COMMENTARY

The Conjunctive Trace

Blake A. Richards[1] and Paul W. Frankland[1,2,3,4]*

ABSTRACT: Memories serve to establish some permanence to our inner lives despite the fleeting nature of subjective experience. Most neurobiological theories of memory assume that this mental permanence reflects an underlying cellular permanence. Namely, it is assumed that the cellular changes which first occur to store a memory are perpetuated for as long as the memory is stored. But is that really the case? In an opinion piece in this issue of Hippocampus, Aryeh Routtenberg raises the provocative idea that the subjective sense of memory persistence is not in fact a result of persistence at the cellular level, rather, that "supple synapses" and multiple "evanescent networks" that are forever changing are responsible for our memories. On one level, his proposal could be construed as a radical challenge to some of our most fundamental theories of the neurobiology of memory, including Donald Hebb's proposal that memories are stored by networks that strengthen their connections to increase the likelihood of the same activity patterns being recreated at a later date. However, it could also be seen as a moderating call, a call for a greater acknowledgement of the dynamic, stochastic, and distributed nature of neural networks. In this response to Routtenberg's article, we attempt to provide a clarification of the dividing line between these two interpretations of his argument, and in doing so, we provide some overview of the empirical evidence that bears on this subject. We argue that the data that exists to date favors the more moderate interpretation: that memory storage involves a process in which activity patterns are made more likely to reoccur, but that an under-appreciated reality is that mnemonic traces may continue to change and evolve over time. © 2013 Wiley Periodicals, Inc.

KEY WORDS: memory storage; consolidation; amygdala; hippocampus; conjunctive encoding; disjunctive encoding; stable trace; dynamic trace

"You cannot step into the same river twice." (Heraclitus)

Memories must be stored somewhere. For us to remember anything there must be a process whereby alterations to our nervous system ensure that later activity patterns reflect earlier activity patterns (by not recognizing this, we would veer off into dualism). However, the exact manner in which later activity patterns reflect earlier ones is not obvious. Memories could be stored by increasing the likelihood that a particular activity pattern would reappear. Alternatively, memories could be stored by decreasing the likelihood that a particular activity pattern would reappear. Neuroscientists often assume the former, but this is not a given, as Routtenberg points out.

To distinguish between these two possibilities, we will develop a simple mathematical formulation. Consider the concrete example of a contextual fear memory experiment wherein the circuits of the amygdala are involved in encoding fearful memories (Fanselow and LeDoux, 1999). When the animal is first placed in the conditioning chamber, it experiences the pairing of the cues with a shock, which we will refer to as experience A_1. During A_1, a particular set of neurons in the amygdala will be active, which we will call α_1. What happens now if the animal is placed back in the chamber? Specifically, what set of neurons in the amygdala, α_2, are activated during the experience A_2 (Fig. 1)? In the case where no memory has been stored, the probability of any given cell, i, being a member of α_2 will be independent of its being active during A_1:

$$P(i \in \alpha_2 | i \in \alpha_1) = P(i \in \alpha_2) \text{ (No memory)}$$

It should be recognized that, as given, this is not strictly speaking a guarantee of no memory storage. When considering a distributed code the probability of individual cells being activated is not enough to completely capture the situation. No memory storage occurs when $P(\alpha 2 \mid \alpha 1) = P(\alpha 2)$, because the mutual information between $\alpha 1$ and $\alpha 2$ is equal to zero in this circumstance (Brunel and Nadal, 1999). Although this would imply that $P(i \in \alpha 2 \mid i \in \alpha 1) = P(i \in \alpha 2)$, the converse is not always true. Nonetheless, we will utilize this simplified, cell-by-cell framework to aid the discussion while recognizing the mathematical limitations of doing so.

