Object expectations alter information use during visual recognition

Laurent Caplette*, Frédéric Gosselin & Greg L. West

Department of Psychology, Université de Montréal,
C.P. 6128, succ. Centre-Ville, Montréal, Qc,
H3C 3J7, Canada

*Corresponding author

Email: laurent.caplette@yale.edu

Present address: Department of Psychology, Yale University, 2 Hillhouse Avenue, New Haven, CT, 06520, U.S.A.

Highlights

- We reveal how expectations influence information use during object recognition
- Expectations accelerate the successful use of coarse information
- Prior expectations lead to a more specialized sampling of information

Abstract

Prior expectations influence how we perceive and recognize objects. However, how they do so remains unclear, especially in the case of real-world complex objects. Expectations of objects may affect which features are used to recognize them subsequently. In this study, we used reverse correlation in neurotypical participants to reveal with high precision how the use of information across time is modulated by everyday object expectations in a recognition task. We show that coarse information leads to accurate responses earlier when an object is expected, indicating that observers use diagnostic features earlier in this situation. We also demonstrate an increased variability in the use of coarse information depending on the expected object, indicating that observers adopt a more specialized recognition strategy when they expect a specific object. In summary, our results reveal how expectations of real-world complex objects affect the use of visual information across time.

Keywords: object recognition, expectation, prediction, representation, information use.

1 Introduction

Every day, we encounter a myriad of objects and we recognize them effortlessly within milliseconds. Understanding how we do so remains one of the greatest challenges of cognitive neuroscience. Since the seminal work of Helmholtz in the 19th century, prior expectations are believed to be responsible for a key part of this feat. Valid expectations increase our accuracy and our speed at recognizing objects (Pinto, van Gaal, de Lange, Lamme & Seth, 2015; Stein & Peelen, 2015; Wyart, Nobre & Summerfield, 2012) and they can bias our perception of ambiguous or noisy sensory input in specific directions (Kok, Brouwer, van Gerven & de Lange, 2013; Sterzer, Frith & Petrovic, 2008).

The mechanisms responsible for the modulation of object recognition by expectations remain unclear. Previous studies have reported that expectations can induce an increase in sensitivity to visual features (Cheadle, Egner, Wyart, Wu & Summerfield, 2015; Wyart et al., 2012), and that they can preactivate internal representations (Cheadle et al., 2015; Kok, Failing & de Lange, 2012; 2014; Wyart et al., 2012). During the recognition of complex objects comprised of multiple features, expectations might alter the use of a subset of features. On the one hand, when no specific object is expected, observers must apply a generic recognition strategy, at least at first: for example, observers may use the same large set of relatively coarse features to recognize the object. On the other hand, when a specific object is expected, observers might adapt and use a more specialized recognition strategy: some features might get more or less used depending on the expected object. They may also use finer and more detailed features to confirm and refine their hypothesis (Bar, 2003; Bar et al., 2006). How different features are affected by expectations has not yet been investigated, since studies assessing the effect of expectations on internal

representations have only used stimuli made of one feature (simple sinusoidal gratings; Bang & Rahnev, 2017; Cheadle et al., 2015; Wyart et al., 2012).

Expectations might also alter the time course of object recognition. Different features, including spatial frequencies (SFs), are typically sampled at different moments during a fixation (Blais, Roy, Fiset, Arguin & Gosselin, 2012; Caplette, Wicker & Gosselin, 2016; Hughes, Nozawa & Kitterle, 1996; Neri & Levi, 2007; Schyns & Oliva, 1994). In a recognition task with everyday man-made objects, subjects typically start using coarse low SF information before they start using finer higher SF information (Caplette et al., 2016; Caplette, Wicker, Gosselin & West, 2017). As such, if expectations alter specific features, their effect may vary across time. Continuously incoming information might also be used to iteratively update internal hypotheses about the fixated object's identity and redirect sampling to more relevant features (Bar, 2003). In that case, it is possible that a prior expectation would reduce the need for these updates, since observers would already be focusing on the right features, and expectations would thus accelerate successful information acquisition on average.

