

Is memory a natural kind?

Kourken Michaelian

Institut Jean-Nicod, France

Memory Studies

4(2) 170–189

© The Author(s) 2010

Reprints and permission: sagepub.

co.uk/journalsPermissions.nav

DOI: 10.1177/1750698010374287

<http://mss.sagepub.com>



Abstract

Though researchers often refer to memory as if it were a unitary phenomenon, a natural kind, the apparent heterogeneity of the various ‘kinds’ of memory casts doubt on this default view. This article argues, first, that kinds of memory are individuated by memory systems. It argues, second, for a view of the nature of kinds of memory informed by the tri-level hypothesis. If this approach to kinds of memory is right, then memory is not in fact a natural kind.

Keywords

memory systems, metaphysics of memory, natural kinds, nondeclarative memory, tri-level hypothesis

Memory in philosophy and psychology

Patricia Smith Churchland once pointed out that we do not know ‘whether searching for the neural substrate for “memory” is like looking for the “principle” that unites jewels, such as amethysts, diamonds, amber, and pearls’ (1986: 152), that is, that we do not know whether memory is a natural kind or, instead, a set of disparate phenomena only superficially similar to each other. Though the question whether memory is a natural kind has otherwise rarely been posed explicitly, once we have grasped the extent of the differences among the various ‘kinds of memory’,¹ it becomes inevitable that we will ask whether the kinds of memory have enough in common to qualify as bona fide subkinds of a single natural kind; whether, that is, memory might not turn out to be a merely nominal kind. Once we have grasped the depth of the differences between declarative and nondeclarative memory, in particular, it becomes inevitable that we will ask whether declarative and nondeclarative memory are subkinds of a single natural kind. That memory is a natural kind is perhaps the default hypothesis; I argue, however, that memory is not in fact a natural kind.²

As I approach it here, the question whether memory is a natural kind belongs to both the philosophy of mind and the philosophy of psychology.³ Philosophers of mind are interested in the nature of memory in general, in the distinctions among different types of memory and in the relations of those types to each other (e.g. Bernecker 2008). They have, however, usually tended to approach these questions in a largely a priori manner, rather than drawing on the empirical psychology of memory. But given a naturalistic orientation in the philosophy of mind, the philosophy of memory should at minimum be consistent with the psychology of memory. (Though Bernecker’s [2008] approach is for the most part a priori, he is to some extent concerned that his account is consistent with the psychology.) And given the sort of thoroughgoing naturalism that I presuppose

Corresponding author:

Kourken Michaelian, Institut Jean-Nicod, Ecole Normale Supérieure, Pavillon Jardin, 29, rue d’Ulm, Paris F-75005, France.

Email: kourken.michaelian@ens.fr

here, a naturalism according to which the philosophy of memory should be not only consistent with but systematically informed by the psychology of memory, the question whether memory is a natural kind becomes also a question in the philosophy of psychology, a question about the nature of memory understood as a theoretical entity (or a set of entities) investigated by psychology.

My intent is by no means to legislate a view of memory for psychology; it is rather to draw out the view of memory already implicit in psychology. But an answer to the question whether memory is a natural kind will nevertheless have normative implications for the practice of the psychology of memory, for such an answer will tell us something about what sorts of empirical approaches to memory are likely to be fruitful; the question whether memory is a natural kind, in other words, is not only of theoretical interest. A natural kind is one determined by the world, by the way properties cluster together in the world, rather than by our ideas about the world (Michaelian, 2008). Thus natural kinds, unlike merely nominal kinds, can support inductions (Quine, 1969). If memory is a natural kind, then it will often be profitable to investigate the various memory phenomena together, as if they constitute a coherent whole, for then investigation of one kind of memory will often tell us something about features of other kinds of memory. But if memory is not a natural kind, then it will not in general be profitable to investigate the various memory phenomena together, for then investigation of one kind of memory will not typically tell us anything about the other kinds of memory; in other words, if the answer to the question whether memory is a natural kind is negative, then we should not be aiming for a general theory of memory.

Anderson's 'rational analysis' of memory (Anderson, 1990; Chater and Oaksford, 1999; Schooler and Anderson, 1997), for example, can perhaps be seen as an attempt to construct such a general theory. If the main thesis of this article – that memory is not a natural kind – is right, then, whatever the merits of the theory, it cannot be successful as an account of memory as a whole. If my more specific thesis about the nature of nondeclarative memory – that many forms of nondeclarative memory are non-cognitive – is right, in particular, then Anderson's theory, which is stated in terms of search for and retrieval of stored representations, can apply at most to declarative memory. This implication of my argument is not unwelcome: this is not the place for a review of Anderson's theory,⁴ but it is significant that he conceives of memory in terms of managing 'a huge database of millions of complex facts and experiences' (Anderson, 1990: 42), a description which does not fit nondeclarative memory.⁵

In order to forestall a basic objection to the project of the article, I point out that the question whether memory is a natural kind is distinct from the question whether the multiple memory systems hypothesis – see next section – is true, so that an answer to the latter question need not automatically settle the former question for us. In light of the near-consensus among memory researchers on the multiple memory systems hypothesis, I take the truth of the hypothesis to be highly probable (and indeed I draw on the multiple memory systems literature throughout the article); but if the various hypothesized memory systems of the brain have enough in common with each other, then memory might nevertheless turn out to be a natural kind. Nor does the fact of widespread agreement on the multiple memory systems hypothesis mean that the view that memory is a natural kind is no longer the default: while most researchers agree that there are multiple memory systems, these systems tend to be discussed together, as if they are all involved in different realizations of a single general phenomenon.

The multiple memory systems hypothesis and kinds of memory

Memory systems

Before determining whether the kinds of memory are subkinds of a single natural kind, it is necessary to establish what are the kinds of memory, to formulate a working list of kinds of memory. Of

the many dozens (perhaps hundreds) of types of memory distinguished by psychologists, only a small number have a chance of qualifying as natural kinds; we thus require a procedure for identifying kinds of memory. My suggestion is that the multiple memory systems hypothesis provides a means of identifying kinds of memory: the suggestion is that kinds of memory are individuated by memory systems, so that we can identify kinds of memory by identifying memory systems. This might seem to amount to a suggestion to replace one difficult question with another, equally difficult one; but it does not, for we know a great deal about the memory systems of the brain.

It is important to distinguish the mundane thought that there are different types of memory from the view of memory embodied in the multiple memory systems hypothesis, the hypothesis that there exist 'a number of different learning and memory systems that possess the capabilities of operating independently as well as in conjunction with one another in the production of the large variety of phenomena of learning and memory with which we are already familiar, and an even larger variety of phenomena still to be discovered' (Tulving, 1984: 165). As Schacter and Tulving (1994) point out, the concept of a memory system includes the concept of a type of memory, but the latter concept does not include the former, and there is indeed an obvious technique for reconciling a view of memory as fundamentally unitary with an admission that there are multiple types of memory: we might simply suppose that the different types of memory are distinct products of a unitary memory system. We might, that is, draw on a unitary memory system hypothesis, a view according to which 'a single learning and memory system mediates the behavioral and experiential plasticity of human beings, and other higher animals, and . . . all phenomena of learning and memory in a species reflect the many ways in which that single system works' (Tulving, 1984: 165).