The situation of no memory allows us to then distinguish two possible manners in which a memory could be stored. In the first, which we will refer to as "conjunctive encoding," the probability of a neuron in the amygdala being active during exposure A_2 increases if it was active during A_1:

$$P(i \in \alpha_2 | i \in \alpha_1) > P(i \in \alpha_2) \text{ (Conjunctive encoding)}$$

Note that a conjunctive code is ultimately what is implied by the Hebbian aphorism that "neurons that fire together wire together (Hebb, 1949)." Therefore, it is worth noting that conjunctive encoding lies at the heart of most modern theorizing about the neurobiology of memory. The extreme version of this is the

[1] Program in Neurosciences and Mental Health, The Hospital for Sick Children, Toronto, Ontario, Canada, M5G 1X8; [2] Department of Psychology, University of Toronto, Toronto, Ontario, Canada M5S 3GM; [3] Department of Physiology, University of Toronto, Toronto, Ontario, Canada M5S 1A8; [4] Institute of Medical Science, University of Toronto, Toronto, Ontario, Canada M5S 1A8

Grant sponsor: Wellcome Trust Senior Research Fellowship in Basic Biomedical Science

*Correspondence to: Paul Frankland, Program in Neurosciences and Mental Health, The Hospital for Sick Children, Toronto, Ontario, Canada, M5G 1X8. E-mail: paul.frankland@sickkids.ca

Accepted for publication 12 December 2012
DOI 10.1002/hipo.22089
Published online in Wiley Online Library (wileyonlinelibrary.com).

