Philosophical Psychology

ISSN: 0951-5089 (Print) 1465-394X (Online) Journal homepage: www.tandfonline.com/journals/cphp20

Varieties of memory, varieties of reconstruction, varieties of memory trace

Simon Alexander Burns Brown

To cite this article: Simon Alexander Burns Brown (16 Oct 2025): Varieties of memory, varieties of reconstruction, varieties of memory trace, Philosophical Psychology, DOI: 10.1080/09515089.2025.2574486

To link to this article: https://doi.org/10.1080/09515089.2025.2574486

Published online: 16 Oct 2025.

PHILOSOPHICAL PSYCHOLOGY
https://doi.org/10.1080/09515089.2025.2574486

Varieties of memory, varieties of reconstruction, varieties of memory trace

Simon Alexander Burns Brown

Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, London, UK

ABSTRACT

The biological world is rich in variation, both in bodies and in minds. A particularly clear case is memory, where traditional taxonomies increasingly face challenges capturing the full extent of variation. Meanwhile, a central debate within philosophy of memory has focused on whether episodic memory requires memory traces, given the role of simulation in episodic remembering. Do other forms of memory involve traces and simulation in the same way as episodic memory, and if they do, does this undermine episodic memory's claim to be a natural kind? Trace minimalist approaches see the role of simulation in memory retrieval as a means of extracting a rich, reliable yet fallible representation from a trace storing compressed information. This insight can be generalised to a broader class of theories, Minimal Trace Minimalism. Within this broader class, 'Trace Pluralist' theories suggest, surprisingly, that the basic computation of compressed storage and potentially fallible reconstruction shows up not just in episodic memory, but equally in paradigmatic cases of semantic memory. This motivates replacing the episodic/semantic distinction with a richer, more systematic array of categories, and reevaluating our understanding of semantic memory and related notions such as belief.

ARTICLE HISTORY
Received 31 October 2024
Accepted 8 October 2025

KEYWORDS
Episodic memory; semantic memory; compression; trace minimalism; simulationism; memory traces

1. Introduction

This paper will suggest new ways of thinking about kinds of memory. It will suggest that semantic memory is likely to involve reconstruction in much the same sense – and for much the same reasons – as episodic memory. It will suggest that neither episodic nor semantic memory need constitute well-defined, clean kinds, underpinned by distinctive mechanisms. Rather, each is likely to come in many overlapping varieties, across different cases within individuals but especially across species.

CONTACT Simon Alexander Burns Brown simonabbrown@gmail.com The London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

Such a conclusion may be exasperating to a certain kind of theorist, who expects science to traffic in well-defined categories. However, before turning to more detailed arguments favoring such views, one can start to develop more favorable intuitions by turning to memories of school science classes.

In secondary school, I preferred physics to biology. This was partly thanks to idiosyncratic features of the classes, but there was also an element of teenaged laziness involved. I did not like biology's making me memorize all those *facts*—the parts of a cell, the classes and orders and families. I liked being able to memorize a handful of equations and derive everything else in the exam.

These differences in what is taught in biology and physics classes are not purely accidental. Deep differences in subject matter support these different approaches. Biology is characterized by variety and by exceptions. The best candidates for laws in biology relate to evolution by natural selection; and natural selection *implies* variety – there needs to be variation to be selected on, and intermediate forms to evolve through. As Murray et al. (2017, p. 14) put it, "natural selection produces advantages, not parsimony." Understanding the biological requires allowing for imperfect patterns, abundant fuzzy-boundaried kinds, and special cases.

All this is doubly true for cognition, characterized by even more complexity and variation than other biological capacities, thanks to learning (hence determination by one's unique series of experiences), flexibility, and creativity. Neat, definitive taxonomies are *possible* here, but should not be expected: adequate taxonomies are *likely* to need many categories, to be riven with exceptions and intermediate cases; and to sit alongside equally useful cross-cutting taxonomies, especially when including other species.

Psychology nonetheless often posits clean, simple taxonomies. Especially in memory science, dichotomies have ruled: long-term/short-term; declarative/procedural; episodic/semantic. Yet any categorization of memory systems must reckon with variation across species. We face the *Problem of Interspecific Variation* (Boyle & Brown, 2025; cf. related discussion in; Boyle, 2022, 2023; Keven, 2022; Malanowski, 2016; Pan, 2022): on any way of characterizing episodic memory, there are almost certainly species with systems, processes or states with some of the relevant features but not others. Under what conditions should we categorize these as genuine episodic memory, or as merely *episodic-like*? Closely related is the *Problem of Alternatives* (Boyle & Brown, 2025): insofar as other species might have episodic-*like* memory which is not episodic memory proper, should this be thought of as semantic memory, associations, or some other category? And how should *these* types be individuated across species in light of variation? Murray et al. (2017) identify no fewer than seven memory systems operating just in humans, each with its own distinctive evolutionary history, without even attempting to search for systems unique to non-human species. Others have rethought the episodic-semantic distinction even in the human case (Addis & Szpunar, 2024; Aronowitz, 2022; De Brigard et al., 2022; Henke, 2010; Rubin & Umanath, 2015; Sherman et al., 2024).

Much of philosophy of memory has focused on a somewhat independent question: it has taken for granted episodic memory's distinctness from other forms of memory, and asked what if anything sets it apart from *imagination*. Roughly, *causalists* claim that episodic memory involves a memory trace, while *simulationists* claim that episodic memory can occur without a memory trace, and is therefore not fundamentally distinct from imagination. *Trace Minimalism* (TM) offers an intermediate position, on which simulation is indeed central to episodic memory, but a distinctive kind of simulation that draws on memory traces.[1]

As Michaelian and Robins (2018, p. 14) aptly characterize the debate since the classic presentation of causalism in Martin & Deutscher (1966), "Like most subsequent causal theorists, Martin and Deutscher focus on episodic memory, memory for past events." This "focus on episodic memory" can take two forms: explicitly offering a theory designed to only apply to episodic memory; or implicitly assuming that one's theory will generalize to memory in general. Martin and Deutscher explicitly claim that their account generalizes with a few modifications, although they do little to substantiate this. The subsequent literature has often had less to say on non-episodic memory.

But if episodic memory as a category is likely to be blurred, we should ask about the role of memory traces and simulation in other forms of memory. Attempting to generalize accounts of the role of imagination in memory can give rise to new ways of thinking about memory traces and their role in all forms of memory. And these ways of thinking about memory traces suggest ways of approaching the question of the taxonomy of memory.

I will illustrate this through showing how different views within one family of approaches to the issue of memory traces and simulation – which I call *Minimal Trace Minimalism* (MTM) – suggest very different accounts of semantic memory and its relationship to episodic memory. All versions of MTM claim that episodic memory involves both traces and simulation, because they interact in a specific way to gain distinctive advantages for cognition. Traces *compress* large amounts of information. Given that simulation can be used to reconstruct full representations on the basis of these compressed traces in a reliable way, this gives subjects access to sufficiently accurate representations of a much wider array of events than would otherwise be possible. MTM constitutes a plausible, distinctive position in the debate on traces and simulation. However, it is consistent with a wide range of views about the reconstruction processes and traces involved. On some of these, such as Werning and Cheng's *Trace Minimalism* (TM), episodic memory involves a highly distinctive kind of trace and of reconstruction. However, other views are plausible, including *Trace Pluralism*, on which multiple kinds of trace and of reconstruction are involved in different cases of episodic remembering. Thinking through Trace Pluralism will point to a further radical possibility: even paradigmatic cases of semantic memory may involve the same kinds of simulation, reconstruction, and memory traces as episodic memory.

I begin by summarizing TM's key insight concerning the role for compression in memory (Sec. 2). Sec. 3 shows how this insight can be generalized beyond the specific commitments of TM, articulating MTM and Trace Pluralism and providing reasons to think the latter is plausible. Sec. 4 shows we can fruitfully apply a very similar framework to semantic memory, with major implications for our understanding of even paradigm cases of semantic memory. Sec. 5 defends the idea that allowing for reconstruction in semantic memory does not illegitimately confuse reconstruction with inference. Sec. 6 considers the upshot for our taxonomy of memory: while we could distinguish kinds of memory on the basis of distinct varieties of compression and reconstruction, the result would be rather different to the standard episodic-semantic distinction.

