Neither Noise nor Signal
The Role of Context in Memory Models

Ian O'Loughlin

Pacific University, Forest Grove, USA
ian.oloughlin@pacificu.edu

Abstract. Context plays a crucial role in learning and memory, but a satisfactory characterization of this role in models of memory remains elusive. Classical and recent studies show that context cannot be meaningfully treated as either the figure or the ground, the noise or the signal, in memory models. This impasse belies certain cognitivist assumptions common to traditional cognitive science. A number of postcognitivist movements in philosophy and cognitive science have offered effective critiques of the basic framework, often borrowed in memory science, that depicts the human cognitive system as a dimensionless executive control unit receiving and transforming signals as input from the environment. These revisionary movements have also offered up alternative dynamic approaches to cognitive modeling and explanation, which can and should be deployed in memory science in order to resolve the impasse surrounding the modeling of context and memory.

1 Context and Memory

Context plays a crucial role in learning and memory. Although researchers across several domains agree that this is so, a satisfactory characterization of this role in models of memory remains elusive [1, pp. 101],[2]. In early conditioning studies, the context of a conditioning experiment was treated as a necessary evil: as a subject is conditioned to respond to a stimulus, the apparatus of the experiment also engenders responses similar to a conditional stimulus [1]. This noise in conditioning experiments eventually came to be treated as a modular stimulus in its own right, spawning a vast, albeit somewhat heterogeneous, collection of methodologies surrounding context conditioning [3, 4]. In memory models, a representation of context came to play a very similar role to representations of any other target stimuli [5, 6]. Recently, however, a number of studies have once more demonstrated incongruities in treating context as just another represented stimulus [7, 8]. This leaves modelers at an impasse: context cannot be adequately modeled as another signal received by the learning and remembering system, but nor can it be adequately modeled as merely the noise in contrast to which such signals proceed.

This impasse, however, is predicated on an assumption about the nature of cognitive systems. This is the assumption that there is a static baseline of awareness which receives signals in a void. Everything must be either signals or noise only if remembering systems are essentially zero-dimensional receivers of input, if the framework for our models is a basically computational picture of discrete phases of input, storage, and output. In cognitive science, these cognitivist assumptions are beginning to give way under their own weight. Especially in approaches within the dynamic systems family of modeling methodologies [7],[9, 10], as well as research trajectories informed by embodied and situated cognition [11–13], the subject as a zero-point receiver and transmitter of representations is quickly becoming obsolete [14–16]. Philosophers of mind have been debating the merits and demerits of just this framework for a century, and recent cognitive science has begun to attend to these lessons in earnest.

Although researchers in the cognitive science of human memory have not commonly been among the first to question cognitivist assumptions, whether models based on the computational architecture of encoding, storage, and retrieval can adequately serve has lately come under fire from both cognitive psychologists and philosophers[17–19]. In the emerging foundation for new models of human memory, the line between data and algorithm has been erased[20], the phases of memory processes have become continuous rather than discrete[21], and memory construction and reconsolidation have become the rule rather than the exception [22]. These developments taken together point to a plausible, albeit revisionist, resolution to the context modelers' impasse: context has resisted memory models that rely on fundamentally computationalist architectures because these frameworks insist that a stimulus is either a signal being received as "input" or the mere noise against which such signals are discerned. Relinquishing this demand may entail the dissolution of modularity and of any notion of a dimensionless "executive unit" which receives input and directs output–revisions which have been independently motivated in the philosophy of cognitive science–but the resulting frameworks will allow models to faithfully characterize the complex role that context plays in human memory.

2 Modeling Figure and Ground in the Cognitive Psychology of Memory

Although many strands still connect the two, research into *learning* and research into *memory* have been long divided into two parallel (and sometimes less parallel) trajectories in cognitive psychology. Although the primary focus of this paper is memory research, the story of the way context has come to be modeled begins with learning and conditioning experiments in the middle of the twentieth century. Early researchers in conditioning quickly discovered that the context of any given successful case of conditioning became something like its own conditioned stimulus, in addition to whatever target stimulus to which the subject had been conditioned in the context in question. Experimenters began treating the context as another CS in their models, and thus set the stage for what would become the standard treatment of context in models of human memory [1, pp. 101]. Furthermore the revisions needed to make memory models adequate for context are almost certainly analogous to needed revisions in learning models, especially since these treatments have remained more or less parallel in the divergent research traditions. Researchers who are concerned to draw conceptual connections between these, such as Mark Bouton, have made this evident:

> "In contemporary learning theory, *context* is typically defined as the background stimuli provided by the apparatus (e.g., Skinner box) in which the experiment is conducted. This definition is similar to the one used in the human memory literature, where *context* is often defined as stimuli provided by cues emanating from the room in which the experiment occurs (e.g., Smith, 1988; Smith and Vela, 2001). In either case, the context is a relatively long-duration stimulus that surrounds or embeds the target stimuli that are to be learned or remembered (whether they are conditioned or discriminative stimuli, or items in a list)." [2, pp. 233]

Memory researchers have even perhaps tended to be a bit more cautious than those modeling learning in their treatment of context. While in the conditioning literature context is often treated as a sort of necessary evil, an additional CS that always accompanies conditioning, context in the memory literature is often characterized as a *cue*, a signal that aids, perhaps crucially, in retrieval of memory. In an article from 1995–in which the cracks in the computationalist edifice are already beginning to show, and to which we shall return presently–Wesley Spencer and Raz Naftali sensibly and sensitively wrote that "The term *context* denotes all conditions and circumstances under which memory for an item was acquired." [7, pp. 528] This characterization is distinct from Bouton's in a subtle but important aspect: those researchers who treat context as "emanating cues" are essentially treating context as a signal, following the CS assignation in learning experiments, while those who follow Spencer and Naftali in treating it as the whole of the "conditions and circumstances" surrounding an encoded or retrieved memory are essentially treating context as the ground of such signals. As each of these researchers is well aware, *experimentally*, context has failed to behave well in either of these roles.

The Spencer and Naftali study demonstrated one such incongruity. It has long been known that aging has adverse effects on the memory abilities of human subjects. In their study, Spencer and Naftali showed that memory ability for context deteriorates much earlier and more severely in aging humans than does memory ability for *content*. [7] That is, aging adults show much more susceptibility to memory failure for context than would be expected were context just another stimulus that has been stored and retrieved. While semantic memory (remembering facts) and episodic memory (remembering past episodes) remain largely intact quite late in the life of the healthy remembering subject, our ability to remember the contexts associated with these retrieved memories declines dramatically long before this. It is perhaps this poor fit between context and signal that induces the authors to characterize context as "circumstances" rather than "cues", although they do not explicitly recommend any particular theoretical revisions.

Michael Fanselow has more recently demonstrated another difficulty in treating context as a memory item that is stored and retrieved. The habituation and extinction rates for remembering and forgetting in a large variety of circumstances have been thoroughly studied and documented, including basic conditioning experiments. Fanselow has found that habituation and extinction rates for *context* conditioning are not only *divergent* from what would be expected, but indeed in some cases they are *inverse* [8][1]. Fanselow's study showed that although, as is well known, the effectiveness of an instance of classical conditioning is inversely proportional to the delay between US and CS (at least above a very brief threshold, the role of which is controversial), the effectivenss of context conditioning is emphdirectly proportional, for a number of cases, to the delay between the introduction of the context "stimulus" and unconditioned stimulus. Rats introduced to a context "stimulus" *simultaneous* with, or directly preceding, unconditioned footshock stimilus, exhibit no conditioning whatsoever to the context "stimulus", despite the fact that simultaneous and directly preceding instances of CS are standardly what produce the most reliable conditioning. Once more, treating the context as a stored, remembered item in the fashion of a conditioned stimulus or environmental cue seems directly contravened by empirical results. These are not the only difficulties faced by researchers who have tried to model context within a framework of stored and retrieved items; context is sometimes assigned multiple roles, or multiple neural bases, to try to resolve these seeming inconsistencies [5, pp. 2431].