While the unitary memory system hypothesis was perhaps the default view in the past, it is no longer so today. This is not the place for a review of the process by which researchers have converged on the multiple memory systems hypothesis; and even a superficial review of the evidence in favour of the hypothesis is not feasible here.⁶ What matters for present purposes is simply that the hypothesis is today well established and thus constitutes a legitimate (though not an infallible) starting-point for philosophical enquiry. Disagreement over the details of the correct taxonomy of memory systems persists, but there are relatively few today who reject the multiple memory systems hypothesis itself (Willingham, 1997); the declarative/nondeclarative distinction, in particular, is acknowledged basically universally (Tulving, 2000). Admittedly, while the evidence strongly seems to favour the hypothesis, the consensus on the hypothesis is imperfect (Foster and Jelicic, 1999), and alternatives exist (e.g. Roediger et al., 1999). Strictly speaking, then, my conclusion here has the form of a conditional: if the multiple memory systems hypothesis is correct, then memory is not a natural kind.

Kinds of memory

Given the memory systems approach, what kinds of memory must we acknowledge? The basic distinction is between declarative and nondeclarative memory.⁷ Declarative memory is subserved by the medial temporal lobe and employs a specialized store in the neocortex, includes processes of encoding, storage and retrieval, and is a matter of conscious recollection of information (Milner et al., 1998). Declarative memory divides into episodic memory, which is concerned with the recollection of personal experiences and semantic memory, which is concerned with the recollection of general knowledge (Schacter et al., 2000): episodic memory, unlike semantic memory, involves auto-noetic consciousness (awareness of subjective time) (Tulving, 2002) and relies on the frontal lobes (Squire and Zola-Morgan, 1998; Wiggs et al., 1999). Nondeclarative memory, on the other

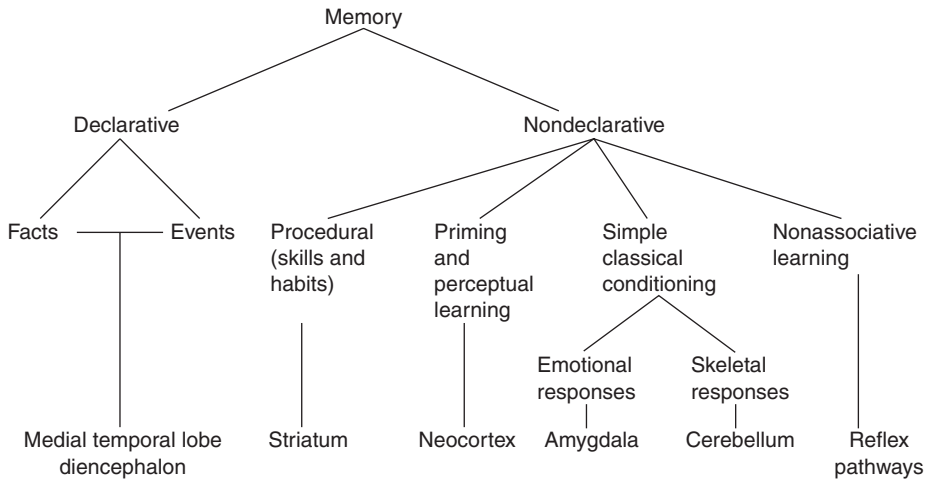


Figure 1. Taxonomy of mammalian long-term memory systems. Reprinted with permission from Elsevier, from L.R. Squire (2004) 'Memory Systems of the Brain: A Brief History and Current Perspective', *Neurobiology of Learning and Memory* 82: 171–7.

hand, is a more varied collection of phenomena: as Schacter and Tulving note, while declarative memory 'refers to a specific memory system with reasonably well characterized functional and neural properties', nondeclarative memory 'is used in a more descriptive sense to refer to a class or collection of memory functions that share certain features in common but also differ from one another in various ways' (Schacter and Tulving, 1994: 22). In general, while declarative memory is concerned with the storage and retrieval of representations,⁸ nondeclarative memory is not, being a matter instead of changes in behaviour and responses to stimuli (e.g. conditioning, priming) (Milner et al., 1998). Consequently, the concepts of encoding, storage and retrieval get little traction here (Tulving, 2000).⁹ Squire's diagram (see Figure 1) provides an overview of the brain structures subserving the various types of nondeclarative memory; note that, in general, nondeclarative memory is a matter of changes to neurons that are themselves part of performance systems (Milner et al., 1998). Thus given the memory systems approach to identifying kinds of memory, we must distinguish at least the following kinds of memory: semantic, episodic and the various non-declarative memories.

The tri-level hypothesis and kinds of memory

The question is whether these kinds of memory are, strictly speaking, subkinds of a single natural kind. In certain other domains, there are established procedures for answering analogous questions about putative natural kinds. Suppose that we want to know, for example, whether the dingo and the domesticated dog, which seem on the face of it to be subkinds of a single kind, are in fact such subkinds. The answer is not self-evident – if it seems to be, we have only to remind ourselves of similar cases (e.g. Smith Churchland's 'jewel' case) in which the 'self-evident' turned out to be incorrect. But we can (after sufficient investigation) identify (behavioural, morphological and other) properties characteristic of the putative kind (*Canis lupus*). And we can identify a set of common (genetic, environmental and perhaps other) mechanisms responsible for the clustering of those properties. The results of these investigations justify an affirmative answer to the question.

This is not to say that it is easy to construct a zoological taxonomy – this is in fact often extremely difficult to do – but only to point out that since we have a well worked-out understanding of the nature of the relevant kind of kind (the kind species), we know how to answer the question in principle. There is no similarly uncontroversial procedure for answering the question whether memory is a natural kind, but the memory systems approach to identifying kinds of memory suggests one plausible procedure. What we need is a general understanding of the nature of kinds of memory: I have suggested that memory systems individuate kinds of memory; I now argue that an understanding of the nature of memory systems grounds a general understanding of kinds of memory of the sort required to answer the question whether a given two kinds of memory are subkinds of a single natural kind.¹⁰

The concept of a memory system

While I drew earlier on our knowledge of particular memory systems, I have so far said little about the concept of a memory system as such. I want now to suggest that we can understand the concept of a memory system via the tri-level hypothesis (Marr, 1982; Pylyshyn, 1984). According to the tri-level hypothesis, a cognitive system has reality at (and so requires description at) three levels (Dawson, 1998: 288):¹¹

The computational level: The system must be described in terms of the information-processing task it performs.

The algorithmic level: The system must be described in terms of the procedure it uses to perform that task.

The implementational level: The system must be described in terms of the way in which that procedure is implemented in its wetware or hardware (i.e. in terms of the physical properties and processes used by the system to realize the procedure).

If the tri-level hypothesis is right, then memory systems require descriptions at each of the computational, algorithmic and implementational levels. My suggestion is that the concept of a memory system, as it is used in memory systems research, is precisely a tri-level concept, a concept of a system having reality at (and so requiring description at) each of these three levels.¹²

Though it is somewhat more explicit than them, the proposal to understand the concept of a memory system as a tri-level concept accords well with many characterizations of the concept by memory researchers.¹³ Summarizing their concept of a memory system, Schacter and Tulving write: ‘We [suggest] that particular memory systems be specified in terms of the nature of their rules of operation, the type of information or contents, and the neural pathways of mechanisms subserving them’ (1994: 31). This clearly suggests a tri-level approach to memory systems. Similarly, Kim and Baxter write: ‘On a biological level, a memory system is usually defined as a neural structure (or network of structures) and its interconnections, which together operate on a certain type of information and then participate in the storage of that information, either within the structure itself or elsewhere’ (2001: 324). This, again, clearly suggests the tri-level approach. The examples easily could be multiplied.

