

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://jneurosci.org/content/preparing-manuscript#journalclub>.

Expectation Suppression Dampens Sensory Representations of Predicted Stimuli

Kevin S. Walsh and David P. McGovern

Trinity College Institute of Neuroscience, Trinity College Dublin, Dublin 2, Ireland
Review of Richter et al.

The historic debate surrounding the role of inferential mechanisms in perception has been stoked by recent technological and theoretical advances that have allowed researchers to formalize models of predictive processing and investigate their corollaries at every level from behavioral to neurobiological. These models contend that perceptual systems exploit statistical regularities in the environment to reconstruct a coherent and stable perceptual experience based on the noisy and ambiguous data received at sensory receptors (Hohwy, 2013). According to “predictive coding,” a cascade of descending predictions distills raw sensory data, scything through expected sensory input so that the remaining, unaccounted-for, information (termed prediction error) can be used to update stimulus representations and refine sensory predictions (Clark, 2013). According to this theory, the continuous process of prediction error minimization is played out across the levels of the neural processing hierarchy, dynamically tuning the relative influence of error and expectation to best account

for the full flood of stimulation at a given moment.

One of the key pieces of evidence cited as support for predictive processing models is the diminished neural responses to sensory stimuli rendered predictable by recent stimulus history, termed expectation suppression (Summerfield and de Lange, 2014). Expectation suppression is temporally distinguishable from repetition suppression and adaptation effects (e.g., Todorovic and de Lange, 2012) and has been demonstrated across sensory modalities in EEG, MEG, and BOLD signals, as well as spiking activity in nonhuman primates (de Lange et al., 2018). The debate about the neural circuitry underlying this process has spawned a number of proposed implementations (e.g., Spratling, 2017), with most suggesting a key role of feedback connections, terminating in infragranular and supragranular layers of lower cortical areas, in suppressing expected lower-level activity (Friston, 2009). However, there is still some debate about the specific nature of the mechanism supporting this effect.

Two primary accounts of expectation suppression have emerged in the literature (de Lange et al., 2018). According to the “dampening” account, activity representing expected input is diminished, leaving the information carried in the residual prediction error for selective processing (Friston, 2005). Others have argued for a “sharpening” account of ex-

pectation suppression, where this pattern of activity is reversed (e.g., Kok et al., 2012). Here, sensory responses to predicted stimuli are enhanced as discrepant information is attenuated, reducing the overall amplitude of activity and sharpening the representation of the anticipated input (de Lange et al., 2018). For example, Kok and de Lange (2014) reported enhanced activity in regions of primary visual cortex (V1) representing a perceived Kanizsa illusion (an illusory shape perceived to occlude veridical shape stimuli), whereas activity associated with the Pacman figures inducing the illusion was suppressed.

The difficulty in testing these competing hypotheses is that they both predict the reduction in activity associated with expectation suppression. However, their accounts of the information that is suppressed differ: “dampening” requires expectation suppression to disproportionately inhibit neurons tuned to the expected stimulus; “sharpening” requires it to primarily suppress activity in neurons not tuned to the expected stimulus.

Recently, Richter et al. (2018) used fMRI to investigate the characteristics of expectation suppression throughout the ventral visual stream and to adjudicate between the sharpening and dampening accounts, defining V1 and object-selective lateral occipital cortex (LOC) as *a priori* ROIs. Twenty-four participants were trained to associate eight pairs of object

Received Aug. 20, 2018; revised Oct. 20, 2018; accepted Oct. 27, 2018.

K.S.W. was supported by an IRC Government of Ireland Postgraduate Scholarship. We thank Dr. Redmond O'Connell for providing feedback on the manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. David P. McGovern, Trinity College Institute of Neuroscience, Trinity College Dublin, Dublin 2, Ireland. E-mail: mcgoved1@tcd.ie.

<https://doi.org/10.1523/JNEUROSCI.2133-18.2018>

Copyright © 2018 the authors 0270-6474/18/3810592-03\$15.00/0




Figure 1. An illustration of the three conditional probability conditions, where the leading image perfectly predicts the trailing image (1:1), two leading images predict the same trailing image (2:1), or a single leading image predicts two possible trailing images (1:2). In the main experiment, the images were each presented for 500 ms with no interstimulus interval and an intertrial interval of 4110–6300 ms. These were not the actual stimuli used in the experiment.

images in a statistical learning paradigm, where the first image predicted the second image according to one of three conditional probabilities: 1:1 (first image is perfectly predictive of the second); 1:2 (first image predicts one of two possible images); and 2:1 (two images predict a single second image; Fig. 1). The subsequent fMRI experiment was identical to the training conditions, except that unexpected images were occasionally presented following the first image, violating the predictive relationships described above. In addition, throughout the training and the main experiment, participants were asked to report the presence of inverted images, which occurred on 11% of trials. This ensured that participants paid attention to the stimuli but did not notice the relationships between images as the conditional probabilities were rendered task-irrelevant.

