

corresponding to Bayesian conditionalization. A learning rule carries out gradient ascent in the ‘likelihood’ of the data presented at the visible units. All of this is achieved with no explicit representation of probability, but merely via simple and distributed ‘neural’ computations.

The Boltzmann machine does not scale-up well. However, related ideas have evolved in a variety of directions. One approach focuses on representing complex probability distributions through sparse and structured ‘graphical models’ which implicitly capture dependencies between variables (e.g., [6]). Indeed, general-purpose programming languages for compositionally specifying and sampling from arbitrary probability distributions have been created (e.g., [7]).

A different development de-emphasizes compositional representation, and focuses on learning, typically with richly connected networks without a transparent interpretation. For example, ‘restricted’ Boltzmann machines can be ‘stacked’ into multiple layers (e.g., in deep belief networks [8]). More broadly, deep learning has been scaled up to achieve state-of-the-art machine-learning performance [9].

More neurobiologically realistic implementations of sampling algorithms have recently been developed, some of which implement sampling for discrete variables on networks of spiking neurons (e.g., [10]). Other schemes for sampling continuous variables build on the link between energy and probability, producing dynamics in networks of excitatory and inhibitory neurons that implement an advanced sampling algorithm (e.g., [11]).

In contrast to our sampling proposal, Friston’s free-energy approach (e.g., [12]) does not treat the entire state of the brain as a single sample from a posterior probability distribution. The free-energy approach also does not

accurately represent the probability of every possible hypothesis – far from it. The true posterior distribution is approximated by a simpler distribution, and minimizing free energy brings this simpler distribution into approximate correspondence with the true posterior. In Friston’s model, neurons encode the parameters of this approximating distribution (cf [13]), often a simple Gaussian distribution, which yields an elegant neurobiological implementation of the free-energy approach.

We argued [2] that sampling will produce reasoning errors, such as the unpacking effect and the conjunction fallacy, if the sampler only samples a single mode in a multimodal distribution. Perhaps approximating a multimodal posterior distribution with a single (e.g., Gaussian) mode may be a different route to producing these same errors. Thus these various approximations to Bayesian inference may provide competing explanations of the fallacies and biases observed in explicit reasoning with probabilities.

<sup>1</sup>University of Warwick, Coventry, UK

<sup>2</sup>Warwick Business School, Coventry, UK

\*Correspondence:

a.n.sanborn@warwick.ac.uk (A.N. Sanborn).

<http://dx.doi.org/10.1016/j.tics.2017.04.009>

#### References

1. Alday, P.M. *et al.* (2017) Commentary on Sanborn and Chater: Posterior modes are attractor basins. *Trends Cogn. Sci.* 21, 491–492
2. Sanborn, A.N. and Chater, N. (2016) Bayesian brains without probabilities. *Trends Cogn. Sci.* 20, 883–893
3. Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U. S. A.* 79, 2554–2558
4. Ackley, D.H. *et al.* (1985) A learning algorithm for Boltzmann machines. *Cogn. Sci.* 9, 147–169
5. Geman, S. and Geman, D. (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* 6, 721–741
6. Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann
7. Piantadosi, S.T. *et al.* (2016) The logical primitives of thought: empirical foundations for compositional cognitive models. *Psychol. Rev.* 123, 392–424
8. Hinton, G.E. and Salakhutdinov, R.R. (2006) Reducing the dimensionality of data with neural networks. *Science* 313, 504–507

9. LeCun, Y. *et al.* (2015) Deep learning. *Nature* 521, 436–444
10. Pecevski, D. *et al.* (2011) Probabilistic inference in general graphical models through sampling in stochastic networks of spiking neurons. *PLoS Comput. Biol.* 7, e1002294
11. Aitchison, L. and Lengyel, M. (2016) The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Comput. Biol.* 12, e1005186
12. Friston, K. (2010) The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138
13. Ma, W.J. *et al.* (2006) Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438

## Forum

### Serial Dependence across Perception, Attention, and Memory

Anastasia Kiyonaga,<sup>1,\*</sup> Jason M. Scimeca,<sup>1</sup> Daniel P. Bliss,<sup>1</sup> and David Whitney<sup>1</sup>

**Information that has been recently perceived or remembered can bias current processing. This has been viewed as both a corrupting (e.g., proactive interference in short-term memory) and stabilizing (e.g., serial dependence in perception) phenomenon. We hypothesize that this bias is a generally adaptive aspect of brain function that leads to occasionally maladaptive outcomes.**