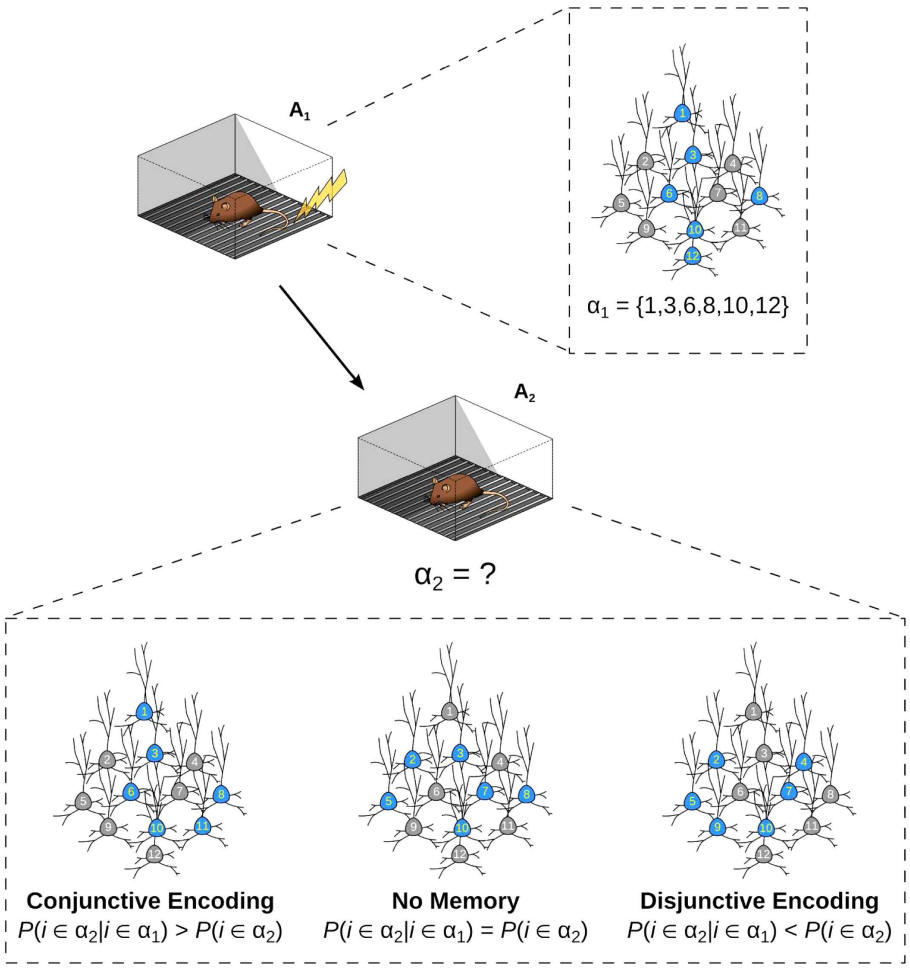

FIGURE 1. Conjunctive versus disjunctive encoding of a fear memory. During a fear conditioning experience, A_1, a particular subset of cells, α_1, is activated in the amygdala. Here, for illustrative purposes we show a schematic of the circuit with a total of twelve cells, six of which are contained in α_1 (top right). When the animal is placed back in the conditioning chamber, this experience A_2 will lead to the activation of a new subset of cells, α_2. If no memory was stored then α_2 will be independent of α_1, and the new subset of cells that is activated will randomly overlap with the subset that was active during conditioning (bottom middle). In contrast, if a memory is stored then α_2 must be dependent on α_1, but this can be in one of two ways. Either the cells that were active during conditioning are more likely to be reactivated during recall (conjunctive encoding, bottom left) or the cells that were active during conditioning are less likely to be reactivated during recall (disjunctive encoding, bottom right). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com]

situation in which the exact same activity pattern is recreated during retrieval [i.e., $P(i \in \alpha_2 \mid i \in \alpha_1) = 1$].

In contrast, another way in which a memory could be stored would be by "disjunctive encoding." In this scenario, the probability that a neuron will be active during A_2 is actually lowered as a result of its being active during A_1:

$$P(i \in \alpha_2 | i \in \alpha_1) < P(i \in \alpha_2) \text{ (Disjunctive encoding)}$$

The extreme case of disjunctive encoding is when the cells that were active during encoding are never reactivated [i.e., $P(i \in \alpha_2 \mid i \in \alpha_1) = 0$]. Although disjunctive encoding is often not considered by researchers when thinking about the neurobiology of memory, it is in fact a perfectly legitimate method for information storage.

Some initial observations are important. First, this basic mathematical framework can, of course, be applied to other tasks beyond fear conditioning and to other circuits beyond the amygdala. Second, in describing these two encoding methods we are assuming that roughly the same number of cells get activated during A_1 and A_2, which means that we are not considering the degenerative cases where the entire population's activity is ramped up or down (i.e., a situation where either all cells become more likely to be in α_2 or all cells become less likely to be in α_2). This assumption is based on evidence that retrieval and encoding activate similar numbers of neurons (Guzowski et al., 1999; Han et al., 2007). Third, given that neural networks exhibit a great deal of stochasticity (Faisal et al., 2008; Shadlen and Newsome, 1994; Shadlen and Newsome, 1998; Faisal et al., 2008), the two extreme scenarios (exact reactivation or complete exclusion of an activity pattern) are almost certainly not how brain circuits operate and also not something which any neuroscientist would advocate. As such, we suggest that any argument against them is ultimately an argument against a straw man. Fourth, there is nothing preventing the brain from combin-

ing conjunctive and disjunctive encoding strategies, both within and across different networks; these encoding strategies are not mutually exclusive though they do lead to very different experimental predictions.

We can now begin to clarify some of Routtenberg's arguments by identifying two important questions he raises:

- **Q1**. To what extent are conjunctive versus disjunctive encoding strategies used by the brain?
- **Q2**. If we consider additional recall experiences beyond a single event (i.e., experiences A_3, A_4, . . ., A_N) is there a consistent set of cells engaged in encoding the memory across all experiences or does the memory trace change over time?

A critical point of clarification is that these two questions are distinct. We could have a scenario wherein the brain utilizes conjunctive codes to store memories, but over time and across multiple retrieval events the traces are gradually altered, such that there is no overlap between the initial activity pattern and the memory traces underlying retrieval at distant time-points.

