



HHS Public Access

Author manuscript

Curr Opin Behav Sci. Author manuscript; available in PMC 2018 October 01.

Published in final edited form as:

Curr Opin Behav Sci. 2017 October ; 17: 141–146.

Does mental context drift or shift?

Sarah DuBrow^a, Nina Rouhani^{a,b}, Yael Niv^{a,b}, and Kenneth A Norman^{a,b}

^aPrinceton Neuroscience Institute, Princeton University, Princeton, NJ 08544

^bDepartment of Psychology, Princeton University, Princeton, NJ 08544

Abstract

Theories of episodic memory have generally proposed that individual memory traces are linked together by a representation of context that drifts slowly over time. Recent data challenge the notion that contextual drift is always slow and passive. In particular, changes in one's external environment or internal model induce discontinuities in memory that are reflected in sudden changes in neural activity, suggesting that context can shift abruptly. Furthermore, context change effects are sensitive to top-down goals, suggesting that contextual drift may be an active process. These findings call for revising models of the role of context in memory, in order to account for abrupt contextual shifts and the controllable nature of context change.

Keywords

episodic memory; temporal context; event boundaries; situation models; latent cause models

Introduction

Much of what we experience is transient. Yet, many of our internal representations tend to show inertia, drifting more gradually than our external environment. Slowly drifting context can act as a bridge between more transient item representations. That is, if two items are associated in memory to the same slowly-changing context, this provides an indirect route for one of those items to activate the other [1]. This property has been captured in computational models of temporal context [2–5], which posit that *gradual integration of sensory inputs* creates a slowly drifting context representation that can then be linked to representations of individual items. These models have been very successful in explaining psychological and neural data, but they have recently been challenged by data suggesting that mental context can rapidly shift in response to both surprising events and changing task demands. Below, we first review data consistent with the context drift framework. Then, we

Corresponding author: Kenneth A Norman, 137 Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, knorman@princeton.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

We have no conflicts of interest.

describe three lines of challenging findings, as well as theories that have the potential to reconcile the “drifting” and “shifting” nature of mental context.

Evidence for slow, passive contextual drift

When freely recalling random lists of words, two effects have been identified as particularly strong and reliable – the recency effect and the contiguity effect – both of which can be explained by a slowly drifting mental context representation bound to items in memory [2]. The recency effect, or the tendency to recall end-of-list items especially well, can be explained by the idea that context in the recall period is most similar to context at the end (vs. beginning) of the list. The contiguity effect, or the tendency to transition in recall between items that were close to each other at study, can be explained by the idea that neighboring items from the study period share context, and therefore prime each other for recall. Neural data, particularly in the hippocampus, have also shown properties consistent with slow and automatic drift. Hippocampal “time cells” show gradually changing activity patterns on the order of seconds during unfilled delays even when the animal’s location is fixed [6–9], consistent with drift. Sensitivity to longer timescales on the order of days to weeks has also been observed in hippocampal place cells [10–12] and independently of place coding [13,14]. This neural drift may influence memory by enabling distinct events encountered in close temporal proximity to be linked according to the overlap in their neural activity profiles [15]. Furthermore, the degree to which neural activity patterns change across events has been associated with memory for temporal order and distance [16–19]. Critically, this relationship has been observed in the same regions that show slow drift [20*, 21]. A link between slow drift and spontaneous organization of memory was recently established in an fMRI study that tracked lingering activation of recently-experienced items [22]. When neural activity associated with the previous item’s category (celebrity, location, or object) persisted into the current item’s encoding period, this led to items being clustered at recall according to the previous item’s category (as would be expected if this slowly-drifting category activity served to contextualize the current item’s memory trace).

Challenging data: Can context shift abruptly?

Recent work has challenged the notion that contextual drift is always gradual. In particular, inducing abrupt changes in stimulus features and/or task goals creates separation in memory [19,23–25]. Similar effects can also be observed with naturalistic stimuli such as written stories or films, in which changes (“event boundaries”) occur at the narrative level (e.g., going from cooking a meal to eating the meal) [26]. Event boundaries can exert a sharp disruptive effect on memory, such that accessing information across a boundary is impaired even when controlling for the time elapsed [27–29]. Cognitively, these findings have been explained in terms of the idea that participants form *situation models* that describe the properties of the event [30,31]. One’s currently-active situation model may be a particularly strong component of mental context, and thus changes in that model may serve as powerful context shifts.

