

Editorial Focus. Machines and Mentality, Knowledge and Its Representation, Epistemic Aspects of Computer Programming, Connectionist Conceptions, Artificial Intelligence and Epistemology, Computer Methodology, Computational Approaches to Philosophical Issues, Philosophy of Computer Science, Simulation and Modeling, and Ethical Aspects of Artificial Intelligence.

Aims and Scope. *Minds and Machines* affords an international forum for discussion and debate of important and controversial issues concerning significant developments within its areas of editorial focus. Well-reasoned contributions from diverse theoretical perspectives are welcome and every effort will be made to ensure their prompt publication. Among the features that are intended to make this journal distinctive within the field are these:

- Strong stands on controversial issues are especially encouraged;
- Important articles exceeding normal journal length may appear;
- Special issues devoted to specific topics will be a regular feature;
- Review essays discussing current problem situations will appear;
- Critical responses to previously published pieces are also invited.

This journal is intended to foster a tradition of criticism within the AI and philosophical communities on problems and issues of common concern. Its scope explicitly encompasses philosophical aspects of computer science. All submissions will be subject to review.

Photocopying. *In the U.S.A.:* This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by Kluwer Academic Publishers for users registered with the Copyright Clearance Center (CCC) Transactional Reporting Service, provided that the fee of USD 8.00 per copy is paid directly to CCC. For those organisations that have been granted a photocopy licence by CCC, a separate system of payment has been arranged. The fee code for users of the Transactional Reporting Service is 0924-6495/95 USD 8.00.

Authorisation does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale.

In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to Kluwer Academic Publishers, P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

Minds and Machines is published quarterly.

Subscription price, per volume: Institutions USD 288.00.

Second-class postage paid at Rahway, N.J. USPS No. 007824.

U.S. Mailing Agent: Mercury Airfreight International Ltd., 2323 Randolph Ave., Avenel, NJ 07001.

Published by Kluwer Academic Publishers, Spuiboulevard 50, P.O. Box 17, 3300 AA Dordrecht, The Netherlands, and 101 Philip Drive, Norwell, MA 02061 U.S.A.

Postmaster: Please send all address corrections to: *Minds and Machines*, c/o Mercury Airfreight International Ltd., 2323 Randolph Ave., Avenel, NJ 07001, U.S.A.

Printed on acid-free paper

Why Cognitive Science Is Not Formalized Folk Psychology

MARTIN PICKERING

University of Glasgow, Human Communication Research Centre, Department of Psychology, 56 Hillhead Street, Glasgow G12 9YR, Scotland (email: martin@psy.gla.ac.uk)

and

NICK CHATER¹

University of Edinburgh, Department of Psychology, 7 George Square, Edinburgh EH8 9JZ, Scotland (email: nicholas@cogsci.ed.ac.uk)

Abstract. It is often assumed that cognitive science is built upon folk psychology, and that challenges to folk psychology are therefore challenges to cognitive science itself. We argue that, in practice, cognitive science and folk psychology treat entirely non-overlapping domains: cognitive science considers aspects of mental life which do not depend on general knowledge, whereas folk psychology considers aspects of mental life which do depend on general knowledge. We back up our argument on theoretical grounds, and also illustrate the separation between cognitive scientific and folk psychological phenomena in a number of cognitive domains. We consider the methodological and theoretical significance of our arguments for cognitive science research.

Key words. Folk psychology, modularity, defeasible reasoning, knowledge representation, propositional attitudes, language, cognition, perception, functional architecture.

Introduction

There is an influential account of the relationship between cognitive science and folk psychology, which we shall call the "received view". The story is that cognitive science, via the computer metaphor, shows how folk psychological explanation can serve as the basis of a rigorous and scientifically respectable account of mind: the meaning of mental representations is given by the contents of propositional attitudes; these representations are manipulated by computational processes which respect that meaning (e.g. Fodor 1975, 1980, 1987; Fodor & Pylyshyn 1988; Pylyshyn 1984). So, if it is true that John believes that fish have toes, then John must have a mental representation which has the content that fish have toes, and the computational processes which act upon that representation must respect that meaning. These ideas give a clear sense in which folk psychology could be vindicated by scientific psychology – talk about beliefs and desires can no longer be necessarily condemned on the grounds of being hopelessly unscientific, since they may play a part in a scientific, computational explanation of mind.

This story is attractive for defenders of folk psychology, since it purports to show that folk psychology can be put on a scientific footing. It is also potentially

Minds and Machines 5: 309–337, 1995.

© 1995 Kluwer Academic Publishers. Printed in the Netherlands.

attractive for cognitive science. The possibility that folk psychology can be incorporated into cognitive scientific explanation gives the field a vast head-start, by recruiting everyday intuitions concerning the contents of mental states, the transitions between such states and how they result in behaviour. Indeed, cognitive science is seen as folk psychology made precise, with appropriate additions and extensions. In other words, cognitive science is taken to be formalized folk psychology.

Opposed to this harmonious picture are those who argue that cognitive science should be seen as replacing or overthrowing folk psychology, rather than building upon it. Stich (1983) cites data from neuropsychology and attribution theory to argue that much of cognitive science undermines folk psychological explanation. Similarly, P. S. Churchland (1986) argues that developments in the neurosciences overturn folk psychological wisdom (see also P. M. Churchland 1989). Furthermore, Stich and Churchland also present general arguments for their position: Stich argues that cognitive scientific explanations should supervene on brain states, but that folk psychological explanation is "broad", i.e. it depends on the state of the external world; and Churchland argues that folk psychology is no more likely to provide a good foundation for science than did folk astronomy or folk medicine.

Both sides of this debate therefore consider folk psychology to be importantly related to cognitive science; they differ on whether the relationship is one of cooperation or competition. In the light of this debate, it is not surprising that a great deal of attention has been given to the nature of the folk psychological explanation and whether it is true, in the philosophical and foundational literature in cognitive science (e.g. Devitt 1990; Horgan & Woodward 1985; Fodor 1986).

But this debate seems to be strangely removed from the practice of cognitive science. If the received view is correct, one would expect theorists to make frequent, productive use of propositional attitude explanation; if the sceptics are correct, one would expect cognitive science to be gradually removing such explanation from the discipline, as the phenomena under consideration become better understood. Indeed, according to both views, one would expect cognitive scientists to be daily confronted with their reliance on folk psychology, for good or ill, in their theoretical, experimental and computational work. In fact, we suggest, talk of beliefs and desire is noticeable by its absence throughout the discipline. On the surface, at least, cognitive science seems to be independent from, rather than deeply enmeshed in, folk psychology.

In this paper, we suggest that appearances are not deceptive, and argue that the separateness of cognitive science and folk psychology stems from the fact that the aspects of mental life with which they are concerned do not overlap: folk psychology applies to "knowledge-rich" aspects of cognition, which have proved completely unamenable to cognitive science, whereas cognitive science has been limited to "knowledge-free" aspects of cognition, about which folk psychology is silent. So, for example, cognitive science is successful at understanding relatively low level perception and language processing, motor control, the structure of

memory stores, the limits of attention, the structure of concepts, and so on. Theories in these areas make no reference to beliefs or desires, because they fall into one of two categories. Either they are concerned with special purpose, modular, psychological processes, which are isolated from world knowledge, or they are concerned with structural aspects of the mind, rather than the content of what is represented. By contrast, the domain of folk psychology is the contents of everyday, common-sense thought, which is the paradigmatically knowledge-rich. We shall argue that cognitive science has been conspicuously unsuccessful in unravelling this aspect of mental life, and illustrate the reasons for this by considering the problems encountered in attempts within artificial intelligence (AI) to model everyday thinking computationally.

According to our view, folk psychology and cognitive science should be seen neither as cooperative nor as standing in competition. They are, *de facto*, independent enterprises, since cognitive science simply cannot handle the knowledge-rich phenomena which are the subject matter of folk psychology. This implies that the discussions of folk psychology in the foundations of cognitive science are irrelevant to the current practice in cognitive science, and also has wider significance for the explanation of human behaviour.

Our argument is organized as follows. First, we briefly review the received view, explaining how cognitive science is supposed to be formalized folk psychology. We then consider a field in which the formalization of human knowledge, and its computational implementation, has been attempted: traditional AI. We consider in detail the problems that led to the failure of the AI experiment, and argue that these problems are extremely serious and difficult to solve. The intractability of formalizing folk psychology explains why there has been a singular lack of success in attempting to provide a cognitive science of knowledge-rich mental processes; and why cognitive science has succeeded only by focusing on knowledge-free aspects of cognition, with which folk psychology is not concerned. We then consider how the boundary between knowledge-free and knowledge-rich processes, i.e. between phenomena which can be addressed by cognitive science and those which can only be handled by folk psychology, is reflected in the practice of cognitive science. Firstly, we consider how the maxim that cognitive science should be restricted to knowledge-free processes is embodied in two methodological principles: that cognitive science focus on "functional architecture" and "modularity". Secondly, we show how the knowledge-free/knowledge-rich distinction is important in specific areas of cognitive science, taking the examples of vision, memory and language processing. Finally, we consider the implications of this distinction for cognitive science and folk psychology, and for the explanation of human behavior in general.

Folk Psychology and Cognitive Science

In this section, we give a brief sketch of how we understand cognitive science and folk psychology. Both terms have been used in a variety of rather different ways,

and it is not of course our purpose to argue about terminology. Rather we are concerned to show that there are two projects in the explanation of human behavior which are, in practice, non-overlapping: the first involves computational explanation, the second involves explanation in terms of beliefs, desires and other propositional attitudes. With this in mind, let us discuss cognitive science and folk psychology in turn.

