



Predictions drive neural representations of visual events ahead of incoming sensory information

Tessel Blom^{a,1}, Daniel Feuerriegel^a, Philippa Johnson^a, Stefan Bode^{a,b}, and Hinze Hogendoorn^{a,c}

^aMelbourne School of Psychological Sciences, The University of Melbourne, Melbourne, VIC 3010, Australia; ^bDepartment of Psychology, University of Cologne, 50923 Cologne, Germany; and ^cHelmholtz Institute, Department of Experimental Psychology, Utrecht University, 3584 CS Utrecht, The Netherlands

Edited by Randolph Blake, Vanderbilt University, Nashville, TN, and approved February 17, 2020 (received for review October 10, 2019)

The transmission of sensory information through the visual system takes time. As a result of these delays, the visual information available to the brain always lags behind the timing of events in the present moment. Compensating for these delays is crucial for functioning within dynamic environments, since interacting with a moving object (e.g., catching a ball) requires real-time localization of the object. One way the brain might achieve this is via prediction of anticipated events. Using time-resolved decoding of electroencephalographic (EEG) data, we demonstrate that the visual system represents the anticipated future position of a moving object, showing that predictive mechanisms activate the same neural representations as afferent sensory input. Importantly, this activation is evident before sensory input corresponding to the stimulus position is able to arrive. Finally, we demonstrate that, when predicted events do not eventuate, sensory information arrives too late to prevent the visual system from representing what was expected but never presented. Taken together, we demonstrate how the visual system can implement predictive mechanisms to preactivate sensory representations, and argue that this might allow it to compensate for its own temporal constraints, allowing us to interact with dynamic visual environments in real time.

prediction | neural delays | time-resolved decoding | visual system

Visual information takes time to travel from the retina and through the visual system (1), such that visual information becomes progressively outdated as it flows through the visual system. This neural transmission delay presents a computational challenge for the visual system, because, without compensation, neural representations of visual information (such as the position of a moving object) would lag behind the present moment. As a result, a moving object's perceived position would consistently lag its true position (2, 3), making accurate interaction with that object impossible.

One way the brain might compensate for neural delays is through prediction (3–6). In support of this idea, predictable visual stimuli have been found to be represented in the visual system with a shorter latency than unpredictable stimuli in cats (7), macaques (8), and humans (9). Furthermore, prior expectations about an expected stimulus have been found to induce sensory templates of that stimulus in both the visual and the auditory cortex (10–12). The early visual cortex has even been shown to represent expected spatiotemporal sequences: After subjects have been familiarized with a spatial sequence, the presentation of the starting stimulus of that sequence can set off a series of neural representations that resembles the neural representations triggered by the presentation of the full sequence (13).

Together, these studies show, convincingly, that predictive mechanisms are sufficient to activate stimulus-specific neural representations in multiple modalities. However, for predictive mechanisms to compensate for neural delays, they must activate their associated representations before those representations would normally be activated. In other words, predictions would need to preactivate neural representations that would typically be driven by sensory input, before the actual arrival of that input.

Only then would the prediction actually serve to reduce the delay that visual information accumulates as it flows through the visual system. Whether predictive mechanisms are indeed able to activate neural representations ahead of transmission delays in this way is unknown.

In the present study, we investigated whether predictive activation of visual representations that are ordinarily driven by afferent sensory information occurs before the arrival of that afferent sensory information, that is, ahead of neural transmission delays. Using time-resolved electroencephalography (EEG) decoding, we show that predictions can indeed drive sensory-like representations of visual events well ahead of the usual arrival of afferent sensory input.

An important consequence of such an anticipatory architecture is that, when predictions contradict sensory events, the visual cortex briefly represents the expected, but not presented, stimulus. This is precisely what we observe in our data: When predictions are violated, sensory information arrives too late to prevent the visual system from representing what was expected but never presented. Taken together, we provide direct evidence that predictive neural mechanisms are able to activate neural representations ahead of incoming sensory information, providing a mechanism by which the visual system might compensate for neural transmission delays.

Results and Discussion

Thirty-two observers viewed apparent motion sequences while their brain activity was recorded using EEG (average of 4,168 trials per observer). Each motion sequence consisted of between

Significance

Visual information takes time to travel from the retina and through the visual system, such that the sensory information available to the brain lags behind events in the present moment. Prediction has long been considered a fundamental principle in neuroscience. Using time-resolved EEG decoding, we show that predictive mechanisms are sufficient to activate sensory-like neural representations of anticipated future events, and that these representations are activated before the arrival of afferent sensory information. This reveals that predictive neural mechanisms might allow the visual system to overcome its neural processing delays and interact with our environment in real time.

Author contributions: T.B., P.J., and H.H. designed research; T.B. performed research; T.B., D.F., and H.H. analyzed data; and T.B., D.F., P.J., S.B., and H.H. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

Data deposition: All data and code required to reproduce the analyses and figures are uploaded onto the Open Science Framework and can be found at <http://doi.org/10.17605/OSF.IO/V9YX8>.