Understanding how situation models might be implemented in the brain has been a major focus of recent work. Researchers have argued that a posterior medial network (PMN),

including the parahippocampal cortex, retrosplenial cortex, and other regions of the default network, integrates internal and external information in order to represent the features of the current event [32]. This proposal is consistent with data showing that these regions are capable of integrating information across long time scales on the order of minutes [33] – an important prerequisite for constructing models of events that unfold at this time scale. This same network has been found to represent information about specific movie scenes during both movie viewing and memory recall, even across individuals [34]. A key property of brain regions involved in representing situation models is that neural patterns should be relatively stable within a temporally-extended event and change abruptly at event boundaries. By fitting a model with this property to movie-viewing fMRI data, Baldassano and colleagues [35**] showed that regions in the PMN are well-described by shifts between relatively stable patterns, lasting on the order of minutes. The “shift” moments identified in the PMN corresponded to moments that human observers identified as event boundaries, further suggesting that these regions are sensitive to high-level event content. Interestingly, perceptual regions tended to identify a larger number of short events, presumably due to a greater sensitivity to rapidly changing environmental features. In a different study, the hippocampus was found to be sensitive to narrative context at an even coarser level than the PMN, separately representing two interleaved story lines [36*]. Thus, the hippocampus and PMN may support high-level structure, integrating within stable events and rapidly shifting when a new situation is encountered. Together with lower-level regions, multiple timescales or hierarchies of event structure could be represented simultaneously.

Inferring change in the state of the world

How can computational models account for these sharp shifts in mental context? Within the temporal context framework, an abrupt increase in contextual drift rate can explain boundary effects on memory [37], but what might drive this increase? One possibility is that prediction errors, or discrepancies between expected and actual outcomes, generate event boundaries and ensuing context shifts [38–40], consistent with evidence that prediction errors can drive segmentation in learning [41,42]. According to a recent theory of latent cause inference, people build up a library of statistical models of their environment, and then infer which is the hidden, or “latent,” cause of their current sensory observations [43–46]. Large prediction errors signal that the current latent cause is no longer relevant (i.e., the situation has changed) and a new or stored latent cause should be inferred. Figure 1 illustrates how this inference process might shape mental context during a real-world situation: a surprise party at work. The appearance of a birthday cake during a meeting elicits a prediction error signaling that work is no longer relevant; instead, the new latent cause is a party.

Prediction errors may influence memory by segmenting experience when the inferred latent cause is updated, thereby reducing interference. In support of this idea, Gershman and colleagues [47] found that old items (simple line segments) were better remembered when they were followed by a prediction error and therefore “separated” from similar, interfering, items. Rouhani and colleagues [48] further showed that large reward prediction errors boosted performance on multiple memory measures, consistent with the idea that prediction errors led to increased segmentation and reduced interference. By contrast, when changes are small and predictable, a single latent cause is inferred and memories are more malleable

and susceptible to interference [49,50]. Neurally, the orbitofrontal cortex, which is sensitive to unobservable states [51], has been shown to represent the probability that each latent cause is currently active during the inference process [52]. Thus, one intriguing possibility is that the orbitofrontal cortex may draw on memory representations from the hippocampus and temporally extended situational information from the PMN to represent the current context for goals and actions.

Goal-directed control of context

While latent-cause models can account for abrupt context shifts, they do not (in their most basic form) speak to the strong influence of people's goals on the extent to which context lingers. For example, in much of the event segmentation literature, participants are asked to indicate when an event boundary occurs. This focus on detecting changes is associated with substantial improvements in memory performance [53–55*]. One potential explanation for this improvement is that the task of segmenting events boosts one's ability to notice changes (and thus experience prediction errors), leading to an increased number of inferred event boundaries, which reduce interference and enhance memory. In contrast, performing an associative task with the goal of remembering the temporal order of items might slow contextual drift so that items can be better remembered together in sequence [24]. In this case, hippocampal phase coding may support within-context sequential binding [25]. When sensory cues change abruptly but the goal stays the same, the hippocampus and lateral prefrontal cortex may help bind information that spans the change [56], perhaps by maintaining a stable internal context [57].