Generating the hierarchy of kinds of memory

My suggestion above was that memory systems individuate kinds of memory, so that we can use the identification of the memory systems of the brain as a means of identifying kinds of memory – by

distinguishing among memory systems, we indirectly distinguish among kinds of memory. This does not yet give us a description of the hierarchy of kinds of memory, for there need not be a one-to-one correspondence between memory systems and kinds of memory: there might be a kind of memory with various subkinds; each of the subkinds might correspond to one of the memory systems of the brain; but there need not be a memory system corresponding to the 'higher' kind.

The tri-level approach to memory systems suggests a simple solution to this problem. First: suppose that if two memory systems satisfy a single tri-level description, then they are members of a single natural kind. The idea is that the implementational- and algorithmic-level descriptions specify the mechanisms responsible for the clustering of the properties mentioned in the computational-level description. This is sufficient to generate a hierarchy of kinds of memory system, with more general descriptions specifying 'higher' kinds and more specific descriptions specifying 'lower' kinds. Next: replace the claim that memory systems individuate kinds of memory with the claim that kinds of memory system individuate kinds of memory. This is sufficient to generate the hierarchy of kinds of memory, with 'higher' kinds of memory being individuated by higher kinds of system and 'lower' kinds of memory being individuated by lower kinds of system.

Given this general understanding of the nature of kinds of memory, the procedure for answering the question whether two kinds of memory are subkinds of a single natural kind involves answering the following three questions:¹⁴

- Is there an information-processing task common to the relevant memory systems?
- Is there a procedure for performing that task common to the systems?
- Is there an implementation of that procedure common to the systems?

A negative answer to one of these questions means that the relevant memory systems are not members of a single kind of memory system; and it thus means that the kinds of memory they individuate are not subkinds of a single kind of memory.

The tri-level approach and declarative memory. Before proceeding, I pause to clarify the relevance of the tri-level hypothesis to the question whether memory is a natural kind. I have proposed that kinds of memory system individuate kinds of memory, and I have argued for a tri-level understanding of memory systems. Given these starting-points, the tri-level hypothesis is indirectly relevant to the question whether memory is a natural kind: in order to know whether two kinds of memory are subkinds of a common kind, we need to know whether there is a corresponding kind of memory system; and in order to determine whether there is a corresponding kind of memory system, we need to investigate the commonalities between the relevant systems at the computational, algorithmic and implementational levels. No doubt other starting-points are available: one might argue for an alternative understanding of the concept of a memory system, or one might propose another means entirely of individuating kinds of memory. My aim here is the modest one of developing one plausible approach to the question whether memory is a natural kind rather than the more ambitious one of canvassing every plausible approach to the question.

In order to show both that posing the question in terms of the tri-level approach does not settle it in advance and that doing so generates plausible results, I briefly consider two relatively simple examples. Consider, first, the question whether digital computers have memories. It is doubtful that we can even specify a single information-processing task performed by both human memory and computer 'memory' – for one thing, human memory is designed to solve the problem of deciding what to forget, while computer memory is not. But even if there is sufficient commonality at the

computational level, it is clear that there are no significant similarities between human memory and computer memory at the algorithmic or implementational levels. Thus no natural tri-level description (no matter how general) of a kind including both human memory and computer memory can be given. Given the proposed criterion, then, human memory and computer memory are not subkinds of a single natural kind. And this is the result that is wanted, for a 'kind' that includes both human memory and computer memory does not support inductive generalizations: attempt, for example, to predict human memory distortions on the basis of computer memory 'distortions', and you will most often fail.¹⁵

Consider, second, the contrast between episodic and semantic memory. There are, as noted earlier, differences between the episodic and semantic systems, especially at the computational level; in particular, only episodic memory involves autoegetic awareness. But it is nevertheless possible to give a natural general description of an information-processing problem solved by both systems; both are involved (roughly) in making information acquired in the past available again for present use, in manageable quantities and in a timely manner.¹⁶ Since the similarities at the algorithmic and implementational levels are profound, a coherent tri-level description of a kind including both episodic and semantic memory presumably can be given (though it is not feasible to give precise algorithmic- and implementational-level descriptions here, since these will involve detailed descriptions of the workings of encoding, consolidation, storage and retrieval, as well as detailed descriptions of the brain-level mechanisms producing these processes). Hence given the proposed criterion, episodic and semantic memory are subkinds of a single natural kind, i.e. declarative memory. And this, again, is the desired result, for the kind declarative memory does support inductive generalizations: episodic and semantic memory are, for example, affected similarly in amnesia.

Avoiding overgeneration. We can, of course, always formulate a more specific or a more general description of a system; the question, then, is how we can know which descriptions of memory systems to take seriously. There are two dangers here. First, there is the danger that we will be forced to acknowledge an endless proliferation of kinds of memory because we are forced to acknowledge an endless proliferation of memory systems – this, because we can always generate a narrower description of a system. We do not want to be forced to acknowledge the existence of a kind of memory devoted, for instance, to 'early morning episodic memories', simply because we can describe such a system by refining the computational description of the episodic system. Second, there is the danger that the claim that memory is a natural kind will turn into a triviality because we can easily give a description of a memory system so highly general that the various memory systems of the brain will all automatically satisfy it.

I deal with the second worry below. A problem related to the first worry has been dealt with convincingly by Schacter and Tulving. Responding to charges that the multiple memory systems hypothesis licenses researchers to posit an endless number of memory systems, they propose that three criteria should be met before we posit a new memory system. The core idea of the first criterion is that we should not individuate the information-processing problem solved by the system too narrowly; this allows us to avoid positing a new memory system for every type of information processed (Schacter and Tulving, 1994: 15). The second criterion amounts roughly to a statement of the need for a multi-level description of a system (Schacter et al., 2000: 629). The third criterion is that we should be able to observe dissociations between the purported systems on a wide range of tasks (since it is relatively easy to produce a dissociation on some task or other) (Schacter and Tulving, 1994: 17) (Schacter et al., 2000: 629). Only when these three criteria are met, they argue, are we in a position to say that a hypothetical memory system has biological reality.

Schacter and Tulving's criteria are not directly useful here, for we are interested in criteria for positing kinds of memory system, abstract entities which group together biologically real memory systems. But their first criterion provides a clue: classes of information-processing problems can be more or less natural, more or less arbitrary; the suggestion seems to be that we should be concerned with natural classes of information-processing problems. Similarly, I suggest, we should be concerned with natural classes of procedures and natural classes of implementations. The notion of naturalness at work here will have to be cashed out eventually. But assuming that some suitable notion is available – and philosophers have developed a range of accounts of naturalness – then we can avoid having to posit endlessly many kinds of memory system: we are to take seriously only the natural tri-level descriptions.

The reliance on natural classes of information-processing problems, in particular, raises additional worries, though not, I think, fatal ones. For if we adopt a more restrictive notion of naturalness, the tri-level approach will rule out certain higher or more general kinds of memory, including declarative memory, which, I argued earlier, should be counted a natural kind. The worry is that as semantic and episodic memory have different functions (solve different information-processing problems), they will receive different computational-level descriptions, and thus they will come out, on the tri-level approach, simply as distinct kinds.