¹To whom correspondence may be addressed. Email: tblom@student.unimelb.edu.au.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.191777117/-DCSupplemental>.

1 and 44 black wedges presented sequentially at 100-ms intervals in eight positions on a circular trajectory around fixation (Fig. 1, *Inset*). Using the EEG responses to the initial stimulus in each sequence, we trained multivariate pattern classifiers to distinguish patterns of EEG activity evoked by the stimulus in each of the eight positions (9). Across all pairwise comparisons, classifiers were able to decode the presented stimulus position starting 50 ms after stimulus onset (Fig. 1; $P < 0.01$).

We then investigated whether the same neural representations could also be driven by predictive mechanisms alone. To do so, we used the classifiers trained to distinguish neural responses to pairs of stimuli presented in positions two steps apart (e.g., position 1 vs. 3, 2 vs. 4, etc.) from time points between 50 ms and 150 ms after stimulus onset; 50 ms is the first time point when position information was evident in the EEG signal, and, by extension, 150 ms is the earliest time point at which the neural response to the second stimulus in the sequence (presented 100 ms later) might have contaminated the training data. These

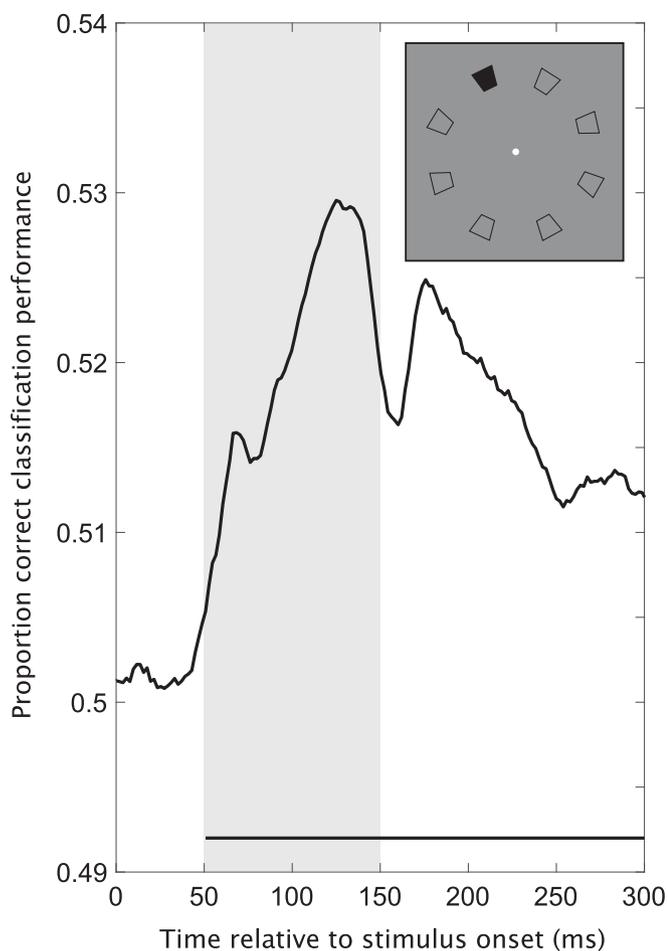


Fig. 1. Mean classification performance over time for multivariate classifiers trained to distinguish neural patterns evoked by the presentation of a black wedge in each of eight possible positions, as depicted in *Inset*. Stimuli were presented in sequences, and classifiers were trained and tested only on the first stimulus of each sequence. Plotted performance reflects the average classification performance at individual time points over all possible pairwise comparisons between positions. Time points at which classification performance significantly exceeded a permuted labels null distribution ($P < 0.01$) are indicated by the line below the curve. Classification performance first rises above chance at 50 ms poststimulus. The shaded area indicates the training period subsequently used for analyses of neural activity after the final stimulus in the motion sequences.

classifiers were then tested on the time period directly following the final stimulus in sequences of at least four consecutive stimuli. We used the classifiers to evaluate the evidence in the neural signal coding for the stimulus in the two positions adjacent to the final stimulus. One of these is the position ahead of the final stimulus, and the other is the position behind the final stimulus (Fig. 2A). We tested the proportion of trials that the classifier assigned to the “Ahead” position rather than the “Behind” position. As the stimulus is never presented in the position ahead of the final stimulus, any classification assignment to this category can be seen as evidence of predictive mechanisms activating sensory-like neural representations of the stimulus in the anticipated position.

We use temporal generalization matrices to reveal classification assignments at each combination of training and test time points. Importantly, temporal generalization matrices reveal which neural representations become active when (14). In our case, they reveal when particular representations from the training set (isolated by the choice of training time point) become active in the test set. Fig. 2B shows classification assignments plotted as a temporal generalization matrix. Diagonal lines connect points where training and test times are equal, indicating the usual time course of sensory information as it evolves through the visual system. From top to bottom, separate lines indicate the usual time course of sensory information corresponding to the penultimate, final, and expected stimulus presentations.

