

are independent of PKA. This is again distinct from the CA1 synapse where direct phosphorylation of GluA1 is important to alter the channel conductance and/or trafficking of AMPARs (Huganir and Nicoll, 2013). Interestingly, the GluA3-dependent LTP requires Epac, but how this protein modulates the channel properties of GluA3 containing AMPARs remains to be investigated. Finally, the study reveals that GluA3, not GluA1, plays a key role in motor learning, in contrast to the acquisition of declarative memories such as fear memory where GluA1 is essential (Kessels and Malinow, 2009).

In summary, the study by Gutierrez-Castellanos et al. (2017) has provided compelling evidence that GluA3 plays an essential role in LTP at the PF-PC synapse and motor learning (Figure 1). This GluA3-dependent synaptic plasticity is distinct from the GluA1-driven LTP in the hippocampus in that it involves exclusively enhanced chan-

nel conductance, but not trafficking, of AMPARs. It would be interesting to know whether this form of plasticity also exists in other regions of the brain. Another outstanding question is how the activation of cAMP-Epac signaling leads to changes in the channel properties of AMPARs. It would also be interesting to know whether this mechanism interacts with GluA2-mediated receptor trafficking shown to occur during LTD at this synapse. Clearly there is much still to do, but the findings of Gutierrez-Castellanos et al. (2017) add to the diversity of mechanisms of synaptic plasticity that underlies forms of learning and memory in the CNS.

REFERENCES

- Benke, T.A., Lüthi, A., Isaac, J.T., and Collingridge, G.L. (1998). *Nature* 393, 793–797.
- Gao, Z., van Beugen, B.J., and De Zeeuw, C.I. (2012). *Nat. Rev. Neurosci.* 13, 619–635.
- Gutierrez-Castellanos, N., Da Silva-Matos, C.M., Zhou, K., Canto, C.B., Renner, M.C., Koene, L.M.C., Ozyildirim, O., Sprengel, R., Helmut, W., Kessels, H.W., and De Zeeuw, C.I. (2017). *Neuron* 93, this issue, 409–424.
- Henley, J.M., and Wilkinson, K.A. (2016). *Nat. Rev. Neurosci.* 17, 337–350.
- Huganir, R.L., and Nicoll, R.A. (2013). *Neuron* 80, 704–717.
- Jia, Z., Agopyan, N., Miu, P., Xiong, Z., Henderson, J., Gerlai, R., Taverna, F.A., Velumian, A., MacDonald, J., Carlen, P., et al. (1996). *Neuron* 17, 945–956.
- Kessels, H.W., and Malinow, R. (2009). *Neuron* 61, 340–350.
- Meng, Y., Zhang, Y., and Jia, Z. (2003). *Neuron* 39, 163–176.
- Wu, Y., Arai, A.C., Rumbaugh, G., Srivastava, A.K., Turner, G., Hayashi, T., Suzuki, E., Jiang, Y., Zhang, L., Rodriguez, J., et al. (2007). *Proc. Natl. Acad. Sci. USA* 104, 18163–18168.
- Zamanillo, D., Sprengel, R., Hvalby, O., Jensen, V., Burnashev, N., Rozov, A., Kaiser, K.M., Köster, H.J., Borchardt, T., Worley, P., et al. (1999). *Science* 284, 1805–1811.

A Two-Way Street between Attention and Learning

Tessa Rusch,¹ Christoph W. Korn,¹ and Jan Gläscher^{1,*}

¹Institute for Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Martinistrasse 52, 20246 Hamburg, Germany

*Correspondence: glaescher@uke.de

<http://dx.doi.org/10.1016/j.neuron.2017.01.005>

In an elegant model-based fMRI study, Leong et al. (2017) demonstrate how attention and learning interact to facilitate value-based decision-making. They combine computational modeling with empirical measures of attentional selection derived from eye-tracking data and multivariate pattern analyses.

A fundamental coordinate of behavior is searching for reward while avoiding punishment. Reinforcement learning (RL) describes how agents achieve the goal of maximizing long-term future return in unknown environments (Sutton and Barto, 1998). The framework of RL has proven indispensable for our understanding of learning from a behavioral as well as a neurobiological perspective. Despite its wide-spreading success, RL suffers from an inherent shortcoming: RL maps potential reward to *all* available states and actions—making it computationally intractable when scaled up to high-dimensional learning problems (Sutton and Barto, 1998). Yet such an extensive mapping may be disproportionate for many problems, because even in exceedingly complex environments only some aspects tend to be relevant. Given the success with which humans act within their surroundings, they seem to extract the pertinent aspects of their environments to simplify learning and decision-making. Thus, humans are likely to make use of selection mechanisms to downscale high-dimensional environ-

ments to a manageable size. If selection mechanisms guide learning and decision-making, how do agents decide what is relevant for the problem at hand, and what should be dismissed?