In the present study, we investigated for the very first time how expectations of specific real-world objects modulate the use of SFs across time. Using real-world objects comprised of multiple features allowed us to examine how expectations affect the use of different features. Moreover, there are preliminary indications that expectations of complex objects involve different mechanisms than expectations of simple features such as gratings of different orientations (Denison, Piazza & Silver, 2011; Denison, Sheynin & Silver, 2016; Kok, Rait & Turk-Browne, 2020). Participants were exposed to object images followed by object names (no-expectation condition), and to object names followed by

object images (expectation condition; Figure 1a). They were asked to indicate whether the name matched the object shown: it did on 50% of trials. Given that there was a total of 80 objects, the object most probable to appear for a given object name was the object described by the name (Figure 1b). Object images had their SFs randomly sampled across time on each trial (Figure 1c; Caplette et al., 2016; 2017). This allowed us to recover which SFs at each moment during the object fixation led to more accurate responses, i.e., which SFs were used at each moment for each condition. Note that the influence of expectations over the use of different features might be mediated by attention or other mechanisms: the goal of our study is not to disentangle expectations from attention (see Summerfield & Egner, 2009; 2016) but simply to assess whether subjects adopt different object recognition strategies when they expect the specific object shown vs when they do not. We believe this is an important issue that remains to be addressed.

Figure 1. Experimental paradigm and sampling method. **A)** On each trial, a fixation cross appeared on a mid-gray background for 300 ms, followed by either an object (no-expectation condition) for 333 ms or an object name (expectation condition) for 1000 ms. Then, another fixation cross appeared for 300 ms, followed by either an object name or an object, depending on what had been presented before. Each trial ended with a blank mid-gray screen. Participants had to indicate whether the name and the object matched. Stimuli, names and fixation crosses have been enlarged for display purposes; names have been translated from French. **B)** Objects and object names had a 50% probability of matching. Since there was a total of 80 objects, each other object had a probability of appearance of about 0.6%. **C)** On each trial, we randomly generated a matrix of dimensions 256×40 (representing respectively SFs and frames) in which most elements were zeros and a few were ones. We then convolved this sparse matrix with a 2D Gaussian kernel (a “bubble”). This resulted in the trial’s sampling matrix, shown here as a plane with a number of randomly located bubbles. Every column of this sampling matrix was then rotated around its origin to create isotropic 2D random filters. Finally, these 2D random filters were dot-multiplied by the base image’s spectrum and the result was inverse fast Fourier transformed to create a filtered version of the image for every video frame.

2 Materials and Methods

2.1 Participants

Fifty-nine right-handed neurotypical adult participants (29 males; mean age = 22.07; SD = 2.87) were recruited on the campus of the University of Montreal. Subjects had normal or corrected-to-normal vision and did not suffer from any visual or reading disability. The

study was approved by the ethics board of the University of Montreal's Faculty of Arts and Sciences. Written consent from all participants was obtained after the procedure had been fully explained, and a monetary compensation was provided upon completion of the experiment.

2.2 Materials

The experimental programs ran on Mac Pro (Apple Inc.) computers in the Matlab (Mathworks Inc.) environment, using functions from the Psychophysics Toolbox (Brainard, 1997; Kleiner, Brainard, Pelli, Ingling & Murray, 2007; Pelli, 1997). All stimuli were presented on Asus VG278H monitors (1920×1080 pixels at 120 Hz), calibrated to allow a linear manipulation of luminance. Luminance ranged from 1.6 cd/m^2 to 159 cd/m^2 .

2.3 Stimuli

Eighty grayscale images of man-made objects were selected from the database used in Shenhav, Barrett & Bar (2013) and from Internet searches (Caplette et al., 2016; 2017). Images were 256×256 pixels and median object width was 220 pixels. The objects were cropped manually and pasted on a homogenous mid-gray background. The Fourier amplitude spectrum of each image was set to the mean Fourier amplitude spectrum across images and the mean luminance of each image was set to the global mean luminance across images using the SHINE toolbox (Willenbockel et al., 2010). Resulting images had a root mean square (RMS) contrast of 0.20.