Cognitive science is defined by its style of explanation rather than its subject matter. A cognitive scientific explanation explains mental life in *computational* terms. The range of topics which admit of computational explanation cannot, of course, be determined *a priori*, but it is standard to assume that cognitive science takes in traditional cognitive psychological topics such as memory, reasoning, problem solving, language comprehension and production, perceptual processing, motor control and skill learning, as well as cognitive developmental psychology and cognitive neuropsychology. The methods and tools of cognitive science are various, and are drawn from experimental psychology, computer science and artificial intelligence, together with relevant aspects of linguistics, logic and philosophy, and other disciplines.

By folk psychology we mean to refer to the body of knowledge which underlies everyday explanations of human behavior in terms of beliefs, desires and other propositional attitudes. Such explanation is both extremely widespread and of enormous importance. Propositional attitudes allow us to ascribe meaningful mental states to each other, to explain thought in terms of internal transitions between meaningful states, and to explain how mental life is embedded in the world, by connecting the contents of these internal states with perception and action.

By saying that a person has a particular propositional attitude, we are ascribing to that person a mental state which has the same content as the "proposition" component of a sentence of natural language expressing that propositional attitude. Thus, propositional attitude explanation recruits all the resources of natural language and allows them to be used to characterize mental states and their transitions. The central idea is that the internal mental state transitions correspond to natural language *arguments*, with the relationship between mental states and natural language being given by the ascription of propositional attitudes. Rational thought corresponds to good argument, and irrational thought to flawed or incoherent argument.

Let us illustrate this point with an example. Suppose that John returns from holiday to find that the front door of his house has been forced open. He goes to the nearest public phone and calls the police. Folk psychology enjoins us to explain John's behavior much as he would explain it himself. He sees the door, and comes to *believe* that it has been forced open. He then rapidly runs through an argument about keys, forced locks, illegal entry, rising crime rates and so on. The conclusion of this argument is that a burglary has occurred, and may perhaps still be in progress. On the basis of this conclusion, John decides (this decision can

be cast as an argument concerning the best course of action) that he should not enter (since he does not want to encounter the burglars) but to call the police. We judge John to be rational if his arguments seem to be compelling; and irrational if they make no sense.

Notice that we are not identifying folk psychology with the common-sense understanding of the nature of mental life in general (although this is a common usage, outside of the philosophical debate, e.g. Bruner 1990). There is consensus that our various and inchoate intuitions about memory, language, attention or visual perception are no more to be taken seriously as a basis for scientific inquiry than common-sense intuitions about the workings of the heart, the legs or the acoustic properties of the ear. What is controversial is whether explanation in terms of propositional attitudes can be sustained.

How Cognitive Science Might Be Formalized Folk Psychology

The received view is that a computational account sticks closely to folk psychological explanation. For example, a computational account of John's thoughts on returning home should be built upon a folk psychological story like the one we gave above: the internal argument that folk psychology ascribes to John is viewed as a description of John's internal thought processes. Specifically, internal formulae of a language of thought correspond to the contents of the propositional attitudes; and internal computational operations carry out the steps in the argument. Of course, the received view is not committed to the assumption that our explanation of John's behaviour, or, for that matter, John's own explanation, can be taken entirely for granted and built into a cognitive scientific explanation. The details of the story may be inaccurate – it is the style of explanation, in terms of mental states characterised by propositional attitudes, and in terms of internal argument, that is crucial.

The feasibility of the received view as a basis for cognitive scientific, computational explanation of the mind therefore depends on whether or not it is possible to formalize and computationally implement the types of argument that folk psychology invokes. We shall argue in the next section that these arguments cannot in general be formalized, let alone implemented computationally. But for now let us consider the possibilities for formalization of everyday argument in as positive a light as possible. There are two areas in which mathematically and computationally well-understood normative theories have been developed.

The first area concerns argument about what action should be taken, given certain beliefs and desires, and can be modeled by mathematical decision theory (North 1968). The idea is to model beliefs about what could happen if a given action is chosen, by associating probabilities with each possible outcome, or sequence of outcomes; desires are modeled by real numbers which represent the utility of these possible states. According to decision theory, the rational action to choose is that which has the greatest expected utility, which this can be calculated

from the beliefs and desires using probability theory. Decision theory is attractive from a computational point of view, since it lends itself to a straightforward computational treatment (Schacter 1986) (although in practice, relevant information is not always available, and the range of options may be too large to explore thoroughly). While there has been a considerable amount of work on planning in AI (Allen *et al.* 1990) and decision making in psychology (Kahneman *et al.* 1982), debate in the foundations of cognitive science has concentrated almost entirely not on arguments about how beliefs and utilities should determine action, but about the prior question of what should be believed.

This second kind of argument is the subject of theories of inference, of which by far the most developed part is formal logic. Formalizing a natural language argument, such as is found in folk psychological explanation, in logical terms, requires first translating the argument into a logical language. We can then consider the steps in the argument and consider whether or not they are licensed by the rules of the logic being used. If they are, then the argument is valid – if the premises are true, then the conclusion must be true. The argument has been shown to be rational by being given a logical analysis. The use of logic also appears attractive from a computational point of view, since logical operations can be mechanized in well understood ways, by a theorem prover, which manipulates logical formulae according to the proof theory of the logic.

Logical theories of inference provide a potential explanation of the internal operations of the cognitive system. The idea is that thought is conducted in an internal logical language, a “language of thought” (Fodor 1975). The natural language sentences ascribed to people in folk psychological explanation are viewed as translations of these internal formulae, and internal logical operations determine the passage from premises to conclusions. It is this picture of how cognitive science can formalize folk psychology which inspires Fodor and Pylyshyn (1988) to claim that cognition is mechanized proof theory.²

To summarize: the received view of the relationship between folk psychology and cognitive science is that folk psychology explains behavior in terms of arguments, which are assumed to be causally involved in the production of behavior; cognitive science shows how this is possible, by formalizing these arguments, and realizing them computationally. We have mentioned two kinds of argument, concerning decision making and logical inference, where it is at least clear how this project of formalization could begin. Since, as noted above, there has been an emphasis on the application of logical methods rather than decision making, within the foundations of cognitive science, we focus entirely on this kind of argument in this paper.

The rest of this paper aims to establish two claims. We shall argue that the programme of formalizing folk psychology is unworkable, for reasons relating to the problems of dealing with common-sense knowledge in artificial intelligence. However, we do not take this to affect current cognitive science, because we then argue that current cognitive science is entirely separate from, rather than based upon, folk psychology.

The AI Experiment

The received view claims that cognitive science can formalize and computationally implement folk psychological explanations of human behavior, explanations which specify internal arguments that putatively underlie behavior. From the point of view of developing theories in cognitive science, this suggests a specific research strategy: to understand everyday argument, as embodied in folk psychological explanations, by specifying (i) the everyday knowledge involved, (in verbal form) and (ii) the mechanisms of inference used (e.g. inference in a particular logic). We call these the requirements of “right information” and “right reason”.

The attempt to build computer systems which can cope with knowledge-rich domains within artificial intelligence has followed just this strategy. The problem of right information has been tackled by using verbal reports of domain experts (or the programmer), and the problem of right reason has been addressed by relying on logical inference. If artificial intelligence is unable to formalize folk psychology at all, then the question of whether any particular formalization might be cognitively plausible (i.e. whether there could be a cognitive science of knowledge-rich processes) simply does not arise. The problems that we outline below suggest that this project faces deep and unsolved problems, and hence that a cognitive science based on formalized folk psychology is at best a long way off, and at worst impossible.

Let us consider the problems of right information and right reason in turn.

RIGHT INFORMATION

In AI the problem of finding out the information that should be incorporated into a system is frequently extremely difficult. In highly formalized mathematical or game-like domains, representational questions are often partially solved from the outset, by existing mathematical characterizations of the nature of the domain. So, for example, representing the information sufficient to prove a particular mathematical theorem, or to decide on the best move in chess, may be possible. (This is not to say, of course, that the exact form of the representation, and finding appropriate algorithms to exploit it, may not be extraordinarily difficult.)

In everyday domains, however, the relevant knowledge must be ascertained and formalized from scratch. The approach that AI research into, for example, knowledge representation (Brachman & Levesque 1985), qualitative reasoning (Weld & De Kleer 1990) and common-sense reasoning (Davis 1990), has adopted is precisely that which the received view recommends. The information built into the computational systems is specified by our relevant propositional attitudes; this information is known as “knowledge”, with no connotations of truth or justification. The relevant knowledge, originally in verbal form, is then translated into a logical language, such as the predicate calculus or some variant (e.g. the situation calculus (McCarthy & Hayes 1969); semantic networks (Rumelhart *et al.* 1972);

frame representations (Minsky 1975); or conceptual dependency representations (Schank 1975)); despite the translation of logical terms, all nonlogical terms are, of course, directly borrowed from natural language. This close correspondence with natural language means that the contents of typical AI knowledge databases can be reasonably easily read off by inspection. For the sake of concreteness, we can assume that the relevant information in a domain is elicited, and cast in terms of a list of statements of predicate calculus.

Unfortunately, so far it has proven to be impossible to carry out this program successfully. Perhaps the most important reason is that everyday, common-sense knowledge turns out not to be decomposable into separable domains that can be formalized relatively independently. Indeed, common-sense knowledge appears to have an ineliminably open-ended character – in making or understanding any particular common-sense argument, there seems to be an indefinite, and arbitrarily far-ranging, body of knowledge involved.