2. Trace minimalism

TM has been developed by Cheng, Werning, and collaborators (Cheng & Werning, 2016; Cheng, Werning, & Suddendorf, 2016; Fayyaz et al., 2022; Werning, 2020). Its core idea is that episodic remembering involves simulation and memory traces collaborating in a specific way, inspired by a prominent account of perception. Predictive processing accounts see perception as making uncertain, error-prone, but generally reliable predictions of what is going on around us on the basis of a general model of the world and sparse sensory inputs. TM sees episodic memory as making uncertain, error-prone, but generally reliable predictions about particular past events on the basis of a general model of the world and sparse memory traces. The brain thereby produces a representation of a scenario, a coherent temporal sequence of events, generally including spatial layouts and a setting, participants, and their interactions.[2] Memory traces do not specify all this; instead they specify a limited portion or gist of the pattern of activity that does represent a whole scenario. Producing a full representation of the past event involves filling out the rest of this pattern via computations which, in effect, estimate the most likely properties of an event to have given rise to this memory trace, given the brain's general model of the environment. That is, memory traces are *compressed* representations of scenarios, which can only be turned into full representations of scenarios through reconstruction.

In general, compression involves transforming a representation into a representation (or, more precisely, information-carrying entity) which uses fewer bits. In the context of the brain compressing rich sensory-cognitive-emotional representations of scenarios, the relevant kind of compression will generally be *lossy* compression, where some information is stripped out entirely. A good lossy compression algorithm will strip away only irrelevant aspects of the initial representation, and especially *redundant* information, which can be reliably reconstructed from the information that remains in the compressed representation.

This reconstruction process is a key reason why remembering can be inaccurate: the regularities used to bridge the gap between trace and full representation might have changed since encoding; precise details are not fully determined by trace and regularities in any case, with multiple versions of the event deemed tolerably accurate; and on top of this, depending on how the trace is stored it too may have changed. Yet remembering will be broadly reliable: the relevant regularities will not have changed *that* much, attention will have ensured that the most important aspects of the scenario were encoded in the memory trace, and so on. The perception analogy is again helpful: it too is broadly reliable, especially in naturalistic scenarios and about aspects of the scene most likely to be important (Werning, 2020, p. 316).

A key attraction of TM is its promise of reconciling the attractive elements of both causalism and simulationism, while avoiding their pitfalls.

Causalism accounts for episodic memory's reliability, by positing traces which retain information from past experiences. But it struggles to account for various phenomena which are naturally predicted by trace minimalism. Werning and Cheng give several, including "epistemic generativity" (Werning, 2020), and evidence of simulation, especially distortion at retrieval, which undermines simple forms of Causalism according to which the entire content of an instance of remembering derives directly from a trace. But a particularly interesting argument for TM over Causalism is that the latter is *information-theoretically implausible*:

Insofar as important information about the scenario can be reliably reconstructed on the basis of a much more minimal trace, traces themselves do not need to store more information than is necessary for that reconstruction. Thus, lossy compression can enable a greater amount of information to be available later, to a tolerable degree of accuracy, for a given storage capacity. Werning (2020, p. 217) claims that remembering as much as we in fact can would exceed the informational capacity of our brains if we did not use compression; at the very least, doing so would be expensive evolutionarily.[3] The general idea that compression is important to memory traces back to Barlow (1961/2012), via models of hippocampal function such as those of, Teyler & DiScenna (1986), Treves & Rolls (1994), and McClelland et al. (1995), and is endorsed explicitly by many contemporary views of episodic memory (e.g., Kerrén et al., 2025; Nagy et al., 2025). Hints of the view can even be found in historical philosophy (De Brigard 2014, p. 405).

Like TM, Simulationism also posits more reasonable amounts of storage than Causalism, and can easily explain distortion. However, Simulationism faces a challenge which TM can avoid. Insofar as Simulationism rejects the necessity of a causal link between remembered experience and remembering, and insists that remembering is just a form of imagination, it faces a challenge of accounting for *how* that form of imagination manages to be reliable about the specifics of particular events. Werning (2020) gives an extensive formal argument that reliability of this kind can only be achieved through a causal link to the event in question. Trace Minimalism posits such a link. Relatedly, Cheng (2024) argues that unique features of how a particular event played out are much more likely to be captured by a simulation based on a memory trace, than are unique features predicted of a future event on the basis of merely semantic memory: i.e., episodic memory is likely to be closer to the truth if it uses a trace.

A dedicated simulationist can reply that while there does need to be a causal link carrying information from the original event to reconstruct it reliably, this need not be a trace. In principle, it would be possible to accurately reconstruct some past events on the basis of general semantic information or testimony specifying enough unique features of those events, just as experts can reconstruct specific historical events without having themselves experienced them. However, TM could respond with an argument that relying on such processes alone would be highly atypical for a well-designed system: it would seem highly maladaptive to *not* use memory traces in something like the way posited by TM nearly all of the time. This argument would draw on parallel considerations to the information-theoretical argument given by Werning against Causalism: just as relying entirely on stored information with no reconstruction implies storing excessive amounts of information, relying entirely on reconstruction is likely to require excessive amounts of computation. Traces can be specialized for reconstruction, containing just the information that reconstruction processes will predictably use. Drawing on whatever information happens to be at hand instead of such traces will require much more energy at the reconstruction phase. Think back to the case of learning physics. My preferred strategy was heavy on computation to enable minimal storage: I would only learn a handful of equations and derive the rest in the exam as needed; but this relied on being able to do all of those computations in the exam. Had I instead learned a greater number of equations, with some of them being special cases of the equations I did learn, I could have cut down on the amount of calculating I needed to do in the exam. And for more difficult exams, relying entirely on deriving the answer to every question from first principles would be ill-advised, to say the least. In general, there is a trade-off between amount of storage and amount of computation required for reconstruction, and relying entirely on either without using the other at all will be at best an extreme edge case. The underlying trade-off here provides deep reasons why typical remembering will use both sorts of mechanism, interacting with one another in a distinctive way. Cases not involving both mechanisms will therefore involve a rather different overall mechanism to the normal one.

3. Minimal Trace Minimalism

The previous section detailed how TM offers to reconcile the attractive elements of both causalism and simulationism, and showed how its doing so is not ad hoc, but rather independently motivated by considering how the mind could efficiently make available reliable-enough information about past events. TM as developed by Cheng and Werning involves numerous further commitments than just these, however. Most relate to the idea that episodic memory is underpinned by a particular neural mechanism, and that compression and reconstruction take a certain form supported by that mechanism. Minimal Trace Minimalism (MTM) eschews these further commitments. It is therefore not committed to the idea that episodic memory constitutes a natural kind in any strong sense, even as it adopts the positions described in Sec. 2 that allow for reconciling the attractions of causalism and simulationism.

TM was developed as a view of what would make episodic memory a natural kind (Cheng & Werning, 2016). This idea was tied to thinking that there is a distinctive neural mechanism at play in episodic memory and only in episodic memory: hippocampal replay, a distinctive kind of oscillatory activity in which place cells fire in rapid sequences often corresponding to sequences of locations the animal has previously traversed (Liu et al., 2022; Ólafsdóttir et al., 2018). While such neural events have been linked to various other functions besides episodic memory, such as planning (Ólafsdóttir et al., 2018), Cheng and Werning suggest that this is because these other functions draw on episodic traces (Cheng, 2024, p.7; Cheng and Werning 2016, pp. 1365–1369). The claim that episodic memory is distinctive in involving hippocampal replay drives further commitments both about its neural implementation and psychological distinctiveness. For example, Cheng and Werning insist that episodic memory always represents temporally extended scenarios rather than merely spatially extended scenes. It also contributes to the idea that traces, reconstruction, and compression take a particular form, captured by models developed by Cheng's lab involving artificial neural networks. These models have architectural features inspired by relevant brain areas (Fayyaz et al., 2022; Zeng et al., 2023). A "cortical" part of the model is a convolutional network, with layers organized hierarchically such that patterns of activity in higher layers correspond to more abstract features of the environment. A "hippocampal" part of the network stores memory traces, which act as pointers to portions of activity patterns in the "cortical" network. In these models, reconstruction takes the form of completing these "cortical" patterns. Traces so-understood do not store information compositionally or in the form of explicit representations of external features; they carry a much smaller amount of information, which in the first instance only specifies part of the pattern of "cortical" activity. Given Werning and Cheng's specific understanding of "representation," this implies memory traces are not representations at all.[4]

Minimal Trace Minimalism (MTM) drops most of these commitments. MTM is the commitment only to the theses that episodic memory uses memory traces which carry some information about past events, that these memory traces do not typically fully specify the content represented when the relevant episodic memories are retrieved, and that this is because these memory traces are compressed, with the function of enabling tolerably accurate reconstruction of these events at retrieval. These are the theses that allow TM to capture the distinctive benefits and avoid the costs of both Causalism and Simulationism. Whereas TM takes on further strong commitments about both the neural implementation and psychological-level details of both the (i) reconstruction processes and (ii) traces, MTM remains neutral on these. Thus, TM is a member of the family of views captured by MTM. Another member of that family can be called *Trace Pluralism*. Trace Pluralism adopts MTM's commitments, and takes on the further commitment that there are many different kinds of reconstruction processes and traces. To see why Trace Pluralism ought to be a live option, it will help to consider the many ways that one might go about designing a system to reconstruct a scenario from sparse information.

