes: mental processes replicate worldly processes, mental laws imitate physical laws. There is a kind of isomorphism between the world and the mind.¹⁴

McGinn's suggestion is this: Suppose we divide the world into dynamic and static features. A static feature of the world is this desk in front of me at this time, as well as its perceivable properties at this time. I represent this feature of the world by having a kind of model of the desk in my head. In the event that the desk is moved from one place to another, or is smashed by a sledgehammer, I will want to represent these dynamic features of the world, and I can do this, not by peering in at a model of the desk moving, but by actually doing something with the model that is isomorphic to the desk's moving. Mental processes for manipulating a partial model "replicate worldly processes."

Unfortunately this suggestion is no better equipped to provide a easy solution to the frame problem than the previous suggestion was. If we divide the representation of dynamic and stable features of the world between mechanisms for symbol manipulation and the symbol itself, then the map theory implies that we do not explicitly represent dynamic features of the world, or at least, in the best case, that we represent dynamic features of the world very differently from how we represent stable features. But this seems wrong. When we represent the

world, we do not represent the world as a static entity, rather, we explicitly represent the world as a changing, evolving thing. Static and dynamic properties seem to be packed into the same representational medium. The map approach presently on the table has it that the dynamic properties of the world are mirrored by the activity of mechanisms for manipulating representations. It is hard to see how we can have as the contents of our representations these dynamic properties if this version of the map theory were right.

There is a possible reply to this complaint. We have assumed that the distinction between the internal model and the manipulations of this model is the same as the distinction between symbols and symbol manipulations. But that is wrong. Internal models are not *symbols*, strictly speaking. The point of calling something a symbol is to emphasize the fact that representational properties accrue to that thing. And therein lies the difference between a symbol and an internal model. The map theory we are presently considering will want to deny that representational properties accrue to the internal model by itself. Instead, the representational properties associated with belief accrue to the model *and* cognition's manipulations of that model. Hence, contra the criticism above, our belief contents are not divided between different kinds of mechanisms according to whether a static or dynamic feature of the world is represented. Rather, the complex system composed of the model-cum-manipulations is the fundamental bearer of belief content.

This alternative way of thinking about where to locate content helps the map theorist to avoid the initial criticism. But ultimately it is not a real solution, for it falls prey to some serious difficulties.

Any theory of cognition needs to explain epistemic relations, and it is hard to see how the proposal under consideration is going to do this. Suppose we show Jones a ruby, and he infers that rubies are red. Our theory of cognition needs to explicate this inferential relation in terms of physical states and events – that is, in terms of the machinery of cognition. The sentential model offers us the following sensible suggestion: in order to

account for epistemic relations between beliefs, we should appeal to mechanisms for manipulating internal representations. Thus, according to this sentential model, we need a mechanism that takes the mental sentence “this ruby is red” as input and outputs “all rubies are red.” Thus the sensible route offered by sententialism is that we explain epistemic and inferential relations by appealing to the mechanisms for manipulating inner structures.

The map theory we are presently considering would be hard-pressed to do something similar. According to that theory, belief content *arises out of* the activity of these mechanisms. Thus, it seems we can not appeal to these mechanisms to explain inference, for inference is a move from one content to another. The role of mechanisms to generate content by simulating dynamic properties of the world is at odds with the role of mechanisms to effect transitions between contents. The map theorist cannot have the mechanisms for manipulating inner structures playing both roles. If these mechanisms are used to explain inference, then the appearance of contentful structures – i.e. symbols – should be antecedent to the application of the inferential mechanisms. But this conflicts with using these same mechanisms to explain how contents arise.

One might reply by drawing our attention to the fact that there is nothing incoherent about inferential role semantics. And inferential semantics seems to do what we just said cannot be done: The same thing takes on both the content-giving role and the inference-explaining role – namely, the patterns of transitions from one structure to the next. If inferential role semantics is a coherent proposal (regardless of whether it is true), then so is the map theory presently under consideration.