We will remain agnostic as to exactly how Routtenberg would answer the two questions above phrased in this way. Some of his arguments seem to promote the idea of disjunctive codes over conjunctive codes, while others seem to be a demand for a greater recognition of the dynamic nature of memory traces. We would argue that the existing literature suggests that the primary mode of encoding is conjunctive. Although disjunctive encoding is an interesting possibility that may be employed by the brain sometimes, we are not convinced by Routtenberg's arguments that it is time to abandon this central conception of conjunctive memory storage. However, Routtenberg may well be correct that there is a great deal of dynamism to mnemonic engrams, such that long-term memory traces are ever-changing.

EVIDENCE FOR CONJUNCTIVE MEMORY ENCODING

Perhaps the clearest illustration of conjunctive memory encoding comes from a study by Reijmers et al. (2007). They generated a transgenic mouse model (the TetTag mouse), wherein the expression of tetracycline-transactivator is driven by the activity-regulated c-Fos promoter, allowing doxycycline controlled conditional tagging of active cells via tau-LacZ. Using this approach, the authors conducted a study to examine exactly the situation described by our basic framework. Mice were trained in a contextual fear memory task while temporarily not receiving doxycycline, leading to LacZ expression in those cells active during the initial fear memory encoding (i.e., providing a tag of cells in α_1). Following this, doxycycline was reintroduced and 3 days later the animals were tested in the context. One hour after testing, mice were sacrificed and their brains were stained for the activity-driven gene Zif268/Egr-1(i.e., providing a tag of cells in α_2). By examining double tagged cells, Reijmers et al. (2007) were able to determine whether the number of cells that were incorporated into both α_1 and α_2 was above chance (indicating a conjunctive code) or below chance (indicating a disjunctive code). The result was that animals given the training and the recall experience (but not control groups who were only exposed to the environment without shock) exhibited a number of cells activated by both the initial experience and the retrieval experience that was significantly greater than chance. This provides a very clear indication that a conjunctive code is used in the amygdala for storing fear memories.

Evidence of a conjunctive code for fear memories in the amygdala also comes from a 2009 paper by Han et al. (2009). Previous studies had demonstrated that overexpression of CREB (cAMP response element-binding protein) in lateral amygdala pyramidal cells increased the likelihood of those cells being active during encoding (i.e., α_1) and subsequent recall (i.e., α_2) of a fear memory (Han et al., 2007). One possible interpretation for this observation is that the CREB overexpressing cells are preferentially incorporated into a conjunctive memory trace (Josselyn, 2010). If this interpretation is correct, then selective elimination of cells that express high levels of CREB during encoding should also selectively eliminate the associative memory stored in the amygdala. To determine whether this was the case, Han et al. used an inducible diphtheria toxin receptor strategy to selectively ablate either CREB overexpressing cells or a random, similar-sized, subpopulation of lateral amygdala cells. When CREB overexpressing cells were selectively ablated, the fear memory (measured by freezing levels) was eliminated. In contrast, when randomly selected cells were ablated the memory was intact. Hence, ablation of CREB overexpressing cells, which are active during encoding, leads to forgetting. This provides a clear, causal demonstration of a conjunctive code in the amygdala, because a disjunctive code would in fact predict the opposite effect, that is, ablating cells that are active during encoding (i.e., ablating α_1) should have little to no effect on the memory (since α_2 would not contain the ablated cells).

What about memory systems other than the amygdala? There is also evidence for conjunctive codes in the hippocampus. An early demonstration of conjunctive coding was reported by Guzowski et al. (1999). In this article, the authors described a then new immediate early gene mapping technique that they had developed known as cellular compartment analysis of temporal activity by fluorescent in situ hybridization (catFISH). This method takes advantage of the fact that the immediate early gene *Arc* exhibits an intracellular shift in the location of its expression following cellular activation, enabling comparisons of cells activated at two different time points. Guzowski et al. examined the distribution of active cells in the CA1 region of hippocampus in animals placed into two different contexts or the same context twice. Roughly, the same number of cells were activated in each group, but animals placed in the same context twice had 90% overlap in the activity patterns compared to only 26% for animals placed into different contexts (corresponding to chance levels). Therefore, this greater than chance probability of reactivation is indicative of a conjunctive code for contextual memories in CA1. Similar results were obtained by Vazdarjanova and Guzowski (2004) when they used catFISH to examine this issue in more depth (Vazdarjanova and Guzowski, 2004). In this study,

the percentage of coactivated cells in the animals placed in the same context twice was slightly lower but still accounted for roughly 70% of the total number of active cells in CA1. A similar pattern was observed in CA3, indicating that both regions employ conjunctive encoding for memory storage.