#### Previous Memory Encroaches on Current Memory

One of the most remarkable aspects of our visual experience is that we perceive a stable environment despite the constantly changing image on the retina (e.g., from eye movements, blinks, and occlusions). A core function of visual working memory is to temporarily maintain representations that bridge disruptions and facilitate that stability. However, working memory is a limited system that sometimes fails to properly segregate recently activated

(but now irrelevant) representations from those that should guide current behavior. For example, memory performance for lists of words is slower and more error-prone if a probe item belonged to a previous memory set [1] (Figure 1A). Such proactive interference also occurs during change detection for arrays of visual stimuli, wherein a probe item will more likely be incorrectly identified as part of the current array if it had appeared in the previous array [2] (Figure 1B). In a spatial delayed-response task, moreover, eye movements of monkeys are biased toward the remembered location from the previous trial [3] (Figure 1C). That is, previously relevant information can linger and impact what we currently remember over the short term. The bulk of proactive interference research has focused on control mechanisms to overcome that impact, because it is considered a fundamental obstacle to cognition [1]. Recent research in visual perception, however, may provide a window into the generally adaptive nature of the processes underlying proactive interference.

### Previous Perception Integrates with Current Perception

The perception of visual stimuli often follows a pattern that could be described as proactive interference: new stimuli appear more similar to just-seen stimuli than they truly are. In an orientation judgement task, for instance, perceived orientations are biased toward the angle of previously attended stimuli [4] (Figure 1D). This serial dependence occurs for a range of features, including numerosity [5] and complex stimuli like faces [6,7] (Figure 1E), which suggests that it is a general visual processing principle. This principle may extend beyond visual perception, moreover, and encapsulate the effect of proactive interference in working memory. In fact, many tasks that are used to investigate serial dependence in perception are structured identically to canonical working memory tasks: participants maintain a stimulus over a short delay and then report on its features (Figure 1). Therefore,

serial dependence in perception and proactive interference in memory may stem from similar memory traces that persist across trials.

