pathophysiology and treatment of Alzheimer's disease. Brain *68*, 465–17.

Lowe, D. (2016). Muscarinic agonists come back around. http://blogs.sciencemag.org/pipeline/archives/2016/04/07/muscaranic-agonists-come-back-around.

Major, A.J., Vijayraghavan, S., and Everling, S. (2015). Muscarinic attenuation of mnemonic rule representation in macaque dorsolateral prefrontal cortex during a pro- and anti-saccade task. J. Neurosci. *35*, 16064–16076.

Schmitz, T.W., and Nathan Spreng, R.; Alzheimer's Disease Neuroimaging Initiative (2016). Basal forebrain degeneration precedes and predicts the cortical spread of Alzheimer's pathology. Nat. Commun. *7*, 13249.

Sperling, R., Mormino, E., and Johnson, K. (2014). The evolution of preclinical Alzheimer's disease: implications for prevention trials. Neuron *84*, 608–622.

Vijayraghavan, S., Major, A.J., and Everling, S. (2018). Muscarinic M1 receptor overstimulation disrupts working memory activity for rules in primate prefrontal cortex. Neuron *98*, this issue, 1256–1268.

Zen and the Art of Making a Bayesian Espresso

Lei Zhang,[1] Saša Redžepović,[1] Michael Rose,[1] and Jan Gläscher[1,]*
[1]Institute for Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Martinistrasse 52, 20246 Hamburg, Germany
*Correspondence: glaescher@uke.de

https://doi.org/10.1016/j.neuron.2018.06.023

In this issue of *Neuron*, Konovalov and Krajbich (2018) argue that a Bayesian inference is employed when learning new sequences and identify distinct brain networks that track the uncertainty of both the current state and the underlying pattern structure.

> "You look at where you're going and where you are and it never makes sense, but then you look back at where you've been and a pattern seems to emerge."
>
> Robert Pirsig, *Zen and the Art of Motorcycle Maintenance*

Our lives are shaped by rituals, little sequences of states and actions that give rhythm to our days. Preparing breakfast, driving to work, lunch with colleagues, getting groceries, and preparing dinner—a substantial amount of our behavior is driven by invariant scripts for specific situations. Here's the most treasured ritual of my morning: turning on the espresso machine, grinding the beans, smelling that scent of fresh coffee, tamping the grounds and adjusting the portafilter in the brew head, pulling the shot, seeing the crema rise in the cup, and finally tasting the first sip. I always repeat this sequence in the same way, and for me there is no better way to start the day. However, this ritual is also part of several longer sequences: preparing a latte, a cappuccino, or a flat white all include pulling a shot of espresso. For an outside observer, it is only in the course of the sequence that it becomes clear what this final product will be.

In the current issue of *Neuron*, Konovalov and Krajbich (2018) tackle the question of how our brains learn to detect deterministic sequences like the one above in the stream of sensory inputs. To this aim and in contrast to the classic serial reaction time task, which is commonly used to investigate implicit sequence learning (Nissen and Bullemer, 1987), they employed a novel pattern detection task where subjects were instructed to predict upcoming stimuli in a stream of images. These images were configured as patterns of different lengths or no pattern (Figure 1B). A common finding in probabilistic (e.g., Bornstein and Daw, 2012) and deterministic sequence learning (e.g., Rose et al., 2010) is a sharp drop in reaction time (RT) that a participant needs to detect the next stimulus, indicating that the sequence has been learned. In the current study, RT drops serve as an indicator that a pattern of a certain length has been detected. In addition, each image was shown as an animation, starting with a scrambled version of all images from which the current image gradually emerges. This clever manipulation prolonged reaction times, leading to better detection of RT drops in each condition and hence to more meaningful and interpretable behavioral data. Indeed, Konovalov and Krajbich (2018) observed significant RT drops when they compared pattern against no pattern, and the effect also depended on the pattern lengths, as the RT drop occurred earlier for shorter sequences.

