

Journal Club

Editor's Note: These short reviews of recent *JNeurosci* articles, written exclusively by students or postdoctoral fellows, summarize the important findings of the paper and provide additional insight and commentary. If the authors of the highlighted article have written a response to the Journal Club, the response can be found by viewing the Journal Club at www.jneurosci.org. For more information on the format, review process, and purpose of Journal Club articles, please see <http://www.jneurosci.org/content/jneurosci-journal-club>.

Unraveling the Neural Mechanisms Which Encode Rapid Streams of Visual Input

 William Turner

Melbourne School of Psychological Sciences, University of Melbourne, Melbourne, 3010 Victoria, Australia
Review of King and Wyart

As you read these words, light is striking your retinas, triggering chemical and electrical cascades. For the information contained within these wavelengths to be of use to you, your brain must encode, process, and transform it into more abstract representations (Treisman, 1986). Importantly, this sense-making process takes time. Carefully controlled experiments have shown that a single flashed image can generate reliable cortical responses that last ~ 1 s (Carlson et al., 2013). In everyday situations, our eyes do not have the luxury of receiving distinct packets of information at convenient 1 s intervals, and are instead bombarded with continuous inputs. How then do our brains encode and process rapid streams of visual events without mixing and muddying this information?

In a recent paper published in *The Journal of Neuroscience*, King and Wyart (2021) investigated this question using EEG recordings from 15 healthy adult participants. Neural activity was recorded while participants viewed ~ 5000 Gabor patches (black and white gratings tilted at random angles) flashed every 250 ms.

Importantly, each successive stimulus was uncorrelated from its neighbors, allowing the researchers to examine how distinct stimuli were encoded in the brain over time.

Given that we can function in the face of continuous visual input, our brains must be able to encode stimulus representations, which persist even as new stimuli appear. As a first step, King and Wyart (2021) elegantly showed that it is possible to detect these lasting representations via whole-brain decoding analyses, even for stimuli flashed in rapid succession. To examine the encoding of stimulus-specific information, the authors fit multivariate linear regression models to predict the orientation of each presented stimulus from the voltage of all EEG electrodes. They also examined whether the change in orientation between stimuli could be predicted, as influential theories suggest that unexpected changes in input may be prioritized when encoding information (Rao, 1999; Friston, 2005). With cross-validation, they assessed the accuracy of these predictions and found that, across participants, both stimulus angles and successive angular differences could be decoded between ~ 50 and 950 ms after stimulus onset. Critically, at any given moment, between 2 and 5 stimulus angles and 2 and 4 angular differences could be decoded. This suggests that multiple stimuli were simultaneously encoded in participants' neural activity.

One way that multiple, distinct stimulus representations could be simultaneously encoded in the brain is if information were

routed through different cortical regions over time. It has long been known that visual information cascades along a "visual hierarchy" of different brain areas after leaving the retina. As a second step, King and Wyart (2021) built on this knowledge and demonstrated that stimulus-specific activity appeared to flow along the visual hierarchy, in an apparent "traveling wave." Specifically, they used encoding models and source localization techniques to show that stimulus information traveled from "low-level" regions in the occipital cortex to more "high-level" inferotemporal and dorsoparietal areas. While one must be careful when making inferences about source-localized EEG signals, these findings are consistent with the view that, at any given moment, multiple "snapshots" of the recent visual past are encoded along the visual hierarchy.

The visual system can be thought of as a network of interconnected neural populations set within distinct "layers" (different cortical regions; e.g., retina, lateral geniculate nucleus, primary visual cortex, etc.). Visual information enters at the retina and flows through the network over time. Thinking in this way, questions arise regarding the structure of this network. For example, what neural populations are present at each layer, and what connections exist both within and between layers? As the third and final step in their paper, King and Wyart (2021) considered these questions and developed a computational framework that attempts to model the

Received Oct. 5, 2021; revised Dec. 2, 2021; accepted Dec. 10, 2021.

This work was supported by Australian Research Council FT200100246 and DP180102268 to Hinze Hogendoorn. I thank Daniel Feuerriegel, Hinze Hogendoorn, Caoimhe Moran, Melinda Turner, Bruce Turner, and Morgan Weaving for helpful comments and feedback on the initial draft manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to William Turner at william.turner@unimelb.edu.au.

<https://doi.org/10.1523/JNEUROSCI.2013-21.2021>

Copyright © 2022 the authors

macroscopic structure and dynamics of the visual hierarchy.