On each trial, participants were shown a short video (333 ms) consisting of an object image with random SFs gradually revealed at random time points (e.g., Video S1; Video

S2; Caplette et al., 2016; 2017); that is, on each video frame, there would typically be several SFs shown among all possible SFs, and these would change from frame to frame. To create these dynamic stimuli, we first randomly generated, on each trial, a matrix of dimensions 256 x 40 (representing respectively SFs from 0.5 to 128 cycles per image (cpi), or 0.1 to 21.3 cycles per degree (cpd), and frames, each lasting 8.33 ms; the matrix was in fact larger because of additional padding) in which most elements were zeros and a few were ones. The probability of an element being one was adjusted on a trial-by-trial basis to maintain performance at 75% correct. We then convolved this *sparse matrix* with a 2D Gaussian kernel (a “bubble”; $\sigma_{\text{SF}} = 1.25$ cpi or 0.2 cpd; $\sigma_{\text{time}} = 12.5$ ms). This resulted in the trial’s *sampling matrix*: a SF $\times$ time plane with randomly located bubbles. Every column of this sampling matrix was then rotated around its origin to create isotropic *2D random filters*. Finally, these 2D random filters were dot-multiplied by the base image’s Fourier spectrum and the result was inverse fast Fourier transformed: this had the effect of filtering the original image in the randomly selected SF bands while keeping the original phase spectrum (Figure 1c). To augment luminance resolution, we applied noisy-bit dithering to the final stimuli (Allard & Faubert, 2008).

2.4 Procedure

Participants sat in front of a computer monitor, in a dim-lighted room, over the course of two days. They completed two 500-trial blocks in each experimental session. In half of the trials, an object name was shown before the stimulus (*expectation* condition); in the other half, the stimulus was shown before the object name (*no-expectation* condition). These two conditions were alternated in mini-blocks of 50 trials (the first condition was

counterbalanced across participants) and a short pause occurred between mini-blocks. Specifically, each trial was comprised of the following consecutive events on a mid-gray background: a fixation cross (300 ms), a basic-level object name (1000 ms) or the video stimulus (333 ms), a fixation cross (300 ms), the video stimulus (333 ms) or a basic-level object name (until response, max. 1000 ms), and a blank screen (until response if not yet made + 200 ms). When shown after the stimulus, the object name remained on screen either until a response was provided or for a maximum of 1 s, in which case it was replaced by a blank screen until a response was provided; when shown after the name, the video stimulus was immediately replaced by a blank screen until a response was provided (this is in addition to a final 200 ms blank screen after the response; Figure 1a). The number of bubbles was adjusted on a trial-by-trial basis using a gradient descent algorithm within each condition to maintain performance at 75% correct (actual mean performance was 75.90%). Subjects were asked to indicate whether the name matched the object as accurately and as rapidly as possible. The object name and the object matched on 50% of trials; on the trials in which they didn't match, the name was randomly chosen among the names of all other objects. During the task, chin rests were used to maintain viewing distance at 76 cm; images subtended 6 x 6 degrees of visual angle.

2.5 Data Analysis

Six experimental sessions were excluded prior to the analysis: one because of technical problems, three because the mean numbers of bubbles were over three standard deviations above the mean, and two because the mean response times were over three standard deviations above the mean.

Accuracies and response times were z-scored within each object (to minimize variability due to differences in object recognizability), block (to minimize variability due to task learning) and subject (to minimize variability due to interindividual differences in performance), separately for each condition. Trials associated with z-scored response times or accuracies over 3 or below -3 were discarded from the regressions (1.88% of trials). Sparse matrices were also z-scored within each trial (using the global mean and standard deviation of each sparse matrix). Only spatial frequencies up to 64 cycles per image (cpi), or 10.7 cycles per degree (cpd), were analyzed, since higher SFs typically do not contribute to accurate object recognition (Caplette et al., 2016; Caplette, West, Gomot, Gosselin & Wicker, 2014; Gold, Bennett & Sekuler, 1999).

To uncover which SFs in which time frames led to accurate object recognition in each condition, we performed sums of z-scored sparse matrices weighted by z-scored accuracies, separately for each subject and each condition (since matrices are random, this is equivalent to a least-squares multiple linear regression). The resulting matrices were then summed across subjects and convolved with a Gaussian kernel ($\sigma_{\text{SF}} = 6.5$ cpi or 1.1 cpd; $\sigma_{\text{time}} = 46$ ms); henceforth, we will refer to these matrices as classification images. To compare the conditions, we subtracted the classification image of the no-expectation condition from the classification image of the expectation condition. The same procedure was repeated 500 times while randomly permuting the accuracies (and thus disrupting any potential association between predictors and responses) to obtain an empirical null distribution of classification images. These null classification images were then used to transform the classification images (both observed and null) into z-scores. Finally, statistical thresholds at the pixel and cluster levels ($p < .05$, two-tailed, corrected for the

Family-Wise Error Rate (FWER); $Z > 2.3$ primary cluster threshold) were determined using the maximum statistic method (Holmes, Blair, Watson & Ford, 1996).