Other internal goals may exert an even more dramatic effect on contextual drift than the encoding task. For example, simply imagining oneself in a different context or mentally transforming one's current context can reduce context-dependent forgetting [58]. This is consistent with the idea that the features of context that organize memories may be largely internally generated and controllable. One implication of this idea is that rapidly shifting one's mental context may be helpful for intentional forgetting, by reducing contextual match between study and test. This view is supported by studies using the list-method directed forgetting paradigm, in which participants study two lists and are cued to forget or remember the first list immediately after studying that list. Sahakyan and Kelley [59] found that instructing participants to forget the first list has similar behavioral effects as instructing participants to change their mental context between lists. A recent neuroimaging study tested this idea more directly by interleaving words in the first list with irrelevant scenes, an approach first used by Gershman and colleagues [60] to tag and track context with fMRI. When instructed to forget the first list, fMRI activity associated with the scene context dropped, and the extent of that drop was related to participants' subsequent failure to recall items from that list [61*, Figure 2]. Like the context-maintenance effects described above, goal-directed forgetting has been associated with interactions between the medial temporal lobes and the prefrontal cortex [62].

Conclusions and open questions

We have reviewed evidence that inferring changes in the current situation and deliberately changing our mental context have substantial effects on memory behavior and the brain. We therefore suggest that – rather than being a compulsory, passive and uncontrollable process – contextual drift can be slowed and speeded adaptively in the service of achieving our goals. Indeed, it is possible that the slow, steady drift observed in unstructured passive encoding tasks is the result of an inference process over arbitrary, unrelated stimuli. That is, when studying a list of unrelated words, individual words might nudge your sense of the composition of the list, but this won't strongly change the overall inference that you are “listening to a list of random words”, so overall levels of contextual drift will be slow and steady (compared, e.g., to a situation where the word sequence suddenly switches from random words to words that spell out a coherent story). Having said this, it is important to note that drifting and shifting contexts may exist simultaneously in the brain; there are likely to be some contextual representations that drift passively over time and persist across event boundaries [22,56] and other contextual representations that are more sensitive to the inferred situation and thus more likely to shift suddenly.

The main challenge, moving forward, is to develop a mechanistic model that accounts for both the shifting and drifting aspects of mental context. Latent cause models are a promising way to account for sharp shifts in mental context, but questions remain. For instance, Figure 1 depicts one latent cause switching to another, suggesting that latent causes are mutually exclusive, but it seems more reasonable to posit that multiple latent causes can coexist (i.e., you may continue to think about, or even discuss work-related issues at the surprise party, making both the “work” and “party” latent causes relevant).

From a statistical perspective, superposition of latent causes is straightforward to model; in fact, existing latent cause models infer a continuous probability distribution over latent causes [43]. However, this superposition poses challenging cognitive and neural questions. For example, if multiple stored latent causes are relevant, they might make different predictions about what will happen next. How do we reconcile these predictions when they differ, and might this difference influence how easy it is to keep these causes simultaneously active? Also, if multiple latent causes are active, how do we assign credit so that we update the appropriate latent cause when new information comes in? These issues intersect with the existing literatures on cognitive control, multitasking, and reinforcement learning [44*, 52,63,64]. Drawing on these literatures might provide useful insight into how latent-cause models could be implemented in the brain, and how they interact with neural memory processes.