What accounts for this effect? One of the seminal contributions of their work is the hierarchical Bayesian pattern learning model, whereby individuals update their beliefs about both the structure (i.e., pattern length) and the state (i.e., upcoming image). The structure-level belief informs the state-level belief, guiding participants' prediction of the next image (face, F; house, H; or landscape, L), which differs for patterns of different lengths (Figure 1A). For example, if the observed image sequence F, L, F was part of a length-2 sequence, then L would be the prediction for the next image; if it was part of a length-3 sequence, then F would be the next predicted stimulus (Figure 1B). Depending on the structure-level belief (is it a 2- or 3-length pattern?), the Bayesian learner would compute different state-level beliefs (is next image

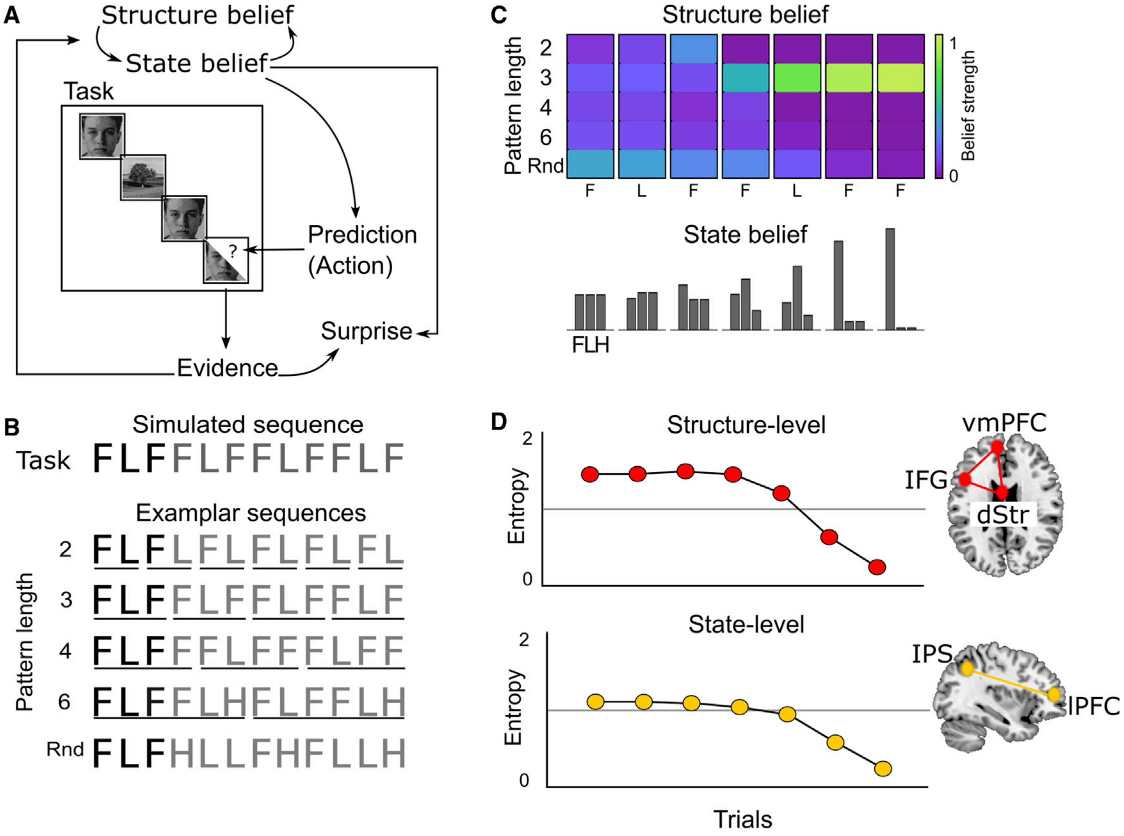

Figure 1. Hierarchical Bayesian Belief Updating in Deterministic Sequence Learning

(A) Hierarchical Bayesian pattern learning model. Structure-level beliefs (about the pattern length) are informing state-level beliefs (about the identity of the next image), which give rise to predictions about the next stimulus in the sequence. Sensory evidence then updates the beliefs at both levels.