To uncover whether different object expectations had a different effect on the use of SFs, we conducted a variance analysis. Given that a lower number of bubbles in one condition could result in an artefactually higher variance (see below), we first randomly subsampled trials so that the mean number of bubbles shown (and the number of trials) was the same in both conditions (83% of the original trials were kept). We then concatenated the z-scored accuracies of all subjects together and did the same for the sparse matrices, and we performed one weighted sum for each object name in each condition. Since we were interested in the facilitation of the recognition of an object by the expectation that this same object would appear, we considered only match trials in which the object name corresponded to the object shown (for both the expectation and the no-expectation conditions). Resulting classification images were convolved with a Gaussian kernel ($\sigma_{\text{SF}} = 6.5$ cpi or 1.1 cpd; $\sigma_{\text{time}} = 46$ ms) and variance across objects was computed for each spatial frequency in each time frame for each condition. To compare the conditions, we computed the ratio of the variance in the expectation condition over the variance in the no-expectation condition (Bailer, 1989). The same analyses were repeated with permuted samples, as explained above; since variance can only be positive, statistical tests were one-tailed, except for the ratio between the conditions.

A potential confound in our analyses is the different signal-to-noise ratios in the conditions resulting from the different average numbers of bubbles. This difference alone predicts, for the “mean” analysis reported in Figure 2, higher Z-scores in the condition with more bubbles when, as is the case here, bubbles reveal less than half of the search space,

which is the opposite of what we observed. In other words, these results cannot be explained by the number-of-bubbles confound. A lower number of bubbles could, however, lead to a higher variance because of the increased noise, as we observed following our “variance” analysis on all trials. For this reason, we repeated the analysis on a subsample of the trials selected to equate the number of bubbles in the two conditions, as explained above. This last analysis, which corroborated our initial results, is the one shown in Figure 3.

3 Results

We first investigated whether subjects performed better on average when they expected the incoming objects. We assessed the amount of information required to do the task in each condition. The no-expectation condition required significantly more bubbles (mean = 45.20; $\sigma = 15.30$) to maintain a 75% accuracy than the expectation condition (mean = 33.71; $\sigma = 10.47$; $t(58) = 8.28$, $p < .001$, $d_z = 1.08$). In accordance with previous studies (Pinto et al., 2015; Stein & Peelen, 2015; Wyart et al., 2012), subjects performed better in the expectation condition. Subjects also responded faster in the expectation condition (mean RT = 598 ms; $\sigma = 183$ ms) compared to the no-expectation condition (mean RT = 729 ms; $\sigma = 141$ ms; $t(58) = 11.27$, $p < .001$, $d_z = 1.47$) although this could be a product of the experimental design rather than neural processing *per se* (in the no-expectation condition, subjects had to read the word just before responding).

3.1 Coarse information is used more and earlier when a specific object is expected

To uncover which SFs presented at which moments led to correct responses in each condition, we reverse correlated accuracies with spatial frequencies presented at each moment (see Methods). Mean results are displayed in Figure 2: higher z-scores indicate a stronger correlation between the presentation of a given SF on a given time frame and recognition accuracy. In both conditions, a large band of relatively low spatial frequencies (LSFs; approx. 0.1–5 cycles per degree, cpd) was significant throughout the stimulus duration; however, the peak use of LSFs occurred at distinct latencies when comparing both conditions. Specifically, a large significant cluster peaked at 1.7 cpd and 150 ms ($Z = 11.70$, $p_{FWER} < .001$, $d_z = 1.60$) and at 1.8 cpd and 233 ms ($Z = 12.11$, $p_{FWER} < .001$, $d_z = 1.50$) in the no-expectation condition (Figure 2a). In the expectation condition, a cluster peaked at 1.8 cpd and 100 ms ($Z = 14.33$, $p_{FWER} < .001$, $d_z = 1.71$) and at 1.5 cpd and 217 ms ($Z = 14.46$, $p_{FWER} < .001$, $d_z = 1.69$; Figure 2b). Contrasting results from both conditions resulted in a significant difference in the same band, peaking at 2 cpd and 67 ms ($Z = 4.01$, $p_{FWER} = .01$, $d_z = .52$): these early low SFs led to more accurate responses in the expectation condition than in the no-expectation condition (Figure 2c). Note that this difference cannot be due to the different signal-to-noise ratio in the two conditions which, in fact, predicts the opposite result (see section 2.5).