But there is an important difference between inferential role semantics and the present variation of the map proposal. According to inferential role semantics, the semantic values of structures need not be settled prior to effecting inferential transitions between those structures. Yet even though the semantic values of those structures have not been settled, it *has* been antecedently settled that the structures in question are the

bearers of content. This leads to the following differences between inferential role semantics and the present map proposal. Inferential roles are not themselves content-bearing. They merely give content to inner structures. The map proposal differs in claiming that belief content supervenes on the whole mess – the structure-manipulating activity is itself part of the primary bearer of belief content. Thus, inferential role semantics take inferential moves as primary and tries to show how these moves confer contents on the things moved. The present map proposal, on the other hand, tries to take isomorphism as primary, and then, in order to avoid the problem of how an inner isomorph will get used by cognition, it looks for cognitive *activities* that are isomorphic to the world. But in doing so, there does not appear to be enough resources left for explaining what the mechanisms of inference are.

There is a final move a map theorist might make along these lines. Suppose that there are further mechanisms in addition to the ones that mirror dynamic properties of the world. These additional mechanisms are the ones that fill the cognitive role needed to explain attributions of inferential activity. But this strategy just takes us back to where we started. According to this strategy, there are (i) internal structures, (ii) mechanisms that manipulate these structures in a worldly like way, and now, (iii) additional mechanisms for manipulating all this. Of course now we find ourselves back at the original problem, the problem of how we use a self-contained inner isomorph.

The last problem I will discuss is, I believe, the most serious. It is the problem of how we are going to give an account of *practical* reasoning using the alleged advantages of the map account. Suppose that Jones is a jungle explorer who finds himself face to face with a tiger. Suppose we grant that Jones can run a simulation of *Me-In-Current-Circumstances-Faced-By-A-Tiger*. He runs this simulation using a partial isomorph of the world and a set of cognitive mechanisms for manipulating the isomorph in worldly like ways. Why should we think that the result of the simulation will be Jones fleeing from the scene? Why would not the result of the simulation be a representation of Jones's being eaten by the tiger? The map approach has not

provided any indication of how a mental model is supposed to be employed in practical reasoning.

The frame problem is of course a practical problem. It concerns how we retrieve the right knowledge at the right time. What counts as the *right* knowledge at a certain time is fixed partly by the current context and states of affairs. But it is also fixed by the goals, interests and desires of the agent. If Jones is a daredevil, searching desperately for new thrills, the appearance of a hungry tiger may elicit a very different plan, intention or reflexive action from the plan, intention, or action elicited in you. Yet you and Jones may represent the world in roughly the same way. Thus, if “running a simulation” on one’s inner isomorph is to offer up a plan, intention or action, a great deal more needs to be said about how the simulation gets used in practical reasoning. And once we start supposing that simulations get *used* in further cognitive activity, then the benefits offered by the alleged economy of running a simulation as opposed to the cumbersome activity of retrieving sentences diminish. Running a simulation was supposed to replace the problems created by trying to find a set of principles for organizing and manipulating sentences. Those problems reappear when we notice that a simulation does not get an agent very far in producing practical conclusions.

5. CONCLUSION

I have argued for two conclusions. First, appealing to the frame problem does not help to generate a special argument against the sentential theory. The best we can do is offer a kind of pessimistic induction against sententialism. Second, a map theory of belief content will not make solving the frame problem easier. The map theory faces serious obstacles of its own when addressing this problem. Solving the frame problem is just hard, no matter what medium of representation you embrace.¹⁵

NOTES

¹ At the end of this section, I briefly discuss the terminological confusion surrounding the ‘frame problem’.

² Churchland (1986, p. 394); I have omitted Churchland’s references to other literature.

³ Haugeland (1987, pp. 92–93).

⁴ Sentential theories of belief content have had numerous defenders. See Harman (1973), Fodor (1975–1981), Field (1981), Evans (1982), Davies (1998).

⁵ Fodor (1975).

⁶ Churchland (1986, pp. 392–394).

⁷ For discussions of the role of context in practical problem solving, see Haugeland (1985) and Dreyfus (1992). It is interesting to note that Fodor makes a related point in his (1983), where he expresses doubts as to whether there could be laws of belief fixation given that the activity of the central system is holistic.

⁸ Again, see Fodor (1983).