We can, given the tri-level approach, classify them as subkinds of a common kind by relying on a description of a more general information-processing problem solved by both episodic and semantic memory; but whether such a description is relevant here turns ultimately on the specific account of naturalness that we adopt. Thus it might be preferable, in order to secure the classification of declarative memory as a natural kind, to rely directly on the point (made earlier) that it supports inductive generalizations. This might seem only to push the problem back a step, for we now face the question of how broad a range of generalizations needs to be supported by something before it counts as a natural kind – declarative memory supports a certain range of generalizations, but each of episodic and semantic memory supports a broader range. But however, exactly, we answer this question, the point secures the right sort of fundamental difference between memory as a whole and declarative memory, for while declarative memory supports some range of generalizations, memory as a whole does not. And this is enough to render it plausible that declarative memory, unlike memory as a whole, is a natural kind.

Accommodating multiple realizability. One might object against the tri-level approach that the requirement of a common implementation is overly restrictive, in that, due to implementational-level differences among the episodic memory systems of individual humans, for example, it rules even episodic memory out as a natural kind. But the requirement of a common implementation is not meant to be understood in such restrictive terms: the requirement is for significant similarity at the implementational level, not for perfect similarity. Minor implementational-level differences among individual episodic systems thus do not imply that they are not members of a common natural kind.

One might similarly object that the tri-level approach is incompatible with functionalism (the view that what makes a mental state a mental state of a given type is a matter exclusively of the role that it plays in a larger system [Block, 1980], the dominant view among philosophers of mind on the nature of mental states – more generally: that the tri-level approach is incompatible with multiple realizability. Again, the intention behind the tri-level approach is not to require perfect similarity at the algorithmic and implementational levels, but rather to require only significant similarity; and this should be sufficient to allow for multiple realizability.

But still one might insist that functionalism at least suggests that only computational-level similarity between two kinds of memory is required for them to qualify as subkinds of a single

natural kind. As noted earlier, the tri-level approach is not the only possible approach to the nature of kinds of memory; less restrictive approaches are certainly available, including the proposed computational-similarity-only approach. But note that if my argument below that nondeclarative memory is non-cognitive succeeds, then my conclusion that memory is not a natural kind follows even given the less restrictive computational-similarity-only approach.

The tri-level approach and nondeclarative memory

Nondeclarative memory as non-cognitive

We have seen that the tri-level approach implies that semantic and episodic memory are subkinds of a single kind (declarative memory); what of declarative and nondeclarative memory? As Schacter and Tulving point out, procedural memory (under which heading they include most¹⁷ of nondeclarative memory):

... is involved in learning various kinds of behaviors and cognitive skills and algorithms, its productions have no truth values, it does not store representations of external states of the world, it operates at an automatic rather than consciously controlled level, its output is noncognitive, and it can operate independently of the hippocampal structures. (Schacter and Tulving, 1994: 26)

The contrast with the other memory systems is vivid: these systems:

... are concerned with cognition. That is, the final productions of all these systems can be, and frequently are, contemplated by the individual introspectively, in conscious awareness. Any conversion of such a product of memory into overt behavior, even symbolic behavior such as speech or writing, represents an optional postretrieval phenomenon, characterized by considerable flexibility regarding the behavioral expression. Such flexibility is absent in nondeclarative forms of memory. (Schacter and Tulving, 1994: 27)

Thus there are radical differences between nondeclarative and declarative memory at the computational and the implementational levels: declarative memory (but not nondeclarative memory) is concerned with propositional knowledge; and declarative memory (but not nondeclarative memory) is controlled by the hippocampus.

There are perhaps also deep dissimilarities between them at the algorithmic level. Though nondeclarative memory researchers often write as if nondeclarative memory involves processes of encoding, storage, etc., it is not clear that this talk should be taken as more than metaphorical. Tulving argues that 'Trying to bring [the concepts of encoding, storage, and retrieval of information] into play in the study of, say, behavioral skills such as reading or writing is awkward at best and silly at worst' (2000: 38). If my argument below is right, then nondeclarative memory indeed cannot involve encoding, etc. in anything like the sense in which declarative memory involves encoding.¹⁸

The depth of the differences between declarative and nondeclarative memory should not be underestimated. Consider in particular the computational question: do nondeclarative memory and declarative memory perform a common information-processing task? What is required here is a natural general description of an information-processing problem solved in different ways by the nondeclarative and declarative memory systems. The problem, to put it bluntly, is simply that we need not mention information-processing at all in order to give a complete description of nondeclarative memory – nondeclarative remembering is in an important sense not a kind of cognition. We need not, in particular, in order to give an adequate description of the workings of nondeclarative memory, say anything about even the storage of information by a nondeclarative memory

system, whereas we do need, in order to give a description of the workings of declarative memory, to say something about the storage of information: nondeclarative memory does involve the modification of the brain of the organism on the basis of its experience; but unlike declarative memory, it does not involve the modification of the brain of the organism as a means of making information acquired in the past available again to the organism in the future.^{19,20}

The point is a fairly straightforward application of Occam's Razor. There is no need to posit the storage of information by nondeclarative memory, for the simple reason that any such stored information would be causally, and hence explanatorily, inert. Information stored by declarative memory is not inert: it is processed in various ways and thereby affects the behaviour of the organism; were we to suppose that declarative memory does not store information, we would be unable to explain certain modifications in the behaviour of the organism.²¹ But no such work is done by any information supposedly stored by nondeclarative memory: explanations of modifications of an organism's behaviour via nondeclarative learning go through just as well if we do not posit the storage of information by nondeclarative memory; this, because any such information would never be processed and therefore would be without effects on the behaviour of the organism. Appeals to changes at the neural level do all the work done by appeals to neural changes combined with appeals to storage of information; we should therefore say that no information is stored.

Objections

Representation in declarative memory. One might object that the appeal to considerations of parsimony cuts both ways: if we can explain modifications of the behaviour of an organism via nondeclarative learning without supposing that nondeclarative memory stores information, it might seem that nothing prevents us from similarly explaining modifications of the behaviour of the organism via declarative learning without supposing that declarative memory stores information. But in fact if we do not suppose that declarative memory involves the reproduction of past representations, there is much behaviour that we will be unable to explain, for the straightforward reason that declarative memories are available to consciousness and thus can affect the activities of other mental systems. For example, I remember that it has been warm every day for the past week; I infer that it is warm today; I therefore leave home without bringing a sweater. An appeal to a stored representation is essential to the explanation of my behaviour in this case, for the memory does not cause the behaviour directly, but only via an intermediate process of reasoning. As Schacter and Tulving point out, the productions of a declarative system can be 'contemplated by the individual introspectively' (1994: 26–7), that is, the subject has access to declarative memory representations. It is this that allows them to enter into the causation of behaviour in various indirect ways, via reasoning, imagination, etc.; if the subject were not to have such access, declarative memory would be unable to provide inputs to other cognitive processes.²²

Rule-based information. One might object that since there are natural ways of describing nondeclarative memory as involving information-processing, we are not strictly forced to characterize it as non-cognitive. Cohen and Squire, for example, speak of procedural learning as if it involves information-processing: the idea seems to be that in procedural learning rules governing action are stored by the system:

Whether a task can or cannot be learned in amnesia seems to depend on the nature of the information and not on the extent of motor involvement demanded by the task. We propose that perceptual-motor and pattern-analyzing skills belong to a class of operations governed by rules or procedures; these operations have information-processing and memory characteristics different from those operations that depend on specific,

declarative, data-based material. . . . This distinction between procedural or rule-based information and declarative or data-based information [is] reminiscent of the classical distinction between 'knowing how' and 'knowing that'. (Cohen and Squire, 1980)²³

So, for instance, they point out that amnesic patients can display eyeblink conditioning.