Figure 2. Use of SFs across time for each condition and difference between the conditions **A)** No-expectation condition. **B)** Expectation condition. **C)** Expectation – No-expectation. Light gray outlines indicate statistical significance at the cluster level and dark gray outlines indicate statistical significance at the pixel level ($p < .05$, two-tailed, FWER-corrected).

Having confirmed an increased early use of low SFs in the expectation condition, we assessed whether low SFs were used earlier irrespective of overall use. To do so, we conducted a latency analysis in individual subjects. We first made the range of values in the classification images (i.e., the maps of regression coefficients) of each subject in each condition vary between 0 and 1 to avoid that signal-to-noise ratio influences our latency estimates. We then verified at which latency spatial frequencies under 5 cpd were first above a threshold of .75 in each transformed classification image. A paired t-test revealed that latencies were significantly smaller in the expectation condition compared to the no-expectation condition (38 ms vs 63 ms; $t(58) = 2.77$, $p = .007$, $d_z = .36$).

3.2 Specific object expectations lead to specific information use

We then investigated whether object recognition strategies were more specialized when specific objects were expected, i.e., whether different object expectations were modulating the use of SF information differently. To do so, we performed one reverse correlation analysis for match trials of each object expectation in each condition and we computed the variance across the resulting classification images for each condition (after subsampling

trials to equalize the numbers of bubbles across the conditions; see section 2.5). There was no significant variance in the no-expectation condition (Figure 3a). This is not surprising, given that subjects did not know which object would appear in advance. In the expectation condition, however, there was significant across-object variance in low SFs around 1.2 cpd and 167 ms ($Z = 5.17$, $p_{FWER} = .007$; Figure 3b). When contrasting the conditions, there was significantly more across-objects variance in the expectation condition than in the no-expectation condition in a low SF cluster peaking at 1.1 cpd and 175 ms ($Z = 4.75$, $p_{FWER} < .05$, cluster-level; Figure 3c). Thus, different object expectations modulated the use of these SFs in that time window differently (see Figure 3d and object examples in Figure 3e).

Figure 3. Variance across objects in the use of SFs across time for **A)** the No-expectation condition and **B)** the Expectation condition. **C)** Ratio of variances of both conditions (Expectation over No-expectation). Light gray outlines indicate statistical significance at the cluster level and dark gray outlines indicate statistical significance at the pixel level ($p < .05$, one-tailed for each condition and two-tailed for the difference, FWER-corrected). **D)** Illustration of the use, for each object and condition, of the 1.1 cpd SFs around 175 ms (the peak activation in the significant cluster in C). Black circles indicate the mean for each condition and error bars indicate the standard deviation. **E)** Illustration of some of the objects for which the 1.1 cpd SFs are most and least used around 175 ms, when filtered in a narrow SF band centered on 1.1 cpd. Upper row: wheelchair, beach bucket and shovel, book. Lower row: knife, gas pump, fan.

4 Discussion

In this study, we investigated how expectations of complex everyday objects modulate the use of SFs for object recognition. When no specific object was expected, a

large band of relatively low SFs (approx. 0.1-5 cycles per degree, cpd) was used (i.e., correlated positively with accuracy), and its use did not vary across objects. Observers seem to apply a generic recognition strategy in this case, and they use a large portion of the feature space. Such low SFs have been associated with object recognition in prior studies (e.g., Caplette et al., 2016; 2017; Hughes et al., 1996). When a specific object was expected, SFs around 2 cpd at the beginning of stimulus presentation led to a higher accuracy, likely because participants now sampled specific low SF features that helped the recognition of the expected object. Indeed, participants now had an idea of what object was likely to appear in front of them so they should have tried to match this object's representation to the incoming sensory information (by focusing on the relevant orientations, image areas, etc.). Several other studies observed heightened sensitivity when stimuli were expected (Cheadle et al., 2015; Stein & Peelen, 2015; Wyart et al., 2012); here, we extend this finding and show how specific SFs of real-world objects are affected.