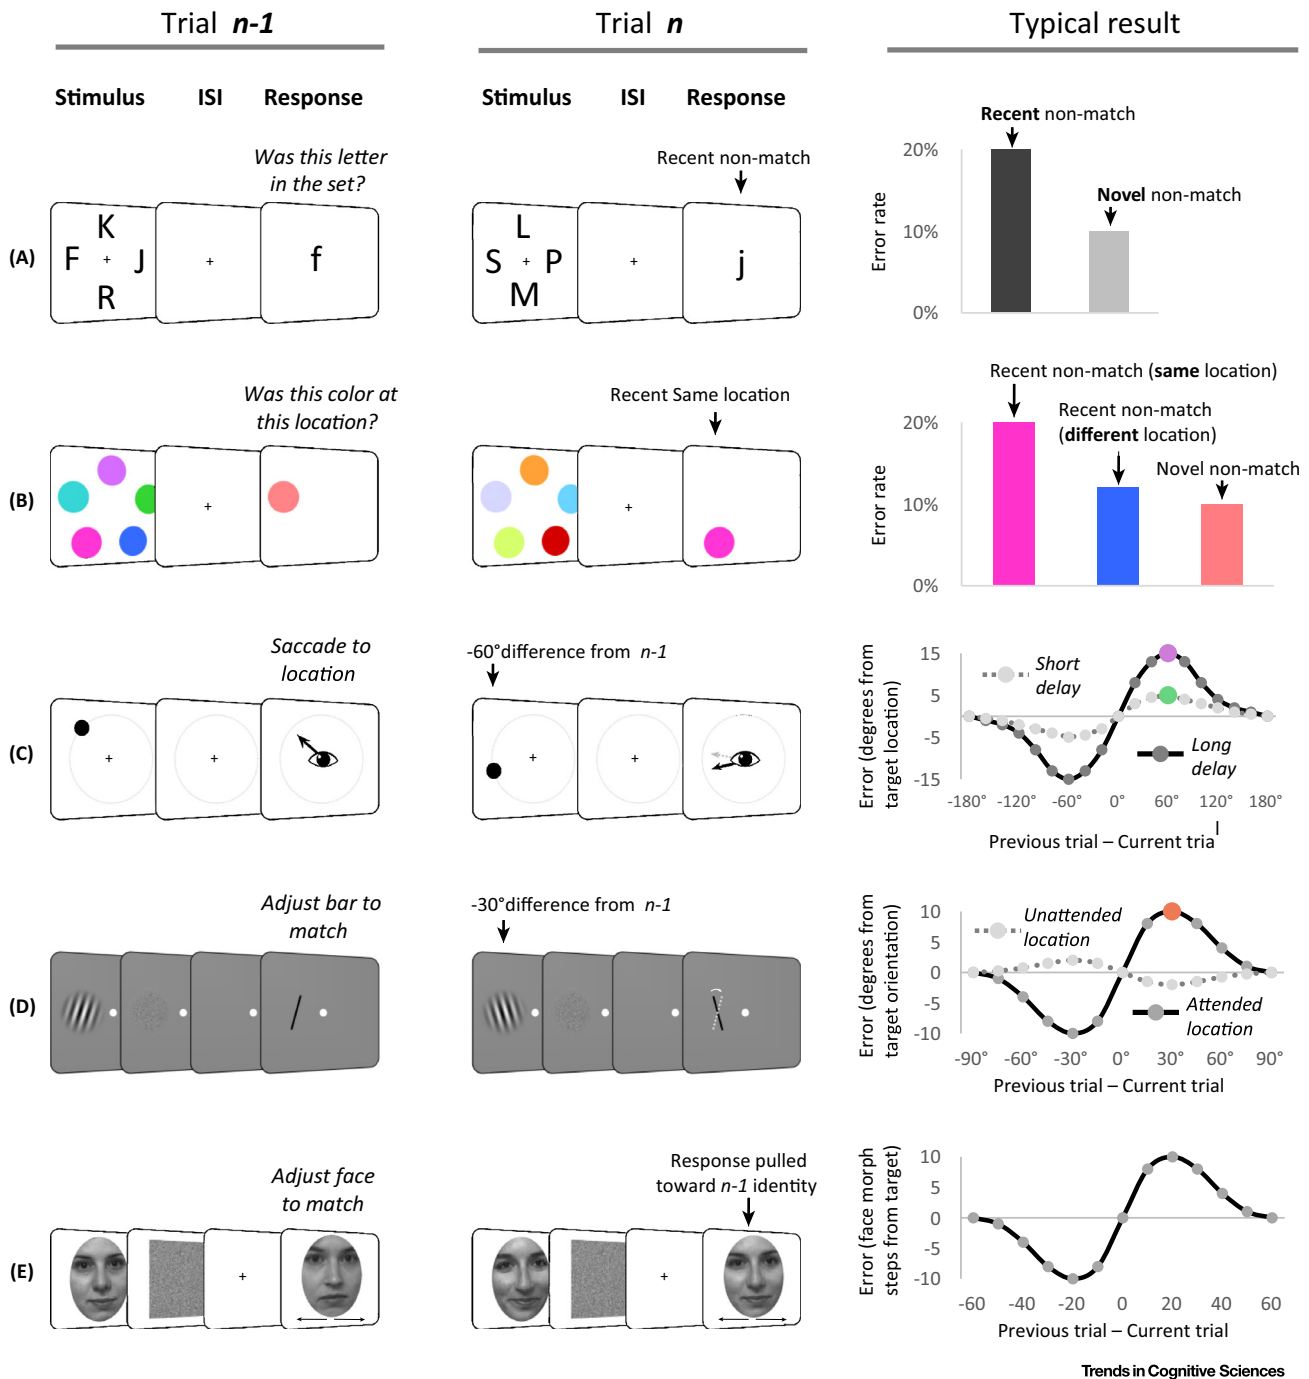
Serial dependence in perception describes a recency-bias much like the proactive interference effect in working memory. Unlike proactive interference, however, serial dependence has been framed as an adaptive process, integrating successive stimuli to improve signal-to-noise in the service of perceptual stability [4–7]. Because the recent visual past is typically a good predictor of the future, this carry-over may bias neural signals to enhance perception and decision making for incoming stimuli (i.e., predictive coding [8]). However, such temporal smoothing could become maladaptive when it sustains representations that are irrelevant to the current situation (i.e., proactive interference). An intriguing hypothesis, therefore, is that proactive interference is a maladaptive form of serial dependence, which arises from a mostly advantageous processing system. That is, the very feature that allows working memory to confer such impressive coherence is also the root of one of its most notorious handicaps.

### What Is the Value of Temporal Integration in Perception and Memory?

Intrusions by concurrent [9] and recent [2] memories have been considered a significant limitation on visual working memory capacity and quality, because they consume a finite representational space that would otherwise be used to maintain more relevant information [1,2]. This framing assumes that recent representations should be discarded, but neglects to consider the potential advantages of using related old information to scaffold new representations. Expectations can preactivate sensory representations and facilitate processing of predicted inputs [8], and active working memory content can improve detection and attentional processing of similar visual stimuli [10].