In his article, Routtenberg suggests that Vazdarjanova and Guzowski's (2004) results are evidence of a disjunctive encoding scheme in the hippocampus, because 30% of the active cells were only activated by one but not both of the experiences. However, this is an attack on the straw man of the extreme conjunctive encoding scenario, wherein the exact same set of cells is reactivated. In both Guzowski et al. (1999) and Vazdarjanova and Guzowski (2004) the number of coactivated cells observed was above chance. Mathematically speaking, this tells us that even though the exact pattern was not recreated, since 30% of the cells were not coactivated, the probability of a given cell being a member of α_2 is higher than chance (and not lower than chance) if it was in α_1; this places us explicitly into a conjunctive encoding framework for contextual memories. There are several reasons why the original pattern of activity is unlikely to be faithfully recreated. First, the stochastic nature of neural activity would make it highly unlikely for the exact same pattern to reappear (Tomko and Crapper, 1974; Shadlen and Newsome, 1994; Faisal et al., 2008). Second, each contextual experience is unique, with slightly different sensory inputs, endogenous circumstances, and so forth, and one of the experiences involves a recall event while the other does not. Instead of asking whether all or most of the cells are activated by both experiences, the critical question that must be asked is this: is the number of coactivated cells above or below chance?

A recent study by Liu et al. (2012) provides causal evidence for the conclusion that conjunctive encoding is used to store contextual memories in the hippocampus (Liu et al., 2012). These authors used viral infection of channelrhodopsin-2 (ChR2) in the TetTag mouse described above to tag and reactivate cells in the dentate gyrus (DG) that were active during fear conditioning. Conditioning was performed in one context, following which the animals were placed in a different, previously encountered, neutral context. What predictions would conjunctive versus disjunctive encoding have for the effects of optical activation of the tagged cells in this neutral context? In a conjunctive encoding situation, one would expect that the DG cells that were active during fear memory encoding would be part of the memory trace for the conditioning context. Therefore, we would predict that activating these cells would in fact artificially induce recall, which would lead to fearful behavior in the neutral context. In contrast, in a disjunctive encoding situation, the neurons in the DG that were active during fear memory encoding would not be a part of the memory trace for the conditioning context. As such, activating these cells would not induce recall, so the animals' behavior in the neutral context would not indicate any sense of fear. The findings from Liu et al. (2012) were very clearly those predicted by conjunctive encoding: when the ChR2 expressing cells were optically activated the animals exhibited significant freezing behaviour in the otherwise neutral context.

Routtenberg cites Liu et al. (2012) in his article, and recognizes that others may construe their results as evidence "...that memory is stored in selected cells..." But, he goes on to suggest that this is not a necessary conclusion and that we can reinterpret Liu et al.'s findings if we instead consider a different possibility. To quote: "Information is ... not actually stored in the sense of an extant residuality. Thus, the physical indeterminancy of the engram is given its biological reality in the retrieval process, as for example the retrieved freezing behavior in the Liu et al. report which is taken for the fear engram." Here, Routtenberg seems to be arguing that Liu et al. were not in fact inducing recall by activating a fear memory engram. Rather, the freezing of the animals in the neutral context during optical stimulation was caused by some unspecified mechanism, one that does not involve information having been stored via persistent modifications to the hippocampal network. Importantly, Liu et al. conducted control experiments to show that in the absence of conditioning or the absence of optical activation the animals did not freeze, which demonstrated that both a fear memory and reactivation of the cells was necessary to induce freezing. We would therefore argue that Routtenberg's assessment of this article is incorrect, and it underlines the problem with his arguments against conjunctive encoding more generally. Although, a disjunctive encoding scheme for memory retrieval is interesting and theoretically possible it leads to specific predictions that can be tested. The evidence that we have reviewed here is clear: in both the amygdala and hippocampus a conjunctive encoding scheme appears to be the primary mode of encoding and the data does not indicate disjunctive encoding. The situation may be different in other mnemonic circuits, or for other tasks, but this phenomenon must be demonstrated first to be worthy of further consideration.