To illustrate this point, let us reconsider John's inference that he has been burgled. Knowledge involved in making this inference includes facts such as: that burglars must enter a house before burgling it; that a door is a potential entrance; that locked doors cannot be opened; that forcing the lock will allow the door to be opened; that locks do not spontaneously break; that people do not generally force locks without criminal intentions; and so on. But it also includes information concerning alternative explanations: for example, that John's wife, who might conceivably have locked herself out and had to break in, is away on business in Greece; that Greece is too far away for a day trip back home; that he spoke to her yesterday, by telephone, and she had no plans to return. There seems to be a never ending stream of relevant information, and this information can come from entirely unexpected quarters (e.g. flight times from Greece). The task of building an AI program to deal with common sense inferences makes it painfully clear that if these pieces of information are ignored, inference is likely to go hopelessly awry. Fodor (1983) makes this point, arguing that non-deductive inference is always what he calls isotropic: relevant information can be drawn from anywhere and everywhere, and cannot be separated into evidentially disconnected domains.

A final, and devastating, difficulty is that each argumentative link that we postulate is itself as rich as complex as John's original inference. So, for example, the inference from John's speaking by phone to his wife the day before, to the fact that she cannot be back already, will involve endless information about airlines, alternative means of travel and their speed, and so on. This phenomenon has been dubbed the "fractal character of common-sense" (Chater & Oaksford 1990).

We have argued that the everyday arguments involved in folk psychological explanations, when examined closely, turn out to be open-ended and draw upon an indefinite amount of knowledge; and that this is one of the reasons that it has not proved to be possible to model such inferences computationally. But this appears to raise a problem: if folk psychological explanation is indefinitely

complex, how is it, then, that succinct folk psychological explanations can be used in everyday life?

In fact, however, there is no puzzle: in everyday folk psychological explanation only the bare bones of the argument need to be specified, since the rest, the indefinitely large body of information which resists formalization, is common to the agent whose mental life is being explained and the audience to whom the explanation is addressed. In building a computational model, no such understanding can be taken for granted: the task in AI is to build a system which can make and understand such arguments completely automatically. And this is a task which, in view of the open-ended character of common-sense knowledge, has not been solved by formalization in terms of logical rules. We shall see later that explanation which relies on shared understanding of human behavior may, nonetheless, be vitally important. Such interpretative explanation is central to many areas of the social sciences and humanities.

RIGHT REASON

The most immediate problem for formalizing common-sense arguments is that such arguments are usually not logically valid: even if the premises are true, the conclusion is at best plausible, but not certain. This creates immediate difficulties given that AI, in line with the standard version of the "received view", is based on logic. But the problem is deeper – most non-logical arguments (arguments for the plausibility, rather than the certain truth of the conclusion, given the premises) have eluded formal characterization. In order to formalize everyday inference, as ascribed to people by folk psychology, AI has been forced to set out on its own, and attempt to develop theories of this kind of uncertain reasoning (Ginsberg 1987; Shafer & Pearl 1990).

A crucial feature of common-sense inferences is that they can be overturned by the addition of later information. John's inference that he has been burgled above will be overturned if he finds that his daughter has come back from college, having forgotten her key. Inferences which can be overturned by the *addition* of premises are called *non-monotonic* inferences. Even the most mundane everyday non-monotonic inferences are extremely difficult to model formally or computationally, despite the intensive study that has been devoted to them in the AI.

There are many formal approaches to non-monotonic reasoning in AI and we shall illustrate the problems encountered by considering the most popular and the most direct: the attempt to develop a non-monotonic logic for everyday reasoning. The most well known example of this approach is Reiter's (1980, 1985) default logic.³

Reiter distinguishes between two kinds of information: certain, hard facts about the world, and default rules which support plausible inferences. The set or sets of beliefs sanctioned by a default theory are obtained by starting with the hard facts,

and repeatedly applying the default rules to derive plausible but not certain conclusions. The basic intuition underlying Reiter's approach is that plausible, default inferences should be allowed when, but only when, their conclusion is consistent with what is already known. So, for example, a default rule from the premise that Tweety has wings to the conclusion that Tweety flies should be assumed to be correct unless the conclusion is not consistent with other known information. This rule would not apply given the additional premise that Tweety is an ostrich (assuming it is known that ostriches can't fly).

The fundamental difficulty for formal approaches such as Reiter's is dealing with cases in which evidence points to conflicting conclusions. Suppose that there is a second default rule, that badly injured creatures cannot fly, and consider what can be concluded from two hard facts: that Tweety has wings and is badly injured. In practice, it is clear that the latter rule should take priority – badly injured birds presumably are unlikely to be able to fly. This follows because of what we know about the *meaning* of the terms involved, what wings, flight and injury refer to, and our general world knowledge about how these are related. But, of course, neither meaning nor general world knowledge can be used by accounts of non-monotonic inference, such as Reiter's, which depend only on the *form* of the predicates mentioned, and hence there is no way of resolving conflicting evidence sensibly.

It is therefore not surprising that Reiter's system cannot handle cases of conflicting defaults successfully. If the first default rule is considered first, the conclusion that Tweety flies can be drawn, and the second default rule is blocked, since its conclusion is inconsistent with what has already been derived. But if the second rule is applied first, the opposite conclusion that Tweety does not fly can be drawn, and then the first rule is blocked. The only way to decide which of these extensions is appropriate is to order the two rules, to ensure that the correct inference is drawn. But, in general, ensuring that rules are correctly ordered is just equivalent to the original problem of deciding which inference should be drawn. Since there is no way of telling which of these extensions is to be preferred, all that can be concluded overall is that the disjunction of these two conclusions holds: that Tweety either flies or does not. Thus, rather than adjudicating between conflicting pieces of evidence, the system can merely draw conclusions which are compatible with both. This "problem of weak conclusions" (McDermott 1987) is endemic to formal approaches to non-monotonic inference: information about form is simply not sufficient to specify how conflict should be resolved (see Oaksford & Chater 1991 for a detailed consideration of various attempts within AI and cognitive science to get around this problem and the difficulties that they run into; see also Garnham 1993; Chater & Oaksford 1993).

The problems of right reason stem from the fact that common-sense reasoning aims to find an explanation that best fits with all the information that is known. Fodor (1983) calls this "inference to the best explanation" and notes that it is what he calls *Quinean* – that is, beliefs cannot be confirmed one by one, since

whatever data are encountered, it is always possible to maintain any particular belief by altering other beliefs appropriately. That is, entire systems of belief, an entire scientific or common sense theory must be collectively be tested against empirical data (Quine 1953). Consider, for example, the situation faced by a jury. Generally, it is not possible to assess each aspect of the defence or the prosecution's case in isolation – it is crucial to consider how well each story fits together as an overall explanation of the evidence. That is, the jury must judge between entire cases presented between the defence and the prosecution, considered as a whole.

The Quinean character of inference to the best explanation gives rise to non-monotonicity – there is always the possibility that additional information will dramatically change the plausibility of a case, story, or theory, and thus upset previous inferences. And isotropy ensures that the relevant information could be of any sort whatever, and might draw on arbitrary background knowledge. It is these factors which make the formalization of common-sense inferences so difficult.

We have so far viewed these difficulties as concerning the requirement of right reason. But equally, they bring out the problem that the relevant information has not been encoded. In the above example of conflicting defaults, the system has not been given enough information to decide which is the most sensible way to resolve conflicting defaults; and encoding such information in a logical form appears to be a literally endless task, since each new piece of knowledge will be as defeasible as the rest, and will require still further knowledge to specify how its defaults should be resolved, and so on. The open-ended character of common-sense, when squeezed into a logical representation (and equally a verbal representation), has suggested to some that such representations are not appropriate for representing common-sense knowledge (Dreyfus & Dreyfus 1986). There is however a resounding silence concerning alternative forms of knowledge representation which might escape these difficulties.

IMPLICATIONS

We have argued that the attempt to devise computational models based on formalized folk psychology has failed. Folk psychological explanations are "knowledge-rich" – they make reference to common-sense argument and knowledge; cognitive-scientific explanations are "knowledge-free" – they eschew common-sense. Thus cognitive science is not (currently at least) formalized folk psychology, as the received view suggests; and cognitive science and folk psychology do not stand in opposition, as the critics of the received view have argued. Rather, (current) cognitive science and folk psychology have entirely separate concerns.

Depending on whether or not the received view is accepted, the significance of

the difficulties that we have considered above will be assessed rather differently. For advocates of the received view, the moral is that formalizing folk psychology is extremely difficult and not currently tractable. A natural consequence of this point of view is that cognitive science should, for the moment at least, turn to areas which skirt the mire of knowledge-rich explanation. We shall see below that both Fodor (1983) and Pylyshyn (1984) advocate methodological principles which keep cognitive science firmly away from knowledge-rich processes, so this seems a reasonable interpretation of their position.

For sceptics regarding the received view, who believe cognitive science should replace folk psychology, the difficulties that we have described are wont to seem more than extraordinarily difficult technical puzzles, but as symptoms of the underlying wrong-headedness of the project of trying to ground cognitive science in folk psychology. Alongside this negative thesis, it would be attractive to be able to offer a positive alternative to the received view, which could successfully model knowledge-rich processes. Needless to say, however, there is no such practical alternative on offer. Despite optimistic discussion in philosophical circles about possible alternative accounts, including neural networks (Ramsey *et al.* 1990; Bechtel & Abramson 1991) and holograms (Haugeland 1978), the practical situation is gloomy. By giving up the received view, theorists no longer have the right to assume what content internal states have, and what transitions can occur between states. In short, the mind, in so far as we think of the mind as constituted by knowledge-rich processes, must be viewed as a black box, only the inputs and outputs of which may be known; the only instrument for looking inside the box, folk psychological intuition, must be rejected as untrustworthy. This situation makes the problem of knowledge-rich processes appallingly difficult.