Likewise, the lingering activation of previously relevant stimuli may increase sensitivity for related information in the environment, and thereby enhance encoding into working memory. In the same way that our visual environment is predictable from moment to moment – rendering serial dependence adaptive in everyday cognition – our mental workspace is often characterized by a sequential narrative (i.e., conversation or problem solving) or by a continuous progression of sensory events. In these structured environments, it would be adaptive to integrate recent stimuli that inform current encoding and retention demands, to increase the cohesion of working memory episodes. The drawback of such an integration process is that it would hamper differentiation between successive stimuli if the priority were to detect changes in the environment.

How might we recognize when past information is likely to remain useful? Serial dependencies are tuned by the similarities between previously and currently relevant stimuli. The extent of influence from previous stimuli depends on their distance from current stimuli in time [4,6,11], space [2,4], and feature similarity [3,4,6], while the bias is virtually nonexistent once that difference exceeds a certain magnitude. This range across which serial dependence occurs is referred to as the continuity field [4]. Serial dependence is also limited to stimuli that were previously attended [4,11] or actively maintained [2], and absent for stimuli that were merely passively viewed. Moreover, the temporal gradient of serial dependence may differ depending on the changeability or autocorrelation of particular features in the natural world [7]. Stimulus attributes with low temporal autocorrelation can even promote repulsion (instead of attraction), which would maximize sensitivity to change [7]. These modulatory influences reinforce the adaptive nature of serial dependence in a constantly changing but statistically predictable environment; it only occurs when stimulus



**Figure 1.** Example Task Schematics and Representative Results Demonstrating Serial Dependence. (A) In a recent-probes recognition test for verbal material, error rate is higher for nonmatch items that were presented in the previous trial memory set (trial  $n-1$ ). (B) In a change-detection test for visual features, error rate is higher for nonmatch items that were presented in the previous trial, especially when the nonmatch item was previously presented at the same spatial location as the current test probe. (C) In an oculomotor delayed-response test for spatial locations, eye movement responses are pulled toward the remembered location in the previous trial. This bias is modulated by the distance between the current and previous location and by the length of the inter-stimulus delay interval (ISI) of the current trial. In the example shown here, the location on the previous trial was  $60^\circ$  clockwise from the current location, and the eye movement is thus biased in the clockwise direction by  $15^\circ$  (on long delay trials, purple dot) or  $5^\circ$  (on short delay trials, green dot). (D) In a continuous-report test for orientations, orientation judgments are biased toward the orientation presented on the previous trial. This bias is modulated by the difference between the current and previous orientation and only occurs for attended (versus unattended) orientations. In the example shown here, the orientation on the previous trial was  $30^\circ$  clockwise from the current orientation, and the orientation judgement is thus biased in the clockwise direction by  $10^\circ$  (orange dot). (E) In a continuous-report test for face identity, face judgements are biased towards the identity of the face presented on the previous trial. This bias is modulated by the number of morph steps (in face-morph space) between the current and previous face.

characteristics signal that the previous information remains potentially relevant.

### Neural Sources of Serial Dependence in Perception and Working Memory

Despite the potential utility of serial dependence for smoothing perception and memory, we know little about how the brain produces this behavior. One possibility is that it manifests at the earliest stages of cortical sensory processing, via increased sensitivity – perhaps due to short-term synaptic plasticity [12] – among neuronal ensembles that are responsive to the feature value of the previously attended stimulus [4]. Distributed functional magnetic resonance imaging (fMRI) patterns across early visual cortex are shifted in the direction of the attended orientation from the previous trial [11], but the underlying source of this bias is unknown. The observation that serial dependence relies on attention [2,4,11] suggests that a signal may originate from higher cortical areas to produce the bias; the attentional signal from recent experience may persist in the form of active firing or a residual synaptic trace. The vast body of knowledge on attentional control and stimulus representation in working memory [12,13] may therefore illuminate the neural mechanisms underlying serial dependence, especially if we examine the bias in the context of similar trial history effects that have been studied from the working memory perspective (i.e., proactive interference).

Like sensory recruitment in working memory [13], serial dependence also seems to manifest in cortical regions that represent the particular task content: proactive interference from verbal material is reflected in fMRI activity patterns in the temporal cortex [14], whereas serial dependence for visual material is evident in the early visual cortex [11], and residual previous trial information in an oculomotor delayed-response task manifests in the frontal eye fields of non-human primates [3]. Moreover, when the delay between