In traditional models of human memory, it can be difficult to know how else context could be modeled. Although the details vary, the basic framework in memory research is to model human remembering on the encoding, storage, and retrieval of memory traces.[1] That is, the human rememberer is taken be fundamentally analogous to an input/output device; memories are stored data, that has been encoded, that can be transformed into various outputs under certain conditions. Stored memory traces are taken to *represent* the fact or episode that will later be recalled, and treating context as a stored item is just to treat it as a stored representation.[2] Many attempts in cognitive psychology and neuroscience to understand the role of context in memory involve an attempt to better understand the location and structure of the stored representation of the context. Fanselow as well, when describing the context conditioning in his footshock experiments, writes:

> "This exploration allows the animal to store an integrated representation of the context and once that integrated representation is formed exploration decreases. This contextual representation is also necessary for contextual fear conditioning, so the animal must be given sufficient time to explore the chamber prior to shock delivery." [8, pp. 79]

But the notion that the human mind is an engine primarily in the business of the storage and transformation of encoded representations is a product of the 20th-century computationalism that has been fraught with basic challenges of late. In certain areas of psycholology, cognitive science, and the philosophy of mind a number of developments have provided good reason to think that this picture of cognition as the computational

[1] It is true that there has been some revisionary *language* in memory science lately: 'persistence' rather than 'storage', reconsolidation instead of 'consolidation', but I have argued elsewhere that these changes–although indicative of an admirable spirit–have so far been little other than cosmetic[19].

[2] Matthew Sanders and colleagues make this explicit, writing: "In order for context conditioning to proceed, the various stimuli in the environment must be associated with one another as the context. Therefore, we have proposed that one function of the hippocampus is to assemble a contextual representation..." [6, pp. 220].

transformations of representations is hopelessy inadequate to the data and conceptually problematic.

3 The Remembering Subject in Postcognitivist Philosophy of Mind

In recent decades, a number of researchers have successfully brought postcognitivist considerations to bear on topics in cognitive science, based on the work of several early-twentieth century philosophers. Some among these are Hubert Dreyfus, who has been applying Wittgensteinian and Heideggerian ideas to work in artificial intelligence [23], Alva Nöe, who has been fruitfully applying ideas from Merleau-Ponty and others to work on the cognitive science of perception [24], and Daniel Hutto, who has been utilizing certain insights from Wittgenstein and others in revisionist critiques of psychological explanation[25], as well Peter Hacker and Maxwell Bennett, who have devoted much, albeit controversial, work to applying Wittgenstein, critically, to cognitive neuroscience [26]. The resulting Heideggerian, Wittgensteinian, Merleau-Pontian, postcomputational, postcognitivist view of mind and cognition is far from complete or uncontroversial. These revisions have gone in different directions, among which there are inconsistencies, conceptual difficulties, and uncharted territories. Nonetheless, there are lessons to be taken from this developing characterization of mind that may help to make sense of the troublesome role that context is proving to play in models of human memory.

First, as others have noted, the postcomputational mind is not amenable to the bright line between data and algorithm that our models once drew [27, pp. 97]. This is a point that is especially emphasized by those postcognitivist researchers who owe at least moderate allegiance to connectionism and neural networks.[3] The jury is still out as to whether it makes sense, in any contexts, to talk about a neural network "representing" or "storing" anything at all. What is certain is that if such networks can rightly be said to store information, this must be in a very different

[3] Connectionism, of the medium-strength variety, is the endorsement of neural network architectures as robustly informative for those who would study minds. That is, there are few who doubt that neural network architectures have *something* to do with cognition, a least at the "subsymbolic" level, and there are few who think that neural network architectures have already solved all of our puzzles about cognition, but many among the postcognitivist vanguard stand on the middle ground here, urging that cognitive scientists and philosophers of mind have real and revisionary lessons to learn by paying close attention to the way that neural networks actually function.

sense (perhaps of both terms: 'store' and 'information') than we are used to employing to describe traditional computational architectures and processes. Since the changing weights in a neural network act as both process and product, any "information" in the network cannot be separated from the processes acting on said information, not even in principle. The notion of representation seems even worse off in this regard. There has not yet been a convincing case presented for how any part or process of a neural network can rightly be said to represent anything at all, but even if such an account were available, it would almost certainly bear little resemblance to the inert, encoded, "stored representations of context" above. A model of human memory that respects the connectionist dissolution of the data/algorithm distinction must not treat context as an item that is stored and accessed in the processes of memory.