(B) Simulated data and possible corresponding exemplar sequences. The top row shows a simulated sequence of pattern length 3 (F, face; L, landscape; H, house). Black and gray letters represent observed and yet-to-be-seen images, respectively. The rows below display corresponding exemplar sequences for each possible pattern length. After observing the first three images (F, L, F), all pattern lengths are still possible. Certain exemplar sequences become less likely with more incoming evidence (e.g., on trial 4 [F, L, F, F], pattern length 2 [F, L, F, L] becomes unlikely).

(C) Belief updates. The top row shows the structure-level belief for the simulated data in (B) for the first 7 trials. Initially, the structure belief is distributed with high uncertainty, but as the evidence accumulates, the specific belief in pattern length 3 is strengthened, while it decreases for all others. The next row shows the corresponding state-level beliefs (summed across all structure patterns), which also become increasingly specific for a pattern length 3.

(D) Entropy and the brain networks. As the high uncertainty in the beginning gets resolved by accumulating evidence, both structure and state entropy derived from (C) declines, which correlates with activation in specific and distinct brain networks.

F or L?). Conceptually, this hierarchical Bayesian model is similar to model-based learning (see Dolan and Dayan, 2013, for a review), where the structure beliefs represent the model of the task, which governs the computation of expected values for each state (i.e., image).

The model assumes a common prior belief distribution for the entire structure space (i.e., how likely each pattern length is), and as more information is gathered, this belief is updated and able to predict whether a structure exists, and if so, which pattern length the sequence has. For example, at the seventh trial (having observed six images in a pattern length-3 sequence F, L, F, F, L, F, previously) a successful Bayesian learner would assign a probability of 0 to pattern lengths 2 and 4 and only update non-zero beliefs about the other patterns (Figure 1C, top). Conditional on this structure-level belief, the Bayesian learner then computes the probability of observing each image on the next trial (i.e., the state-level belief; Figure 1C, bottom). Following previous literature, Konovalov and Krajbich (2018) regressed participants' RT on Shannon entropy (e.g., Bornstein and Daw, 2012) and Shannon surprise (e.g., Harrison et al., 2006) derived from both levels of beliefs. Entropy is a measure of global uncertainty about a belief distribution, which is maximal when the beliefs about the pattern lengths and images are equally likely (Figure 1C, first trial). As some of pattern lengths become more and more unlikely as the sequence unfolds, the entropy declines (Figure 1D). Shannon surprise, on the other hand, measures how much an observed stimulus in the sequence violates the state-level beliefs (Figure 1A) and is conceptually similar to a prediction error computed during reinforcement learning. Konovalov and Krajbich (2018) showed that RTs on their task (including their drops) were well explained by the state-level entropy and surprise, but not by the structure-level entropy. This finding is in itself a bit surprising but can probably be explained by the high correlation between the two levels of entropy: structure-level uncertainty about the pattern length will also result in state-level uncertainty about which image comes next (Figure 1D).

Having the computational model at hand, the next question is straightforward: how are the two levels of entropy implemented in the brain? Using model-based functional neuroimaging, Konovalov and Krajbich (2018) first identified a frontoparietal network, including the intraparietal sulcus (IPS) and bilateral lateral prefrontal cortex (lPFC), that encoded the state-level entropy (Figure 1D, bottom). This largely replicates the previous findings on the neural signatures of state prediction error (cf. IPS; Gläscher et al., 2010) and the arbitration between model-free learning and model-based learning (cf. lPFC; Lee et al., 2014). Notably, activations in both the IPS and the lPFC monotonically increased as the pattern became more complex, demonstrating that the frontoparietal network is not only tied to state-level uncertainty but also reflects the pattern length.