But note that the reference to 'rule-based' information is redundant, for an appeal to changes at the neural level is sufficient to explain, for example, classical conditioning. (It is noteworthy that even in emotional conditioning, conditioning and episodic memory are dissociable: e.g. fear conditioning can occur without memory for the conditioning events and vice versa [Willingham, 1997: 7].) In general, a system need not represent the rules governing a learned behaviour in order for its behaviour to be governed by them. The assumption of storage of rule-based information does no additional explanatory work.

That this sort of reference to information-processing by nondeclarative memory is redundant is brought out nicely by a passage by Milner et al. They write:

. . . nondeclarative memory storage does not depend on specialized memory neurons or systems of neurons whose only function is to *store* rather than process information. Rather, simple nondeclarative memory storage results from changes in neurons that are themselves *components* of the reflex pathway. The storage of nondeclarative memory is embedded in the neural circuit that produces the behavior. . . . the organization and implementation of nondeclarative memory is different from declarative memory where a whole neural system, the medial temporal-lobe memory system, is needed to ensure the remembrance of things past. (Milner et al., 1998: 454)

The reference to the storage of information here is clearly unnecessary: that the reflex pathway has been modified is a sufficient explanation for the emergence of new behaviour; to say that the reflex pathway stores information adds nothing to the explanation. Thus the fact that there are ways of talking about nondeclarative memory as if it were computational by itself settles nothing. Notoriously, we are free to treat even simple systems such as thermometers as information-processors (Dennett, 1987). But parsimony requires that we give the lowest-level explanations of these systems possible: the lowest-level explanation of a thermometer possible confines itself to the physical level; and we therefore should not take the thermometer to be an information-processor. Similarly, while it is possible and even natural to describe nondeclarative memory as cognitive, parsimony dictates that we refrain from doing so: the lowest-level explanation of nondeclarative memory possible confines itself to the neural level; and we therefore should not take nondeclarative memory to be computational.

Implicit representation. One might suggest that though nondeclarative memory does not involve explicit representation of rule-based information, nevertheless it might involve the sort of implicit representation familiar from the literature on connectionism (Smolensky, 1988). But the connectionist literature provides no support for this suggestion. The question about implicit vs explicit representation is about how systems represent information in general rather than about how they represent rule-based information: a system might represent rule-based information, but it will do so in the same manner as it represents data-based information; in the connectionist case, both types of information will be represented implicitly, while in the classical case, both types of information will be represented explicitly. But connectionist systems do not implicitly represent the rules governing their behaviour: a system represents whatever it represents; and we (from our external perspective) can formulate the rules which govern its operation. The situation is no different than in the classical case: in that case, the system will represent both rule-based and data-based

information; but the rule (formulated by us) describing the role played by the represented rules in the operation of the system is not represented by the system.

Thus the notion of implicit representation does not provide us with a means of saving the suggestion that nondeclarative memory involves the representation of rule-based information: a reflex pathway is modified in a certain way, thus producing a new pattern of behaviour, which we can describe by means of a rule; there is no reason to suppose that the pathway implicitly represents the rule in question.

Knowledge-how

The point that nondeclarative memory is non-cognitive will strike some as surprising and others as obvious. It might appear surprising because it is usual for us to think of memory as having a sort of unity, for us to suppose that nondeclarative memory is involved in information-processing. It might appear obvious because the declarative/nondeclarative distinction is intimately related to the old familiar philosophical distinction between knowing that and knowing how, a distinction that philosophers have found it necessary to draw because they have recognized the existence of a deep gap between these two 'kinds' of knowledge.²⁴ Philosophers are not unanimous on the difference in kind between knowledge how and knowledge that. But the view that there is a difference in kind is certainly the default.

The philosophical debate over the relation of knowledge-how to knowledge-that has largely been conducted either in an a priori manner (e.g. Ryle's [1949] argument that the 'intellectualist legend', according to which performing an action requires contemplating the rule governing the action, leads to a vicious regress) or on the basis of linguistic considerations (e.g. Stanley and Williamson's [2001] defence of the view that knowledge-how is a type of knowledge-that, which turns on the linguistic forms of knowledge-how and knowledge-that attributions). My argument here provides a novel, empirically grounded reason for endorsing the default view: coming to know how to do something is sometimes (though not always) a matter of nondeclarative learning; given that nondeclarative memory is non-cognitive, then, at least some knowledge-how is non-cognitive, and this secures the traditional gap between knowledge-how and knowledge that.

Qualifications

The claim that nondeclarative memory is non-cognitive is subject to an important qualification. Priming (a kind of nondeclarative memory) would seem to be straightforwardly cognitive, since it is a matter of changes in the subject's ability to recognize objects. But priming is precisely the exception that proves the rule. Schacter and Tulving, in their discussion of cognitive vs non-cognitive memory systems, explicitly group a special 'perceptual representation system' supposed to subserve priming with the episodic, semantic and working memory systems: the perceptual representation system 'plays an important role in identifying words and objects, it operates at a presemantic level, and it is typically involved in nonconscious or implicit expressions of memory, such as priming' (Schacter and Tulving, 1994: 28; see also Schacter and Buckner, 1998). If something like this approach is right, then priming should not, after all, be classified with procedural memory, conditioning, etc.²⁵ Working memory (which is normally grouped neither with declarative nor with non-declarative memory) is also cognitive. There is thus an interesting question whether working memory and priming can be classified as subkinds of a kind that includes also declarative memory. Limitations of space prevent me from taking this question up here, though it merits further investigation.




Figure 2. Revised taxonomy of memory systems (a dotted line indicates uncertainty about that aspect of the hierarchy)

There is a further question about whether the various kinds of nondeclarative memory (excluding priming) are subkinds of a single natural kind, that is, whether nondeclarative memory itself is merely a nominal kind. There are suggestions in the literature that it might be: researchers are prone to note that nondeclarative memory is a rather heterogeneous collection of phenomena. But the question is complex, and considerations of space thus again dictate that I leave it for future research.

Memory is not a natural kind

If nondeclarative memory is not involved in information-processing, then the question whether memory is a natural kind must be answered in the negative: memory is not a natural kind, for we cannot give a natural computational-level description adequate to both declarative and nondeclarative memory. Alternative approaches to the question whether memory is a natural kind might generate different results. But we have in any case a powerful reason for a negative answer to the question: that some but not all kinds of memory are cognitive is strong evidence that memory is not a natural kind, whatever specific approach to the question we decide to adopt. Nevertheless, the suggestion that memory is not a natural kind is likely to provoke various objections. The remainder of the article is devoted to sketching responses to those that I anticipate.

Memory and 'memory'

The first objection turns on our habitual use of a single term to refer both to nondeclarative and to declarative memory: humans, the objection runs, have a well-developed capacity for detecting natural kinds (Kornblith, 1993); and our use of the same term to refer to both nondeclarative and declarative memory thus provides at least some evidence that memory is a natural kind.

Two points should be made in response to this objection. First: since ordinary language does mark something like the distinction between nondeclarative and declarative memory – we ordinarily distinguish between ‘remembering how’ and ‘remembering that’ – it is not clear that ‘memory’ in ordinary language is a natural kind term. Second: while we do have a capacity for detecting natural kinds, our capacity is far from infallible; while ordinary language might carry some weight

here, it is far from decisive; if it turns out that ordinary language fails to mark a distinction revealed to us by scientific investigation, so much the worse for ordinary language – ‘memory’ might (to refer to a familiar example) be like ‘jade’.