Acknowledgments

This work was supported by the John Templeton Foundation (Grant #36751), grants R01MH112357 (K.A.N) and R01DA042065 (Y.N.) from the National Institute for Mental Health, and the National Science Foundation's Graduate Research Fellowship Program DGE-1656466 (N.R.). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

References

1. Manning JR, Kahana MJ, Norman KA. The role of context in episodic memory. *Cogn Neurosci*. 2014; doi: 10.3758/BF03333552
2. Howard MW, Kahana MJ. A distributed representation of temporal context. *J Math Psychol*. 2002; 46:269–299.
3. Howard MW, Fotedar MS, Datey AV, Hasselmo ME. The temporal context model in spatial navigation and relational learning: toward a common explanation of medial temporal lobe function across domains. *Psychol Rev*. 2005; 112:75–116. [PubMed: 15631589]
4. Howard MW, Shankar KH, Aue WR, Criss AH. A distributed representation of internal time. *Psychol Rev*. 2015; 122:24–53. [PubMed: 25330329]
5. Estes WK. Statistical theory of spontaneous recovery and regression. *Psychol Rev*. 1955; 62:145–154. [PubMed: 14371893]
6. Macdonald CJ, Carrow S, Place R, Eichenbaum H. Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *J Neurosci*. 2013; 33:14607–16. [PubMed: 24005311]
7. Macdonald CJ, Lepage KQ, Eden UT, Eichenbaum H. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*. 2011; 71:737–749. [PubMed: 21867888]
8. Naya Y, Suzuki WA. Integrating what and when across the primate medial temporal lobe. *Science*. 2011; 333:773–776. [PubMed: 21817056]
9. Eichenbaum H. Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci*. 2014; 15:732–44. [PubMed: 25269553]
10. Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. Neuronal code for extended time in the hippocampus. *Proc Natl Acad Sci U S A*. 2012; 109:19462–7. [PubMed: 23132944]
11. Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer MJ. Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci*. 2013; 16:264–6. [PubMed: 23396101]
12. Lever C, Wills T, Cacucci F, Burgess N, O’Keefe J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*. 2002; 416:90–94. [PubMed: 11882899]
13. Rubin A, Geva N, Sheintuch L, Ziv Y. Hippocampal ensemble dynamics timestamp events in long-term memory. *Elife*. 2015; 4
14. Mankin EA, Diehl GW, Sparks FT, Leutgeb S, Leutgeb JK. Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. *Neuron*. 2015; 85:190–201. [PubMed: 25569350]
15. Cai DJ, Aharoni D, Shuman T, Shobe J, Biane J, Song W, Wei B, Veshkini M, La-Vu M, Lou J, et al. A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*. 2016; 534:115–8. [PubMed: 27251287]
16. DuBrow S, Davachi L. Temporal memory is shaped by encoding stability and intervening item reactivation. *J Neurosci*. 2014; 34
17. Jenkins LJ, Ranganath C. Distinct neural mechanisms for remembering when an event occurred. *Hippocampus*. 2016; 26:554–559. [PubMed: 26845069]
18. Jenkins LJ, Ranganath C. Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *J Neurosci*. 2010; 30:15558–15565. [PubMed: 21084610]
19. Ezzyat Y, Davachi L. Similarity Breeds Proximity: Pattern Similarity within and across Contexts Is Related to Later Mnemonic Judgments of Temporal Proximity. *Neuron*. 2014; 81:1179–89. [PubMed: 24607235]
- 20*. Lositsky O, Chen J, Toker D, Honey CJ, Shvartsman M, Poppenk JL, Hasson U, Norman KA, Andersson J, Jenkinson M, et al. Neural pattern change during encoding of a narrative predicts retrospective duration estimates. *Elife*. 2016 Nov.5:e16070. In this fMRI study, judgments of temporal distance were found to correlate both with narrative event boundaries as well as with drift in patterns of entorhinal and ventrolateral cortex activity, consistent with both the “drift” and “shift” accounts of mental context. [PubMed: 27801645]