When tested at early time points following the final stimulus, the temporal generalization matrix reflects the sustained sensory activity coding for the stimulus behind the final stimulus, as indicated by the blue classification assignment in the top of Fig. 2B. Since the stimulus was presented in this position 100 ms prior to the final stimulus, we expected to find neural representations coding for this position activated in the EEG signal (Fig. 2B, upper solid black line).

Following the final stimulus (middle solid black line), we observed a cluster of significant classification assignments in favor of the anticipated position ahead of the final stimulus ($P < 0.001$, two-sided), without the stimulus ever being presented in this position. Since the stimulus was presented in the behind position only 100 ms earlier and the classifier is forced to choose either the Ahead or the Behind position, classification assignments to the Ahead position provide strong evidence that predictive mechanisms in the visual system indeed operate by activating sensory-like representations of anticipated object positions.

For predictive mechanisms to compensate for neural transmission delays, they must be able to produce sensory-like representations in the absence of sensory input. However, this is not sufficient; they must further be able to do so at a shorter latency than the actual afferent sensory input. Only then can neural activity get ahead of transmission delays. This is precisely what we observed: Predictive activation of the representation coding for the stimulus in the position ahead of the final stimulus is evident in the EEG roughly 70 ms to 90 ms earlier than it would have been had it been evoked by sensory input alone (Fig. 2B; the strong dark red activation cluster around training time 90 ms to 150 ms and test time 150 ms occurs above the dashed line denoting the usual time course of sensory information coding for the same position). By preactivating the anticipated future position of the object, the visual system would be able to represent the object’s position more quickly than would be possible on the basis of afferent sensory information.

These results suggest a predictive neural architecture in which sensory-like representations are activated before sensory input can arrive. Importantly, in such an architecture, when predictions are not fulfilled, sensory information about the actual events

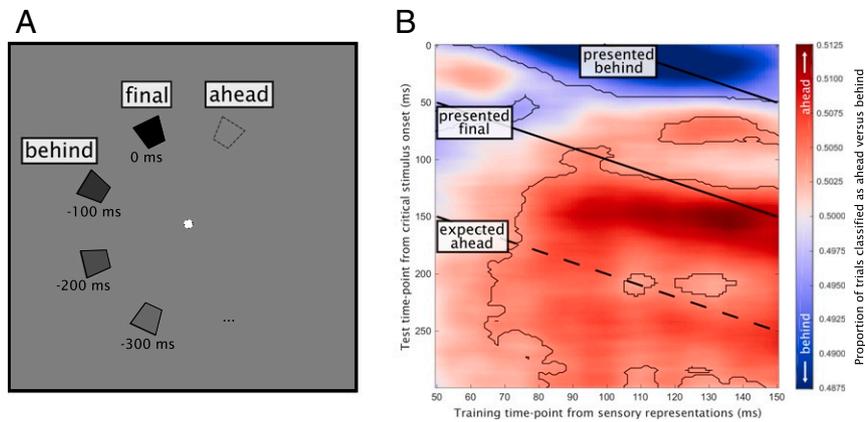


Fig. 2. Activation of sensory representations coding for the anticipated stimulus position at the end of the motion sequence. (A) Motion sequences consisted of at least four sequential stimuli in consecutive positions presented for 66 ms with a 33-ms interstimulus interval. The final stimulus is the last presented stimulus in the sequence (at 0 ms). The stimulus position ahead of the final stimulus could be expected at 100 ms, but was never presented. The stimulus position behind the final stimulus was presented at -100 ms. (B) Temporal generalization matrix revealing the proportion of trials classified as Ahead rather than Behind at each combination of training time points and test time points relative to the onset of the final stimulus. Diagonal lines connect points where training time corresponds to the time elapsed after the (real or expected) presentation of each stimulus. The upper and middle solid black lines indicate the time courses of the stimulus behind the final stimulus and the final stimulus itself, respectively. The dashed black line indicates the time course of the anticipated stimulus (100 ms after the final stimulus), had it been presented as expected. The cluster of above-chance classification performance in favor of the stimulus position ahead of the final stimulus reveals that the neural representations coding for the anticipated position are activated. Outlined clusters represent time points that are significantly different from the permuted null distribution ($P < 0.001$, two-sided).

should arrive too late to prevent the counterfactual anticipatory activation. To test this hypothesis, we also analyzed trials in which the motion sequences unexpectedly reversed direction after a minimum of eight consecutive stimuli. Using the same approach as described above, we probed neural representations in these reversal trials immediately following the last stimulus that adhered to the original motion direction (again termed the final stimulus; Fig. 3A). In contrast to the previous analysis, in these reversal trials, the stimulus in the position behind the final stimulus is not only presented at -100 ms but is also presented again 100 ms after the final stimulus.