Implementational-level similarity

The second objection begins by pointing out that although different brain structures are involved in different kinds of memory, there are nevertheless common neural mechanisms (namely, modifications of the strength of existing synaptic connections and growth of new synaptic connections [Kandel and Pittenger, 1999]) employed by the different memory systems; this provides some hope, the objection continues, that deep similarities between nondeclarative and declarative memory can be found at the implementational level at least.

There are three points to be made in response to this objection. The first is that the use by different memory systems of similar mechanisms at the neural level is not particularly telling: as Weiskrantz (1990: 99–100) points out, ‘neurons are neurons and nothing more than neurons. At the level of cellular mechanism there are limited degrees of freedom. It may be, even if there are independent memory systems, that the cellular events used in all are similar or even identical’.²⁶ The second point is simply that even extensive similarity at the implementational level is insufficient (given the approach adopted here) to ground the claim that memory is a natural kind: what is required is tri-level similarity. The third point is that there are extensive functional architectural differences between declarative and nondeclarative memory; thus their degree of similarity at the implementational level should not be overestimated.

Interacting memory systems

The third objection turns on the fact (a prominent theme in recent memory research) that declarative and nondeclarative memory interact in various ways (Kim and Baxter, 2001; McDonald et al., 2004; Poldrack and Packard, 2003; Poldrack et al., 2001) – the fact of this interaction confers a certain degree of plausibility to the idea of a single complex memory system (with relatively independent specialized nondeclarative and declarative subsystems), and hence on the claim that nondeclarative and declarative memory are subkinds of a single kind. It is not just that, as Squire notes, ‘The memory systems of the brain operate in parallel to support behavior’ (2004: 174). It is that many learning and memory tasks seem to draw upon multiple memory systems: for example, the amygdala (which subserves emotional learning) can exert an influence on the consolidation of declarative memories (McGaugh, 2000).

I respond that to suggest on the basis of interactions among memory systems that there is a single memory system (a single kind of memory) including both nondeclarative and declarative memory is a mistake, for the interaction between nondeclarative and declarative memory is likely not of the right sort. The interaction between systems is at least arguably not computational: it does not appear that information is passed from nondeclarative memory to declarative memory or vice versa. Researchers interested in interaction among memory systems are, rather, concerned with competition (and similar relations) among them. Nor, I think, does the existence of competition among memory systems provide more indirect support for the claim that memory is a natural kind. Poldrack and his colleagues do, for example, argue that ‘competition may serve as a mechanism to arbitrate between two fundamentally incompatible requirements of learning: the need for flexibly accessible knowledge (supported by the [medial temporal lobe]) and the need to learn fast, automatic responses in specific situations (supported by the striatum)’ (Poldrack et al., 2001: 549). But

it will not do to claim that nondeclarative and declarative memory serve the common general function of learning and to argue on this basis that nondeclarative and declarative memory are subkinds of a single natural kind. The question whether learning is a natural kind is not distinct from the question whether memory is a natural kind: the evidence that memory is not a natural kind is at the same time evidence that learning is not a natural kind. Hence the claim that both the nondeclarative and the declarative systems are involved in 'learning' can provide no more support for the view that memory is a natural kind than does the claim that both systems are involved in 'memory'.

Reinforcement learning

This response to the third objection suggests a fourth (and final) objection, which turns on the existence of apparently general theories of learning. Reinforcement learning theory (Kaelbling et al., 1996; Sutton and Barto, 1998), in particular, has been applied in the context of multiple memory systems research, and if a single theory of learning applies to both declarative and non-declarative learning, this suggests that memory might after all be a natural kind.

Reinforcement learning is an approach to the artificial intelligence problem of designing agents that can learn from interactions with their environments. Reinforcement learning contrasts with supervised learning, in which agents are explicitly instructed about which actions are correct: in reinforcement learning, the agent receives feedback from the environment, and must explore (i.e. test different policies) to discover optimal actions (Sutton and Barto, 1998: 3–4). Thus though reinforcement learning is effectively a normative theory, it is nevertheless plausible as a descriptive theory of learning. It is, however, crucial in the present context to note that reinforcement learning refers not to a single learning method but rather to a type of learning problem (Sutton and Barto, 1998: 4); significantly different types of learning thus fall under the general heading of reinforcement learning.

Daw et al. (2005) have recently applied reinforcement learning theory to the question of how the brain chooses between competing systems for behavioural choice in cases of disagreement between systems; they are concerned in particular with the dorsolateral striatum, which supports 'habitual or reflexive control', and with the prefrontal cortex, which supports 'more reflective or cognitive action planning' (Daw et al., 2005: 1704). Their account associates model-free learning (a method in which an action is simply associated with a value) with the dorsolateral striatum, while model-based learning (a method which involves constructing a model of the environment) is associated with the prefrontal cortex. In model-free learning (temporal-difference learning, in particular), the system employs caching, i.e. 'the association of an action or situation with a scalar summary of its long-run future value': 'Working with cached values is computationally simple but comes at the cost of inflexibility: the values are divorced from the outcomes themselves and so do not immediately change with the re-valuation of the outcome. This is also the defining behavioral characteristic of habitual control' (Daw et al., 2005: 1705). In model-based learning, in contrast, the system 'constructs predictions of long-run outcomes, not through cached storage, but rather on the fly, by chaining together short-term predictions about the immediate consequences of each action in a sequence' (tree search): 'Search in deep trees can be expensive in terms of memory and time and can also be error-prone. However, that the predictions are constructed on the fly allow them to react more nimbly to changed circumstances, as when outcomes are re-valued. This . . . is the behavioral hallmark of cognitive (or "goal-directed:") control' (Daw et al., 2005: 1705).

While Daw et al. do indeed employ techniques from reinforcement learning theory to describe the behaviour of the two systems, it is crucial here that the techniques applied differ between the two cases: to describe the system supporting reflective or cognitive action planning (the sort of

system at work, for example, when an action associated with a devalued outcome ceases to be performed), they resort to a technique incorporating an explicit model of the environment; while to describe the system supporting habitual or reflexive action control (the sort of system at work, for example, when an action associated with a devalued outcome continues to be performed), they need not invoke such a model (Daw et al., 2005: 1705). Thus it appears that there is no common computational-level description of the two systems: while one system influences behaviour by relying on information about the environment acquired in the past, the other produces behaviour automatically – its selection of an action is influenced by its past experience with its environment, but does not involve a representation of that environment. Moreover, while the reinforcement learning account is given in terms of computations performed by systems, it is unclear whether the dorsolateral striatal system itself need be given a computational description: while an invocation of stored information would seem to be indispensable to describe the flexible workings of the pre-frontal system, a neural-level description would appear to be sufficient to describe the automatic workings of the dorsolateral system. Thus this use of reinforcement learning theory gives us little additional reason to suppose that memory is a natural kind.

I said at the outset that an answer to the question whether memory is a natural kind is not only of theoretical interest, but that an answer will have implications for what sorts of empirical approaches to memory are likely to be fruitful. But one might wonder whether the very empirical literatures discussed in the course of my argument for the answer that memory is not a natural kind themselves belie that claim, since they might seem to suggest that the contrary default answer poses no barrier to fruitful empirical research on memory.