In the current issue of *Neuron*, Leong et al. (2017) propose that humans use selective attention to cope with the dazzling complexity of their surroundings. Their study suggests that attention acts as a filter on the learning problem by amplifying a subset of environmental states while dampening irrelevant ones. Crucially, they propose a two-way street

ments to a manageable size. If selection mechanisms guide learning and decision-making, how do agents decide what is relevant for the problem at hand, and what should be dismissed?

In the current issue of *Neuron*, Leong et al. (2017) propose that humans use selective attention to cope with the dazzling complexity of their surroundings. Their study suggests that attention acts as a filter on the learning problem by amplifying a subset of environmental states while dampening irrelevant ones. Crucially, they propose a two-way street




Figure 1. Sketch of the Mutual Influence of Attention and Learning as Proposed in Leong et al.

Attention (derived from measures of eye-tracking data and from multivariate decoding of object categories in the ventral visual stream) affects the computation of compound values during choice and the updating of individual feature values. This leads to an amplification of the reward-predicting feature in the attended category. Meanwhile, learning of feature values affects the allocation of attention toward object categories, which gradually increases the attentional focus during learning.

between attention and learning: attention facilitates learning, and learned values in turn inform attentional selection.

To investigate the potential interplay between learning and attention, [Leong et al. \(2017\)](#) presented subjects with a multifaceted probabilistic learning task which bears resemblance to the Wisconsin Card Sorting Task ([Berg, 1948](#)): on each trial, participants chose between three compound stimuli. Compounds were defined by three different features along the object categories faces, landmarks, and tools (e.g., Einstein, Big Ben, and screwdriver, [Figure 1](#)). Within a given block of the task, only one object category (the target category, e.g., faces) indicated a probabilistic reward, and within that category only one feature (the target feature, e.g., Einstein) was highly predictive of the reward, whereas the remaining

two features (e.g., Clooney and Lincoln) would rarely precede reward. Features from the remaining two categories were completely unrelated to reward. Participants were instructed about these regularities in the task.

Previously, [Niv and colleagues \(2015\)](#) demonstrated in a similar experimental setting that a feature RL model outperformed a variety of alternative models in predicting subjects' behavior. Feature RL implies a drastic downscaling of the learning problem in comparison to a naive RL model ([Sutton and Barto, 1998](#)). Within the latter model, the values of *all* possible compound stimuli have to be represented and updated (i.e., all possible combinations of nine features within three compound objects, resulting in $9 \times 3 = 27$ compounds in the current task). In contrast, within the feature RL model, choices and

learning depend on the nine feature values alone. Decisions are based on the values of the multi-categorical compound objects, which are in turn linear combinations of the individual feature values. After each choice, learning takes place by updating the chosen feature values according to a prediction error signal. This earlier work provided evidence that humans structure high-dimensional environments according to task requirement but did not isolate the selection mechanism that guides learning and choice.

Here, [Leong et al. \(2017\)](#) investigate whether selective attention facilitates the learning process by boosting the relevant category. To directly test the interaction between attention and learning, Leong et al. derived two *empirical* measures of selective attention—one from eye-tracking data and another from applying

multivariate pattern analysis (MVPA) in a creative way. Stimulus categories were faces, landmarks, and tools. MVPA on the inferior temporal cortex can easily discriminate these object categories (Haxby et al., 2001). So far, most studies have used MVPA to decode stimulus representations from BOLD activation patterns along the visual processing stream. Leong et al. apply this technique in order to decode selective attention to the different object categories (Figure 1). This demonstrates one of the innovative approaches of this paper. In contrast to earlier studies that had to infer attentional selection from a computational model (Niv et al., 2015; Yu and Dayan, 2005), here the authors have a direct, empirical, trial-by-trial measure of selective attention at their disposal. In different variants of the feature RL model they used this attention measure to bias either the calculation of the compounds' expected values or the reward prediction errors that update the feature values—or both. Model comparison indicated that selective attention influences both choice and learning (Figure 1).