⁹ Haugeland (1987). I should emphasize that Haugeland is only tentatively suggesting the “quasi-pictorial” alternative to sententialism. He does not in that paper (or anywhere else that I know of) commit himself to the view.

¹⁰ The isomorphism here is quite abstract – in particular, the suggestion is not that the models must *resemble* what is represented. See McGinn (1989) for more on the notion of isomorphism relevant to characterizing an inner model.

¹¹ Haugeland (1987, pp. 90–91). See also Braddon-Mitchell and Jackson (1996) for a similar point about the “map” theory of belief content.

¹² A word about terminology. A theory which posits mental models as the medium of belief content is what I shall call a *map theory*. I am following the usage established by Braddon-Mitchell and Jackson (1996). Examples of map theories are McGinn’s “modeling” approach to belief (1989), Lewis’s account of “map-like” representation (1994), and Jackson’s “radical holism” (1998).

¹³ Lewis (1994, p. 422).

¹⁴ McGinn (1989, p. 176).

¹⁵ Thanks go to John Pollock and Daniel Stoljar for their comments on earlier drafts of this paper.

REFERENCES

- Braddon-Mitchell, D. and Jackson, F. (1996): *Philosophy of Mind and Cognition*, Oxford: Blackwell.
 Churchland, P.S. (1986): *Neurophilosophy*, Cambridge: MIT Press.

- Davies, M. (1998): *Language and Thought*, 'Language, Thought and the Language of Thought (Aunty's Own Argument Revisited)', in P. Carruthers and J. Boucher (eds.), New York: Cambridge UP.
- Dreyfus, H. (1992): *What Computers Still Can't Do*, Cambridge: MIT Press.
- Evans, G. (1982): *The Varieties of Reference*, New York: Oxford UP.
- Field, H. (1981): 'Mental Representation', reprinted in N. Block (ed.), *Readings in Philosophy of Psychology*, Vol.2, Cambridge: Harvard UP.
- Fodor, J. (1975): *The Language of Thought*, New York: Crowell.
- Fodor, J. (1981): 'Propositional Attitudes', reprinted in *Representations*, Cambridge: MIT Press.
- Fodor, J. (1983): *Modularity of Mind*, Cambridge: MIT Press.
- Glymour, C. (1987): 'Android Epistemology and the Frame Problem' in Z. Pylyshyn (ed.), *The Robot's Dilemma*, Norwood, NJ: Ablex.
- Harman, G. (1973): *Thought*, Princeton: Princeton UP.
- Haugeland, J. (1985): *Artificial Intelligence: The Very Idea*, Cambridge: MIT Press.
- Haugeland, J. (1987): 'An Overview of the Frame Problem', in Z. Pylyshyn (ed.), *The Robot's Dilemma*, Norwood, NJ: Ablex.
- Hayes, P. (1987): 'What the Frame Problem Is and Isn't', in Z. Pylyshyn (ed.), *The Robot's Dilemma*, Norwood, NJ: Ablex.
- Jackson, F. (1998): 'Mental Causation Without the Language of Thought', in *Mind, Method and Conditionals: Selected Essays*, New York: Routledge.
- Lewis, D. (1986): *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, D. (1994): 'Lewis, David: Reduction of Mind', in S. Guttenplan (ed.), *A Companion to the Philosophy of Mind*, Oxford: Blackwell.
- Lormand, E. (1990): 'Framing the Frame Problem', *Synthese* 82, 353–374.
- Lormand, E. (1996): 'The Holorobophobe's Dilemma' in Z. Pylyshyn and K. Ford (eds.), *The Robot's Dilemma Revisited*, Norwood, NJ: Ablex.
- McCarthy, J. and Hayes, P. (1969): 'Some Philosophical Problems from the Standpoint of Artificial Intelligence', in B. Meltzer and D. Michie (eds.), *Machine Intelligence 4*, Edinburgh: Edinburgh UP.
- McGinn, C. (1989): *Mental Content* Oxford: Blackwell.
- Pollock, J. and Cruz, J. (1999): *Contemporary Theories of Knowledge*, 2nd edn., Lanham, MD: Rowman & Littlefield.

Department of Philosophy

Clark University

950 Main Street

Worcester, MA 01610

USA

E-mail: shendricks@clarku.edu