The temporal generalization matrix for reversal trials is shown in Fig. 3B. This reveals that the pattern of EEG activity immediately following the presentation of the final stimulus initially reflects the Behind stimulus position, which had been presented earlier at -100 ms (Fig. 3B, blue classification assignment cluster

around the upper black line; $P < 0.001$, two-sided). This was followed by a cluster of classification assignment in favor of the Ahead position (the horizontal dark red cluster corresponding to training times 110 ms to 150 ms and test time 150 ms; $P < 0.001$, two-sided). The timing of this preactivation was comparable to the previous analysis in Fig. 2B, with position representations preactivated between 70 ms and 90 ms earlier than they would have been on the basis of sensory input alone. This suggests that the incoming afferent sensory information signaling the reversal was subject to neural delays, and arrived too late to prevent the visual system from transiently representing the anticipated position of the stimulus ahead of the final stimulus. When the new sensory input did finally arrive (lower black line), the neural position representation was corrected, and the visual system then represented the stimulus which was actually presented, as indicated by significant classification assignments

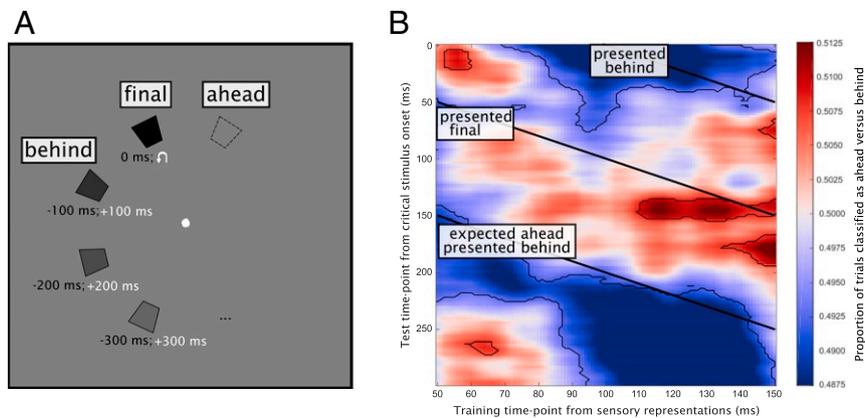


Fig. 3. Activation of sensory representations coding for the anticipated stimulus following a motion reversal. (A) Reversal sequences consisted of at least eight sequential stimuli in consecutive positions presented for 66 ms each, with a 33-ms interstimulus interval. The final stimulus was the last presented stimulus in the sequence (at 0 ms) before the reversal. The stimulus position ahead of the final stimulus was expected at 100 ms, but was not presented. Instead, the stimulus position behind the final stimulus that was presented at -100 ms was presented again at $+100$ ms. (B) Temporal generalization matrix revealing the proportion of trials classified as Ahead rather than Behind at each combination of training and test time points after the final stimulus. Diagonal lines connect points where training time corresponds to the time elapsed after the presentation of each stimulus. Outlined clusters represent time points that are significantly different from the permuted null distribution ($P < 0.001$, two-sided).

for the Behind position (Fig. 3B, blue cluster in the bottom right; $P < 0.001$, two-sided). This pattern of results provides further evidence that predictive mechanisms in the visual system are driving sensory representations ahead of the arrival of corresponding sensory information, and confirms our hypothesis that, when sensory events are counterfactual to predictions, sensory input arrives too late to prevent the visual system erroneously representing what is expected but never comes to pass.

We ran two control analyses to evaluate alternative explanations for the observed results. First, we demonstrated that the results cannot be accounted for by eye movements. Three participants completed the experiment while eye position was recorded using an eye tracker, and we performed an identical classification analysis on the basis of eye position time courses. Eye movement patterns contained no information regarding the position of the stimulus (SI Appendix, Fig. S1). Additionally, we investigated the possibility that the pattern of classification performance in favor of the stimulus position ahead of the final stimulus might reflect delayed negative responses to the stimulus behind the final stimulus. Because the classifier was forced to choose between Ahead and Behind, deactivation of the representation coding for the stimulus behind the final stimulus could, in theory, yield the same results as preactivation of the position ahead of the stimulus. To evaluate this possibility, we analyzed EEG data from shorter sequences in which only two stimuli were presented in a sequence. Strong expectations regarding the next stimulus position could not yet have built up, yet the theorized deactivation would still have occurred. These sequences revealed no decodable information in favor of the position ahead of the final stimulus (SI Appendix, Fig. S2). This indicates that the observed pattern of predictive activation is not an artifact of a deactivation of the preceding stimulus representation but is indeed related to prediction of the next stimulus position.

To evaluate which brain regions are involved with the prediction of the next stimulus position, we plotted the scalp topographies of the event-related potentials (ERPs) evoked by the physical presentation of the stimulus in the training data at the time of peak classification performance in the test data (120 ms). These scalp topographies show that the ERPs are dominated by contralateral occipital responses, as would be expected by the presentation of lateralized visual stimuli (SI Appendix, Fig. S3).

Taken together, these results demonstrate that predictive mechanisms are sufficient to drive sensory-like stimulus representations even in the absence of sensory input, and activate stimulus representations with a shorter latency than could be achieved by afferent sensory input, that is, ahead of neural transmission delays.