Another distinction that has become notoriously blurry in postcognitivist frameworks for cognition is that between subject and environment. This erasure is most well known in its guises as embodied, embedded, enactive, and extended cognition, but these branches stem, in part, from philosophical considerations about the nature of the situated subject. Classical computationalism models the human cognitive system as a dimensionless, floating, executive function that stands apart from the environment it contingently inhabits, receiving input and directing output at a distance, as it were. There have now been many arguments, from many different traditions, deployed against the possibility of such distance between subject and world, but one well known example comes from Wittgenstein. Wittgenstein seeks to convince us that once we frame the problem in terms of a gap between mind and world, *we will never be able to cross it.* Wittgenstein's grocer, upon being asked for five red apples, could not successfully recognize which apples were red on the basis of an inner mental image of redness unless he had first recognized which inner mental image was red on the basis of something else [28][17]. I take it that at least some instances of the "4E" movement in cognitive science (especially among those under the rubrics of the first three E's) are explicit attempts to resolve the difficulty framed by Wittgenstein. Notably, these are of course not attempts to resolve the difficulty by *bridging* the gap in question, but rather by obviating the apparent need for an explanation that bridges this gap. The embedded, enactive mind is already in and of the world, and so there is neither a narrow nor wide gap to be bridged in our explanations. This is especially important for models of context, since the traditional treatment of context in memory models *just is* the attempt to separate out the environment from the subject and salient stimulus in question.[4]

Both of these developments, the erasure of the dividing line between data and algorithm, as well as the stretching, complicating, and blurring of the lines that separate subject from environment, are indicative of a more general tendency in postcognitivist philosophy of mind and cognitive science: a cautious and critical attitude toward modularity and informational encapsulation. This can be expressed as skepticism about whether cognitive phenomena can or should be explained in terms of "information processing"–once the unchallenged currency of inquiry in cognitive science. Even setting aside the notorious difficulties in getting from syntax to semantics (or, set in another key, from Shannon information to meaningful content), information processing explanations are fundamentally modular ordeals. A common theme among the approaches that seek to revise or reject the basic tenets of classic computationalism is the rejection of just this fundamental, encapsulated modularity: whether we are taking the Watt governor, the charts of expert navigators, or perceptrons as a starting point for understanding cognitive systems, our explanations and models will be relentlessly continuous, dynamic, and interdependent.[5]

This leaves those of us who would model cognitive phenomena in a new and difficult position: these models have traditionally been cast in terms of processes and (conceptually isolated) products, and they have taken the target to be a cognitive subject in an environment, receiving stimuli and performing operations on them as such. If these elements have been proscribed, one might wonder whether whatever is left can even go under the name *model*. This is a reasonable concern, but fortunately the business of dynamic, non-modular, fundamentally embedded and interdependent models has already gotten itself off to a solid start.

4 Modeling without Modularity in Learning and Memory

The role played by context in contemporary models has been fraught with internal and external inconsistencies. Treating the "context" of a mem-

[4] Although I am here taking my philosophical cue from Wittgenstein, this point brings to mind some particularly *Heideggerian* critiques of standard cognitive science, given Heidegger's insistence that we are always already involved in the world, as Dreyfus often points out [23]. An excellent treatment of this Heideggerian/Wittgensteinian point is a recent monograph by Lee Braver[16].

[5] This is not to say that informational encapsulation has no place in models, or that a given explanans may not fruitfully be broken up into component "systems"–rather it is something like a shift in the burden of proof, a consistently suspicious attitude toward over-modularizing that which can be done in continuous and dynamic terms.

ory as a stored, discrete item that can be tranformed by computation-like processes has not worked. Treating the context of a memory as though *it is not* one of the stored stimuli affecting the memory process has not worked any better. That is, although many elements of memory models demonstrate empirical tendencies that only imperfectly fit the encoding-storage-retrieval framework for these models, context, in particular, is the problem child. The resources traditionally available to cognitive modeling just seem inapt for whatever role context is playing. A mark of the recent postcognitivist movements in cognitive science and philosophy is the rejection of the computation-inspired modeling techniques that seem to frame exactly this problem. These movements have been spurred forward by similarly ill-behaved empirical cases from other research domains already, and it is reasonable to think that letting go of cognitivist assumptions about modularity, processes and products, and subjects and environments, may resolve the otherwise intractable difficulties that memory researchers continue to encounter in trying to fit the square nature of context into the round holes of their explanatory models. New models and explanatory resources have begun to resolve impasses in other areas of cognitive science; there are good reasons, as evinced above, to think that context is a *particularly* problematic candidate for classically cognitivist models. Yet it can seem that without the usual toolkit of information, modular systems, and an executive subject that performs processes on stimuli received from the environment, we are simply not left with a model–or, for that matter, perhaps not even with an explanation.