This modeling framework consists of a hierarchical system, all candidate configurations of which share the same core structure: 10 layers with one observable excitatory unit (x) and one unobservable inhibitory unit (y) per layer. The activity of these units represents the postsynaptic potentials of aligned pyramidal neurons (x) and unaligned (and therefore unobservable with EEG) interneurons (y). To explore different network architectures, the authors systematically searched across configurations of within- and between-layer connections, connection weights, and unit activation functions. For each candidate network, they considered the activity patterns that would be observed if excitatory unit activity was recorded with EEG.

To distinguish between the activity patterns of different networks, the authors used temporal generalization analyses (King and Dehaene, 2014). This involves training decoders on data from specific time points and examining the accuracy of these “temporal decoders” when tested at different time points. This can reveal rich information about the dynamics underlying time-resolved neural data. For example, if a decoder accurately generalizes from one time point to another, this suggests that similar populations of neurons are active at each time point. However, if a decoder does not generalize (or generalizes with below-chance accuracy), then this suggests that the patterns of neural activity significantly differ (or are opposing).

The authors first performed temporal generalization analyses on participants’ EEG recordings, revealing three key details. First, stimulus onset evoked transient waves of activity, evidenced by a diagonal pattern of temporal generalization within a training by testing time matrix (consistent with hierarchical processing, King and Wyart, 2021, their Fig. 3). Second, stimulus offset (and the concurrent onset of subsequent stimuli) evoked transient waves of opposite amplitude, evidenced by flanking diagonals of below-chance decoding. Finally, maintenance of stimulus-related information increased over time, evidenced by increasing peri-diagonal generalization.

To find the simplest model that could account for these patterns, the authors searched through increasingly complex network architectures. After searching through ~1.5 billion models, two valid models were found. Both were feedforward, locally recurrent networks in which stimulus information was maintained in inhibitory units and only briefly observable in

excitatory units following changes in input. Somewhat counterintuitively, this suggests that stimulus-specific information was maintained in populations of neurons undetectable with EEG.

In appraising the work of King and Wyart (2021), it is important to consider the generality of their findings. A series of recent studies have shown that it is possible to decode multiple representations of distinct object-images at rapid presentation rates (5 and 20 Hz) (Grootswagers et al., 2019a,b; Robinson et al., 2019). Therefore, the broad observation that multiple visual events are simultaneously encoded along the visual hierarchy holds for different stimuli and presentation rates.

Interestingly, Grootswagers et al. (2019a,b) found that late-stage stimulus processing was impaired at ultra-rapid presentation rates (20 Hz). This processing impairment might be explained by the finding of King and Wyart (2021) that information is increasingly maintained at higher levels of the visual hierarchy, as prolonged maintenance may increase the likelihood of interference from subsequent stimuli. However, since this previous work predates the development of the current modeling framework, it remains to be seen whether the network architectures uncovered by King and Wyart (2021) can capture the neural dynamics observed in these past studies, as well as those associated with the processing of visual features other than orientation, more generally.

King and Wyart (2021) used random stimulus sequences to disentangle neighboring stimulus representations. However, in real-world scenes, visual inputs are often temporally correlated or contextually predictable. Because it takes time for the brain to process information, there is a lag between when an event happens in the “outside world” and when it is registered in the brain. To overcome this, the brain may attempt to exploit spatiotemporal correlations or contextual expectations to predict future inputs and better align neural representations with the outside world (e.g., Hogendoorn and Burkitt, 2019, their Fig. 2). Supporting this theory, studies have found evidence of predictive mechanisms within human and nonhuman visual systems.

For example, studies of neural responses to predictable motion have found that retinal ganglion cells in salamanders, rabbits, and monkeys (e.g., Berry et al., 1999; Liu et al., 2021), as well as neurons in the primary visual cortex (V1) of cats and monkeys (e.g., Jancke et al., 2004; Benvenuti et al.,

2020), display anticipatory firing patterns. That is, they appear to preempt the arrival of stimuli that are moving along predictable trajectories. On a more macroscopic level, recent human EEG and MEG experiments have found evidence of anticipatory “pre-activation” of stimulus representations over the visual cortex, for predictable stimulus sequences (e.g., Kok et al., 2017; Blom et al., 2020). Finally, human fMRI and transcranial magnetic stimulation studies have also shown that feedback between layers of the visual system (e.g., between areas V5/MT and V1) appears to mediate the formation of predictive visual representations (Sterzer et al., 2006; Vetter et al., 2015).