Moreover, we observe that the effect of expectations is largely constrained to an early time window: the difference between the conditions occurred around a latency of 67 ms and our latency analysis of individual results indicates that the use of low SFs started earlier on average when a specific object was expected. These results suggest that expectations accelerated the effective use of low SFs. One possibility is that, when there is no specific expectation, subjects don't know what features (e.g., orientations) are optimal to sample right at the stimulus onset, but they can gradually adapt to the incoming sensory information: this would lead for all objects to SFs correlating more strongly with accuracy later during recognition, as is observed. When there is a specific expectation however, subjects already know what information to look for: as such, they sample the most useful

features right from the start, leading to a high correlation with accuracy throughout stimulus presentation and a significant difference with the no-expectation condition at the beginning of stimulus presentation. Although this supports the general idea that hypotheses are generated and iteratively updated during object recognition (Bar, 2003; Bullier, 2001; Di Lollo et al., 2000), we did not find any evidence for the more specific proposal that high SFs are used to refine the original hypothesis (Bar, 2003; Bar et al., 2006).

Expectations further increased the across-objects variability in the use of similar SFs around 1.1 cpd, at a peak latency of 175 ms. This indicates that subjects were using these low SFs differently depending on the expected object. This is the first demonstration, to our knowledge, that prior object expectations modulate the degree to which some features are represented and used for object recognition. It indicates that object recognition strategies are different and more specialized when we expect a particular object, varying depending on this expected object. It is difficult to investigate precisely why these SFs were used more for specific objects and less for other objects given the characteristics of our experiment. Our stimuli were not chosen for their intrinsic properties and we do not have any data associated to these stimuli that we could use to correlate with this SF use. Less use of these low SFs might co-occur with more use of other, possibly higher, SFs that are different for every object (and therefore not observable in our results). It might also correlate with some other object properties such as their complexity or their familiarity. We investigated a related possibility that low SFs around 175 ms were less used for some objects because these objects were associated with smaller RTs, but we did not find any support for this hypothesis (results not shown). Future studies should try to correlate intrinsic object properties with information use.

Few other studies used reverse correlation in the context of expectations and all that did employed grating orientation discrimination tasks (Bang & Rahnev, 2017; Cheadle et al., 2015; Wyart et al., 2012). However, semantic object-based predictions might be explained by different mechanisms than predictions based on simple correlations between such lower-level features. Moreover, expectations based on associations learned in a distant past – like the ones our subjects made – and ‘mnemonic’ expectations based on associations learned recently might rely on different neural processes (Caplette, Gosselin, Mermillod & Wicker, 2020; Hindy, Ng & Turk-Browne, 2016). Using real-world objects varying on multiple feature dimensions also allowed us to investigate whether the features sampled are different when objects are expected.

Our results indicate clearly for the first time that object expectations do influence how object features are sampled. These effects may very well be mediated by top-down attention. Our goal was not to disentangle the respective contributions of attention and “pure” expectation effects in sampling of features — the focus of much of the research efforts in the past 10 years or so (e.g., Kok, Rahnev, Jehee, Lau & de Lange, 2012; Summerfield & Egner, 2009; 2016) — but rather to examine this sampling when objects are either expected or not.

In summary, we have shown how expectations influence the recognition of complex everyday objects across time. We showed that expectations overall accelerate the successful use of low SFs and that specific expectations will influence how much low SFs are used, indicating that prior object expectations lead to more specialized object recognition strategies. These results reveal for the first time the mechanisms underlying the effects of expectations on the recognition of real-world objects.

References

- Allard, R., & Faubert, J. (2008). The noisy-bit method for digital displays: Converting a 256 luminance resolution into a continuous resolution. *Behavior Research Methods*, 40, 735–743.
- Bailer, A. J. (1989). Testing variance equality with randomization tests. *Journal of Statistical Computation and Simulation*, 31:1, 1–8.
- Bang, J. W., & Rahnev, D. (2017). Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Scientific Reports*, 7, 17072.
- Bar, M. (2003). A Cortical Mechanism for Triggering Top-Down Facilitation in Visual Object Recognition. *Journal of Cognitive Neuroscience*, 15, 600–609.
- Bar, M. et al. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences U.S.A.*, 103, 449–454.
- Blais, C., Roy, C., Fiset, D., Arguin, M., & Gosselin, F. (2012). The eyes are not the window to basic emotions. *Neuropsychologia*, 50, 2830–2838.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.
- Bullier, J. (2001). Integrated model of visual processing. *Brain Research Reviews*, 36(2–3), 96–107.