STABILITY VERSUS DYNAMISM IN LONG-TERM MEMORY TRACES

Despite the fact that some persistent changes must occur in order to store information there is nothing preventing the brain from continually inducing new modifications in order to alter and update the substrates of memory storage. To make this concrete, we consider again the mathematical framework introduced above, but we now expand the experiences and activity patterns beyond a single encoding and recall event, and instead consider a set of N experiences $A_1, \ldots, A_N$ and their corresponding activity patters $\alpha_1, \ldots, \alpha_N$. Even though the existing evidence in the amygdala and hippocampus suggests that α_1 and α_2 typically overlap significantly more than chance, the information contained in these altered probabilities could be propagated to later experiences in two different manners. In one case, which we will call a "stable trace", a core subset of cells would continue to have a very high probability of being reactivated during each experience, such that these cells would continue to be members of α_3, α_4, and so forth, even though other cells may be randomly active or inactive. The other possibility, which we will call a "dynamic

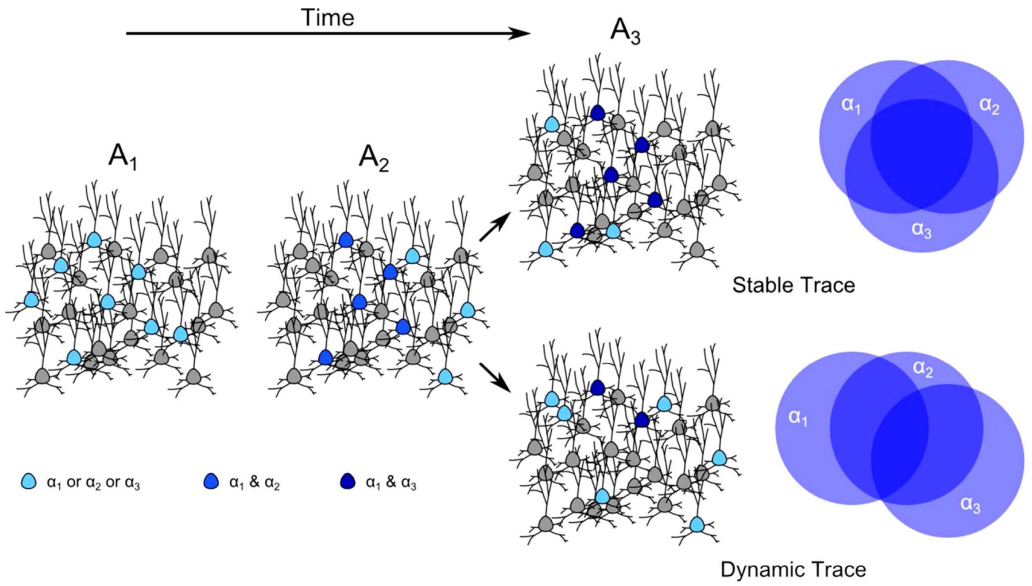

FIGURE 2. Stable versus dynamic long-term memory traces. Over extended periods of time a conjunctive memory trace may either stabilize or continue to evolve. In this example, we show the progression of a memory trace in a schematic circuit over three experiences, A_1, A_2, and A_3. If conjunctive encoding is used to store the memory, then the overlap in the activity patterns for consecutive experiences must be greater than chance, that is, the overlap between α_1 and α_2, or α_2 and α_3, must be above chance. In this schematic, there are twenty-six cells, and eight are activated during each experience, making the number of coactive cells expected by chance approximately two to three cells (8/26 × 8/26 × 26 cells = 2.46 cells of overlap on average). However, whether the overlap between non-consecutive activity patterns is greater than chance depends on whether the memory trace is stable or dynamic. In the case of a stable trace (top right), the heightened probability of being activated propagates within the same subset of cells that was originally active during A_1, leading to a similar degree of overlap in the activity patterns for consecutive and non-consecutive traces (as illustrated by the Venn diagram). The other possible scenario is a dynamic trace (bottom right), wherein the heightened probability of being activated gradually propagates to a new subset of cells. The result is that although A_2 and A_3 still exhibit greater than chance overlap (thereby preserving the memory) the overlap between A_1 and A_3 will diminish (as illustrated by the Venn diagram). Over extended periods of time this could result in a completely altered trace that does not resemble the original trace that was laid down during encoding. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com]

trace", is for the information to be maintained via an updating of the memory trace over time, such that α_2 may overlap with α_1 to a large extent, and α_3 may overlap with α_2 to a large extent, but α_3 and α_1 could be quite different (Fig. 2).