In practice, models using alternative computational approaches have also focused largely on knowledge-free processes. Knowledge-free models have been provided of reading (e.g. Seidenberg & McClelland 1989), learning the past tense (Rumelhart & McClelland 1986), aspects of visual processing (Linsker 1988; Zennel 1989), speech perception (McClelland & Elman 1986; Abu-Bakar & Chater 1993) and motor control (Sanger 1993; Uno *et al.* 1993), and reinforcement learning mechanisms (Sutton & Barto 1991). In these domains, knowledge has not been provided by the programmer, but has been derived by learning mechanisms of the system itself. In knowledge-rich domains, it has not been possible to apply learning mechanisms, since knowledge-rich processes are so ill-specified and complex.

There have been some attempts to model knowledge-rich processes in connectionist terms. But these have generally amounted to non-standard implementations of symbolic architectures, and have been just as committed to the received view as standard artificial intelligence (Chater & Oaksford 1990). Thus, neural networks have been used to implement semantic networks (Hinton 1981; Shastri 1985; Smolensky 1987), production systems (Touretzky & Hinton 1985), schemata (Rumelhart *et al.* 1986) and specialist knowledge representation

formalisms such as μ -clone (Derthick 1987). All of these approaches simply inherit all the difficulties that we raised above for traditional approaches to knowledge representation and common-sense reasoning in artificial intelligence. In particular, then, the skeptic of the received view cannot use our arguments that the received view is not correct to show that neural networks, or some other non-standard alternative computational model, must be correct. Whatever computational formalism is used, only knowledge-free processes can be modeled successfully. The modeling of knowledge-rich processes must start from the assumption that the received view is correct, and immediately runs into the range of problems that we have outlined. So, in conclusion, even if the received view is not correct, the skeptic can offer no workable alternative route to a cognitive science of knowledge-rich processes; cognitive science cannot, therefore, encroach upon the territory of folk psychology.

If our arguments are correct, the debate over the received view, which has loomed so large in the philosophical foundations of cognitive science is very remote from actual practice. Both advocates and sceptics regarding the received view must concede that the territory of folk psychology, knowledge-rich mental processes, is beyond the scope of current cognitive science, and hence folk psychology and cognitive science will, in practice, not overlap. As we shall see, cognitive scientists have in practice taken this view, and concentrated on studying knowledge-free aspects of cognition.

Steering Cognitive Science Away from Knowledge-Rich Processes

How should cognitive science proceed in the face of the problems with accounting for knowledge-rich explanation? In this section, we consider two methodological principles which warn cognitive science away from just those areas in which folk psychology is operative.

The first is Pylyshyn's (1984) call for concentration on aspects of the cognitive system which can be understood in sufficiently abstract terms to be entirely neutral about the content of mental representations; cognitive science should focus on what he calls functional architecture. The second is Fodor's (1983) call for concentration on modular cognitive processes, i.e. processes which do not involve common-sense and hence do not slip into isotropy (and where, of course, folk psychology does not apply). By following such recommendations, cognitive science can studiously avoid contact with folk psychology, rather than be founded upon it. It is ironic that these principles, which keep cognitive science and folk psychology apart, are advocated by the main advocates of the received view that cognitive science is formalized folk psychology: in different contexts, they argue that cognitive science is based on folk psychology, and call for cognitive science to eschew just those areas of cognition with which folk psychology is concerned.

Pylyshyn proposes a criterion of "cognitive penetrability": a cognitively penetrable process is one that can be influenced by a person's beliefs and desires,

that is, those aspects of cognition to which folk psychological explanation is relevant. In our terms, the criterion of cognitive penetrability precisely distinguishes knowledge-rich from knowledge-free cognitive processes. He notes that since this *functional architecture* is the basis for cognitive computation, it cannot itself be dependent on the results of that computation. In particular, the operations of the functional architecture will be independent of what knowledge is being deployed, what inferences are being made and so on. The functional architecture will be explicable without reference to the particular mental processes it supports, and hence without reference to particular beliefs and desires. Pylyshyn commends uncovering the functional architecture as a tractable research goal for cognitive science.

Functional architecture is not the only aspect of the cognitive system which can be cognitively impenetrable, however. Specialized modular processes, which do not have access to common-sense information, will also be unaffected by beliefs and desires, and hence are a possible subject for cognitive-scientific inquiry.

It is such modular systems which Fodor (1983) argues should be the focus of cognitive science (and Fodor also recognizes that it may be possible to study some of the mechanisms underlying non-modular processes – what he calls horizontal faculties appear to match up roughly with Pylyshyn's functional architecture). Fodor assumes that the processes underlying perception and language comprehension on the one hand, and motor control and language production on the other, will be modular. According to this view, whether or not some particular cognitive process is comprehensible to cognitive-scientific explanation depends on whether or not it is a module. This has important methodological ramifications. For example, to a first approximation at least, in perception, the division between processes which receive top-down input, and those which are driven entirely by the perceptual stimulus distinguishes those processes which can be tractably explained and those that cannot. Fodor notes: "...the limits of modularity [i.e. cognitive processes which do *not* make reference to general world knowledge] are likely to be the limits of what we are going to be able to understand about the mind, given anything like the theoretical apparatus currently available." (Fodor 1983: 126)

Putting Fodor's and Pylyshyn's discussions together, the two kinds of processes that should be studied by cognitive science are both knowledge-free: one depends on a body of information isolated from world knowledge (such as Fodor's modules); the other is concerned with computationally primitive processes. So, while Fodor and Pylyshyn explicitly suggest that cognitive science, or part of it, should be formalized folk psychology, they recommend that cognitive science in practice shy away from any area in which folk psychology is relevant.

We wholeheartedly agree that cognitive science will only succeed by shying away from the territory of folk psychology; but if this is correct, current cognitive science cannot be formalized folk psychology, as Fodor and Pylyshyn's statement of the received view suggests. Cognitive science should not, and does not,

formalize folk psychology, and hence does not, after all, presuppose its truth. But this position also undercuts the position of the sceptics regarding the received view, such as Stich and the Churchlands, who argue that cognitive science displaces folk psychology. In line with the arguments we have made, and in accordance with Fodor and Pylyshyn's methodological canons, practical cognitive science and folk psychology are concerned with separate issues.

THE BOUNDARY BETWEEN COGNITIVE SCIENCE AND FOLK PSYCHOLOGY: SPECIFIC CASES

It is now clearly a key practical question to determine where the boundary between cognitive science and folk psychology lies. How can we find out which aspects of mental life are knowledge-rich and which are knowledge-free? This is a far from straightforward question, and it can in fact be thought of as a central aim of cognitive science itself. In this section, we shall briefly consider three areas of current cognitive-scientific research, namely vision, memory and language comprehension, in order to show how this division arises and why it is important. We also illustrate how folk psychological explanation is independent of cognitive science in each of these areas.

Vision. Marr (1982) stresses that computational models should focus on the bottom-up aspects of visual perception, aspects of visual perception which are determined by the perceptual stimulus alone, rather than those that depend on beliefs and expectations about the world. The programme of research which he began has proved to be one of the most productive in cognitive science, integrating computational models, psychophysical experiments and neurophysiological data (see, e.g., Richards 1988).

Given this focus on bottom-up processing, Marr viewed the separation of distinct aspects of perception that did not involve general knowledge as an important research goal. Only once a process has been isolated can a computational account of its operation be formulated. One way to factor out the influence of world knowledge is to conduct psychophysical demonstrations with stimuli which have no real world interpretation. For example, Julesz's (e.g. 1960, 1971) demonstration that depth perception can be induced by stereopsis over random dot patterns provides, for Marr, an excellent demonstration that there must be purely bottom-up mechanisms for stereo depth perception (interestingly, Marr is explicit that he does not rule out the possibility of some top-down input from world knowledge, but he considers its study to be intractable). Similarly, Braddick's (1979) demonstration of apparent motion using successive frames of random dot patterns suggests that there are bottom-up mechanisms for motion perception, and Julesz's (1975) studies of immediate preattentive texture discrimination, using meaningless texture elements, shows that there are segmentation

processes in vision which do not depend on breaking up the scene into cognitively meaningful chunks. Quite generally, this line of psychophysical inquiry uses meaningless stimuli, so that knowledge-free aspects of perception can be distinguished and studied.

It is interesting to note that prior to Marr, the dominant information processing approach to perception, both in the "New Look" approach to perception in psychology (Bruner 1957; Neisser 1967) and AI (Freuder 1974; Tenenbaum & Barrow 1976), focused on the influence of the subject's *expectations* (i.e. beliefs) on perceptual processing, and sometimes even suggested that there are no separable bottom-up aspects of perception. This approach proved to be unworkable – computational models were hopelessly weak and inflexible and psychological theory and experiment soon became bogged down. Only when the focus switched to studying perceptual processes which are independent of expectations, and thus separable from the processes studied by folk psychology, did a genuine cognitive science of vision begin to emerge.

Memory. The psychology of memory has been dominated by the study of nonsense materials (such as strings of numbers or letters) or decontextualized materials (such as word lists or pairs of words), where the influence of prior knowledge can be minimized. This tradition stems from Ebbinghaus's (1885) experiments on rates of forgetting and runs through to the present day (see, e.g., Baddeley 1990). The bulk of theoretical work on memory stands firmly in this knowledge-free tradition: distinctions in modern memory theory between Central Executive, Visual-Spatial Scratch Pad, Articulatory Loop, Semantic, Episodic and Procedural Memory stores and the account of putative relations between them are all characterized independently of the specific contents of these stores.