The pattern completion processes in Cheng's artificial networks are broadly about different layers of those networks filling out patterns of activity in line with regularities in how these layers respond to "perceptual" inputs, where each layer works more or less similarly. Such artificial networks are simplified models of the biological neocortex. The simplifications Cheng makes are reasonable and productive modeling choices for many purposes. Nonetheless, they obscure the possibility of there being multiple varieties of processing supporting distinct varieties of reconstruction. The cortex supports many kinds of computation; indeed, for higher forms of cognition as opposed to perceptual categorization, we have little idea of exactly how they are implemented, so we do not even know *how* such a model distorts them. It will be helpful to sketch some of the potential for variety in reconstruction processes.

One dimension of variation of possible reconstruction processes is generality. We could imagine processes that use very general cognitive processes to elaborate on memory traces, which can also be used for tasks such as for determining whether certain states of affairs are likely under different conditions. At one extreme lie processes for highly general functions, such as performing Bayesian inference. At the other, memory reconstruction heavily relies on processes specialized for reading memory traces and nothing else.

Another dimension might be how perceptual or cognitive the relevant processes are. Exactly how to draw the perception-cognition boundary (indeed, whether there is a robust boundary at all) is controversial (Block, 2023; Clarke & Beck, 2023; Green, 2023; Siegel, 2017; Watzl et al., 2021). However, for illustrative purposes, suppose the following captures two sorts of process at different ends of the spectrum: perceptual systems are honed for constructing a structured, image-like representation of the overall likely layout of specific features (textures, shapes, colors, spatial relations of mid-sized objects, faces, pitch etc.) in one's immediate surroundings, in a fast, automatic and inflexible way driven by sensory input. They do this through embedding (either in their rules of processing or in their formats) a great deal of information about environmental regularities: that certain patterns of light tend to mean a certain 3D shape; that it is likely that those two protruding shapes with an occluding object separating them (but aligned just so) are really parts of one object; and so on. Cognitive systems use less image-like, more abstract or language-like representations, processing them according to quite different rules and often considerably more flexibly but often more slowly. They use more explicit representations of generalities than implicitly encoded regularities in their rules of operation, and can revise these more easily. Their rules of operation are often flexible, as are the tasks they can perform – thus they rely on inference to the best explanation, various kinds of modeling etc.

Both sorts of process could potentially be used for the simulation involved in episodic memory, in rather different ways.

Using perceptual systems offline (i.e., in a non-stimulus-driven, non-immediately-action-guiding way)[5] might allow a simulation process to tap into the environmental regularities they implicitly encode. This could enable the compression of a huge amount of sensory information. For example, for Bob to store the details of how Alice's face looked on a particular occasion, there would be no need to explicitly record, as it were, every pixel. Instead, suppose that Bob's visual face-recognition subsystem has developed a schema for Alice's face, primarily for recognizing

Alice, and for filling out the low-level features when seeing her under poor lighting conditions etc.; and a similar schema for emotional facial expressions. Then all that need to be stored, to be able to later recall a good approximation of Alice's face on a particular occasion (including the space between her eyes, the way her mouth went taught and her eyes widened – down to details that are difficult to capture in language), is the information that it was Alice, and that she looked angry. Additional storage capacity could be used for any event-specific details that do not fit such patterns – Alice's wearing sunglasses, say. In turn, the idea that episodic remembering draws on offline perceptual systems broadly suggests an explanation of why it is often associated with *imagery*—which is often understood in precisely the terms of offline use of perceptual systems – and why these same systems show up neurally in online perception, mental imagery, and episodic memory.[6]

Yet episodic memories do not always focus on the kinds of properties perceptual systems are specialized for. Aphantasics likely have some form of episodic memory without imagery, although it is impaired, reflecting lower connectivity between hippocampus and perceptual areas (Milton et al., 2021; Monzel et al., 2024). Neurotypical humans often have *some* episodic memories with minimal perceptual details, focusing more on internal mental states and feelings (e.g., remembering coming to some realization via a chain of reasoning) or social details of a scenario (the gist of who said what to whom) without including details of what they looked or sounded like. More common still are episodic memories which combine perceptual imagery with both memories of emotions and thoughts and with broader context such as the location of the event and its role in one's overall life, information about how different elements of an event were causally related to one another and to other events one remembers, abstract features of the event in question, and their significance within a particular social and cultural moment. The way Pablo Casals recounts discovering Bach's Cello Suites brings out how such different elements are often bound together in his recollection, and in this respect is plausibly illustrative of how many instances of episodic remembering are experienced:

> One day I told my father I needed especially to find some new solo music ... we stopped at an old music shop near the harbor [*event context including location and his intentions*]. I began browsing through a bundle of musical scores. Suddenly I came upon a sheaf of pages, crumbled and discolored with age [*sensory details*]. They were unaccompanied suites by Johann Sebastian Bach—for the cello only! I looked at them with wonder ... What magic and mystery, I thought, were hidden in those words? [*cultural context and feelings/thoughts*] ... I forgot our reason for being at the shop [non-sensory mental events]. All I could do was stare at the pages and caress them. That scene has never grown dim. Even today ... I am back again in the old musty shop with its faint smell of the sea. [*vivid sensory details*] (Casals & Kahn, 1974, p. 46)

In cases like this, it is highly likely that multiple cognitive and neural systems are working in tandem, processing very different kinds of information in distinctive ways (e.g., with information about Bach, feelings of wonder, and the smell of the shop being handled quite differently, even if they are all also being related to one another in a coherent overall scenario). To the extent that psychological systems use very different kinds of computation and inference, formats, prior probabilities etc. specialized for particular tasks, we can expect simulation (including in episodic memory) to draw on these for compressing and reconstructing different kinds of information. It is unlikely that modeling all these different processes of compression and reconstruction as simple pattern completion will capture the relevant variety at either the psychological or neural levels. Indeed, theoretical treatments of episodic memory in psychology and neuroscience – from Tulving (1983) to Conway & Pleydell-Pearce (2000) to Murray et al. (2017) to Addis & Szpunar (2024) – have typically in one way or another incorporated the assumption that it typically draws on and coordinates multiple different perceptual and cognitive systems to represent different aspects of any given episode.

The reason TM is restrictive here is not that it cannot account for episodic memory including non-perceptual details. Cheng's model explicitly allows that the stored patterns and reconstruction may take place at different levels of the sensory hierarchy, including more or less abstract and detailed contents depending on context and attention, with some episodic memories allowing relatively little access to sensory contents (Cheng, 2024, p. 3). Rather, the issue is that it seems to rule out the possibility posited by Trace Pluralism – which should be a live one – of contrasting mechanisms being involved in remembering.

With variety in the kinds of reconstruction available comes variety in memory traces. Insofar as episodic memory involves compressed memory traces, it involves traces which, when fed to specific systems, will be likely to result in a content within some range of acceptable values. The forms taken by traces will vary with the systems they are to be fed to. For example, traces designed to be read by connectionist pattern completion might simply be pointers to fragments of an activation pattern in a network. Yet traces designed to be read by more classical kinds of computation might have explicit, compositional content about the external world, consisting of explicit summaries of the gist of an event plus ways in which it diverged from that overall gist (or representations warning of risks of specific errors when reconstructing based on the other parts of the trace). Furthermore, it is not obvious they would need to take the specific form of sequences, or be tied to a particular neural implementation in hippocampal place cells – let alone that hippocampal replay always encodes for particular episodic memories.[7]

Advocates of TM can point to evidence that hippocampal traces based around sequences are often implicated in episodic memory. However, this does not establish that memory traces *only* take this form. It could be that some traces take this form (particularly for perceptually based memories) and others don't; and it could be that many cases of episodic memory involve *hybrid* traces, where sequences are bound to further compressed cortical information that is then read in a completely different way by other cortical systems. Indeed, it is plausible, given the plethora of different varieties of information that can be bound into one episodic memory, that such hybrid traces do sometimes occur. Thus, there are reasons to either remain neutral about the form traces might take, both in terms of neural implementation and in terms of their computational properties (as per MTM), or even to lend some credence to Trace Pluralism.