#### Box 1. Predictions for a Unified Principle of Serial Dependence across Perception, Attention, and Memory

Although serial dependence may adaptively support visual stability over time, it fundamentally involves a biased representation of the current stimulus. Thus, it should be downregulated in certain situations (i.e., those that demand differentiation [7]). This regulation may rely on the same mechanisms that detect and dampen proactive interference in working memory. If this is the case, cues for differentiation (like a context shift) should initiate control processes to segregate representations and curtail the impact of previous information. Conversely, if proactive interference arises from the same adaptive processes as perceptual serial dependence, the impact of the previous information on current memory should be amplified when context cues (like spatial or semantic similarity) signal its continued relevance for the current circumstances. While there is a compelling rationale for the adaptive role of serial dependence, the argument is mostly theoretical. Strong empirical evidence for the value of serial dependence would emerge from manipulations of signal-to-noise ratio in a predictable visual environment, wherein endogenous representations of previous stimuli should actually boost perception and encoding of incoming stimuli to improve performance. Such investigations would clarify the context-dependent flexibility of serial dependence and confirm its utility in multiple domains of processing.

stimulus and response is lengthened, the magnitude of attraction toward the previously remembered stimulus grows [3,15]. This implies that the effect of integration with previous stimuli can continue to evolve while the current stimulus is maintained in working memory, which may reflect the contribution of postperceptual processes. In fact, judgments that rely more on perceptual comparisons – and are less influenced by postperceptual and attentional processes – can be repelled away from preceding stimuli (i.e., tilt after-effects; [15]), and repulsive adaptation may also increase with the duration of exposure to the previous stimulus [4]. The ultimate consequence of serial dependence – that is, whether it segregates to promote change detection or integrates to promote stability – likely depends on the balance of perceptual and memory demands.

### Concluding Remarks

Serial dependencies in perception and working memory have been studied separately, but the parallels between them highlight the value of synthesizing these research areas moving forward [10]. We can harness the careful psychophysics and modeling data arising from perceptual serial dependence research to illuminate principles of stability in working memory. Conversely, we can exploit many decades of working memory research to inform the study of stimulus representations and adaptive control in

serial dependence. Several outstanding questions and testable hypotheses stem from the integration of these concepts (Box 1). If we appreciate proactive interference as a consequence of a system that typically supports cohesive internal representations across time, this perspective should clarify the principles of efficient temporal dependence in brain function, and why they sometimes lead us astray.

### Acknowledgments

We thank Mark D'Esposito for insightful feedback, as well as Derek Nee and Rob White for helpful discussions and comments on an earlier draft of the manuscript.

<sup>1</sup>Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

\*Correspondence: [kiyonaga@berkeley.edu](mailto:kiyonaga@berkeley.edu) (A. Kiyonaga).  
<http://dx.doi.org/10.1016/j.tics.2017.04.011>

### References

1. Jonides, J. and Nee, D.E. (2006) Brain mechanisms of proactive interference in working memory. *Neuroscience* 139, 181–193
2. Makovski, T. and Jiang, Y.V. (2008) Proactive interference from items previously stored in visual working memory. *Mem. Cognit.* 36, 43–52
3. Papadimitriou, C. et al. (2017) Ghosts in the machine II: neural correlates of memory interference from the previous trial. *Cereb. Cortex* 27, 2513–2527
4. Fischer, J. and Whitney, D. (2014) Serial dependence in visual perception. *Nat. Neurosci.* 17, 738–743
5. Cicchini, G.M. (2014) Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proc. Natl. Acad. Sci. U. S. A.* 111, 7867–7872
6. Liberman, A. (2014) Serial dependence in the perception of faces. *Curr. Biol.* 24, 2569–2574
7. Taubert, J. et al. (2016) Different coding strategies for the perception of stable and changeable facial attributes. *Sci. Rep.* 6, 32239

8. Summerfield, C. and de Lange, F.P. (2014) Expectation in perceptual decision making: neural and computational mechanisms. *Nat. Rev. Neurosci.* 15, 745–756
9. Oberauer, K. and Lin, H.-Y. (2017) An interference model of visual working memory. *Psychol. Rev.* 124, 21–59
10. Kiyonaga, A. and Egner, T. (2013) Working memory as internal attention: toward an integrative account of internal and external selection processes. *Psychon. Bull. Rev.* 20, 228–242
11. John-Saaltink, E.S. (2016) Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex. *J. Neurosci.* 36, 6186–6192
12. Stokes, M.G. (2015) 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn. Sci.* 19, 394–405
13. D'Esposito, M. and Postle, B.R. (2015) The cognitive neuroscience of working memory. *Annu. Rev. Psychol.* 66, 115–142
14. Oztekin, I. and Badre, D. (2011) Distributed patterns of brain activity that lead to forgetting. *Front. Hum. Neurosci.* 5, 86
15. Fritsche, M. *et al.* (2017) Opposite effects of recent history on perception and decision. *Curr. Biol.* 27, 590–595