- Caplette, L., West, G., Gomot, M., Gosselin, F., & Wicker, B. (2014). Affective and contextual values modulate spatial frequency use in object recognition. *Frontiers in Psychology, 14*, 512.
- Caplette, L., Wicker, B., & Gosselin, F. (2016). Atypical Time Course of Object Recognition in Autism Spectrum Disorder. *Scientific Reports, 6*, 35494.
- Caplette, L., Wicker, B., Gosselin, F., & West, G. L. (2017). Hand position alters visual recognition by modulating the time course of spatial frequency use. *Journal of Experimental Psychology: General, 146*, 917–923.
- Caplette, L., Gosselin, F., Mermillod, M., & Wicker, B. (2020). Real-world expectations and their affective value modulate object processing. *NeuroImage*.
- Cheadle, S., Egner, T., Wyart, V., Wu, C., & Summerfield, C. (2015). Feature expectation heightens visual sensitivity during fine orientation discrimination. *Journal of Vision, 15*, 14.
- Denison, R. N., Piazza, E. A., & Silver, M. A. (2011). Predictive context influences perceptual selection during binocular rivalry. *Frontiers in Human Neuroscience, 5*, 166.
- Denison, R. N., Sheynin, J., & Silver, M. A. (2016). Perceptual suppression of predicted natural images. *Journal of Vision, 16*, 6–15.

- Di Lollo, V., Enns, J. T., & Rensick, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General*, 129(4), 481–507.
- Dupuis-Roy, N., Faghel-Soubeyrand, S., & Gosselin, F. (2019). Time course of the use of chromatic and achromatic facial information for sex categorization. *Vision Research*, 157, 36–43.
- Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research*, 39, 3537–3560.
- Hindy, N. C., Ng, F. Y., & Turk-Browne, N. B. (2016). Linking pattern completion in the hippocampus to predictive coding in visual cortex. *Nature Neuroscience*, 19, 665–667.
- Holmes, A. P., Blair, R. C., Watson, J. D. G., & Ford, I. (1996). Nonparametric Analysis of Statistic Images from Functional Mapping Experiments. *Journal of Cerebral Blood Flow and Metabolism*, 16, 7–22.
- Hughes, H. C., Nozawa, G., & Kitterle, G. F. (1996). Global precedence, spatial frequency channels, and the statistics of natural images. *Journal of Cognitive Neuroscience*, 8, 197–230.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., & Murray, R. (2007). What's new in Psychtoolbox-3. *Perception 36 ECVF Abstract Supplement*.

- Kok, P., Brouwer, G. J., van Gerven, M. A. J., & de Lange, F. P. (2013). Prior Expectations Bias Sensory Representations in Visual Cortex. *Journal of Neuroscience*, 33, 16275–16284.
- Kok, K., Rait, L. I., & Turk-Browne, N. B. (2020). Content-based dissociation of hippocampal involvement in prediction. *Journal of Cognitive Neuroscience*, 32(3), 527–545.
- Neri, P., & Levi, D. M. (2007). Temporal Dynamics of Figure-Ground Segregation in Human Vision. *Journal of Neurophysiology*, 97, 951–957.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F., & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, 15, 13.
- Schyns, P. G., & Oliva, A. (1994). From Blobs to Boundary Edges: Evidence for Time- and Spatial-Scale-Dependent Scene Recognition. *Psychological Science*, 5, 195–200.
- Shenhav, A., Barrett, L., & Bar, M. (2013). Affective value and associative processing share a cortical substrate. *Cognitive, Affective and Behavioral Neuroscience*, 13, 46–59.

- Stein, T., & Peelen, M. V. (2015). Content-specific expectations enhance stimulus detectability by increasing perceptual sensitivity. *Journal of Experimental Psychology: General*, 144, 1089–1104.
- Sterzer, P., Frith, C., & Petrovic, P. (2008). Believing is seeing: expectations alter visual awareness. *Current Biology*, 16, R697-R698.
- Vinette, C., Gosselin, F., & Schyns, P. G. (2004). Spatio-temporal dynamics of face recognition in a flash: it's in the eyes. *Cognitive Science*, 28, 289–301.
- Willenbockel, V., et al., (2010). Controlling low-level image properties: The SHINE toolbox. *Behavior Research Methods*, 42, 671–684.
- Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences U.S.A.*, 109, 6354–6354.

Supplementary Material

Movie S1. Example of a random stimulus.

Movie S2. Same as S1, slowed down 10 times.

Link to data: <https://osf.io/qj2ka/>

Competing Interests Statement

The authors have no competing interests to declare.