Trace Pluralism, especially, marks an important difference with TM on the issue of whether episodic memory is a natural kind, cleanly mechanistically distinguishable from other forms of memory or information storage (even if they share some components). If various compounds of different brain areas are involved in storing and "reading" memory traces depending on the precise contents of the memory, it becomes much less likely that such mechanisms are unique to storing particular past events, and more likely that some of them are involved in memories for *kinds of events, or for particular objects, or for locations,* and so on as the case may be, especially when we consider variations on these mechanisms across species. Thus, it is vitally important to the Problem of Interspecific Variation that Trace Pluralism is a live view which captures many of the advantages of TM – all those which, as explained above, can be gained by any view within the MTM family. Are there nonetheless reasons to adopt TM?

One point to note is that none of what is said above disproves TM: there are ways of explaining away findings like apparently divergent functions of hippocampal "replay" within the view. Instances in which the sequence of activity maps to actions an animal *subsequently* takes might reflect purely forward-looking simulation. But they might instead reflect the use of memory of a similar sequence of actions in the past as part of planning. Or they might reflect an episodic memory of the predicted experience being stored for the first time. And some similarities in brain activity between memory and processes like future planning may simply reflect methods with insufficient spatial and/or temporal resolution or analytical tools to detect subtle but important differences (Cheng, 2024, p. 9).

Furthermore, there is some evidence *favouring* TM proper, including its neural commitments, which Cheng and Werning take to make it plausible episodic memory constitutes a natural kind with a distinctive mechanism underpinning a cluster of psychological-level features. Making the hippocampus so central to all episodic memory traces and cortex central to reconstruction from those traces correctly predicts the order of activation of different brain areas – in encoding, cortex then hippocampus; in retrieval, cortex (representation of a cue) then hippocampus then (a more elaborate representation in) cortex— (Staresina & Wimber, 2019). Cheng's model can also reproduce specific experimental observations from Zöllner et al. (2023), relating to distinctions between retrieval of where an object was in a specific episode, retrieval of where it is usually found, and guessing (Fayyaz et al., 2022). And it can explain why hippocampal lesions seem to affect episodic memory even for events long before the lesion (although note that there has been controversy in the past about the exact nature of these effects, as subjects are still able to recall some kinds of information about those events, and some models of episodic memory still treat such evidence as showing that traces are indeed transferred to cortex (McClelland et al., 1995; Spens & Burgess, 2024)). However, such evidence is not conclusive: it is particularly difficult for such evidence to rule out versions of Trace Pluralism which allow for Cheng's hippocampal sequence-based traces, *alongside other kinds of trace*, as these make similar broad qualitative predictions about the hippocampus being important.

There are perhaps two broadly methodological reasons why one might nonetheless prefer TM. One is a general preference for stronger, more easily falsifiable hypotheses, perhaps on the grounds that they are likely to lead to more scientific progress. Another, especially important to Cheng and Werning, is the idea that episodic memory should be a *natural kind*, and that therefore the more different properties can be associated with episodic memory and clustered homeostatically the better.

Yet, while these considerations make TM an attractive working hypothesis for a research programme, they do not establish its truth – especially for debates about just how far episodic really is separate from other forms of memory and how many varieties of memory there are across species, i.e., debates in which whether it is a natural kind is one of the issues at stake. It is sometimes suggested that science should *only* traffic in natural kinds, which might be thought to foreclose such debates. Yet it is far from clear that this is true in any sense that entails that science requires "episodic memory" to pick out a *single* natural kind with a unique, well-defined mechanism and no borderline cases: Boyle (2022) and Schwartz (2025) defend views on which there may be multiple natural kinds in the vicinity of episodic memory, each of which could be usefully appealed to in science, while Gomez-Lavin (2025) develops a similarly pluralist view with respect to working memory. Furthermore, Trace Pluralism seems like a working hypothesis for a research programme just as attractive as TM's, a research programme focused on exploring different varieties of trace and reconstruction.

4. Reconstruction in semantic memory

TM seems to set episodic memory apart from other forms of memory by picking out a highly distinctive mechanism — one so distinctive it qualifies as a natural kind (Cheng, 2024; Cheng & Werning, 2016, pp. 1364–1365; Cheng, Werning, & Suddendorf, 2016; Werning, 2020, p. 311). MTM, and especially Trace Pluralism, muddies these waters: for example, hippocampal traces and replay events no longer seem so strongly associated with episodic memory alone. One way in which MTM might still mark a distinction would be if non-episodic forms of memory do not use compression and reconstruction at all. However, this section argues that, perhaps surprisingly, they probably do.

To begin, it is worth saying that compression and closely related processes often have benefits in the absence of reconstruction, unrelated to increasing effective memory capacity, such as producing models of the world that are easier to operate with and avoid overfitting (Kinney & Lombrozo, 2024; Rosch, 1988), cutting out unimportant aspects of remembered episodes (Jeunehomme & D'Argembeau, 2019, 2020), and enabling very constrained nervous systems to represent complex stimuli with only a limited number of neurons (Roper et al., 2017). So compression of some form is almost certainly at play in most varieties of memory. However, while they have some formal similarities to the compression of episodic memory, these uses of compression will involve different mechanisms, and need not be associated with any inverse reconstruction process – so there would still be a marked difference from episodic memory. More central to our case here is the role of compression of semantic memory into traces for reconstruction. Arguably, this is much more widespread than might first appear.

Unfortunately, it is difficult to say anything general about semantic memory. Even more so than episodic memory, it is not sharply defined, or at least not in a uniform way across memory science, and is associated with a number of disparate properties which come apart.[8] Such properties include: being associated with language and definitions, with Tulving (1972) describing it as a "mental thesaurus"; having a language-like rather than image-like format; generalizing over many cases; abstraction in the sense of trafficking in non-perceptual properties; slow, incremental learning; and being realized neocortically rather than hippocampally. Examples where these properties come apart include: general, abstract memories with map-like rather than sentence-like contents having nothing to do with language, depicting anything from where Paris is to social relations (Camp, 2009) – with at least some maps depending on hippocampal place cells and parahippocampal structures (Moser et al., 2008); remembering propositions with limited generality, including idiosyncratic personal facts, such as remembering one's date and location of birth; and remembering generalizations over non-abstract, perceptually representable properties (*fire is hot*), even using mental imagery to remember highly specific kinds of experience one has undergone many times (like one's morning routine), or once underwent many times (such as one's old route to work). It is understandably common to suspect that "semantic memory" in fact does not pick out a unique memory system, even amongst those who think episodic memory does (Cheng, 2024, p. 4; see also Renoult et al., 2019; Renoult & Rugg, 2020).

In this context, we need to consider several cases to understand whether compression sets episodic memory apart from such systems. We can start with the most paradigmatically non-episodic case: remembering that Paris is the capital of France.

Can this sort of memory involve compression and reconstruction? Initially, it might seem hard to even get a grip on what this might mean: *Paris is the capital of France* seems a bare, atomic fact, and remembering it appears like an indivisible, incompressible atomic achievement, which one either simply succeeds or simply fails at. There is no gist to extract, no parts that would allow you to guess the rest based on background knowledge about capitals in general.

However, consider seeming to remember that Toronto is the capital of Canada (this is not a unique case: Sydney, Istanbul and Rio de Janeiro occasion very similar errors). Here we have a non-accidental error, with a similar flavor – of systematic, understandable mis-stepping rather than arbitrary blundering – to many of the errors motivating the view that episodic memory involves reconstruction. Perhaps some people making this error even confabulate associated episodic memories of hearing from trusted sources that Toronto is the capital. Thinking that Toronto is the capital is a common error, and often comes with a feeling of confidence in remembering rather than merely guessing. Sometimes this error is due to being misled, or to intentionally guessing on the sole basis of "Toronto" being a familiar name associated with Canada. But often the error is due to something else: reconstruction of a broader compressed structure of information about Toronto – a general sense of the city – or a broader structure bringing together and summarizing a range of different facts, about capitals, about Canada, and so on. And *correctly* remembering that, say Dublin is the capital of Ireland may well rest on the very same processes: a general sense of the city and country in question, knowledge of capitals in general, and other associated knowledge. It would make sense for semantic memory to function in this way, rarely storing every individual atomic fact as such, but instead storing a more limited cluster of information that allows for sufficiently reliable reconstruction of such a network of atomic facts. Much the same information-theoretical argument as given by Werning for compression and reconstruction in the case of episodic memory can be given for the vast array of facts we need to be able to access: we can compress much of this information, and it would seem incredibly inefficient not to.