The existence of predictive mechanisms driving sensory representations ahead of neural processing delays suggests that such mechanisms might allow the brain to compensate for those delays. As long as events unfold predictably, they would allow the visual system to reduce the impact of neural delays. Interestingly, our results suggest that prediction takes place at multiple stages in the visual system, as evidenced by the fact that the red preactivation clusters in Figs. 2B and 3B are nearly horizontal. This shows that successive neural representations that are ordinarily activated sequentially (i.e., on a diagonal in a temporal generalization matrix) are instead activated nearly concurrently, meaning that the corresponding brain areas may have become aligned in time. This would potentially allow the brain to function in close to real time, as has been proposed previously (5). However, a consequence of such a predictive architecture is that, when predicted events are not fulfilled, sensory information about what actually happened arrives too late to prevent the system from temporarily representing the predicted event. We show that, for visual motion, this indeed happens: The visual system preactivates position representations for locations where the

object is not presented, even in situations whereby the object is actually presented in a different location.

Although the current paradigm allowed us to reveal predictive activation of neural representations ahead of neural delays, it provided no evidence for correlates of these representations in conscious awareness. Whether the brief activation of incorrect representations leads to a conscious percept of the predicted event, even if only short-lived, is therefore an interesting open question. If it does not lead to (sustained) awareness, which intuitively seems to be the case since we do not typically report “seeing” events that never happened, it would mean that the predictive activation of sensory-like representations is not sufficient for perception. Instead, it would potentially serve as input for regions farther down the processing stream, for example, to guide action (15, 16). On the other hand, perhaps we do very briefly become aware of these incorrect predictions. If so, then our conscious recollection of that experience must be corrected in the same way that the sensory representation is overwritten when new information becomes available. In addition to implementing predictive mechanisms, the visual system would then also need to implement retroactive “postdictive” mechanisms to prevent these events from impacting our (later) behavior and conscious experience. In order to create a coherent flow of events in conscious perception without violating causality at the neural level, visual events would have to be reorganized and corrected in a postdictive fashion (17, 18).

This interplay between predictive and postdictive mechanisms has a perceptual parallel in the flash-lag effect, a visual illusion in which a static object flashed in alignment with a moving object is perceived to lag behind the moving object (2). Although convergent evidence suggests the flash-lag effect is due to predictive motion extrapolation mechanisms (3, 6, 19), the flash-lag effect reverses when motion after the flash reverses direction, prompting Eagleman and Sejnowski (17) to propose an account of this illusion in terms of postdiction. In the related flash-grab illusion (20), in which an object flashed on a reversing background is mislocalized, the flashed object is similarly mislocalized in the direction of motion after the flash. It has previously been argued that a postdictive correction of failed extrapolation also contributes to this perceptual illusion (21, 22). Motion-position illusions such as the flash-lag effect and related illusions therefore provide a potential conscious correlate to the predictive mechanisms that we reveal here.

Other authors have used different paradigms to demonstrate conscious correlates of predictive mechanisms. Particularly relevant is a study by Roach et al. (23), who studied the detectability of threshold-level targets at the leading and trailing edges of sinusoidally modulated motion patches. They showed that the detectability of a sinusoidal target was dependent on its phase congruency with the motion patch at the leading edge, but not at the trailing edge. As an explanation for this phenomenon, they proposed that the representation of a future pattern activated by predictive mechanisms is added to the incoming sensory information. When the target and the motion patch are in phase, this boosts the sensory signal and reduces the detection threshold; conversely, when they are out of phase, the sensory signal is attenuated, and the detection threshold rises. This notion of summation presumes that the neural representations activated by predictive mechanisms are the same as those activated by corresponding sensory input, which is consistent with our findings.

Another study showing conscious consequences of interference between motion prediction and sensory input was reported by Hogendoorn et al. (24). They showed targets in between possible positions of an apparent motion display with opposite luminance polarity, and observed interference (as measured using reaction time [RT]) on detection and discrimination

performance when the target was presented in the predicted position of the apparent motion stimulus (i.e., directly between two successive presentations). Interestingly, they observed a dissociation between predictive and postdictive mechanisms when the apparent motion stimulus reversed direction. When using a simple detection task (with low mean RT), interference remained evident in the original direction of motion, consistent with predictive activation. However, when using a discrimination task (with higher mean RT), interference was found in the postreversal direction of motion. This behavioral finding directly parallels the pattern of neural activation we observe here in our reversal trials, whereby the neural signal initially codes for the position ahead of the moving object, and then subsequently codes for the opposite direction once incoming sensory information becomes available.

Although we have not demonstrated a conscious correlate to the neural preactivation, the pattern of results is highly consistent with predictive and postdictive perceptual phenomena reported in previous behavioral experiments. One way to probe the possible conscious correlate of the preactivation in our paradigm would be to apply the same summation logic used by Roach et al. (23) and Hogendoorn et al. (24). This could be achieved by presenting a threshold-level stimulus at the end of an apparent motion sequence, timed such that its sensory representation coincides with the preactivation we demonstrate here. Our current results would predict that preactivation of that neural representation would lower the detection threshold for that stimulus.