This is not a trivial or shallow concern. For many cognitive phenomena, we have a deeply felt need for the revelation of an underlying process made up of component parts and processes, if we are to have an explanation of these phenomena at all. The ability to *model* a phenomenon at all is sometimes taken *to be* the ability to underwrite a phenomenon by charting out the involved information processing in terms of modules and processes. This identification, however, should be resisted. There are certainly explananda that merit mechanical, subcomponent, information-processing explanantia, but it is carelessly myopic synecdoche to treat modeling *per se* as identical to this. Explanations come in many forms.[6] Fortunately, we need not unravel the ultimate nature of explanation before moving forward with revised frameworks for capturing context in a model of human memory. Explanations and models that fit our desiderata

[6] Indeed, adequately characterizing just what explanation consists in is a notorious enough difficulty that I take it that any identification of explanation and information processing-style models can only be a tacit, uncritical move.

are abundant nearby, and these can serve as existence proofs that what we seek to do in constructing these postcognitivist models can, in fact, be done.

The models and explanations that I have in mind, of course, are predominantly those that hail from the "dynamic systems" family of methods and resources. Researchers in the dynamic systems tradition are interested in nothing if they are not interested in models–and often, in models of cognitive phenomena–yet these models persistently eschew the computational trappings that postcognitivist philosophy and science have rejected. That is, the dynamic systems model is built on continuous, interdepenent mathematical tools that frame elements which co-vary, co-consist, and develop continuously in time. Patterns are studied even without being conceptually isolated, elements of the environment are modeled without separating them from a dimensionless executive control unit; any process/product separation questions are simply not well-formed queries of a dynamic model. Nor are stimuli standardly treated as "signals" that have been received or retrieved. These models' primary business is to avoid such things and yet maintain explantory power and rigor.

John Spencer, Michael Thomas, and James McClelland compiled and edited an excellent volume of such models recently, the target domain for which was human development [10]. The dynamic systems and connectionist researchers represented in this volume tackle a number of complex phenomena associated with development, and in each case succeed in building a dynamic and explanatorily powerful model without running afoul of postcognitivist insights. The special role allotted, explicitly and implicitly, to context in dynamic models of child development is nicely characterized by Paul van Geert and Kurt W. Fischer in a summary article within the Spencer et al. volume, who encourage us to follow the dynamic models and treat person and context together as a quasi-unified soft assembly:

> "If we then speak about the development of the child, we actually refer to long-term changes in the child and automatically also to the corresponding long-term changes in the child's contexts. These changes refer primarily to patterns of correlation over time (that is, in a single person's life trajectory) and not to identifiable entities. That is, development refers to a person-context assembly throughout the life span; contexts are no longer to be seen as independent variables or circumstances in which the person can be placed at pleasure, according to the whim of a researcher who

treats such contexts as independent variables, the effect of which has to be estimated over many subjects" [29, pp. 327].

The dynamic development models summarized by van Geert and Fischer tend toward either the very narrow time scale or the very broad, with target phenomena usually taking after either real-time development and reaction to ongoing stimuli or taking after longitudinal, long-term studies of development in the sense of childhood development. These success stories thus flank the usual time scales of memory studies, which tend to fall somewhere in between. This is a challenge for the memory scientist who would model context using dynamic systems resources, but there is no reason to think that the challenge cannot be met. After all, the long-term development studies modeled by the researchers in the Spencer et al. volume are already operating at a remove from real-time interactions. Memory scientists will, in most cases, need to follow suit, building dynamic models based on data extracted from large numbers of remembering cases across time. In such models, *context* will be neither a stimulus nor the unassuming background against which stimuli are discerned, but will follow the soft assembly framework van Geert and Fischer describe. The line one draws between a system and its environment, in a dynamic model, is as specific or nonspecific as the line-drawer pleases. This also obviates the need for memory scientists to decide whether 'context' refers solely to external stimuli, solely to *internal* stimuli, both, or whether the term needs subscripts: memory scientists are well aware that context is presently used, sometimes imperiously, in each of these ways [30].