The authors first confirmed an expectation suppression effect, demonstrating that BOLD responses were larger for unexpected than for expected images in both V1 and area LOC. Furthermore, a whole-brain analysis showed widespread expectation suppression across the ventral visual stream, including bilateral fusiform gyrus, bilateral inferior parietal cortex, and right posterior parahippocampal gyrus. Although the authors reported that the expectation suppression in LOC was significantly stronger in voxels activated by the image stimuli than in nonactivated voxels, this difference was not significant in V1. The authors suggest that the observation that expectation suppression did not selectively suppress voxels underlying the image's sensory representation in V1 may be related to the complexity of the stimuli used in the study. That is, the expectation of a complex object image is not clearly defined at the level of processing characteristic of V1. In support of this explanation, previous fMRI studies using complex stimuli have failed to demonstrate expectation suppression in V1 (e.g., Pajani et al., 2017), whereas more basic

stimuli, such as sinusoidal gratings, have elicited this effect in V1 (Alink et al., 2010; e.g., Kok et al., 2012).

When investigating the effect of conditional probability on expectation suppression, the authors found no evidence of a modulation of activity across conditional probabilities in V1 or LOC and no interaction between conditional probability and expectation in either area. This is surprising given that expectations are only meaningful insofar as they are sensitive to conditional probabilities. In Bayesian models, a prior (expectation) associated with a cue that is 50% predictive (1:2 condition in Fig. 1) should have a distinguishable effect from one that is 100% predictive (1:1 condition) (Bowers and Davis, 2012). Furthermore, Egner et al. (2010) previously demonstrated sensitivity to the conditional probabilities of face stimuli in fusiform face area, reporting greater responses to low and medium probability faces (25% and 50%, respectively) than house stimuli, but equivalent responses to house stimuli and high probability (75%) faces. Richter et al. (2018) point out that sensitivity to conditional probability has also been previously demonstrated in monkey inferotemporal cortex (Ramachandran et al., 2016) and suggested that their analysis may have lacked the requisite sensitivity to detect such an effect. Given this result, they decided to collapse across conditional probabilities for the remaining analyses.

The authors sought to determine whether expected representations were dampened or sharpened based on three predictions: (1) the dampening account predicts that voxels responding preferentially to specific images will exhibit the strongest expectation suppression in response to those images, whereas the sharpening account predicts the opposite; (2) the dampening account predicts that expectation suppression will have the greatest effect for highly selective voxels,

whereas the sharpening hypothesis suggests that activity in these voxels should be enhanced, not suppressed; and (3) the sharpening account predicts that a multivariate pattern analysis (MVPA), used to decode sensory representations based on patterns of neural activity, will exhibit improved classification accuracy for sharpened sensory representations. Alternatively, the dampening of this neural activity will obscure these sensory representations, such that little discrepancy in classification accuracy will be observed between expected and unexpected trials.

Tests of these predictions provided support for the dampening account of expectation suppression. First, a regression of voxel image preference ranks, calculated based on expectation-neutral presentations, against BOLD responses from the main task confirmed that the strength of expectation suppression was greatest in LOC voxels that responded preferentially to the image. No significant differences were observed in V1. Second, expectation suppression was significantly stronger in highly selective LOC voxels. While there was also a trend toward this dampening pattern in V1, the result was not statistically significant. Third, although MVPA classification accuracy was above chance for both expected and unexpected stimuli in V1 (27.9% and 30.2%, respectively) and LOC (18.5% and 19.5%, respectively), there was no significant difference in accuracy based on predictability. Together, these analyses provide evidence for dampened sensory representations underlying expectation suppression consistent with recent research in monkey inferotemporal cortex (Meyer and Olson, 2011; Kumar et al., 2017). However, the classification accuracy result differs from multivariate analyses in previous fMRI research (Kok et al., 2012) and single-unit recordings in monkey inferotemporal cortex (Bell et al., 2016), which indicated that expectation suppression is associated

with improved decoding accuracy for expected stimuli.