21. Manns JR, Howard MW, Eichenbaum H. Gradual changes in hippocampal activity support remembering the order of events. *Neuron*. 2007; 56:530–540. [PubMed: 17988635]
22. Chan SCY, Applegate MC, Morton NW. Lingering representations of stimuli influence recall organization. *Neuropsychologia*. 2017; 97:72–82. [PubMed: 28132858]
23. Polyn SM, Norman KA, Kahana MJ. Task context and organization in free recall. *Neuropsychologia*. 2009; 47:2158–2163. [PubMed: 19524086]
24. DuBrow S, Davachi L. The influence of context boundaries on memory for the sequential order of events. *J Exp Psychol Gen*. 2013; 142:1277–1286. [PubMed: 23957281]
25. Heusser AC, Poeppel D, Ezzyat Y, Davachi L. Episodic sequence memory is supported by a theta-gamma phase code. *Nat Neurosci*. 2016; 19:1374–1380. [PubMed: 27571010]
26. Zacks J, Tversky B, Iyer G. Perceiving, remembering, and communicating structure in events. *J Exp Psychol Gen*. 2001; 130:29–58. [PubMed: 11293458]
27. Swallow KM, Zacks JM, Abrams RA. Event Boundaries in Perception Affect Memory Encoding and Updating. *J Exp Psychol Gen*. 2009; 138:236–257. [PubMed: 19397382]
28. Ezzyat Y, Davachi L. What constitutes an episode in episodic memory? *Psychol Sci*. 2011; 22:243–52. [PubMed: 21178116]
29. Horner AJ, Bisby JA, Wang A, Bogus K, Burgess N. The role of spatial boundaries in shaping long-term event representations. *Cognition*. 2016; 154:151–164. [PubMed: 27295330]
30. Zwaan RA, Langston MC, Graesser AC. The Construction of Situation Models in Narrative Comprehension: An Event-Indexing Model. *Psychol Sci*. 1995; 6:292–297.
31. Radvansky GA. Across the Event Horizon. *Curr Dir Psychol Sci*. 2012; 21:269–272.
32. Ranganath C, Ritchey M. Two cortical systems for memory-guided behaviour. *Nat Rev Neurosci*. 2012; 13:713–26. [PubMed: 22992647]
33. Hasson U, Chen J, Honey CJ. Hierarchical process memory: memory as an integral component of information processing. *Trends Cogn Sci*. 2015; 19:304–313. [PubMed: 25980649]
34. Chen J, Leong YC, Honey CJ, Yong CH, Norman KA, Hasson U. Shared memories reveal shared structure in neural activity across individuals. *Nat Neurosci*. 2016; 20:115–125. [PubMed: 27918531]
- 35**. Baldassano C, Chen J, Zadbood A, Pillow JW, Hasson U, Norman KA. Discovering event structure in continuous narrative perception and memory. *Neuron*. 2017; 95:709–721. This study uses a data-driven Hidden Markov Modeling technique to identify moments in which patterns of fMRI activity shift during narrative processing and recall. Regions in the posterior medial network show neural pattern shifts that line up with human-labeled event boundaries, and hippocampal responses at boundaries identified by the model predict later reinstatement of the just-completed event. [PubMed: 28772125]
- 36*. Milivojevic B, Varadinov M, Vicente Grabovetsky A, Collin SHP, Doeller CF. Coding of Event Nodes and Narrative Context in the Hippocampus. *J Neurosci*. 2016; 36:12412–12424. This fMRI study uses a movie with interwoven narratives to provide evidence that the hippocampus is sensitive to event-level features, such as characters and locations, as well as high-level narrative context that differentiates over time. [PubMed: 27927958]
37. Polyn SM, Norman KA, Kahana MJ. A context maintenance and retrieval model of organizational processes in free recall. *Psychol Rev*. 2009; 116:129–156. [PubMed: 19159151]
38. Reynolds JR, Zacks JM, Braver TS. A computational model of event segmentation from perceptual prediction. *Cogn Sci*. 2007; 31:613–643. [PubMed: 21635310]
39. Zacks JM, Kurby CA, Eisenberg ML, Haroutunian N. Prediction Error Associated with the Perceptual Segmentation of Naturalistic Events. *J Cogn Neurosci*. 2011; 23:4057–4066. [PubMed: 21671745]
40. Zacks JM, Speer NK, Swallow KM, Braver TS, Reynolds JR. Event Perception: A Mind/Brain Perspective. *Psychol Bull*. 2007; 133:273–293. [PubMed: 17338600]
41. Redish AD, Jensen S, Johnson A, Kurth-Nelson Z. Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychol Rev*. 2007; 114:784–805. [PubMed: 17638506]
42. Gershman SJ, Niv Y. Learning latent structure: Carving nature at its joints. *Curr Opin Neurobiol*. 2010; 20:251–256. [PubMed: 20227271]