An apparent counterexample to this claim is that memory stores are often distinguished by the way in which they code information (as ascertained, for example, by examining recall errors). Notice, however, that discussion of content at this level does not encroach on the specifics of folk psychology. Either the upshot is that information is encoded in some specialized code (e.g. in phonological or articulatory terms), or that it is stored in a "semantic" code. What is meant by a semantic code is just that the information is integrated into the store of general knowledge; memory theory takes no stand regarding whether the contents of the this store correspond to the contents of propositional attitudes. That is, it does not encroach into the domain of folk psychology.

There are, however, important determinants of memory performance which do appear to depend crucially on general knowledge. The study of such effects goes back to Bartlett (1932), who presented subjects with passages containing culturally unfamiliar material, and found that their errors in recall were likely to be consistent with their own cultural assumptions. Research based on Bartlett's approach (see, e.g., Neisser 1982; Grunberg *et al.* 1988) is concerned with knowledge-rich explanation, in that its assumptions and conclusions are dependent on information specific to cultures and individuals. These aspects of memory,

while of considerable importance, can not be assimilated into cognitive science, since they depend crucially on knowledge-rich processes.

The interaction between knowledge-free and knowledge-rich aspects of memory can actually be rather intricate. For example, it is a putative structural constraint on short term memory performance that only around seven chunks of information can be remembered (Miller 1956). But what counts as a chunk is dependent on the extent to which the subject is able to make sense of something as a single entity (a number, a word, a well known phrase), which is a knowledge-rich process par excellence. Folk psychology would seek to explain why a person employs *particular* chunks of information in remembering: obviously this is not a question which the psychology of memory can address. Similarly, the exponential form of forgetting curves in long-term memory may be knowledge independent, but how well information is encoded in the first instance is strongly affected by how well it can be integrated into world knowledge (Bransford *et al.* 1972). We claim that cognitive science can only successfully model the knowledge-free aspects of such interactions, and the boundary between knowledge-free and knowledge-rich.

Language and language comprehension. Most modern linguistics has deliberately and explicitly eschewed folk psychology considerations in its theorizing. One obvious example is Chomsky's (1965) distinction between competence and performance, which aims to insulate the study of the structure of language from the factors affecting actual language production (which will, of course, include the beliefs and desires which lead to a speaker's utterance). Chomsky (1980) amplifies this point by stressing that the linguist aims to describe the speaker/hearer's *capacity* to produce/understand sentences, but is not concerned with the choices involved in choosing to make a particular utterance (for which folk psychology might be an appropriate description). Furthermore, Chomsky insists that linguistics can be separated from the study of mental life at large by arguing for the autonomy of linguistic levels, and argues that knowledge of the language is carried in an innate and special purpose language acquisition device, which does not interact with general knowledge. Not everyone would accept all the details of Chomsky's account, but it is generally accepted that there are aspects of language and its mental representation which can be studied in a manner independent of general knowledge.

It is important to remember that the structuralist tradition in linguistics (e.g. Bloomfield 1933) and behaviorist psychology (e.g. Skinner 1957), against which Chomsky was reacting, took actual linguistic behavior as its starting point. Chomsky proposed instead that rigorous formal linguistics could only develop by splitting apart the contribution of knowledge of language from the contribution of general cognition. Note that there has been some reaction against this assumption (e.g. the cognitive linguistics of Lakoff (1986)), but most other "non-transformational" approaches to linguistics such as Lexical-Functional Grammar (Kaplan and Bresnan 1982) and Generalized Phrase Structure Grammar (Gazdar *et al.*

1985) fundamentally retain the assumption that common-sense, knowledge is irrelevant to the goals of linguistics.

Work in language processing has in general assumed that comprehension as a whole employs a component of autonomous linguistic knowledge which interacts with general knowledge. However, there are aspects of comprehension which are independent of expectations and beliefs. There is much debate about exactly what these aspects are and how interaction occurs. This is exactly as we would expect: determining the boundary between knowledge-free and knowledge-rich processing is a fundamental question in cognitive science. Let us consider word recognition and sentence processing in turn.

Swinney (1979) gives evidence that there are context-independent processes in word recognition. Subjects listened to sentences which included an ambiguous word, but where the appropriate sense of the word was clear from the (prior) context. In the standard example, *bugs* can mean (roughly) "insects" or "eavesdropping devices". By using a cross-modal priming paradigm (where subjects listen to a sentence and at the same time make a lexical decision to a written word), he found evidence that the "eavesdropping devices" sense of *bugs* was made momentarily active even when the context made it clear that the "insects" sense was the appropriate one. However, very shortly afterwards, the inappropriate sense was no longer active and only the contextually appropriate sense remained. The analysis given to this experiment by Fodor (1983) is that activation of word senses is modular, but that context can rapidly choose between the available senses. In our terms, the experiment suggests that the initial processing of a word is knowledge-free: these mechanisms can be understood by cognitive science. (Note that this interpretation is somewhat controversial; see Taft (1991) for discussion).

Indeed, the most obvious account is a computationally very simple one: that all senses of an ambiguous word are immediately activated. The cognitive system then attempts to integrate these potential senses with the context and general knowledge; at this point it is generally said that "the most contextually appropriate sense is chosen". But it is not possible to analyse the expression "contextually appropriate" any further in cognitive-scientific terms, because we would then make direct reference to general knowledge. In this particular case, we would need to be able to determine the contexts in which insects are likely to occur. Making this assessment, in general, requires the deployment of arbitrary amounts of world knowledge, and we suggest, cannot therefore be modeled by cognitive science.⁴

In sentence processing, the point at which general knowledge becomes relevant is very controversial (see for instance the papers in Altmann 1989). Again, the usual question regards the resolution of ambiguity: how does the processor make the initial decision to opt for one analysis rather than the other (assuming it does not simply work in parallel)? For instance, the sentence *the man told the girl that Bill knew about the events* has two possible interpretations, one where what the man told the girl was that Bill knew about the events, and the other where the

man told a particular girl, the one that Bill knew, about the events. The first interpretation is generally regarded as much easier to obtain, at least when the sentence is presented in isolation from context (as here). On one standard account (e.g. Frazier 1979), this reading is obtained for purely linguistic reasons (the syntactic analysis is taken to be simpler in a clearly defined sense). But other accounts hold that general knowledge plays an important role during the initial processes of analysis.

One type of account is very similar to Swinney's account for word recognition. All possible interpretations are proposed together, and the most contextually appropriate one is chosen. This model proposes a clear boundary between knowledge-free and knowledge-rich processes. But it is still necessary to determine exactly when general knowledge is consulted. In our example, does the processor choose between the analyses at the end of the sentence or at some earlier point, such as after *knew* or after *that*? There are many such unresolved questions concerning knowledge-free aspects of the account.

A currently favoured version of this model (Crain & Steedman 1985; Altmann & Steedman 1988) is rather more complicated. It assumes that the choice between the readings is made on referential grounds. If the context to our example sentence has mentioned more than one girl, then it makes sense to assume that the speaker wishes to indicate which girl is being talked about by the expression *the girl that Bill knew*. Hence the reading consistent with this analysis is chosen. More strictly, the speaker has set up a "discourse model" containing more than one girl-token, and makes reference to this model to choose the appropriate analysis. This may appear to be an account making explanatory use of a component of general knowledge. But in fact the account breaks into two parts, one knowledge-rich and one knowledge-free. The knowledge-rich component is involved in extracting certain aspects of the (knowledge-rich) interpretation of the context and placing it in the discourse model.

The nature of this process is not specified, and, we would argue, this is because such knowledge-rich processes are beyond the scope of current cognitive science. The second process, which is the focus of cognitive-scientific interest, uses the current state of the discourse model in processing new linguistic material. Such processes can be described in knowledge free terms, and hence lie within the scope of cognitive-scientific theory.

Whatever the merits of the particular accounts of word recognition and sentence processing that we have considered, they are good illustrations of the way in which cognitive scientists have attempted to factor knowledge-free aspects of cognition apart from processes involving general world knowledge, and have attempted to study only the former.

Discussion

We have argued that folk psychology and cognitive science are in practice concerned with entirely separate domains: *a fortiori*, cognitive science is not, and

should not be, folk psychology formalized, as the received view suggests. But equally, the critics of the received view (e.g. Stich 1983; P. S. Churchland 1986), who argue that cognitive science will gradually replace folk psychology, are also mistaken. Cognitive science will not replace folk psychology, since it does not encroach upon its territory. We cannot rule out the possibility that cognitive science and folk psychology may one day co-operate or compete to explain the same mental processes; but for now, and for the foreseeable future, they have entirely separate concerns.

The distinction between knowledge-free and knowledge-rich explanation splits the traditional subject matter of psychology in two; we have argued that a crucial practical question in the development of cognitive science is to determine which aspects of the subject fall on the knowledge-free side of the divide, and hence may be tractably studied by cognitive science. In fact, though, this division can be seen as having a much wider significance in the study of human behavior. The knowledge-free aspect includes cognitive science, and perhaps ultimately the neurosciences. In these areas, it may be possible to provide detailed explanations of the mechanisms underlying human thought.