This also presents a challenge for dynamic systems modelers. As Karl Newell and colleagues point out, differences in time scales provide real, but not insurmountable, conceptual difficulties for dynamic system modelers [31]. The dynamic systems research program is still very much a new endeavor. Given also the fact that it is interdisciplinary in the in the truest sense–it exists *between* disciplines, in methods and approaches that do not properly belong to any of the many disciplines to which these have been applied–one can expect that applying these methods to new cases will produce some puzzlement. For example, *declarative* memory studies rely on linguistic data, which in turn require clever and still controversial techniques to model using dynamic systems resources. Human language *really is*, of course, a largely discrete, modular, information-laden endeavor, and hence experiments dependent on linguistic data can sit uneasily with the emerging dynamic systems paradigm.[7] This obstacle is easily avoided–or

[7] This is not to suggest that language-involving cognitive phenomena are necessarily beyond the ken of dynamic systems modeling techniques, only that there is good

at least delayed–though, by choosing non-linguistic data as explananda to be modeled dynamically within memory science. These are easy enough to come by, even in studies of human memory: non-declarative memory has a rich array of experimental data, and even "declarative" varieties of memory like semantic and episodic memory are often the subject of studies that do not furnish any linguistic *data*. Indeed, dynamic and connectionist models of neuronal function have been common for some time, but even when researches treat hippocampal networks in terms of attractor dynamics, these very same researchers still tend to speak in terms of *stored representations* when they ascend to the cognitive level to discuss context memory [32][33]. One exception is the recent work of Guoqi Ling and colleagues, who have designed an attractor network model by working backward from a novel energy function *designed* for cognitive plausibility, rather than basing the the model on neural dynamics and then encountering an impasse when confronting cognitive-level phenomena [34]. The modeling framework presented by Ling *et al.* has not yet been applied to context in particular, but their characterization of reconstructable (*and nonbinary*) "memory patterns" as stable equilibrium points of dynamical attractor networks already invites a context-rememberer soft assembly.

So a way forward for capturing context in our explanations of human memory has become clear and equipped with reasonably successful resources. Context does not fit the standard computational models that cognitivist frameworks have traditionally offered to memory scientists. Researchers have attempted to tweak these models, but it is clear that tweaking is insufficient. The problematic bad fit between target phenomena and modeling framework makes a significant revision seem all but inevitable, and happily the postcognitivist front in philosophy and cognitive science provides significant resources for the endeavor. There are obstacles and difficulties on all sides–not just because of the intermediate time scale involved in most memory studies, but also because memory science itself does not yet have the language to characterize its phenomena without implicitly invoking computationalist elements like encoding, storage, and retrieval. Furthermore, an interesting subsection of human memory revolves around linguistic abilities, which are resistant to dynamic modeling. Nonetheless, we have good reason to try to construct significantly revised modeling frameworks for human memory and the clearly crucial role that context plays in these processes according to these dynamic modeling techniques. Memory models at the cognitive level need not build on

---

reason to expect successful models of these to be among the late-comers to the dynamic systems table.

computationalist foundations. This can avoid the problematic treatment of context as "content" that we find in work like that of Spencer and Naftali above. We already have many of the needed resources available, even if the initial dynamic models of human remembering phenomena are constrained to non-declarative cases, or to non-linguistic studies of semantic or episodic memory.[8] In thus relinquishing the trappings of the cognitivist framework, context–which is so clearly fundamental to the operations of human remembering–may be adequately modeled by those who seek to explain human memory.