In conclusion, we demonstrate that predictive mechanisms are able to activate sensory-like neural representations ahead of the arrival of afferent sensory information. This provides a candidate mechanism by which the brain would be able to act in the present, while only having access to delayed sensory information. Prediction has long been considered a fundamental principle in neuroscience (25–28), and here we show how predictive mechanisms could allow the brain to function in close to real time despite its internal temporal constraints.

Materials and Methods

The current experiment includes data collected using two slightly different protocols. Unless otherwise stated, the following details pertain to both protocols.

Human Subjects. Twelve observers (six female, average age 25 y) with normal or corrected-to-normal vision participated under the first protocol, and 20 observers (12 female, average age 23 y) with normal or corrected-to-normal vision participated under the second protocol. Both protocols were approved by the human research ethics committee of the University of Melbourne (Ethics ID 1954628), Australia, and conducted in accordance with the Declaration of Helsinki. All observers signed informed consent before participating in the experiment and were reimbursed AUD15 per hour for their time.

Stimuli. The stimulus was a black, truncated wedge presented on a uniform 50% gray background. The stimulus could be presented in one of eight equally spaced locations around a white central fixation point, at 22.5°, 67.5°, 112.5°, 157.5°, 202.5°, 247.5°, 292.5°, and 337.5° of polar angle from the vertical (Fig. 1, *Inset*). Inner and outer edges of the wedge were 6.3° and 7.7° of visual angle away from fixation, respectively. The wedge covered 11° of polar angle, with 1.3° of visual angle at the inner and 1.5° of visual angle at the outer edge. The stimulus was presented for 66 ms, with an interstimulus interval of 33 ms and an intertrial interval of 400 ms between sequences. Stimuli were presented on an ASUS ROG PG258 monitor with a resolution of 1,920 × 1,080 running at 120 Hz. The monitor was controlled by an HP EliteDesk 800 G3 TWR PC running Matlab R2017b with PsychToolbox 3.0.14 extensions (29). Observers viewed the stimuli from a headrest at a distance of 60 cm.

Behavioral Task. Observers were instructed to respond with a button press on the keyboard whenever a wedge was presented in red instead of black. This occurred a total of 32 times per block under protocol 1 and 50 times

per block under protocol 2. The purpose of the task was to keep observers engaged with the stimulus, and behavioral data were not analyzed. Under protocol 1, trials with targets were discarded, and target trials were shown again at the end of each block. Under protocol 2, trials with targets were simply discarded.

Experimental Design. Twelve observers completed six blocks of sequences across three testing sessions (protocol 1). Twenty observers completed two blocks across two testing sessions (protocol 2).

Under protocol 1, each block contained the following types of trials, randomly interleaved:

- 1) Sequences with one, two, or three consecutive presentations starting at each position and moving in both directions were presented 10 times (3 sequence lengths × 8 starting positions × 2 directions × 10 repetitions = 480 trials). These trials are not analyzed in the main results here, but sequences with length 2 are analyzed and presented in *SI Appendix, Fig. S2* as a control.
- 2) Sequences with four, five, six, seven, or eight consecutive presentations starting at each position and moving in both directions were presented twice (5 sequence lengths × 8 starting positions × 2 directions × 2 repetitions = 160 trials). These trials are included in the analysis presented in Fig. 2B.
- 3) Sequences with 16, 20, 24, 28, 32, 36, 40, or 44 consecutive presentations starting at each position and moving in both directions were presented once (8 sequence lengths × 8 starting positions × 2 directions = 128 trials). These trials are included in the analysis presented in Fig. 2B.
- 4) Sequences with 16, 20, 24, 28, 32, 36, 40, or 44 consecutive presentations starting at each position and moving in both directions followed by a reversal and continuation in the opposite direction for 8 to 16 (randomly determined) additional presentations were presented once (8 sequence lengths × 8 starting positions × 2 directions = 128 trials). These trials are included in the analysis presented in Fig. 3B.

A target stimulus was randomly presented in 32 trials per block. Because these trials were appended to the trial list, each block encompassed 928 trials (split up into 16 sets of 58 trials). Each new set was initiated with a button press by the observer. Each observer completed two blocks per session, and a block lasted ~30 min. Taken together, each observer completed 5,568 trials.

Under protocol 2, all types of trials were combined in a single block, randomly interleaved:

- 1) Sequences with four, five, six, seven, or eight consecutive presentations starting at each position and moving in both directions were presented eight times (5 sequence lengths × 8 starting positions × 2 directions × 8 repetitions = 640 trials). These trials are included in the analysis presented in Fig. 2B.
- 2) Sequences with 9, 10, 11, 12, 13, 14, 15, or 16 consecutive presentations starting at each position and moving in both directions were presented four times (8 sequence lengths × 8 starting positions × 2 directions × 4 repetitions = 512 trials). These trials are included in the analysis presented in Fig. 2B.
- 3) Sequences with 9, 10, 11, 12, 13, 14, 15, or 16 consecutive presentations starting at each position and moving in both directions followed by a reversal and continuation in the opposite direction for one to eight (randomly determined) additional presentations were presented four times (8 sequence lengths × 8 starting positions × 2 directions × 4 repetitions = 512 trials). These trials are included in the analysis presented in Fig. 3B.