Even if the default view in practice poses no barrier to fruitful research, it might nevertheless be incorrect, so the question retains its theoretical interest in any case. But I suggest that the default view does do some damage in practice, and that it does less than it might largely because it is already to a certain extent disregarded in practice. On the one hand, I suggest, researchers often disregard the view in practice, focusing on specific memory systems and refraining from posing overarching accounts of memory as a whole; this tends to neutralize the effects of the default view, preventing it from leading researchers on a fruitless quest for a theory of memory as a whole. But on the other hand, the default view is not without consequences, for the accounts of specific kinds of memory that are produced are often presented as if they apply to memory more broadly; this leads to confusion over the scope and implications of accounts of kinds of memory. As noted in the first section, for example, Anderson sometimes advertises his rational analysis as applying simply to memory, whereas the account seems applicable rather to declarative memory (and priming). At minimum, then, the default view tends to generate confusion over the scope of accounts that are best understood as applying to specific kinds of memory; acceptance of the view that memory is not a natural kind should decrease this sort of confusion.

Acknowledgements

Thanks to Kristoffer Ahlström, Louise Antony, Charles Clifton, Jeremy Cushing, Jeff Dunn, Hilary Kornblith, Jonathan Schaffer, John Sutton and three anonymous referees for comments on earlier drafts of this article. The preparation of the article was supported in part by the Agence Nationale de la Recherche, under the contract ANR-08.BLAN-0205-01.

Notes

- 1 I drop the scare-quotes in what follows.
- 2 I will not take up the distinct question whether ‘memory’ is a natural kind term.
- 3 My answer to the question also has implications for epistemology; for a discussion of these, see Michaelian (2010a).

- 4 For discussion of the philosophical implications of the theory, see Michaelian (2010b).
- 5 Anderson does apply the theory to priming; but see my remarks on priming in a later section.
- 6 For relevant history and reviews of evidence, see Gabrieli (1998), Kandel and Pittenger (1999), Schacter and Tulving (1994) and Milner et al. (1998).
- 7 I largely disregard working memory, which seems to be a basically distinct phenomenon (Baddeley, 2003; Baddeley and Hitch, 1974; Repovš and Baddeley, 2006), in what follows.
- 8 Which does not mean that it is not a thoroughly constructive process (Schacter and Addis, 2007; Schacter et al., 1998).
- 9 Nor is there room here for the operation of metacognitive processes analogous to reality monitoring or source monitoring (Johnson and Raye, 1981; Johnson et al., 1993).
- 10 The species example is only intended to be illustrative of the relative difficulty of the question whether memory is a natural kind: biologists and philosophers of biology who hold that species are not kinds (but rather sets or individuals), or that there are no species, or who are simply worried about the availability of multiple accounts of the nature of the kind species can make the same point using, for example, the less problematic case of chemical kinds.
- 11 The tri-level hypothesis obviously makes sense only given the prior assumption that cognition is computation or information-processing; indeed, I use 'computation' and 'information-processing' interchangeably with 'cognition'. Note that this assumption is compatible with connectionism: for the connectionist, cognition is after all still computation, albeit not classical computation. The assumption does, however, rule out anti-representationist forms of dynamicism (Eliasmith, 1996). Since such views remains a small minority tendency within cognitive science, this does not count heavily against my use of the assumption.
- 12 Note that I take the implementational level to include both the brain structures involved in the system (e.g. hippocampal structures) and the cell- (and lower-) level mechanisms used by the system (e.g. changes in synaptic strength).
- 13 Though the tri-level view is well-supported, other, more restrictive understandings of the concept of a memory system have been proposed. In a relatively early paper (1984), Tulving suggested that we distinguish memory systems not only in terms of their computational, algorithmic and implementational properties, but also in terms of 'differences in ontogenetic and phylogenetic development' (Schacter et al., 2000, 628). It is not clear what would be gained by moving to this more restrictive understanding of the concept of a system, for even on the tri-level approach we can grant that considerations of origins can shed light on the nature of given memory systems. For another early discussion of the concept of a memory system, see Tulving (1985).
- 14 These questions (especially the question about commonality at the algorithmic level) might be made easier to answer by relativizing the discussion to a common architecture; I believe that the gain in generality secured by not so relativizing outweighs the cost of the additional difficulty that it entails.
- 15 This is not to say that machines could not have memories but only that current digital computers do not.
- 16 For an extended discussion of the function of declarative memory, see Michaelian (2010b).
- 17 See the discussion of priming later.
- 18 Tulving's choice of examples might mislead: reading would seem to involve encoding, storage, etc., for a subject will normally be able to recall what she or he has read and to make use of the recalled information. But Tulving's point here concerns not memory for what is read but rather the skill of reading: one learns to read, but acquiring that skill is not a matter of encoding and storing information about how to read.
- 19 I discuss exceptions to this claim later.
- 20 If nondeclarative remembering is not information-processing, then the claim that the concept of a memory system is a concept of a system which requires a tri-level description becomes slightly misleading, for it now appears that there is a crucial ambiguity in the concept of a memory system. We can of course speak sensibly of nondeclarative systems. It is just that these are not memory systems in the sense in which, for example, the episodic system is a memory system.
- 21 I take no stand here on whether the fact of causation by stored content requires that we accept the existence of narrow content (Fodor, 1987).
- 22 Arguably, what ultimately matters here is accessibility to a wide range of cognitive processes, rather than conscious accessibility as such. And granted this, one might wonder whether it is necessary to suppose

that a structure must store information in order to be able to influence a wide range of processes – the worry, again, will be that a neural-level story will suffice to explain the effects of declarative memory, just as such a story suffices to explain the effects of non-declarative memory. But while a structure need not in general store information in order to be able to influence a wide range of processes, certain sorts of influence – including that at issue in the example given above – appear to be explicable only if we suppose that the relevant structure stores information.

- 23 I return to the relation between the nondeclarative/declarative distinction and the knowing-how/knowing-that distinction later.
- 24 The cognitive psychologist Jerome Bruner used the suggestive terms ‘memory with record’ and ‘memory without record’ to draw a similar distinction (Milner et al., 1998: 449).
- 25 Perhaps some additional exceptions should be acknowledged here: Bayley and Squire (2002) argue that very limited factual learning is possible even without declarative memory; the suggestion is that this is similar to perceptual learning in that it is a case of learning directly by neocortex. Sensory memory might be an additional exception.
- 26 Though it remains unclear just how much commonality there is at the neural level, considerable progress on the question has been made; see Kandel and Pittenger (1999) and Milner et al. (1998).