43. Gershman SJ, Blei DM, Niv Y. Context, learning, and extinction. *Psychol Rev.* 2010; 117:197–209. [PubMed: 20063968]
- 44*. Gershman SJ, Monfils M-H, Norman KA, Niv Y. The computational nature of memory modification. *Elife.* 2017; 6:e23763. This paper describes a computational account of memory updating in Pavlovian conditioning using latent cause inference. Prediction errors create new memory traces, but when the environment is stable, memories are updated. [PubMed: 28294944]
45. Soto FA, Gershman SJ, Niv Y. Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychol Rev.* 2014; 121:526–558. [PubMed: 25090430]
46. Gershman SJ, Niv Y. Exploring a latent cause theory of classical conditioning. *Learn Behav.* 2012; 40:255–268. [PubMed: 22927000]
47. Gershman SJ, Radulescu A, Norman KA, Niv Y. Statistical Computations Underlying the Dynamics of Memory Updating. *PLoS Comput Biol.* 2014; 10
48. Rouhani N, Norman KA, Niv Y. Dissociable effects of surprising rewards on learning and memory. *bioRxiv.* 2017; doi: 10.1101/111070
49. Gershman SJ, Jones CE, Norman KA, Monfils M-H, Niv Y. Gradual extinction prevents the return of fear: implications for the discovery of state. *Front Behav Neurosci.* 2013; 7:164. [PubMed: 24302899]
50. Gershman SJ, Norman KA, Niv Y. Discovering latent causes in reinforcement learning. *Curr Opin Behav Sci.* 2015; 5:43–50.
51. Schuck NW, Cai MB, Wilson RC, Niv Y. Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron.* 2016; 91:1402–1412. [PubMed: 27657452]
52. Chan SCY, Niv Y, Norman KA. A Probability Distribution over Latent Causes, in the Orbitofrontal Cortex. *J Neurosci.* 2016; 36:7817–28. [PubMed: 27466328]
53. Gold DA, Zacks JM, Flores S. Effects of cues to event segmentation on subsequent memory. *Cogn Res Princ Implic.* 2017; 2:1. [PubMed: 28203629]
54. Pettijohn KA, Thompson AN, Tamplin AK, Krawietz SA, Radvansky GA. Event boundaries and memory improvement. *Cognition.* 2016; 148:136–144. [PubMed: 26780472]
- 55*. Flores S, Bailey HR, Eisenberg ML, Zacks JM. Event Segmentation Improves Event Memory up to One Month Later. *J Exp Psychol Learn Mem Cogn.* 2017; This series of behavioral studies showed that when viewing movies, simply attending to meaningful changes in the narrative in order to perform a segmentation task improved long-term memory for up to a month later. doi: 10.1037/xlm0000367
56. DuBrow S, Davachi L. Temporal binding within and across events. *Neurobiol Learn Mem.* 2016; 134:107–114. [PubMed: 27422018]
57. Polyn SM, Kahana MJ. Memory search and the neural representation of context. *Trends Cogn Sci.* 2008; 12:24–30. [PubMed: 18069046]
58. Masicampo EJ, Sahakyan L. Imagining Another Context During Encoding Offsets Context-Dependent Forgetting. *J Exp Psychol Learn Mem Cogn.* 2014; 40:1772–1777. [PubMed: 24611435]
59. Sahakyan L, Kelley CM. A contextual change account of the directed forgetting effect. *J Exp Psychol Learn Mem Cogn.* 2002; 28:1064–1072. [PubMed: 12450332]
60. Gershman SJ, Schapiro AC, Hupbach A, Norman KA. Neural context reinstatement predicts memory misattribution. *J Neurosci.* 2013; 33:8590–5. [PubMed: 23678104]
- 61*. Manning JR, Hulbert JC, Williams J, Piloto L, Sahakyan L, Norman KA. A neural signature of contextually mediated intentional forgetting. *Psychon Bull Rev.* 2016; 23:1534–1542. This study used fMRI to demonstrate that forget cues in list-method directed forgetting lead to neural deactivation of the previously-active mental context, with greater deactivation leading to greater forgetting. [PubMed: 27150815]
62. Anderson MC, Hanslmayr S. Neural mechanisms of motivated forgetting. *Trends Cogn Sci.* 2014; 18:279–92. [PubMed: 24747000]
63. Feng SF, Schwemmer M, Gershman SJ, Cohen JD. Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cogn Affect Behav Neurosci.* 2014; 14:129–146. [PubMed: 24481850]