Knowledge-rich explanation, by contrast, involves the *interpretation* of human behavior, which is related to folk-psychological explanation. Social psychologists, anthropologists, sociologists, historians and literary theorists consider the way in which people see themselves, others and their culture: such interpretation involves spelling out the beliefs and desires in terms of which people see the world. This project may be very difficult, since the relevant beliefs and desires may not directly be introspectively available to people (this might be particularly true, no doubt, of beliefs and values which are so fundamental to a person's world view that they are unaware of them); and furthermore, it may require great sophistication to interpret and unravel views very different from one's own. While such interpretative explanations are rooted in folk psychological explanation – they ascribe meaning to people's mental states, practices and products – they may enormously extend and enrich our common-sense folk psychology conceptions. However, we currently have no conception of how such explanation could be studied in computational terms, and hence it remains entirely separate from cognitive science.

The failure to distinguish these two projects has led to each being criticized for not doing the job of the other. So, cognitive science has been attacked for failing to integrate with social and cultural explanation, and folk psychology has been attacked as a poor basis for computational explanation. We shall now consider these points in turn.

SCEPTICISM ABOUT COGNITIVE SCIENCE

A common criticism of cognitive science is that it is concerned with the minutiae of cognition, and makes no attempt to integrate with social and cultural studies of

people. Hence it will never be able to explain the “really important” aspects of mind and behavior: how and why we think and act as we do. In order to explain such things, we should turn away from a computational, information-processing, approach to cognition and turn to one where the cultural meaning of what is being thought is fundamental. Some examples of this perspective are Shotter (1975), Coulter (1979; 1983), and Bruner (1990). Here we shall concentrate on Bruner's discussion.

Bruner (1990) attacks current cognitive science for concentrating on the minutiae of cognition, and failing to integrate with social and cultural studies of man. He argues that the initial impetus for information-processing psychology in reaction to behaviorism was the thesis that the cognitive system imposes meaning on the world, and must be understood in these terms. This study of the meaning we invest in the world was to tie up with what he calls the *interpretative* study of man as a social and cultural entity, study which involves attempting to interpret the actions, artefacts and words of individuals and societies. (In particular, this involves the ascription of beliefs, desires, fears and other propositional attitudes – in short, social and cultural explanation involves, though it enriches and elaborates, common-sense folk psychological explanation.) Bruner bemoans the fact that, in practice, this initial direction was lost, as computational ideas pushed psychologists to study just those processes which are independent of such considerations, and calls for psychology to take up the wider project once more. In particular, he advocates the naturalistic study of narrative, and gives some examples of his approach.

We do not disagree with Bruner that the integration of cognitive processes with cultural issues is important. But we do not accept that the study of knowledge-free processes is a trivial side-issue in the scientific study of human beings. In other words, we do not accept that cognitive science has neglected its duty by focusing on knowledge-free, and hence socially and culturally inert, aspects of mind. In fact, we would argue that the fact that computational explanation can only tractably be applied to knowledge-free processes (at least currently) was a fundamental *discovery* of the information processing revolution, and that cognitive science has only succeeded where this discovery has been taken into account.

We believe that the relationship between cognitive science on the one hand, and actual mental activity in particular cultural contexts on the other, is similar to the relationship between human anatomy and physiology, and behavior in particular contexts. For example, we may provide a theory of how running can occur in terms of the anatomy and physiology of the leg. We will also be able to make behavioral predictions on the basis of this theory, for instance how fast it is possible to run. This clearly provides biological constraints on how running may occur in any particular cultural context. But we cannot use this information to determine whether a person will choose to run in any given circumstance. No-one would regard the investigation of the biology of running, independent of any cultural considerations, as irrelevant to the science of human beings. Exactly the

same argument holds for cognitive science: for instance, the study of memory is not rendered irrelevant when it pays no attention to cultural considerations. Of course, the cultural study of how memory is used may be valuable as well, just as the cultural study of running may be.

Clearly, cultural issues are exactly what we have termed knowledge-rich issues, except that the person's general knowledge is the knowledge of culture (among other things). It appears that scepticism towards cognitive science should disappear, once the knowledge-free/knowledge-rich split is recognized. Cognitive science should no longer look like a bare and unconvincing model of human thought and behavior, since it should be clear that providing such a model is not a goal of cognitive science. In fact, its goal is far more modest than this, yet it is far from trivial. Bruner's proposed "cultural psychology" is simply a different enterprise from cognitive science.

SCEPTICISM TOWARDS FOLK PSYCHOLOGY

As we have already seen, opponents of the received view, such as Stich and the Churchlands, argue that folk psychology is being overthrown by cognitive science, perhaps in close alliance the neurosciences. We have already noted that this position presupposes that cognitive science can successfully encroach on the subject matter of folk psychology, which we have argued it cannot.

Here, we consider a number of examples used by sceptics towards folk psychology, which purport to show how it is superseded and thrown into doubt by cognitive science. We argue that none of these examples challenges our thesis that folk psychology and cognitive science are entirely separate.

P. M. Churchland (1981) argues that folk psychology is inadequate as a basis for cognitive science since it fails to address mental illness, the nature of sleep, creativity, individual differences, perception, motor control and memory. But in the light of the discussion above, it seems rather clear that, as a criticism of folk psychology, this misses the mark. First, it is inappropriate to criticize folk psychology for failing to cover processes, such as perception, motor control and memory about which there are good knowledge-free accounts. Folk psychology is only concerned with explaining behavior in terms of knowledge and argument. Second, and more importantly, in knowledge-rich domains in which folk psychology is found wanting, such as the basis of mental illness, cognitive science fares no better. Indeed, while folk psychology can at least provide a description of the phenomena in interpretative terms, cognitive science is entirely unable to begin. Cognitive science cannot begin to explain rational thought in knowledge-rich domains, and hence it has no chance of explaining how such thought breaks down. So Churchland's argument that an inadequate folk psychology will be replaced by cognitive science and the neurosciences cannot be sustained, precisely because folk psychology and cognitive science concern separate domains.

Stich (1983) also considers phenomena which he takes to be beyond the scope of folk psychology, but which, it turns out, are also beyond the scope of cognitive science. So, for example, he considers how folk psychological explanation breaks down in attempting to describe the cognitive state of a person with a severe acquired memory disorder, who attests that McKinley was assassinated, but does not have any idea who McKinley was, or what assassination or even dying involves. Certainly folk psychology does have difficulty making sense of such cases – as we noted above real folk psychological explanation is only possible because the open-ended and seemingly infinite set of relevant background assumptions can be assumed as common to both the individual being explained and the audience to whom the explanation is addressed. When the mental state of the person to be explained is very abnormal, such as in the case of some severe neuropsychological impairments, in cases of severe psychiatric disorders, or for that matter in the interpretation of the mental states of children (see, e.g., Carey 1988 for discussion of how drastically different the child's account of the world can be from that of the adult), then the background assumptions made by the audience may no longer be appropriate. That is, the knowledge which allows the audience to fill out the argument, and judge its coherence, may have little relation to the knowledge of the agent whose mental life is to be explained. Hence, of course, folk-psychological explanation is inevitably less successful. But again, the most important point here does not concern the adequacy or otherwise of folk psychological explanation of such cases: it is simply that cognitive science provides no alternative which can overthrow the folk-psychological analysis.

This point also applies to a rather different case with which Stich attempts to cast doubt on folk psychology. He cites a well-known study by Storms and Nisbett (1970), who gave two groups of insomniac subjects a placebo pill before going to bed. One group was told that the pill would produce the symptoms of insomnia, the other that it would have a relaxing effect. In fact, the first group slept better than the second. The phenomenon can be explained by assuming that the first group attribute their symptoms to the pill rather than to their own states, and thus they do not suffer the anxiety associated with the insomniac state. This explanation is sensitive to world knowledge in many ways: for instance, if they were told that the drug was a placebo or produced different symptoms, they would behave differently. In addition, their behavior requires the belief that the symptoms of insomnia that they experience are indeed associated with the inability to sleep. The reason that Stich invokes the example is to show that folk psychology is inadequate (because it cannot explain such phenomena) and that it is being replaced by more adequate cognitive-science treatments. But, in fact, the explanations offered for this kind of effect, and other similar effects throughout social psychology, are knowledge-rich through and through, and quite beyond the grasp of computational, cognitive scientific explanation.

We draw a very different lesson from considering cases in which folk psychology is inadequate, and the theories within clinical, social, or developmental

psychology that attempt to account for such cases. This is that such accounts, in so far as they succeed at all, do so by using interpretative methods, which ascribe mental states and state transitions (though perhaps rather unusual ones) to the agents under study, firmly in the style of propositional attitude explanation. Thus, such explanations extend and enrich folk psychology, rather than overthrow it. Such knowledge-rich, interpretative explanation is, we have argued, outside the scope of current cognitive science.

Implications and Conclusions

If we are right, and the domains of folk psychology and cognitive science do not overlap, then the folk psychology debate should assume a much less central role in the discussion of the foundations of cognitive science. Practical cognitive scientists can, perhaps, sleep easier in the knowledge that their work does not rely on the truth of a folk theory. But more importantly the distinction between knowledge-free and knowledge-rich processes provides a clear methodological pointer for organizing cognitive scientific research, a pointer which makes sense both of existing methodological canons, and issues which have been seen as central in a range of areas of the discipline.

We have argued that the distinction between knowledge-free and knowledge-rich has a wider significance across the behavioral sciences: specifically that the domains in which cognitive science and neuroscientific explanation will be appropriate will not overlap with the domains in which interpretative explanations of human behavior are appropriate. If this view is correct, while the truth of folk psychology is irrelevant to cognitive science, it seems that it may be highly relevant to the social sciences and humanities, whose interpretative style of explanation is intimately bound up with folk psychology. However, it is possible that the truth of folk psychology, while an interesting theoretical question, may have little more practical significance for the interpretative disciplines than it does not cognitive science. The intractability of the study of knowledge-rich processes means that no systematization of, or replacement for, folk psychology will be forthcoming from cognitive science, at least in the near future. Therefore, whether right or wrong, folk psychology is the only theory of the explanation of knowledge-rich processes that we have available, and will remain so for some time, perhaps for ever.