Collectively, these studies provide strong evidence that predictive mechanisms do exist within human and nonhuman visual systems. However, since King and Wyart (2021) used unpredictable stimuli, the dynamics of such mechanisms will not have been detectable. Future research should therefore investigate the neural dynamics associated with predictable visual sequences, using the current modeling approach (for a discussion of how the effects of adaptation, expectation, and additional confounding factors could be disentangled in such research, see Feuerriegel et al., 2021). By harnessing King and Wyart’s (2021) modeling framework, this research could provide a unique, intermediary view of the neural architecture underlying predictive visual processing, to complement and begin bridging the coarser views given by typical EEG and MEG experiments, and the more fine-grained views given by invasive single-cell neurophysiology experiments.

When visual inputs are predictable, different neural dynamics may be observed and alternative processing architectures may be revealed. For example, it may be necessary to consider network architectures that include feedback and/or lateral connections (and the requisite additional spatial dimension/s), to fully account for participants’ neural activity dynamics. Empirically, if stimulus representations are predictively “pre-activated” (Blom et al., 2020), one may see increased “off-diagonal” generalization for temporal decoders trained on data from unpredictable visual sequences and tested on data from predictable sequences. That is, predictable stimulus representations may propagate more rapidly along the visual hierarchy than unpredictable ones, as predictive mechanisms attempt to compensate for processing delays (Nijhawan and Wu, 2009). Importantly, where novel dynamics are observed, King and Wyart’s (2021) modeling framework will provide a means of

revealing the neural architectures which likely underlie them.

In conclusion, the work of King and Wyart (2021) elegantly reveals how multiple “snapshots” of recent visual events are simultaneously encoded along the visual hierarchy. Moreover, their modeling framework provides neuroimaging researchers with a powerful tool for characterizing the dynamical system architectures underlying time-resolved neural recordings.

References

- Benvenuti G, Chemla S, Boonman A, Perrinet L, Masson GS, Chavane F (2020) Anticipatory responses along motion trajectories in awake monkey area V1. *BioRxiv*. doi: 10.1101/2020.03.26.010017.
- Berry MJ, Brivanlou IH, Jordan TA, Meister M (1999) Anticipation of moving stimuli by the retina. *Nature* 398:334–338.
- Blom T, Feuerriegel D, Johnson P, Bode S, Hogendoorn H (2020) Predictions drive neural representations of visual events ahead of incoming sensory information. *Proc Natl Acad Sci USA* 117:7510–7515.
- Carlson T, Tovar DA, Alink A, Kriegeskorte N (2013) Representational dynamics of object vision: the first 1000 ms. *J Vis* 13:1.
- Feuerriegel D, Vogels R, Kovács G (2021) Evaluating the evidence for expectation suppression in the visual system. *Neurosci Biobehav Rev* 126:368–381.
- Friston K (2005) A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360:815–836.
- Grootswagers T, Robinson AK, Carlson T (2019a) The representational dynamics of visual objects in rapid serial visual processing streams. *Neuroimage* 188:668–679.
- Grootswagers T, Robinson AK, Shatek SM, Carlson T (2019b) Untangling featural and conceptual object representations. *Neuroimage* 202:116083.
- Hogendoorn H, Burkitt AN (2019) Predictive coding with neural transmission delays: a real-time temporal alignment hypothesis. *ENeuro* 6:ENEURO.0412-18.2019.
- Jancke D, Erlhagen W, Schöner G, Dinse HR (2004) Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex. *J Physiol* 556:971–982.
- King JR, Dehaene S (2014) Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn Sci* 18:203–210.
- King JR, Wyart V (2021) The human brain encodes a chronicle of visual events at each instant of time through the multiplexing of traveling waves. *J Neurosci* 41:7224–7233.
- Kok P, Mostert P, Lange FP (2017) Prior expectations induce prestimulus sensory templates. *Proc Natl Acad Sci USA* 114:10473–10478.
- Liu B, Hong A, Rieke F, Manookin MB (2021) Predictive encoding of motion begins in the primate retina. *Nat Neurosci* 24:1280–1290.
- Nijhawan R, Wu S (2009) Compensating time delays with neural predictions: are predictions sensory or motor? *Philos Trans A Math Phys Eng Sci* 367:1063–1078.
- Rao RP (1999) An optimal estimation approach to visual perception and learning. *Vision Res* 39:1963–1989.
- Robinson AK, Grootswagers T, Carlson T (2019) The influence of image masking on object representations during rapid serial visual presentation. *Neuroimage* 197:224–231.
- Sterzer P, Haynes J, Rees G (2006) Primary visual cortex activation on the path of apparent motion is mediated by feedback from hMT+/V5. *Neuroimage* 32:1308–1316.
- Treisman A (1986) Features and objects in visual processing. *Sci Am* 255:114–125.
- Vetter P, Grosbras MH, Muckli L (2015) TMS over V5 disrupts motion prediction. *Cereb Cortex* 25:1052–1059.