64. Botvinick MM, Cohen JD. The computational and neural basis of cognitive control: Charted territory and new frontiers. *Cogn Sci.* 2014; 38:1249–1285. [PubMed: 25079472]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Highlights

- Recent behavioral and neuroscience work in memory has challenged the notion that context always drifts slowly and passively.
- Changes in context, or “event boundaries,” shift context abruptly, creating discontinuity in memory.
- These shifts may be triggered by prediction errors that lead to inferring a new model of the environment.
- Internal goals influence context effects on memory, further suggesting that shifting contexts is under top-down control.

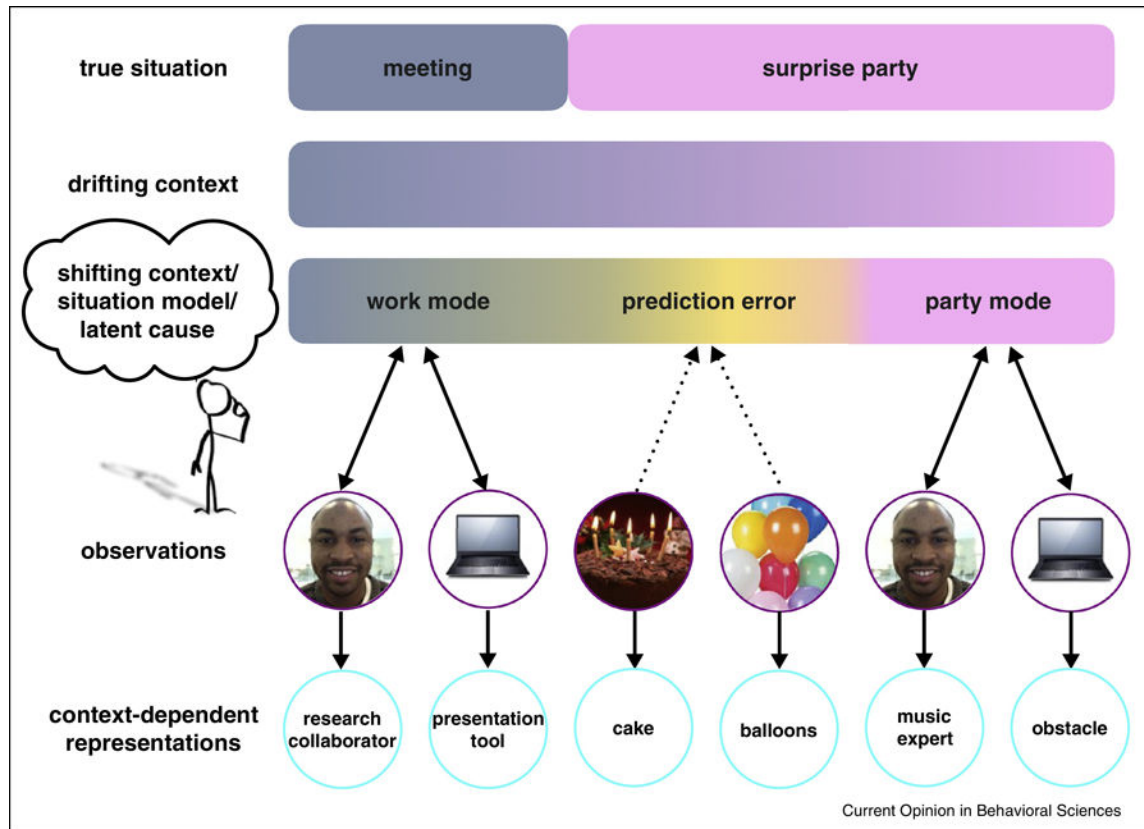


Figure 1. Real world inference process and how it shapes mental context

While meeting with a colleague on a joint project, your mental context reflects a “work mode” that allows you to generate appropriate predictions about what you expect to see (e.g., colleagues) and how you will interact with them (e.g., as a research collaborator). When you see your colleague pull out a cake, this elicits a prediction error because this no longer meets your “work mode” expectations. Because there is still ambiguity about what the underlying situation is, you may need to sample a few more observations (e.g., other party decorations