In accordance with the behavioral findings, Leong et al. (2017) found the winning model to best reflect their neural data. Expected value representations in the ventromedial prefrontal cortex (vmPFC, see Clithero and Rangel, 2014, for a review; Figure 1) and prediction error representations in the ventral striatum (see Garrison et al., 2013, for a review) co-varied with the value and prediction error estimates retrieved from the model that included an attention bias on choice and learning, but not from any of the other models.

Having demonstrated the influence of selective attention on learning, Leong et al. (2017) closed the circle by investigating how attention itself evolves over time. To obtain a direct proxy of attention bias, they computed the standard deviation over the attention weights of the three object categories. A low standard deviation points toward a fairly even distribution of attention over object categories (i.e., a low attention bias), whereas a high value indicates focused attention on a subset of the object categories (i.e., high attention bias). Over the course of an experimental block, the standard deviation increased, suggest-

ing that the attentional focus sharpened in the course of learning (Figure 1). Furthermore, the most attended category correlated with the category containing the highest feature value, pointing toward an interaction between feature value and attention bias. These findings were corroborated by a model-based analysis of attention: the comparison of different models that predict dynamic allocation of attention revealed that the distribution of attention over categories depended on learned feature values (and not on recent choice or reward history; Figure 1).

An interaction between value learning and attention allocation was affirmed by an analysis of the neural data during attention switches. An attention switch was defined as a change in the maximally attended object category. Such switches were more abundant when the standard deviation over the attention weights was low. Accordingly, the BOLD signal in areas attributed to the network of top-down attentional control (Corbetta and Shulman, 2002) was greater during switch than during stay trials. Connectivity analyses revealed an anti-correlation between vmPFC and attention-related areas on stay trials. This suggests that increased value signals in the vmPFC lead to decreased activity in the attention network, which may result in a reduced tendency to switch the attentional focus.

Taken together, these results persuasively show that selective attention is a key mechanism for managing complex learning problems. Furthermore, attention does not act in a single-sided top-down manner but is itself informed by the learning process.

The proposed interaction between attention and learning offers new perspectives on the framework of RL. In the study by Leong et al. (2017), stimuli were carefully controlled. Attentional selection arose purely from the learned values and fed back into learning process in a “top-down fashion.” In many (real-life) scenarios, bottom-up processes (e.g., a colored item among gray items) or top-down processes due to previous experiences (e.g., having seen the iconic photograph of Einstein with the tongue stretched out countless times) are likely to contribute to the formation of attention biases in various manners.

Furthermore, in many situations regularities that may reduce the state space are not known a priori. For instance, rather than an individual feature a specific combination of features could predict reward (e.g., Clooney with a screwdriver but neither Clooney nor screwdriver on their own). In such cases, participants would have to concurrently track several competing hypotheses about the informative aspects of the environment. Over time, they would need to narrow down the hypothesis space and exploit the best option (Donoso et al., 2014). Hypothesis testing could be dynamically interconnected with attentional processes in a similar fashion as feature learning and selective attention.

From a methodological perspective, the current study introduces MVPA as a novel measure of attention. Future studies will likely refine this metric to further specify the attentional selection process. For example, instead of deriving attention measures for the entire object categories, attention on individual features within an object category could be examined.

In sum, the study by Leong et al. (2017) provides an elegant model and innovative methods to gain further insight into the interplay between top-down and bottom-up attentional selection and value learning in complex environments.

REFERENCES

- Berg, E.A. (1948). *J. Gen. Psychol.* 39, 15–22.
- Clithero, J.A., and Rangel, A. (2014). *Soc. Cogn. Affect. Neurosci.* 9, 1289–1302.
- Corbetta, M., and Shulman, G.L. (2002). *Nat. Rev. Neurosci.* 3, 201–215.
- Donoso, M., Collins, A.G., and Koehlin, E. (2014). *Science* 344, 1481–1486.
- Garrison, J., Erdeniz, B., and Done, J. (2013). *Neurosci. Biobehav. Rev.* 37, 1297–1310.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). *Science* 293, 2425–2430.
- Leong, Y.C., Radulescu, A., Daniel, R., DeWoskin, V., and Niv, Y. (2017). *Neuron* 93, this issue, 451–463.
- Niv, Y., Daniel, R., Geana, A., Gershman, S.J., Leong, Y.C., Radulescu, A., and Wilson, R.C. (2015). *J. Neurosci.* 35, 8145–8157.
- Sutton, R.S., and Barto, A.G. (1998). *Reinforcement Learning: An Introduction* (MIT Press).
- Yu, A.J., and Dayan, P. (2005). *Neuron* 46, 681–692.