The demonstration by Richter et al. (2018) of the specificity of expectation suppression in LOC to activity associated with the predicted stimulus is an important finding, explaining the effect of expectation at a neural level and bolstering arguments about the role of inference in perception. However, the nonspecific suppression observed in V1 also raises important questions. This finding appears inconsistent with the theorized cascade of descending predictions, specifying expectations at each level of sensory processing, characteristic of predictive processing (Clark, 2013). Rather, these findings suggest a more modular role of expectation, acting primarily at the level of processing associated with the stimulus (in this case, object-selective LOC), in line with more traditional, feedforward models of perception. It may also be the case that the specificity of suppression is graded along the processing hierarchy, being tightly defined in regions closely tuned to the stimulus in question, but becoming less selective as one descends through lower-level processing areas (e.g., V1). The superior accuracy of V1 MVPA decoding in Richter et al. (2018) may hint at this graded approach, where the representation of the stimulus in question is most suppressed in LOC. This question could be addressed in future research by characterizing the specificity of expectation suppression in intermediate regions between V1 and LOC (e.g., V4), akin to research with motion stimuli (Alink et al., 2010), and across various levels of stimulus complexity. In addition, using MVPA decoding techniques to trace the accuracy of representations across intermediate re-

gions in the visual stream may offer an opportunity to provide greater spatial resolution to accounts of expectation suppression and inform predictive processing models.

The remarkable surge in research on sensory processing as fundamentally inferential is producing new and exciting insights into the mechanisms underlying perceptual experience (Friston, 2018). Richter et al. (2018) have made a significant contribution in characterizing the nature of one of the core phenomena associated with predictive processing. Furthermore, they have raised important questions about the functional architecture of the inferential hierarchy, providing avenues for future research.

References

- Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L (2010) Stimulus predictability reduces responses in primary visual cortex. *J Neurosci* 30:2960–2966. [CrossRef Medline](#)
- Bell AH, Summerfield C, Morin EL, Malecek NJ, Ungerleider LG (2016) Encoding of stimulus probability in macaque inferior temporal cortex. *Curr Biol* 26:2280–2290. [CrossRef Medline](#)
- Bowers JS, Davis CJ (2012) Bayesian just-so stories in psychology and neuroscience. *Psychol Bull* 138:89–414. [CrossRef Medline](#)
- Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci* 36:181–204. [CrossRef Medline](#)
- de Lange FP, Heilbron M, Kok P (2018) How do expectations shape perception? *Trends Cogn Sci* 22:764–779. [CrossRef Medline](#)
- Egner T, Monti JM, Summerfield C (2010) Expectation and surprise determine neural population responses in the ventral visual stream. *J Neurosci* 30:16601–16608. [CrossRef Medline](#)
- Friston K (2005) A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360:815–836. [CrossRef Medline](#)
- Friston K (2009) The free-energy principle: a rough guide to the brain? *Trends Cogn Sci* 13:293–301. [CrossRef Medline](#)
- Friston K (2018) Does predictive coding have a future? *Nat Neurosci* 21:1019–1021. [CrossRef Medline](#)
- Hohwy J (2013) *The predictive mind*. New York: Oxford UP.
- Kok P, de Lange FP (2014) Shape perception simultaneously up- and downregulates neural activity in the primary visual cortex. *Curr Biol* 24:1531–1535. [CrossRef Medline](#)
- Kok P, Jehee JF, de Lange FP (2012) Less is more: expectation sharpens representations in the primary visual cortex. *Neuron* 75:265–270. [CrossRef Medline](#)
- Kumar S, Kaposvari P, Vogels R (2017) Encoding of predictable and unpredictable stimuli by inferior temporal cortical neurons. *J Cogn Neurosci* 29:1445–1454. [CrossRef Medline](#)
- Meyer T, Olson CR (2011) Statistical learning of visual transitions in monkey inferotemporal cortex. *Proc Natl Acad Sci U S A* 108:19401–19406. [CrossRef Medline](#)
- Pajani A, Kouider S, Roux P, de Gardelle V (2017) Unsuppressible repetition suppression and exemplar-specific expectation suppression in the fusiform face area. *Sci Rep* 7:160. [CrossRef Medline](#)
- Ramachandran S, Meyer T, Olson CR (2016) Prediction suppression in monkey inferotemporal cortex depends on the conditional probability between images. *J Neurophysiol* 115:355–362. [CrossRef Medline](#)
- Richter D, Ekman M, de Lange FP (2018) Suppressed sensory response to predictable object stimuli throughout the ventral visual stream. *J Neurosci* 38:7452–7461. [CrossRef Medline](#)
- Spratling MW (2017) A review of predictive coding algorithms. *Brain Cogn* 112:92–97. [CrossRef Medline](#)
- Summerfield C, de Lange FP (2014) Expectation in perceptual decision making: neural and computational mechanisms. *Nat Rev Neurosci* 15:745–756. [CrossRef Medline](#)
- Todorovic A, de Lange FP (2012) Repetition suppression and expectation suppression are dissociable in time in early auditory evoked fields. *J Neurosci* 32:13389–13395. [CrossRef Medline](#)