In each block, a target was randomly presented in 50 trials, and these trials were discarded. Each block was split up into 13 sets, and each new set was initiated with a button press by the observer. In a session, observers completed one block, taking ~90 min. Taken together, each observer completed 3,328 trials.

EEG Acquisition and Preprocessing. The 64-channel EEG data, and data from six EOG and two mastoid channels, were acquired using a BioSemi ActiveTwo EEG system sampling at 2,048 Hz. EEG data were rereferenced offline to the average of the two mastoid electrodes and resampled to 512 Hz. Eleven observers had one bad channel during one of the sessions. This channel was spherically interpolated using EEGlab (30).

All data were epoched relative to stimulus onset. For classification analyses, we make a distinction between two types of epochs: training and test epochs. Training epochs were used to train classifiers to distinguish stimulus positions, whereas the test epochs were used to test the classifiers.

Under both protocols, the training epochs were time-locked to the first presentation in a sequence. The first presentation in each sequence was random and has no history, and therefore its position could not be anticipated. The training epochs were taken from 800 ms before stimulus onset to 500 ms after and were baseline-corrected to the mean of the 200-ms period before stimulus onset.

Epochs in the training sets with large artifacts were removed using an automated rejection procedure. Epochs for which the SD of amplitudes across all electrodes and all time points exceeded 50 μ V were rejected for analysis. Across all observers, 35.2% (SD = 16.0%) of epochs were removed in this way. This conservative rejection procedure was chosen to avoid introducing systematic patterns in the data which the classifiers might exploit.

The test set consisted of all other stimuli, and epochs were again taken from 800 ms before stimulus onset to 500 ms after and were baseline-corrected to the mean of the 800-ms period before stimulus onset. This baseline period was chosen such that it was consistent across all epochs and contained a full cycle of motion on the majority of the epochs, in order to avoid introducing stimulus-specific differences as much as possible.

Multivariate Pattern Analysis. Epochs from the training sets were used to train time-resolved linear discriminant analysis (LDA) classifiers (31). The classifier was first trained to separate neural patterns associated with the presentation of the stimulus in two different positions using the amplitudes from all 64 EEG channels as features. For the classification performance shown in Fig. 1, a set of LDA classifiers was trained on each pairwise combination of eight stimulus positions at each individual time point in the epoch for each individual observer. Each time point was defined as the average of the surrounding 10-ms time window. The classifier was trained on half of the data (to estimate a model) and subsequently tested on the other half (to test the quality of the model, i.e., how well it generalized). This procedure was repeated independently after swapping the roles of the two halves, and both classification performances were averaged. An above-chance classifier performance at a certain time point indicates that the distributed pattern of EEG signal contained information that allowed the classifier to distinguish the two positions of the stimulus. Fig. 1 shows the average classification performance across all of the pairwise comparisons.

To answer the question of whether prediction alone can drive neural representations, we used classifiers trained on sensory representations evoked

by the presentation of the stimulus two positions apart in the training set, and analyzed to which position on either side of the final stimulus the classifier assigned the trial, based on EEG data following the final stimulus in the sequence. Classifiers were trained on all trials from the training set to discriminate the stimulus position Ahead of the final stimulus versus the position Behind the final stimulus, and we investigated classification assignments in the time period following the final stimulus. Because no subsequent stimulus was presented after the final stimulus, this allowed us to probe which position representation was “activated” without a stimulus being presented.

Quantification and Statistical Analysis. Statistical analyses were conducted using permutation tests. For the mean classification performance in Fig. 1, a null distribution was created by generating 1,000 datasets with shuffled position labels (i.e., the position within the stimulus array), and each permuted dataset was analyzed in the exact same way as the original dataset. Time points at which the observed classification performance fell in the highest or lowest 0.5% of the null distribution were taken as significant ($P < 0.01$, two-tailed, uncorrected). For the temporal generalization matrices in Figs. 2B and 3B, a null distribution was created by generating 1,000 datasets with shuffled condition labels (in relation to whether a neighboring location was Ahead or Behind), and each permuted dataset was analyzed in the exact same way as the original dataset. Individual time points at which the absolute value of the difference between the observed classification assignment and chance (0.5) fell in the highest 0.1% of the distribution at that time point were taken as significant ($P < 0.001$, two-tailed, uncorrected). Significant points in the training time–test time space are outlined in their respective figures.

Materials and Data Availability. All data and code required to reproduce the analyses and figures are uploaded onto the Open Science Framework and can be found at <http://doi.org/10.17605/OSF.IO/V9YX8>.