References

1. Michael S Fanselow. Context: Whats so special about it. *Science of memory: Concepts*, pages 101–105, 2007.
2. Mark E Bouton. The multiple forms of context in associative learning theory. *The mind in context*, pages 233–258, 2010.
3. SM Smith. Context: A reference for focal experience. *Science of memory: Concepts*, pages 111–114, 2007.
4. Catherine E Myers and Mark A Gluck. Context, conditioning, and hippocampal rerepresentation in animal learning. *Behavioral neuroscience*, 108(5):835, 1994.
5. Patricia Matus-Amat, Emily A Higgins, Ruth M Barrientos, and Jerry W Rudy. The role of the dorsal hippocampus in the acquisition and retrieval of context memory representations. *The Journal of neuroscience*, 24(10):2431–2439, 2004.
6. Matthew J Sanders, Brian J Wiltgen, and Michael S Fanselow. The place of the hippocampus in fear conditioning. *European journal of pharmacology*, 463(1):217–223, 2003.
7. Wesley D Spencer and Naftali Raz. Differential effects of aging on memory for content and context: a meta-analysis. *Psychology and aging*, 10(4):527, 1995.
8. Michael S Fanselow. Contextual fear, gestalt memories, and the hippocampus. *Behavioural brain research*, 110(1):73–81, 2000.
9. Kees De Bot, Wander Lowie, and Marjolijn Verspoor. A dynamic systems theory approach to second language acquisition. *Bilingualism language and cognition*, 10(1):7, 2007.
10. John P Spencer, Michael SC Thomas, and JL McClelland. Toward a unified theory of development. *JP Spencer, MSC Thomas, & JL McClelland (Eds.)*, pages 86–118, 2009.
11. Edwin Hutchins. *Cognition in the Wild*. MIT press, 1995.
12. Lawrence Shapiro. *Embodied cognition*. Routledge, 2010.
13. Kourken Michaelian and John Sutton. Distributed cognition and memory research: History and current directions. *Review of philosophy and psychology*, 4(1):1–24, 2013.

[8] As characterized above, much of these experiments take the form of tracking the help and hindrance of "context cues" for declarative memory–simple context cue experiments such as those that study positive or negative intereference in semantic memory may be amenable, for example, in the same way that "cognitive control" was found to be in a recent study by Todd Braver [35].

14. Selmer Bringsjord. Computationalism is dead; now what? *Journal of Experimental & Theoretical Artificial Intelligence*, 10(4):393–402, 1998.
15. Anthony Chemero. *Radical embodied cognitive science*. MIT press, 2011.
16. Lee Braver. *Groundless grounds: a study of Wittgenstein and Heidegger*. Mit Press, 2012.
17. David G Stern. Models of memory: Wittgenstein and cognitive science. *Philosophical psychology*, 4(2):203–218, 1991.
18. Jens Brockmeier. After the archive: remapping memory. *Culture & Psychology*, 16(1):5–35, 2010.
19. Ian OLoughlin. The persistence of storage: Language and concepts in memory research. *European Perspectives on Cognitive Science*, 2011.
20. George Houghton, Tom Hartley, and David W Glasspool. The representation of words and nonwords in short-term memory: Serial order and syllable structure. *Models of short-term memory*, pages 101–127, 1996.
21. Scott D Slotnick and Chad S Dodson. Support for a continuous (single-process) model of recognition memory and source memory. *Memory & Cognition*, 33(1):151–170, 2005.
22. Kourken Michaelian. Generative memory. *Philosophical Psychology*, 24(3):323–342, 2011.
23. H.L. Dreyfus. *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press, 1992.
24. Alva Noë. *Action in perception*. MIT press, 2004.
25. Daniel Hutto. Folk psychology as narrative practice. *Journal of Consciousness Studies*, 16(6-8):9–39, 2009.
26. Max R.. Bennett and Peter Michael Stephan Hacker. *Philosophical foundations of neuroscience*. Blackwell Publishing, 2003.
27. Andy Clark. *Mindware: An introduction to the philosophy of cognitive science*. Oxford University Press, Inc., 2013.
28. Ludwig Wittgenstein. *Philosophical investigations*. John Wiley & Sons, 2010.
29. PLC Van Geert, KW Fischer, JP Spencer, MSC Thomas, and J McClelland. Dynamic systems and the quest for individual-based models of change and development. *Toward a Unified Theory of Development: Connectionism and Dynamic Systems Theory Reconsidered*, pages 313–336, 2009.
30. E Eich. Context: Mood, memory, and the concept of context. *Science of memory: Concepts*, pages 107–110, 2007.
31. Karl M Newell, Yeou-Teh Liu, and Gottfried Mayer-Kress. Time scales in motor learning and development. *Psychological review*, 108(1):57, 2001.
32. Nicolas Brunel. Hebbian learning of context in recurrent neural networks. *Neural computation*, 8(8):1677–1710, 1996.
33. Tom J Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O'Keefe. Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876, 2005.
34. Guoqi Li, Kiruthika Ramanathan, Ning Ning, Luping Shi, and Changyun Wen. Memory dynamics in attractor networks. *Computational intelligence and neuroscience*, 2015, 2015.
35. Todd S Braver. The variable nature of cognitive control: a dual mechanisms framework. *Trends in cognitive sciences*, 16(2):106–113, 2012.