References

- Anderson, J.R. (1990) *The Adaptive Character of Thought*. Hillsdale, NJ: Lawrence Erlbaum.
- Baddeley, A. (2003) ‘Working Memory: Looking Back and Looking Forward’, *Nature Reviews Neuroscience* 4: 829–39.
- Baddeley, A. and G.J. Hitch (1974) ‘Working Memory’, *Recent Advances in Learning and Motivation* 8: 47–90.
- Bayley, P.J. and L.R. Squire (2002) ‘Medial Temporal Lobe Amnesia: Gradual Acquisition of Factual Information by Nondeclarative Memory’, *Journal of Neuroscience* 22: 5741–8.
- Bernecker, S. (2008) *The Metaphysics of Memory*. New York: Springer.
- Block, N. (1980) ‘Introduction: What is Functionalism?’, in N. Block (ed.) *Readings in Philosophy of Psychology*, pp. 171–84. Cambridge, MA: Harvard University Press.
- Chater, N. and M. Oaksford (1999) ‘Ten Years of the Rational Analysis of Cognition’, *Trends in Cognitive Sciences* 3: 57–65.
- Cohen, N. and L. Squire (1980) ‘Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociations of Knowing How and Knowing That’, *Science* 210: 207–10.
- Daw, N.D., Y. Niv and P. Dayan (2005) ‘Uncertainty-Based Competition between Prefrontal and Dorsolateral Striatal Systems for Behavioral Control’, *Nature Neuroscience* 8: 1704–11.
- Dawson, M.R.W. (1998) *Understanding Cognitive Science*. Oxford: Blackwell.
- Dennett, D. (1987) *The Intentional Stance*. Cambridge, MA: MIT Press.
- Eliasmith, C. (1996) ‘The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition’, *Journal of Philosophical Psychology* 9: 441–63.
- Fodor, J. (1987) *Psychosemantics*. Cambridge, MA: MIT Press.
- Foster, J.K. and M. Jelicic (eds) (1999) *Memory: Systems, Process, or Function*. Oxford: Oxford University Press.
- Gabrieli, J.D.E. (1998) ‘Cognitive Neuroscience of Human Memory’, *Annual Review of Psychology* 49: 87–115.
- Johnson, M.K. and C.L. Raye (1981) ‘Reality Monitoring’, *Psychological Review* 88: 67–85.
- Johnson, M.K., S. Hashtroudi and D.S. Lindsay (1993) ‘Source Monitoring’, *Psychological Bulletin* 114: 3–28.
- Kaelbling, L.P., M.L. Littman and A.W. Moore (1996) ‘Reinforcement Learning: A Survey’, *Journal of Artificial Intelligence Research* 4: 237–85.
- Kandel, E.R. and C. Pittenger (1999) ‘The Past, the Future and the Biology of Memory Storage’, *Philosophical Transactions of the Royal Society of London B* 354: 2027–52.
- Kim, J.J. and M.G. Baxter (2001) ‘Multiple Brain-Memory Systems: The Whole Does not Equal the Sum of its Parts’, *Trends in Neurosciences* 24: 324–30.

- Kornblith, H. (1993) *Inductive Inference and its Natural Ground*. Cambridge, MA: MIT Press.
- McDonald, R.J., B.D. Devan and N.S. Hong (2004) 'Multiple Memory Systems: The Power of Interactions', *Neurobiology of Learning and Memory* 82: 333–46.
- McGaugh, J.L. (2000) 'Memory: A Century of Consolidation', *Science* 287: 248–51.
- Marr, D. (1982) *Vision*. San Francisco: W.H. Freeman.
- Michaelian, K. (2008) 'Testimony as a Natural Kind', *Episteme* 5: 180–202.
- Michaelian, K. (2010a) 'Generative Memory', *Philosophical Psychology*, forthcoming.
- Michaelian, K. (2010b) 'The Epistemology of Forgetting', *Erkenntnis*, forthcoming.
- Milner, B., L.R. Squire and E.R. Kandel (1998) 'Cognitive Neuroscience and the Study of Memory', *Neuron* 20: 445–68.
- Poldrack, R.A. and M.G. Packard (2003) 'Competition among Multiple Memory Systems: Converging Evidence from Animal and Human Brain Studies', *Neuropsychologia* 41: 245–51.
- Poldrack, R.A., J. Clark, E.J. Paré-Blagoev, D. Shohamy, J. Creco Moyano, C. Myers and M.A. Gluck (2001) 'Interactive Memory Systems in the Human Brain', *Nature* 414: 546–50.
- Pylshyn, Z.W. (1984) *Computation and Cognition*. Cambridge, MA: MIT Press.
- Quine, W.V.O. (1969) 'Natural Kinds', in *Ontological Relativity and Other Essays*, pp. 114–38. New York: Columbia University Press.
- Roediger, H.L., R.L. Buckner and K.B. McDermott (1999) 'Components of Processing', in J.K. Foster and M. Jelicic (eds) *Memory: Systems, Process, or Function*, pp. 31–65. Oxford: Oxford University Press.
- Repoš, G. and A. Baddeley (2006) 'The Multi-Component Model of Working Memory: Explorations in Experimental Cognitive Psychology', *Neuroscience* 139: 5–21.
- Ryle, G. (1949) *The Concept of Mind*. Chicago, IL: University Press.
- Schacter, D.L. and D.R. Addis (2007) 'The Cognitive Neuroscience of Constructive Memory: Remembering the Past and Imagining the Future', *Philosophical Transactions of the Royal Society of London B* 362: 773–86.
- Schacter, D.L. and R.L. Buckner (1998) 'Priming and the Brain', *Neuron* 20: 185–95.
- Schacter, D.L. and E. Tulving (1994) 'What Are the Memory Systems of 1994?', in D.L. Schacter and E. Tulving (eds) *Memory Systems*, pp. 1–38. Cambridge, MA: MIT Press.
- Schacter, D.L., K.A. Norman and W. Koutstaal (1998) 'The Cognitive Neuroscience of Constructive Memory', *Annual Review of Psychology* 49: 289–318.
- Schacter, D.L., A.D. Wagner and R.L. Buckner (2000) 'Memory Systems of 1999', in E. Tulving and F.I.M. Craik (eds) *Oxford Handbook of Memory*, pp. 627–43. Oxford: Oxford University Press.
- Schooler, L.J. and J.R. Anderson (1997) 'The Role of Process in the Rational Analysis of Memory', *Cognitive Psychology* 32: 219–50.
- Smith Churchland, P. (1986) *Neurophilosophy: Toward a Unified Science of the Mind/Brain*. Cambridge, MA: MIT Press.
- Smolensky, P. (1988) 'On the Proper Treatment of Connectionism', *Behavioral and Brain Sciences* 11: 1–74.
- Squire, L.R. (2004) 'Memory Systems of the Brain: A Brief History and Current Perspective', *Neurobiology of Learning and Memory* 82: 171–7.
- Squire, L.R. and S. Zola-Morgan (1998) 'Episodic Memory, Semantic Memory, and Amnesia', *Hippocampus* 8: 205–11.
- Stanley, J. and T. Williamson (2001) 'Knowing How', *Journal of Philosophy* 98: 411–44.
- Sutton, R.S. and A.G. Barto (1998) *Reinforcement Learning*. Cambridge, MA: MIT Press.
- Tulving, E. (1984) 'Multiple Learning and Memory Systems', in K.M.J. Lagerspetz and P. Niemi (eds) *Psychology in the 1990s*, pp. 163–84. Amsterdam: Elsevier.
- Tulving, E. (1985) 'How Many Memory Systems Are There?', *American Psychologist* 40: 385–98.
- Tulving, E. (2000) 'Concepts of Memory', in E. Tulving and F.I.M. Craik (eds) *Oxford Handbook of Memory*, pp. 33–43. Oxford: Oxford University Press.
- Tulving, E. (2002) 'Episodic Memory: From Mind to Brain', *Annual Review of Psychology* 53: 1–25.

- Weiskrantz, L. (1990) 'Problems of Learning and Memory: One or Multiple Memory Systems?', *Philosophical Transactions of the Royal Society of London B* 329: 99–108.
- Wiggs, C.L., J. Weisberg and A. Martin (1999) 'Neural Correlates of Semantic and Episodic Memory Retrieval', *Neuropsychologia* 37: 103–18.
- Willingham, D.B. (1997) 'Systems of Memory in the Human Brain', *Neuron* 18: 5–8.

Author biography

Kourken Michaelian is currently a postdoc on a project on epistemology and metacognition at the Institut Jean-Nicod (CNRS/EHESS/ENS). He recently completed his PhD in philosophy at the University of Massachusetts (Amherst), with a thesis on the epistemologies of memory and testimony. His recent articles include 'Generative Memory', in *Philosophical Psychology*, and 'The Epistemology of Forgetting', in *Erkenntnis*.