ACKNOWLEDGMENTS. This work was supported by the Australian Government through the Australian Research Council’s Discovery Projects funding scheme (Project DP180102268) awarded to H.H. We would like to thank Jane Yook for her help in the data collection, and the members of both the Time in Brain and Behavior Laboratory and the Decision Neuroscience Laboratory at the University of Melbourne for their insightful comments on the work.

1. J. H. Maunsell, J. R. Gibson, Visual response latencies in striate cortex of the macaque monkey. *J. Neurophysiol.* **68**, 1332–1344 (1992).
2. R. Nijhawan, Motion extrapolation in catching. *Nature* **370**, 256–257 (1994).
3. R. Nijhawan, Visual prediction: Psychophysics and neurophysiology of compensation for time delays. *Behav. Brain Sci.* **31**, 179–198 (2008).
4. M. J. Berry, I. H. Brivanlou, T. A. Jordan, M. Meister, Anticipation of moving stimuli by the retina. *Nature* **398**, 334–338 (1999).
5. H. Hogendoorn, A. N. Burkitt, Predictive coding with neural transmission delays: A real-time temporal alignment hypothesis. *eNeuro* **6**, ENEURO.0412-18.2019 (2019).
6. P. A. White, Is the perceived present a predictive model of the objective present?. *Vis. Cognit.* **26**, 624–654 (2018).
7. D. Jancke, W. Erlhagen, G. Schöner, H. R. Dinse, Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex. *J. Physiol.* **556**, 971–982 (2004).
8. M. Subramanian *et al.*, Faster processing of moving compared with flashed bars in awake macaque V1 provides a neural correlate of the flash lag illusion. *J. Neurophysiol.* **120**, 2430–2452 (2018).
9. H. Hogendoorn, A. N. Burkitt, Predictive coding of visual object position ahead of moving objects revealed by time-resolved EEG decoding. *NeuroImage* **171**, 55–61 (2018).
10. P. Kok, M. F. Failing, F. P. de Lange, Prior expectations evoke stimulus templates in the primary visual cortex. *J. Cognit. Neurosci.* **26**, 1546–1554 (2014).
11. F. P. de Lange, M. Heilbron, P. Kok, How do expectations shape perception?. *Trends Cognit. Sci.* **22**, 764–779 (2018).
12. G. Demarchi, G. Sanchez, N. Weisz, Automatic and feature-specific prediction-related neural activity in the human auditory system. *Nat. Commun.* **10**, 3440 (2019).
13. M. Ekman, P. Kok, F. P. De Lange, Time-compressed preplay of anticipated events in human primary visual cortex. *Nat. Commun.* **8**, 15276 (2017).
14. J. R. King, S. Dehaene, Characterizing the dynamics of mental representations: The temporal generalization method. *Trends Cognit. Sci.* **18**, 203–210 (2014).
15. E. Brenner, J. B. Smeets, Fast responses of the human hand to changes in target position. *J. Mot. Behav.* **29**, 297–310 (1997).
16. E. Brenner, J. B. Smeets, How people achieve their amazing temporal precision in interception. *J. Vis.* **15**, 8–8 (2015).
17. D. M. Eagleman, T. J. Sejnowski, Motion integration and postdiction in visual awareness. *Science* **287**, 2036–2038 (2000).
18. S. Shimojo, Postdiction: Its implications on visual awareness, hindsight, and sense of agency. *Front. Psychol.* **5**, 196 (2014).
19. R. Nijhawan, Neural delays, visual motion and the flash-lag effect. *Trends Cognit. Sci.* **6**, 387–393 (2002).
20. P. Cavanagh, S. Anstis, The flash grab effect. *Vis. Res.* **91**, 8–20 (2013).
21. T. Blom, Q. Liang, H. Hogendoorn, When predictions fail: Correction for extrapolation in the flash-grab effect. *J. Vis.* **19**, 3 (2019).
22. K. M. Coffey *et al.*, Expecting the unexpected: Temporal expectation increases the flash-grab effect. *J. Vis.* **19**, 9 (2019).
23. N. W. Roach, P. V. McGraw, A. Johnston, Visual motion induces a forward prediction of spatial pattern. *Curr. Biol.* **21**, 740–745 (2011).
24. H. Hogendoorn, T. A. Carlson, F. A. Verstraten, Interpolation and extrapolation on the path of apparent motion. *Vis. Res.* **48**, 872–881 (2008).
25. R. P. N. Rao, D. H. Ballard, Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
26. H. Feldman, K. J. Friston, Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* **4**, 215 (2010).
27. R. A. Adams, S. Shipp, K. J. Friston, Predictions not commands: Active inference in the motor system. *Brain Struct. Funct.* **218**, 611–643 (2013).
28. J. Hohwy, The predictive mind. *Choice Rev. Online* **52**, 52–1343 (2014).
29. D. H. Brainard, The psychophysics toolbox. *Spatial Vis.* **10**, 433–436 (1997).
30. A. Delorme, S. Makeig, EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* **134**, 9–21 (2004).
31. T. Grootswagers, S. G. Wardle, T. A. Carlson, Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time-series neuroimaging data. *J. Cognit. Neurosci.* **29**, 677–697 (2017).