It may seem obvious that insisting on a strong knowledge-free/knowledge-rich distinction is necessarily divisive regarding different approaches to understanding human behavior. In fact, however, we would argue that a better understanding of the difference between the two kinds of explanation may actually diminish misunderstanding and friction between different methodologies for studying human behavior. On the one hand, it may help persuade those working within a cognitive scientific or neurobiological framework that the patterns of "scientific" explanation that are appropriate to understanding the individual mechanisms of

cognition cannot be applied to study human behavior in its entirety, in its social and cultural context. On the other hand, researchers in the social sciences and humanities who use an interpretative approach may be more convinced that there are certain aspects of cognition which can be studied independently of interpretative factors.

Notes

¹The order of authorship is arbitrary. Martin Pickering acknowledges the support of a British Academy Post-Doctoral Research Fellowship. Nick Chater is now at the University of Oxford, Department of Experimental Psychology, South Parks Road, Oxford, OX1 3UD, England.

²Of course, both decision theory (and the probability theory that underlies it) and logic have been criticized as poor psychological models (Kahneman *et al.* 1982; Johnson-Laird & Byrne 1991; Oaksford *et al.* 1990) although both have their defenders (Braine 1990; Gigerenzer 1993; O'Brien 1993; Pearl 1988). We ignore such issues here.

³Other approaches include McCarthy's (1980) circumscription, McDermott & Doyle's (1980) Non-Monotonic Logic I, Moore's (1985) Autoepistemic Logic, McDermott's (1982) Non-Monotonic Logic II, and Clark's (1978) Predicate Completion. It has been shown that these approaches are closely related (Hanks & McDermott 1985; Shoham 1987), and furthermore, relations have been made (Ginsberg 1986) to a very different approaches to non-monotonic reasoning, the Truth Maintenance System (Doyle 1979). The difficulties that we discuss apply equally to each of these approaches (Oaksford & Chater 1991; and see Israel 1980, Perlis 1987, McDermott 1987 for other general critiques).

⁴It is possible, however, that knowledge-free approximations to plausibility might be computed by specialist processes, based on, for example, word association. Such processes could be studied by cognitive science. This possibility is considered by Fodor (1983).

References

- Abu-Bakar, M. & Chater, N. (1993), 'Processing time-warped sequences using recurrent neural networks: Modelling rate-dependent factors in speech perception' *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Lawrence Erlbaum, 191-197.
- Allen, J., Hendler, J. & Tate, A. (Eds.) (1990), *Readings in Planning*, San Mateo, California: Morgan Kaufman.
- Altmann, G. (Ed.) (1989), *Language and Cognitive Processes: Special Issue on Parsing and Interpretation*, London and Utrecht: Lawrence Erlbaum Associates.
- Altmann, G. & Steedman, M. 1988 'Interaction with context during human sentence processing', *Cognition* 30, 191-238.
- Baddeley, A. D. (1990), *Human Memory: Theory and Practice*, Boston: Allyn and Bacon.
- Bartlett, F. C. (1932), *Remembering: A Study in Experimental and Social Psychology*, Oxford: Cambridge University Press.
- Bechtel, W. & Abramson, A. (Eds.) (1991), *Connectionism and the Mind: An Introduction to Parallel Distributed Processing in Networks*, Oxford: Blackwell.
- Brachman, R. & Levesque, H. (Eds.) (1985), *Readings in Knowledge Representation*, Los Altos, California: Morgan Kaufman.
- Braddick, O. J. (1979) 'Low- and high-level processes in apparent motion', *Philosophical Transactions of the Royal Society of London*. B 290, 137-151.
- Braine, M. D. S. (1990), 'The "natural logic" approach to reasoning', In W. F. Overton (Ed.), *Reasoning, Necessity and Logic: Developmental Perspectives*, Hillsdale, NJ: Erlbaum.
- Bransford, J. D., Barclay, J. R. & Franks, J.J. (1972), 'Sentence memory: A constructive versus interpretative approach', *Cognitive Psychology* 3, 193-209.

- Bruner, J. (1957), 'On perceptual readiness', *Psychological Review* 65, 14-21.
- Bruner, J. (1990), *Acts of Meaning*, Cambridge, Mass: Harvard University Press.
- Carey, S. (1988), 'Conceptual differences between children and adults', *Mind and Language*, 3, 167-181.
- Chater, N. & Oaksford, M. (1990), Logicist cognitive science and the falsity of common sense theories. Technical Report UWB-CNU-TR-90-4, Cognitive Neurocomputation Unit, Department of Psychology, University of Wales, Bangor.
- Chater, N. & Oaksford, M. R. (1993), 'Logicism, mental models and everyday reasoning: Reply to Garnham', *Mind & Language* 8, 72-89.
- Chomsky, N. (1965), *Aspects of a Theory of Syntax*, Cambridge, Mass: MIT Press.
- Chomsky, N. (1980), *Rules and Representations*, Oxford: Blackwell.
- Churchland, P. M. (1981), 'Eliminative materialism and the propositional attitudes', *Journal of Philosophy* 78, 67-90.
- Churchland, P. M. (1989), *A Neurocomputational Perspective*, Cambridge, Mass: Bradford Books/MIT Press.
- Churchland, P. S. (1986), *Neurophilosophy*, Cambridge Mass: Bradford Books/MIT Press.
- Clark, K. (1978), 'Negation as failure', in H. Gallaire & J. Minker (Eds.), *Logic and Databases*, New York: Plenum Press, pp 293-322.
- Coulter, J. (1979), *The Social Construction of Mind*, London: Macmillan.
- Coulter, J. (1983), *Rethinking Cognitive Theory*, London: Macmillan.
- Crain, S. & Steedman, M. (1985), 'On not being led up the garden path: The use of context by the psychological parser', in Dowty, D., Karttunen, L. and Zwicky, A. (Eds.), *Natural Language Parsing: Psychological, Computational and Theoretical Perspectives*. Cambridge: Cambridge University Press.
- Davis, E. (Ed.) (1990), *Representations of Commonsense Knowledge*, Los Altos, California: Morgan Kaufman.
- Derthick, M. (1987), A connectionist architecture for representing and reasoning about structured knowledge. Technical Report CMU-BOLTZ-29, Department of Computer Science, Carnegie-Mellon University.
- Devitt, M. (1990), 'A narrow representational theory of mind', in W. Lycan (Ed.), *Mind and Cognition*, Oxford: Blackwell, pp. 371-398.
- Doyle, J. (1979), 'A truth maintenance system', *Artificial Intelligence* 12, 231-272.
- Dreyfus, H. L. & Dreyfus, S. E. (1986), *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, New York: The Free Press.
- Ebbinghaus, H. (1885, translated 1913), *Memory*, New York: Teachers College Press.
- Fodor, J. A. (1975), *The Language of Thought*, New York: Thomas Crowell.
- Fodor, J. A. (1980), 'Methodological solipsism considered as a research strategy in cognitive psychology', *Behavioral and Brain Sciences* 3, 63-109.
- Fodor, J. A. (1983), *Modularity of Mind*, Cambridge, Mass: MIT Press.
- Fodor, J. A. (1986), 'Banish DisContent', in J. Butterfield (Ed.), *Language, Mind and Logic*, Cambridge University Press, reprinted in W. Lycan (Ed.) (1990), *Mind and Cognition*, Oxford: Blackwell, pp. 420-438.
- Fodor, J. A. (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, Mass: MIT Press.
- Fodor, J. A. & Pylyshyn, Z. W. (1988), 'Connectionism and cognitive architecture: A critical analysis', *Cognition* 28, 3-71.
- Frazier, L. (1979), On Comprehending Sentences: Syntactic Parsing Strategies. PhD Dissertation, Indiana University Linguistics Club.
- Freuder, E. C. (1974), A computer vision system for visual recognition using active knowledge. MIT A.I. Laboratory, Technical Report 345.
- Garnham, A. (1993), 'Is logicist cognitive science possible?', *Mind and Language* 8, 49-71.
- Gazdar, G., Klein, E., Pullum, G. and Sag, I. (1985), *Generalized Phrase Structure Grammar*, Cambridge, Mass: Harvard University Press.
- Gigerenzer, G. (1993), 'The bounded rationality of probabilistic mental models', in K. I. Manktelow & D. E. Over (Eds.), *Rationality: Psychological and Philosophical Perspectives*, London: Routledge, pp 284-313.