I am not claiming that semantic memory's systematically involving simulation *must* be the explanation of such errors. Many explanations are possible, including those appealing to associations and heuristic inference (although the line between these and reconstruction of a compressed trace gets fuzzy – see Sec. 5). Rather, the point is the plausibility that such reconstruction is (a) possible, and (b) frequent.

The idea that reconstruction is involved in non-episodic memory is even more plausible for other cases: cognitive maps of areas one only visits occasionally, one's model or broad understanding of a subject domain, detailed sensory memory of an oft-repeated experience, and so on. For example, if we think of semantic memory as a generative model of the environment, as do many computational approaches (Nagy et al., 2025 are typical in this respect), we could easily imagine that detailed parameters of such a model relating to specific sorts of scenario are not explicitly stored, but rather reconstructed as and when scenarios of that kind come up, using other parts of the model and, indeed, available and relevant episodic memories; or there might be a hierarchy of "levels" to the model, with shallower levels being reconstructed on the basis of deeper levels rather than being stored explicitly. Or return to the case about learning physics. In high-school physics, you just had to remember 3 or 4 equations and you could do the rest on the spot, and many students do learn to engage with the material in this way, whether they explicitly try to store the equations in this way or not.[9] Biologists have to learn many facts, but *some* generalizations cut down on learning individual instances, particularly if one learns the important exceptions along with the generalizations. One does not need to store explicitly that giraffes give birth to live young, provided one knows they are mammals and that the exceptions to this rule about mammals – monotremes – are very unlike giraffes.

We can also add a further set of cases from Aronowitz (2023, pp. 6–7) suggesting a relevant kind of continuum here (though Aronowitz's purposes are slightly different). Aronowitz's cases revolve around Clelia, who believes the post office is closed on Saturdays. Clelia formed this belief a long time ago, but is not currently thinking about it. Aronowitz's different versions of the case flesh out further details differently. At one end of a spectrum, Clelia has an explicit representation with the content "the post office is closed on Saturdays," even when not currently thinking about the topic. At the other end, she has no stored representation with this content, but builds a new representation *ad hoc* each time the question comes up, and may answer the question differently in different contexts. Aronowitz presents two cases intermediate between these. Clelia might have the *outline of a belief*: a representation of a way the world could be, which is "somehow coarser, impoverished, or less vivid than the original," but can be filled in based on context (e.g., when thinking about a specific branch). A different case would be if Clelia has *mere pieces* of a belief, such as some imagery or propositions corresponding to parts of the belief, but where there may be many ways of putting the pieces together with no clear disposition to do so in any one way, like representing "the post office has the same hours as the bank" and "the bank's hours are such-and-such," where inference is needed to make the further step. Aronowitz also gives extended examples of other cases with a similar structure: exemplar-based concepts, and intuitive physics (a stored model which is accessed via running simulations of particular cases). One might wish to distinguish Clelia's *remembering* from her *believing*. But for our purposes, the point is that we have another quite mundane case where someone's mechanism for reliably forming certain thoughts is structured remarkably like compression and reconstruction.

In general, insofar as the mind is at least somewhat well-designed by natural selection and shaped in a helpful direction by learning, it will involve some compression and reconstruction rather than uncompressed storage for most explicit memory. The precise amount of explicit storage vs. compression that is optimal in a given case will depend on factors like the costs of storage, the costs of the relevant kind of processing, the compressibility of the information structure (how elegantly and accurately this information be captured by such processes – with biology scoring lower than physics), and so on. So cases with no compression at all, or with extreme amounts where barely any information is explicitly stored relative to that which is reconstructed, might both be rare; and if information about scenarios is especially compressible with lots of exploitable structure, episodic memory might have more compression and reconstruction than its rivals. But this would have to be shown: as it stands, we should not think that compression and reconstruction *per se* are distinctive of episodic memory. It is worth emphasizing that thinking about semantic memory as likely to involve reconstruction in this way deeply challenges common assumptions about semantic memory as a fixed, stable model or set of representations of the world which is used to reconstruct episodic memory and for general-purpose reasoning and which only changes when revised in light of new evidence: it may turn out that no kind of memory constitutes a completely fixed point, with *all* our memory being at least somewhat subject to construction, compression and reconstruction in different contexts.

MTM does not support a sharp distinction between episodic and other forms of memory. The many forms of explicit memory are likely to share common broad computational approaches of compressing, storing and reconstructing information. So only further commitments about the precise way in which this happens (such as TM's account of hippocampus-specific compression in episodic memory), or some very different approach to individuation, could carve episodic memory out as a distinct kind.

5. Reconstruction and inference

One response to this conclusion is that in the Toronto and physics cases, we do not really have *reconstruction* from a compressed *trace* resulting in a *memory* with enriched content; we have *inference* from a limited number of *memories* (like *Toronto is a big important city in Canada* and $F = ma$) to new *thoughts*. A full discussion of this worry is beyond the scope of this paper, as it would require addressing such controversial issues as the nature of inference (Boghossian, 2014; Marcus, 2021; Quilty-Dunn & Mandelbaum, 2018; Varga et al., 2024). However, we can sketch some considerations suggesting how future work can alleviate this worry. Firstly, insofar as cases of semantic reconstruction like for "Toronto is the capital of Canada" involve inference, it is often not *conscious* inference, and is so fluent and effortless (in part because we have performed these inferences many times before[10]) that subjectively the result may feel more like "aha, I remember now!," than careful inference. Secondly, the "inference" involved may consist of the very same sorts of processes (e.g., Bayesian updating or use of heuristics) posited to be involved in episodic memory reconstruction – and, indeed, in perception – by adherents of TM. And abstracting from the specifics of TM, some of the classic evidence for episodic memory reconstruction seem just as likely to involve inference as the case of thinking Toronto is the capital of Canada: consider the finding that individuals are more likely to say they saw glass on the floor if asked a question about a car collision describing it with the word "smashed" (Loftus & Palmer, 1974); or the finding that individuals are likely to misremember having seen words which fit the overall subject-matter of a list of words they were presented with, such as medicine-related words (Deese, 1959; Roediger & McDermott, 1995). So, the idea that there is a sharp distinction between reconstruction and inference, such that the processes involved in semantic memory fall on one side and episodic memory on the other, seems dubious.

6. Varieties of long-term memory

Suppose it is true that compression, traces, and reconstruction are involved in many non-episodic kinds of memory. It follows that episodic memory cannot be distinguished solely by involving these. However, a different proposal remains a live possibility: perhaps we can individuate varieties of memory according to specific varieties of reconstruction and trace.

It is certainly possible, as suggested above, that we have several different ways of reconstructing, which pair with quite different forms of memory trace, such as those based specifically around perceptual systems, those based on general mechanisms of pattern completion at the sub-representational neural level, those based around classical forms of computation, those based around sequences or dynamically evolving rather than static representations, and so on.

However, we also saw above that these different kinds of reconstruction may be used across memory types. Episodic memory is often rich in perceptual content and typically includes sequences of sub-events, but often uses other kinds of information – emotional, causal, contextual, autobiographical – and insofar as it is plausible that we have different varieties of reconstruction, it is plausible that they will be in play for these different kinds of information even within episodic memory. For this reason, Trace Pluralism has some attraction relative to a TM view positing a unitary mechanism focused around sequences and general pattern completion. Conversely, many of these distinctive kinds of information are used in non-episodic sorts of memory – memory for oft-repeated but highly specific sequences of events such as one's morning routine or a run through an oft-visited maze; cognitive maps of particular locations; general autobiographical information. As such, we might well see particular mechanisms being used for compression, storage, and reconstruction in both episodic and non-episodic memory alike.