We believe that Routtenberg's proposal can be broadly construed as an objection against trace stability. He presents a series of interesting neurobiological observations that could fit quite well with a dynamic trace scenario. The malleability of dendritic spines (Fischer et al., 1998), the continual renewal of proteins at postsynaptic sites (Gray et al., 2006) and the importance of NMDA receptors beyond the initial encoding event for some types of learning (Shimizu et al., 2000; Santini et al., 2001) all fit with a situation wherein the brain continually updates its memory traces. Routtenberg points to an especially interesting consequence of this model of long-term memory storage: if memory traces are continually altered over time then perhaps consolidation is not in fact the crystallization of the original synaptic architecture used to encode the information, but rather a process of information redistribution and re-encoding. This process could help to protect memories by generating highly distributed traces or by generating multiple traces, as Routtenberg suggests. If true, then consolidation could help to prevent interference or memory loss because it effectively "hedges" the information storage against conflicts arising from other changes induced by later experiences. This proposal represents an important alternative to the idea of a stable trace, and we feel that it should be carefully considered by researchers who study long-term memory. Indeed, recent evidence suggests that CA1 utilizes a dynamic trace mechanism, one that could aid in the encoding of elapsed time (Mankin et al., 2012). Furthermore, there is plenty of evidence that as memories mature they are reorganized at the systems level (Frankland and Bontempi, 2005), with the circuits engaged at the time of encoding (or shortly thereafter) differing from those engaged when the equivalent memory is retrieved at more remote time-points (Wheeler et al., 2013). This process of systems consolidation fits quite nicely with Routtenberg's conceptualization of a dynamic trace. It is conceivable that remote memories are altered, distributed across multiple neural circuits, and possibly even duplicated precisely because the long-term consolidation of memory involves flexible, evanescent traces that encode the information. If so, then synapses must indeed remain supple to provide an effective physical mechanism for propagating memory.

CONCLUSION

The quote from Heraclitus that we included at the beginning of this response captures the transitory nature of any physical system, including the brain. When we recall our past,

we can never replicate the internal neural states that we were originally experiencing; this is, perhaps, a truism. Nonetheless, there are different ways in which memory storage and recall may involve a divergence away from the activity patterns that occurred during encoding. In his article, Routtenberg challenges two components of conventional wisdom: (1) that memory storage involves persistent changes that encourage the reactivation of previous activity patterns (the conjunctive trace), and (2) that long-term memory consolidation is a process that stabilizes the cellular or synaptic changes that were originally instigated by encoding (the stable trace). In general, prompting scientists to reconsider these potentially dogmatic assumptions is something which we applaud. However, we have argued here that the challenge to the first component, although interesting, is not supported by the existing literature—on this point the standard wisdom is likely to be correct. Despite this, Routtenberg's challenge to the second component is an important theoretical contribution that deserves further consideration. The apparent permanence of our cherished memories may not in fact be reflected by any long-lasting permanence at the neurobiological level.