Furthermore, Konovalov and Krajbich (2018) examined the neural correlates of the structure-level belief updating derived from their hierarchical Bayesian model, which has rarely been studied in sequence learning. They show a distinct and separate network of brain regions tracking structure-level entropy (and its inverse) including, but not limited to, the dorsal striatum (dStr), left inferior frontal gyrus (IFG), and ventromedial prefrontal cortex (vmPFC; Figure 1D, top). The latter also tracked the different pattern lengths (including the no pattern condition), suggesting that the role of this area goes beyond simple reward coding. Conceptually, this demonstrates the richness of neuroimaging data and its capacity to uncover subtle but meaningful differences in model-based signals that remain unnoticed in the single-dimensionality of behavioral data (Wilkinson and Halligan, 2004).

Taken together, this study provides a novel perspective on the hierarchical nature of deterministic sequence learning (e.g., making an espresso and turning it into a latte or a flat white), which is no less important to our daily life than detecting transition strength in stochastic sequences. The proposed Bayesian pattern learning model tracks both the current state and the underlying structure of the sequence, by integrating the prior beliefs about all possible structures with the history of the observed states. Intriguingly, the computational signals of both levels belief entropy are encoded in two distinct networks of brain regions, replicating findings on model-based planning, valuation, and expectancy violations. This study also demonstrates the neural evidence for an additional layer of structure-level reasoning in sequence learning and supports the general hypothesis that our brains employ Bayesian inference in complex structure learning (Gershman and Niv, 2010).

It is worth noting that although the entropy accounts for the behavioral effects captured by RT, it does not reflect the exact computations when updating beliefs, hence weakening a precise conclusion of the computational function of each brain region. In addition, detailed connectivity analysis in future studies can unravel the functional interactions of the nodes of each network, thus substantiating how other decision-related variables (e.g., action probabilities) are computed in the brain. Furthermore, in the real world, the pattern space can be seemingly large and even reach infinity; thus, it is worth investigating how initial beliefs are formulated and how individuals effectively downscale irrelevant dimensions and represent solely the appropriate information. Attentional selection could be a key process in accomplishing this goal, and incorporating it into the hierarchical Bayesian pattern learning model may provide invaluable insight into more complex structure learning, for both deterministic sequences and stochastic sequences.

In sum, Konovalov and Krajbich (2018) provide an elegant Bayesian inference account of deterministic sequence learning, shedding light on the potential mechanisms of how our brains uncover meaningful structures in a complex environment using a limited amount of evidence.

REFERENCES

Bornstein, A.M., and Daw, N.D. (2012). Dissociating hippocampal and striatal contributions to sequential prediction learning. Eur. J. Neurosci. *35*, 1011–1023.

Dolan, R.J., and Dayan, P. (2013). Goals and habits in the brain. Neuron *80*, 312–325.

Gershman, S.J., and Niv, Y. (2010). Learning latent structure: carving nature at its joints. Curr. Opin. Neurobiol. *20*, 251–256.

Gläscher, J., Daw, N., Dayan, P., and O'Doherty, J.P. (2010). States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. Neuron *66*, 585–595.

Harrison, L.M., Duggins, A., and Friston, K.J. (2006). Encoding uncertainty in the hippocampus. Neural Netw. *19*, 535–546.

Konovalov, A., and Krajbich, I. (2018). Neuron *98*, this issue, 1282–1293.

Lee, S.W., Shimojo, S., and O'Doherty, J.P. (2014). Neural computations underlying arbitration between model-based and model-free learning. Neuron *81*, 687–699.

Nissen, M.J., and Bullemer, P. (1987). Attentional requirements of learning: evidence from performance measures. Cognitive Psychol. *19*, 1–32.

Rose, M., Haider, H., and Büchel, C. (2010). The emergence of explicit memory during learning. Cereb. Cortex *20*, 2787–2797.

Wilkinson, D., and Halligan, P. (2004). The relevance of behavioural measures for functional-imaging studies of cognition. Nat. Rev. Neurosci. *5*, 67–73.