Empirical investigation may yet reveal episodic memory to be characterized by a distinctive kind of reconstruction/trace, and TM's bold hypotheses about that possibility relating to sequence processing in the hippocampus seems like the most promising route to such an outcome. But we should be open to the possibility that things will not turn out that way. If they do not, there will still be a distinctive kind of compression and reconstruction associated with the hippocampus – it is just that this will not map one-one onto cases of remembering that we ordinarily deem episodic, with other systems typically being involved in such cases, and with this mechanism showing up in, for example, memory for *kinds* of sequence rather than just particular events.

If so, one could simply identify episodic memory with the system for storing, simulating, and reconstructing sequences, and accept that it rarely acts on its own and is not proprietary to any familiar kind of state, with nearly all instances of remembering involving multiple systems. After all, Tulving always accepted that in most instances of remembering multiple underlying systems would be combined (Andonovski, 2023; Pan, 2022).

However, this option might require relinquishing any straightforward link to other oft-cited properties of episodic remembering, such as being specially related to particular events, self-understanding, or autonoesis, plus the idea that episodic memory is a particularly important system, rather than one amongst a multiplicity even within the hippocampus.

Conclusion

Trace Minimalism has many virtues. As a part of the MTM family, it offers an attractive proposal about the relationship between simulation, memory traces, and remembering, explaining both the overall reliability of memory and the role within remembering for simulation despite its tendency to construct details diverging from the initial event. Its claiming that episodic memory is a natural kind, and identifying a candidate cluster of properties, meanwhile, offers a bold hypothesis which may yet turn out to be true, and can guide research. However, it is important to separate out these two aspects of TM, and to recognize that they have independent sources of support. This is important for dealing with the issue of individuating episodic memory; but it is also important for understanding other varieties of memory, where, as we have seen, MTM offers insight through prompting the search for mechanisms of compression and reconstruction beyond episodic memory. Indeed, Trace Pluralist versions of MTM on which most remembering involves multiple kinds of compression, reconstruction and memory trace working together, offer an attractive alternative picture to TM, suggesting a research programme which deserves to be pursued too. Such a research programme would have much to learn from TM, given that it would accept that the mechanisms posited by TM may well operate in EM, but it would study these alongside other varieties of compression and reconstruction, the core of so many different varieties of memory.

Notes

1. This paper will not detail the more subtle versions of simulationism and causalism: such discussion can be found elsewhere (Andonovski & Michaelian, 2024; Andonovski et al., 2024; Langland-Hassan, 2022; Michaelian, 2011, 2016; Robins, 2020; Sant'anna et al., 2023; Schacter & Addis, 2020, 2020). All I need for the purposes of this paper is to capture in broad outline why TM appears to be a useful position in the debate.
2. Cheng (2024) helpfully describes his notion of a scenario as a temporally extended version of Maguire and colleagues' notion of a "scene."
3. Werning also suggests preserving large amounts of information may lead to overfitting or making the information unmanageable (Richards & Frankland, 2017). Brown (2024) also emphasizes the potential for making information unmanageable, but points out that overfitting is not an inevitable consequence of storing rich information, given that regularities could be extracted from the excessively rich

information (i.e., the relevant form of compression could occur) before any specific use rather than before storage.

4. I will nonetheless talk of them as "compressed representations," given that their whole role is to carry information that allows for the reconstruction of representations with a greater amount of information, approximating the original representations from which they were derived. We can talk of them as representations in something like the sense that we can talk of a folded up camping chair as a seat despite the fact that there are defensible views on which this is not literally true, given that it needs to be transformed to be used as such. Note that given the qualification "information-carrying entity" in the definition of compression in Section. 2, these traces do not have to be literal representations to count as the products of compression.
5. Related notions of "simulation" and "offline" use of psychological systems have been developed in more detail for other purposes (Carruthers, 2015; Currie & Ravenscroft, 2002; Goldman, 2006; Nichols et al., 1998).
6. See e.g (Dijkstra & Fleming, 2023; Dijkstra et al., 2019; Gottfried et al., 2004; Kosslyn et al., 2006; Pearson et al., 2015; Vaidya et al., 2002; Wadia et al., 2024; Wheeler et al., 2000). although there are complexities here beyond the scope of this paper (Favila et al., 2020; She et al., 2024).
7. It is known that some forms of memory can be formed, sometimes very rapidly, in areas other than the hippocampus, such as fear conditioning in the basolateral amygdala and prefrontal cortex (Tonegawa et al., 2018; Zhang et al., 2020). Furthermore, the data relating to how memories are encoded in hippocampus are complex, and it is likely that replay is not the whole story: in some paradigms, only a fraction of hippocampal place cells seem to be involved in engrams (Tanaka et al., 2018); there is evidence for hippocampal "bar codes" for specific experiences which are distinct from place cell activity (Chettih et al., 2024); and there is some (controversial) evidence for "preplay," i.e., replay-like activity which does not correspond to a past experience (Ólafsdóttir et al., 2018).
8. See Reilly et al. (2024) for a heroic attempt to herd these felines.
9. For someone flawless at mathematics, the physics case is in fact *lossless* compression. That is, thanks to the relevant equations expressing exceptionless laws (as far as answering exam questions goes), the individual facts can be derived perfectly every time, with no understandable misfirings as in the case of Toronto. This does not affect the overall importance of compression and reconstruction: lossless compression still allows for more efficient storage, generalization etc.
10. Perrin (2018) and Langland-Hassan (2022) both give such fluency-related processing a starring role in *episodic* memory.

Acknowledgements

I would like to thank participants at the Causation in Memory workshop, Grenoble 2023, for helpful discussion of this material, and two anonymous reviewers for this journal for helpful comments.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research is part of a project that has received funding from the H2020 European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, [Grant Number 851145].

ORCID

Simon Alexander Burns Brown http://orcid.org/0000-0002-5606-6979

References

Addis, D. R., & Szpunar, K. K. (2024). Beyond the episodic-semantic continuum: The multidimensional model of mental representations. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *379*(1913). https://doi.org/10.1098/rstb.2023.0408

Andonovski, N. (2023). Autonoesis and the Galilean science of memory: Explanation, idealization, and the role of crucial data. *European Journal for Philosophy of Science*, *13*(3), 42. https://doi.org/10.1007/s13194-023-00548-3

Andonovski, N., & Michaelian, K. (2024). *Naturalism and simulationism in the philosophy of memory*. https://hal.science/hal-04638053/

Andonovski, N., Sutton, J., & McCarroll, C. J. (2024). Eliminating episodic memory? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *379*(1913), 20230413. https://doi.org/10.1098/rstb.2023.0413

Aronowitz, S. (2022). Semanticization challenges the episodic-semantic distinction. *British Journal for the Philosophy of Science*, *721760*. https://doi.org/10.1086/721760

Aronowitz, S. (2023). A planning theory of belief. *Philosophical Perspectives*, *37*(1), 5–17. https://doi.org/10.1111/phpe.12178

Barlow, H. (1961/2012). Possible principles underlying the transformation of sensory messages. In W. Rosenblith (Ed.), Sensory Communication (pp. 217–234). Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/9780262518420.003.0013

Block, N. (2023). *The border between seeing and thinking*. Oxford University Press.

Boghossian, P. (2014). What is inference? *Philosophical Studies*, *169*(1), 1–18. https://doi.org/10.1007/s11098-012-9903-x

Boyle, A. (2022). Episodic memory in animals: Optimism, kind scepticism and pluralism. In A. Sant'Anna, C. McCarroll, K. Michaelian (Eds.), *Current controversies in philosophy of memory*. Routledge.

Boyle, A. (2023). Disagreement & classification in comparative cognitive science. *Noûs*, *58*(3), 825–847. https://doi.org/10.1111/nous.12480

Boyle, A., & Brown, S. A. B. (2025). Episodic memory in animals. *Philosophy Compass*, *20*(5), e70037. https://doi.org/10.1111/phc3.70037

Brown, S. A. B. (2024). Episodic memory and unrestricted learning. *Philosophy of Science*, *91*(1), 90–110. https://doi.org/10.1017/psa.2023.16

Camp, E. (2009). A language of baboon thought. In R. W. Lurz (Ed.), *The philosophy of animal minds* (pp. 108–127). Cambridge University Press.

Carruthers, P. (2015). *The centered mind: What the science of working memory shows us about the nature of human thought*. Oxford University Press.

Casals, P., & Kahn, A. E. (1974). *Joys and sorrows*. Simon and Schuster.