REFERENCES

Brunel N, Nadal J-P. 1998. Mutual information, Fisher information, and population coding. Neural Comput 10:1731–1757.

Faisal AA, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. Nat Rev Neurosci 9:292–303.

Fanselow MS, LeDoux JE. 1999. Why we think plasticity underlying pavlovian fear conditioning occurs in the basolateral amygdala. Neuron 23:229–232.

Fischer M, Kaech S, Knutti D, Matus A. 1998. Rapid actin-based plasticity in dendritic spines. Neuron 20:847–854.

Frankland PW, Bontempi B. 2005. The organization of recent and remote memories. Nat Rev Neurosci 6:119–130.

Gray NW, Weimer RM, Bureau I, Svoboda K. 2006. Rapid redistribution of synaptic PSD-95 in the neocortex in vivo. PLoS Biol 4:e370.

Guzowski JF, McNaughton BL, Barnes CA, Worley PF. 1999. Environment-specific expression of the immediate-early gene Arc in hippocampal neuronal ensembles. Nat Neurosci 2:1120–1124.

Han J-H, Kushner SA, Yiu AP, Cole CJ, Matynia A, Brown RA, Neve RL, Guzowski JF, Silva AJ, Josselyn SA. 2007. Neuronal competition and selection during memory formation. Science 316:457–460.

Han J-H, Kushner SA, Yiu AP, Hsiang H-L, Buch T, Waisman A, Bontempi B, Neve RL, Frankland PW, Josselyn SA. 2009. Selective erasure of a fear memory. Science 323:1492–1496.

Hebb DO. 1949. The Organization of Behavior: A Neuropsychological Theory. New York, Wiley.

Josselyn S. 2010. Continuing the search for the engram: Examining the mechanism of fear memories. J Psychiatry Neurosci 35:221–228.

Liu X, Ramirez S, Pang PT, Puryear CB, Govindarajan A, Deisseroth K, Tonegawa S. 2012. Optogenetic stimulation of a hippocampal engram activates fear memory recall. Nature 484:381–385.

Mankin EA, Sparks FT. Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. 2012. Neuronal code for extended time in the hippocampus. PNAS 109:19462–19467.

McClelland JL, Botvinick MM, Noelle DC, Plaut DC, Rogers TT, Seidenberg MS, Smith LB. 2010. Letting structure emerge: Connectionist and dynamical systems approaches to cognition. Trends Cogn Sci 14:348–356.

Reijmers LG, Perkins BL, Matsuo N, Mayford M. 2007. Localization of a stable neural correlate of associative memory. Science 317:1230–1233.

Santini E, Muller RU, Quirk GJ. 2001. Consolidation of extinction learning involves transfer from NMDA-independent to NMDA-dependent memory. J Neurosci 21:9009–9017.

Shadlen MN, Newsome WT. 1994. Noise, neural codes and cortical organization. Curr Opin Neurobiol 4:569–579.

Shadlen MN, Newsome WT. 1998. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. J Neurosci 18:3870–3896.

Shimizu E, Tang Y-P, Rampon C, Tsien JZ. 2000. NMDA receptor-dependent synaptic reinforcement as a crucial process for memory consolidation. Science 290:1170–1174.

Tomko GJ, Crapper DR. 1974. Neuronal variability: Non-stationary responses to identical visual stimuli. Brain Res 79:405–418.

Vazdarjanova A, Guzowski JF. 2004. Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. J Neurosci 24:6489–6496.

Wheeler AL, Teixeira CM, Wang AH, Xiong X, Kovacevic N, Lerch JP, McIntosh AR, Parkinson J, Frankland PW. 2013. Identification of a Functional Connectome for Long-Term Fear Memory in Mice. PLoS Comput Biol 9: e1002853.