- Ginsberg, M. L. (1986), 'Multi-valued logics', in *Proceedings of the Fifth National Conference on Artificial Intelligence*, pp 243-247.
- Ginsberg, M. L. (Ed.) (1987) *Readings in Nonmonotonic Reasoning*, Los Altos, California: Morgan Kaufman.
- Grunberg, M. M., Morris, P. E. & Sykes, R. N. (Eds.) (1988), *Practical Aspects of Memory. Vol 2*, London: Wiley.
- Hanks, S. & McDermott, D. (1985), 'Default reasoning, nonmonotonic logics, and the frame problem', in *Proceedings of the American Association for Artificial Intelligence*, Philadelphia, PA.
- Haugeland, J. (1978), 'The nature and plausibility of cognitivism', *Behavioral and Brain Sciences* 3, 63-73.
- Hinton, G. E. (1981), 'Implementing semantic networks in parallel hardware', in G. E. Hinton & J. A. Anderson (Eds.), *Parallel Models of Associative Memory*, Hillsdale, NJ: Erlbaum.
- Horgan, T & Woodward, J. (1985), 'Folk psychology is here to stay', *The Philosophical Review* XCIV; reprinted in W. Lycan (Ed.) (1990), *Mind and Cognition*, Oxford: Blackwell, pp. 399-420.
- Israel, D. (1980), 'What's wrong with non-monotonic logic?' in *Proceedings of the First National Conference on Artificial Intelligence*, pp, 99-101.
- Johnson-Laird, P. N. & Byrne, R. M. J. (1991), *Deduction*, Hove, UK: Erlbaum.
- Julesz, B. (1960), 'Binocular depth perception of computer generated patterns', *Bell Systems Technical Journal* 39, 1125-1162.
- Julesz, B. (1971), *Foundations of Cyclopean Perception*, Chicago: University of Chicago Press.
- Julesz, B. (1975), 'Experiments in the visual perception of texture', *Scientific American* 232, 34-43.
- Kahneman, D., Slovic, P. & Tversky, A. (1982), *Judgement under Uncertainty: Heuristics and Biases*, New York: Cambridge University Press.
- Kaplan, R. and Bresnan, J. (1982), 'Lexical-functional grammar: a formal system for grammatical representation', in Bresnan, J. (Ed.), *The Mental Representation of Grammatical Relations*, Cambridge, Mass: MIT Press.
- Lakoff, G. (1986), *Women Fire and Dangerous Things: What Categories Tell Us about the Nature of Thought*, Chicago: University of Chicago Press.
- Linsker, R. (1988), 'Self-Organisation in a Perceptual Network', *Computer*, March 1988, 21, 3, I 105-117.
- Marr, D. (1982), *Vision: A Computational Investigation into the Representation and Processing of Visual Information*, New York: Freeman.
- McCarthy, J. (1980), 'Circumscription - A form of non-monotonic reasoning', *Artificial Intelligence* 13, 27-39.
- McCarthy, J. & Hayes, P. (1969), 'Some philosophical problems from the standpoint of artificial intelligence', in B. Meltzer and D. Michie (Eds.) *Machine Intelligence* 4, Edinburgh University Press.
- McClelland, J. L. & Elman, J. L. (1986), 'Interactive processes in speech perception: The TRACE model', in D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Vol 2*, Cambridge, Mass: Bradford Books/MIT Press, pp. 58-121.
- McDermott, D. (1982), 'Non-monotonic logic II: non-monotonic model theories', *Journal of the ACM* 29, 33-57.
- McDermott, D. (1987), 'A critique of pure reason', *Computational Intelligence*, 3, 151-160.
- McDermott, D. & Doyle, J. (1980), 'Non-monotonic logic I', *Artificial Intelligence* 13, 41-72.
- Miller, G. A. (1956), 'The magical number seven, plus or minus two: Some limits on our capacity for processing information', *Psychological Review* 63, 81-97.
- Minsky, M. (1975), 'A framework for representing knowledge', in P. Winston (Ed.), *The Psychology of Computer Vision*, New York: McGraw Hill.
- Moore, R. (1985), 'Semantical considerations in nonmonotonic logic', *Artificial Intelligence* 25, 75-94.
- Neisser, U. (1967), *Cognitive Psychology*, New York: Appleton-Century-Crofts.
- Neisser, U. (1982), *Memory Observed*, San Francisco: Freeman.
- North, D. W. (1968), 'A tutorial introduction to decision theory', Reprinted in: G. Schafer and J. Pearl (Eds.), *Readings in Uncertain Reasoning*, San Mateo, California: Morgan Kaufman, pp. 68-78.
- Oaksford, M. & Chater, N. (1991), 'Against logicist cognitive science', *Mind and Language* 6, 1-38.

- Oaksford, M., Chater, N. & Stenning, K. (1990), 'Connectionism, classical cognitive science and experimental psychology', *AI and Society* 4, 73-90.
- O'Brien, D. P. (1993), 'Mental logic and human irrationality', in K. I. Manktelow & D. E. Over (Eds.), *Rationality: Psychological and Philosophical Perspectives*, London: Routledge, pp. 110-135.
- Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems*, San Mateo, CA: Morgan Kaufman.
- Perlis, D. (1987), 'On the consistency of commonsense reasoning', *Computational Intelligence* 2, 180-190.
- Pylyshyn, Z. W. (1984), *Computation and Cognition: Toward a Foundation for Cognitive Science*, Montpelier, Vermont: Bradford Books/MIT Press.
- Quine, W. V. O. (1953), 'Two dogmas of empiricism', in W. V. O. Quine, *From a Logical Point of View*, Cambridge, Mass: Harvard University Press.
- Ramsey, W., Stich, S. & Garon, J. (1990), 'Connectionism, eliminativism, and the future of folk psychology', in Tomberlin, J. E. (Ed.), *Philosophical Perspectives, Volume 4: Action Theory and the Philosophy of Mind*, Atascadero, CA: Ridgeview.
- Reiter, R. (1980), 'A logic for default reasoning', *Artificial Intelligence* 13, 81-132.
- Reiter, R. (1985), 'On reasoning by default', in R. Brachman and H. Levesque (Eds.), *Readings in Knowledge Representation*, Los Altos, California: Morgan Kaufman.
- Richards, W. (Ed.) (1988), *Natural Computation*, Cambridge, Mass: Bradford Books/MIT Press.
- Rumelhart, D. E., Lindsey, P. H. & Norman, D. A. (1972), 'A process model for long-term memory', in E. Tulving and W. Donaldson (Eds.), *Organization of Memory*, New York: Academic Press.
- Rumelhart, D. E. & McClelland, J. L. (1986), 'On learning the past tenses of English verbs', in D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Vol 2*, Cambridge, Mass: Bradford Books/MIT Press, pp. 216-271.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L. & Hinton, G. E. (1986), 'Schemata and sequential thought processes in PDP models', in D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Vol 2*, Cambridge, Mass: Bradford Books/MIT Press pp. 7-57.
- Sanger, T. (1993), 'A practice strategy for robot learning control', in S. J. Hanson, J. D. Cowan & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems 5*, San Mateo, CA: Morgan Kaufman, pp. 335-341.
- Schacter, R. D. (1986), 'Evaluating influence diagrams', Reprinted in: G. Schafer and J. Pearl (Eds.), *Readings in Uncertain Reasoning*, San Mateo, California: Morgan Kaufman, pp. 79-90.
- Schafer, G. & Pearl, J. (Eds.) (1990), *Readings in Uncertain Reasoning*, San Mateo, California: Morgan Kaufman.
- Schank, R. C. (1975), *Conceptual Information Processing*, Amsterdam: North Holland.
- Seidenberg, M. & McClelland, J. L. (1989) A distributed, developmental model of word recognition and naming', *Psychological Review* 96, 523-568.
- Shafer, G. & Pearl, J. (1990), *Readings in Uncertain Reasoning*, San Mateo, California: Morgan Kaufman.
- Shastri, L. (1985), Evidential reasoning in semantic networks: A formal theory D and its parallel implementation. TR166, Department of Computer Science, University of Rochester.
- Shoham, Y. (1987), 'Non-monotonic logics', M. L. Ginsberg (Ed.), *Readings in Non-Monotonic Reasoning*, Los Altos, California: Morgan Kaufman.
- Shotter, J. (1975), *Images of Man in Psychological Research*, London: Methuen.
- Skinner, B. F. (1957), *Verbal Behavior*, New York: Appleton-Century-Crofts.
- Smolensky, P. (1987), On variable binding and the representation of symbolic structures in connectionist systems. Technical Report CU-CS-355-87, Department of Computer Science, University of Colorado at Boulder.
- Stich, S. (1983), *From Folk Psychology to Cognitive Science*, Cambridge, Mass: Bradford Books/MIT Press.
- Storms, M. & Nisbett, R. (1970), 'Insomnia and the attribution process', *Journal of Personality and Social Psychology* 16, 319-328.
- Sutton, R. S. & Barto, A. G. (1991), 'Time derivative models of Pavlovian reinforcement', in M. Gabriel and J. W. Moore (Eds.), *Learning and Computational Neuroscience*, Cambridge, Mass: MIT Press.

- Swinney, D. (1979), 'Lexical access during sentence comprehension: (Re)consideration of context effects', *Journal of Verbal Learning and Verbal Behavior* 18, 645-659.
- Taft, M. (1991), *Reading and the Mental Lexicon*, Erlbaum: Hillsdale, NJ.
- Tenenbaum, J. M. & Barrow, H. G. (1976), Experiments in interpretation-guided segmentation. Stanford Research Institute, Technical Note 123.
- Touretzky, D. S. & Hinton G. E. (1985), 'Symbols among the neurons: Details of a connectionist inference architecture', in *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 238-243.
- Uno, Y., Fukumura, N., Suzuki, R & Kawato, M. (1993), 'Integration of visual and somatosensory information for preshaping hand in grasping movements', in S. J. Hanson, J. D. Cowan & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems 5*, San Mateo, CA: Morgan Kaufman, pp. 311-318.
- Weld D. S. & De Kleer, J. (Eds.) (1990), *Readings in Qualitative Reasoning about Physical Systems*, Los Altos, California: Morgan Kaufman.
- Zemel, R.S. (1989), TRAFFIC: A Connectionist Model of Object Recognition. Technical Report CRG-TR-89-2, Department of Computer Science, University of Toronto.