Cheng, S. (2024). Distinct mechanisms and functions of episodic memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *379*(1913), 20230411. https://doi.org/10.1098/rstb.2023.0411

Cheng, S., & Werning, M. (2016). What is episodic memory if it is a natural kind? *Synthese*, *193*(5), 1345–1385. https://doi.org/10.1007/s11229-014-0628-6

Cheng, S., Werning, M., & Suddendorf, T. (2016). Dissociating memory traces and scenario construction in mental time travel. *Neuroscience and Biobehavioral Reviews*, *60*, 82–89. https://doi.org/10.1016/j.neubiorev.2015.11.011

Chettih, S. N., Mackevicius, E. L., Hale, S., & Aronov, D. (2024). Barcoding of episodic memories in the hippocampus of a food-caching bird. *Cell*, *187*(8), 1922–1935.e20. https://doi.org/10.1016/j.cell.2024.02.032

Clarke, S., & Beck, J. (2023). Border disputes: Recent debates along the perception–cognition border. *Philosophy Compass*, *18*(8), e12936. https://doi.org/10.1111/phc3.12936

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, *107*(2), 261–288. https://doi.org/10.1037/0033-295x.107.2.261

Currie, G., & Ravenscroft, I. (2002). *Recreative minds: Imagination in philosophy and psychology*. (C. Hoerl ,Ed.) Oxford University Press.

De Brigard F. (2014). The Nature of Memory Traces. Philosophy Compass, *9*(6), 402–414. https://doi.org/10.1111/phc3.12133

De Brigard, F., Umanath, S., & Irish, M. (2022). Rethinking the distinction between episodic and semantic memory: Insights from the past, present, and future. *Memory & Cognition*, *50*(3), 459–463. https://doi.org/10.3758/s13421-022-01299-x

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, *58*(1), 17–22. https://doi.org/10.1037/h0046671

Dijkstra, N., Bosch, S. E., & van Gerven, M. A. (2019). Shared neural mechanisms of visual perception and imagery. *Trends in Cognitive Sciences*, *23*(5), 423–434. https://www.cell.com/trends/cognitive-sciences/fulltext/S1364-6613(19)30059-2?dgcid=raven_jbs_etoc_email

Dijkstra, N., & Fleming, S. M. (2023). Subjective signal strength distinguishes reality from imagination. *Nature Communications*, *14*(1), 1627. https://doi.org/10.1038/s41467-023-37322-1

Favila, S. E., Lee, H., & Kuhl, B. A. (2020). Transforming the concept of memory reactivation. *Trends in Neurosciences*, *43*(12), 939–950. https://doi.org/10.1016/j.tins.2020.09.006

Fayyaz, Z., Altamimi, A., Zoellner, C., Klein, N., Wolf, O. T., Cheng, S., & Wiskott, L. (2022). A model of semantic completion in generative episodic memory. *Neural Computation*, *34*(9), 1841–1870. https://doi.org/10.1162/neco_a_01520

Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University Press USA.

Gomez-Lavin, J. (2025). Working memory is as working memory does: A pluralist take on the center of the mind. *WIREs Cognitive Science*, *16*(1), e1696. https://doi.org/10.1002/wcs.1696

Gottfried, J. A., Smith, A. P. R., Rugg, M. D., & Dolan, R. J. (2004). Remembrance of odors past: Human olfactory cortex in cross-modal recognition memory. *Neuron*, *42*(4), 687–695. https://doi.org/10.1016/s0896-6273(04)00270-3

Green, E. J. (2023). The perception-cognition border. In B. McLaughlin, J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 469–493). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781394259847.ch26

Henke, K. (2010). A model for memory systems based on processing modes rather than consciousness. *Nature Reviews Neuroscience*, *11*(7), 523–532. https://doi.org/10.1038/nrn2850

Jeunehomme, O., & D'Argembeau, A. (2019). The time to remember: Temporal compression and duration judgements in memory for real-life events. *The Quarterly Journal of Experimental Psychology*, *72*(4), 930–942. https://doi.org/10.1177/1747021818773082

Jeunehomme, O., & D'Argembeau, A. (2020). Event segmentation and the temporal compression of experience in episodic memory. *Psychological Research*, *84*(2), 481–490. https://doi.org/10.1007/s00426-018-1047-y

Kerrén, C., Reznik, D., Doeller, C. F., & Griffiths, B. J. (2025). Exploring the role of dimensionality transformation in episodic memory. *Trends in Cognitive Sciences*, *29*(7), 614–626. https://doi.org/10.1016/j.tics.2025.01.007

Keven, N. (2022). What does it take to remember episodically? In A. Sant'anna, C. J. McCarroll, & K. Michaelian (Eds.), *Current controversies in philosophy of memory* (pp. 206–222). Taylor and Francis.

Kinney, D., & Lombrozo, T. (2024). Building compressed causal models of the world. *Cognitive Psychology*, *155*, 101682. https://doi.org/10.1016/j.cogpsych.2024.101682

Kosslyn, S. M., Thompson, W. L., & Ganis, G. (2006). *The case for mental imagery*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195179088.001.0001

Langland-Hassan, P. (2022). Propping up the causal theory. *Synthese*, *200*(2), 95. https://doi.org/10.1007/s11229-022-03635-9

Liu, A. A., Henin, S., Abbaspoor, S., Bragin, A., Buffalo, E. A., Farrell, J. S., Foster, D. J., Frank, L. M., Gedankien, T., Gotman, J., Guidera, J. A., Hoffman, K. L., Jacobs, J., Kahana, M. J., Li, L., Liao, Z., Lin, J. J., Losonczy, A., Malach, R. . . . Zugaro, M. (2022). A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations. *Nature Communications*, *13*(1), 6000. https://doi.org/10.1038/s41467-022-33536-x

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, *13*(5), 585–589. https://doi.org/10.1016/S0022-5371(74)80011-3

Malanowski, S. (2016). Is episodic memory uniquely human? Evaluating the episodic-like memory research program. *Synthese*, *193*(5), 1433–1455. https://doi.org/10.1007/s11229-015-0966-z

Marcus, E. (2021). *Belief, inference, and the self-conscious mind*. Oxford University Press.

Martin C B and Deutscher M. (1966). Remembering. The Philosophical Review, *75*(2), 161. https://doi.org/10.2307/2183082

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457. https://doi.org/10.1037/0033-295X.102.3.419

Michaelian, K. (2011). Generative memory. *Philosophical Psychology*, *24*(3), 323–342. https://doi.org/10.1080/09515089.2011.559623

Michaelian, K. (2016). *Mental time travel: Episodic memory and our knowledge of the personal past*. MIT Press.

Michaelian, K., & Robins, S. K. (2018). Beyond the causal theory?: Fifty years after Martin and Deutscher 1. In K. Michaelian, D. Debus, D. Perrin (Eds.) *New directions in the philosophy of memory* (pp. 13–32). Routledge. https://www.taylorfrancis.com/chapters/edit/10.4324/9781315159591-2/beyond-causal-theory-kourken-michaelian-sarah-robins

Milton, F., Fulford, J., Dance, C., Gaddum, J., Heuerman-Williamson, B., Jones, K., Knight, K. F., MacKisack, M., Winlove, C., & Zeman, A. (2021). Behavioral and neural signatures of visual imagery vividness extremes: Aphantasia versus hyperphantasia. *Cerebral Cortex Communications*, *2*(2), tgab035. https://doi.org/10.1093/texcom/tgab035

Monzel, M., Leelaarporn, P., Lutz, T., Schultz, J., Brunheim, S., Reuter, M., & McCormick, C. (2024). Hippocampal-occipital connectivity reflects autobiographical memory deficits in aphantasia. *ELife*, *13*, RP94916. https://doi.org/10.7554/eLife.94916

Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, *31*(1), 69–89. https://doi.org/10.1146/annurev.neuro.31.061307.090723

Murray, E. A., Wise, S. P., & Graham, K. S. (2017). *The evolution of memory systems: Ancestors, anatomy, and adaptations*. Oxford University Press.

Nagy, D. G., Orbán, G., & Wu, C. M. (2025). Adaptive compression as a unifying framework for episodic and semantic memory. *Nature Reviews Psychology*, *4*(7), 484–498. https://doi.org/10.1038/s44159-025-00458-6

Nichols, S., Stich, S. P., Leslie, A. M., & Klein, D. B. (1998). Varieties of off-line simulation. In P. Carruthers & J. Boucher (Eds.), *Language and thought: Interdisciplinary themes* (pp. 39–74). Cambridge University Press.