like balloons, other non-work friends appear) before realizing that they are throwing you a surprise party. This realization activates a new “party mode” that generates new predictions and activates situationally-relevant goals (e.g., the colleague you had been working with is also a music expert who might have a perfect playlist for the occasion). Note this model assumes that only a single “mode” can be active at a time.

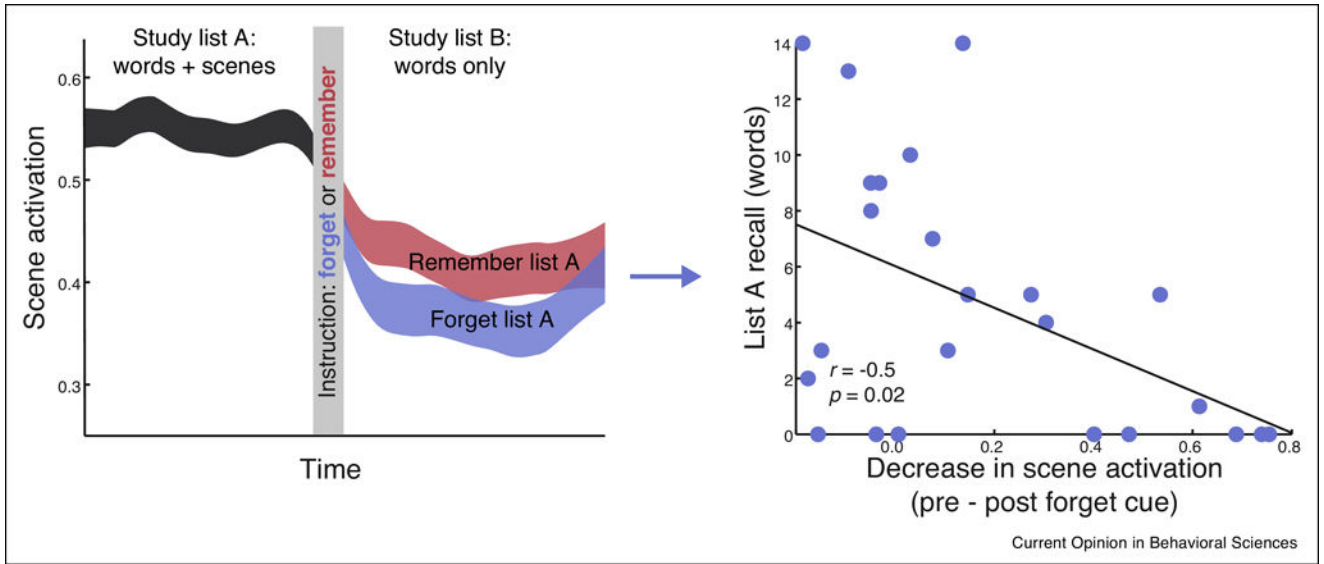


Figure 2. Cued forgetting induces goal-directed shifts in neural context

In a list-method directed forgetting experiment, participants studied a list of words (list A) with intermixed scenes, then they were cued to remember or forget list A, then they studied a second list of words (list B) without intermixed scenes. Lingering scene-related fMRI activation during list B (when scenes were not visible) was used to track persistence of the list A context. Instructing participants to forget (vs. remember) list A led to reduced persistence of scene-related fMRI activation (left), suggesting that participants deliberately shift their context in response to forget cues. The degree to which scene activation decreased in response to the forget cue predicted subsequent forgetting of list A, operationalized as fewer recalled words from that list (right). Adapted from [61*] with permission.