Ólafsdóttir, H. F., Bush, D., & Barry, C. (2018). The role of hippocampal replay in memory and planning. *Current Biology*, *28*(1), R37–R50. https://doi.org/10.1016/j.cub.2017.10.073

Pan, S. (2022). What is so special about episodic memory: Lessons from the system-experience distinction. *Synthese*, *200*(1), 1–26. https://doi.org/10.1007/s11229-022-03500-9

Pearson, J., Naselaris, T., Holmes, E. A., & Kosslyn, S. M. (2015). Mental imagery: Functional mechanisms and clinical applications. *Trends in Cognitive Sciences*, *19*(10), 590–602. https://doi.org/10.1016/j.tics.2015.08.003

Perrin, D. (2018). A case for procedural causality in episodic recollection 1. In K. Michaelian, D. Debus, & D. Perrin (Eds.), *New directions in the philosophy of memory* (pp. 33–51). Routledge.

Quilty-Dunn, J., & Mandelbaum, E. (2018). Inferential transitions. *Australasian Journal of Philosophy*, *96*(3), 532–547. https://doi.org/10.1080/00048402.2017.1358754

Reilly, J., Shain, C., Borghesani, V., Kuhnke, P., Vigliocco, G., Peelle, J. E., Mahon, B. Z., Buxbaum, L. J., Majid, A., Brysbaert, M., Borghi, A. M., De Deyne, S., Dove, G., Papeo, L., Pexman, P. M., Poeppel, D., Lupyan, G., Boggio, P., Hickok, G. . . . Garcea, F. E. (2024). What we mean when we say semantic: Toward a multidisciplinary semantic glossary. *Psychonomic Bulletin & Review*, *32*(1), 243–280. https://doi.org/10.3758/s13423-024-02556-7

Renoult, L., Irish, M., Moscovitch, M., & Rugg, M. D. (2019). From knowing to remembering: The semantic-episodic distinction. *Trends in Cognitive Sciences*, *23*(12), 1041–1057. https://doi.org/10.1016/j.tics.2019.09.008

Renoult, L., & Rugg, M. D. (2020). An historical perspective on Endel Tulving's episodic-semantic distinction. *Neuropsychologia*, *139*, 107366. https://doi.org/10.1016/j.neuropsychologia.2020.107366

Richards, B. A., & Frankland, P. W. (2017). The persistence and transience of memory. *Neuron*, *94*(6), 1071–1084. https://doi.org/10.1016/j.neuron.2017.04.037

Robins, S. (2020). Defending discontinuism, naturally. *Review of Philosophy and Psychology*, *11*(2), 469–486. https://doi.org/10.1007/s13164-020-00462-0

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *21*(4), 803–814. https://doi.org/10.1037//0278-7393.21.4.803

Roper, M., Fernando, C., & Chittka, L. (2017). Insect bio-inspired neural network provides new evidence on how simple feature detectors can enable complex visual generalization and stimulus location invariance in the miniature brain of honeybees. *PLOS Computational Biology*, *13*(2), e1005333. https://doi.org/10.1371/journal.pcbi.1005333

Rosch, E. (1988). Principles of categorization. In A. Collins & E. E. Smith (Eds.), *Readings in cognitive science, a perspective from psychology and artificial intelligence* (pp. 312–322). Morgan Kaufmann Publishers.

Rubin, D. C., & Umanath, S. (2015). Event memory: A theory of memory for laboratory, autobiographical, and fictional events. *Psychological Review*, *122*(1), 1–23. https://doi.org/10.1037/a0037907

Sant'anna, A., McCarroll, C. J., & Michaelian, K. (2023). *Current controversies in philosophy of memory*. Routledge. https://api.taylorfrancis.com/content/books/mono/download?identifierName=doi&identifierValue=10.4324/9781003002277&type=googlepdf

Schacter, D. L., & Addis, D. R. (2020). Memory and imagination: Perspectives on constructive episodic simulation. In A. Abraham (Ed.) *The Cambridge handbook of the imagination* (pp. 111–131). Cambridge University Press. https://doi.org/10.1017/9781108580298.008

Schwartz A. (2025). Are episodic memory and episodic simulation different in kind?. Philosophical Psychology, 1–24. https://doi.org/10.1080/09515089.2025.2451819

She, L., Benna, M. K., Shi, Y., Fusi, S., & Tsao, D. Y. (2024). Temporal multiplexing of perception and memory codes in IT cortex. *Nature*, *629*(8013), 861–868. https://doi.org/10.1038/s41586-024-07349-5

Sherman, B. E., Turk-Browne, N. B., & Goldfarb, E. V. (2024). Multiple memory subsystems: Reconsidering memory in the mind and brain. *Perspectives on Psychological Science*, *19*(1), 103–125. https://doi.org/10.1177/17456916231179146

Siegel, S. (2017). *The rationality of perception*. Oxford University Press.

Spens, E., & Burgess, N. (2024). A generative model of memory construction and consolidation. *Nature Human Behaviour*, 1–18. https://doi.org/10.1038/s41562-023-01799-z

Staresina, B. P., & Wimber, M. (2019). A neural chronometry of memory recall. *Trends in Cognitive Sciences*, *23*(12), 1071–1085. https://doi.org/10.1016/j.tics.2019.09.011

Tanaka, K. Z., He, H., Tomar, A., Niisato, K., Huang, A. J. Y., & McHugh, T. J. (2018). The hippocampal engram maps experience but not place. *Science*, *361*(6400), 392–397. https://doi.org/10.1126/science.aat5397

Teyler T J and DiScenna P. (1986). The hippocampal memory indexing theory. Behavioral Neuroscience, *100*(2), 147–154. https://doi.org/10.1037//0735-7044.100.2.147

Tonegawa, S., Morrissey, M. D., & Kitamura, T. (2018). The role of engram cells in the systems consolidation of memory. *Nature Reviews Neuroscience*, *19*(8), 485–498. https://doi.org/10.1038/s41583-018-0031-2

Treves A and Rolls E T. (1994). Computational analysis of the role of the hippocampus in memory. Hippocampus, *4*(3), 374–391. https://doi.org/10.1002/hipo.450040319

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381–403). Academic Press.

Tulving, E. (1983). *Elements of episodic memory*. Clarendon Press.

Vaidya, C. J., Zhao, M., Desmond, J. E., & Gabrieli, J. D. E. (2002). Evidence for cortical encoding specificity in episodic memory: Memory-induced re-activation of picture processing areas. *Neuropsychologia*, *40*(12), 2136–2143. https://doi.org/10.1016/s0028-3932(02)00053-2

Varga, N. L., Morton, N. W., & Preston, A. R. (2024). Schema, inference, and memory. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford handbook of human memory, two volume pack: Foundations and applications*. (pp. 1400–1425). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190917982.013.47

Wadia, V. S., Reed, C. M., Chung, J. M., Bateman, L. M., Mamelak, A. N., Rutishauser, U., & Tsao, D. Y. (2024). A shared code for perceiving and imagining objects in human ventral temporal cortex. *BioRxiv*, 2024. https://doi.org/10.1101/2024.10.05.616828

Watzl, S., Sundberg, K., & Nes, A. (2021). The perception/cognition distinction. *Inquiry, an Interdisciplinary Journal of Philosophy*, *66*(2), 165–195. https://doi.org/10.1080/0020174x.2021.1926317

Werning, M. (2020). Predicting the past from minimal traces: Episodic memory and its distinction from imagination and preservation. *Review of Philosophy and Psychology*, *11*(2), 301–333. https://doi.org/10.1007/s13164-020-00471-z

Wheeler, M. E., Petersen, S. E., & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. *Proceedings of the National Academy of Sciences*, *97*(20), 11125–11129. https://doi.org/10.1073/pnas.97.20.11125

Zeng, X., Diekmann, N., Wiskott, L., & Cheng, S. (2023). Modeling the function of episodic memory in spatial learning. *Frontiers in Psychology*, *14*. https://doi.org/10.3389/fpsyg.2023.1160648

Zhang, X., Kim, J., & Tonegawa, S. (2020). Amygdala reward neurons form and store fear extinction memory. *Neuron*, *105*(6), 1077–1093.e7. https://doi.org/10.1016/j.neuron.2019.12.025

Zöllner, C., Klein, N., Cheng, S., Schubotz, R. I., Axmacher, N., & Wolf, O. T. (2023). Where was the toaster? A systematic investigation of semantic construction in a new virtual episodic memory paradigm. *The Quarterly Journal of Experimental Psychology*, *76*(7), 1497–1514. https://doi.org/10.1177/17470218221116610