Title: Visual information is predictively encoded in occipital alpha/low-beta oscillations

Abbreviated Title: Occipital rhythms carry predictive information

William Turner[1,2*], Tessel Blom[2], Hinze Hogendoorn[1,2]

[1]Queensland University of Technology, Brisbane, Australia, 4059

[2]Melbourne School of Psychological Sciences, The University of Melbourne, Melbourne, Australia, 3010

* Corresponding Author, Email: w6.turner@qut.edu.au

Number of Figures: 5

Number of Words: 250 (Abstract), 117 (Significance), 650 (Introduction), 1679 (Discussion)

Conflict of interest

The authors declare no competing financial interests.

Acknowledgements

This work was supported by Australian Research Council Grants FT200100246 and DP180102268 awarded to Hinze Hogendoorn. We thank Jane Yook, Milan Andrejević, and Vinay Mepani for help with data collection, and Philippa Johnson for helpful discussions.

Abstract

Hierarchical predictive coding networks are a general model of sensory processing in the brain. Under neural delays, these networks have been suggested to naturally generate oscillatory activity in approximately the alpha frequency range (~8-12 Hz). This suggests that alpha oscillations, a prominent feature of EEG recordings, may be a spectral 'fingerprint' of predictive sensory processing. Here, we probed this possibility by investigating whether oscillations over the visual cortex predictively encode visual information. Specifically, we examined whether their power carries information about the position of a moving stimulus, in a temporally predictive fashion. In two experiments (N = 32, 18 female; N = 34, 17 female), participants viewed an apparent-motion stimulus moving along a circular path, while EEG was recorded. To investigate the encoding of stimulus-position information, we developed a method of deriving probabilistic spatial maps from oscillatory power estimates. With this method, we demonstrate that it is possible to reconstruct the trajectory of a moving stimulus from alpha/low-beta oscillations, tracking its position even across unexpected motion reversals. We also show that future position representations are activated in the absence of direct visual input, demonstrating that temporally predictive mechanisms manifest in alpha/beta-band oscillations. In a second experiment we replicate these findings and show that the encoding of information in this range is not driven by visual entrainment. By demonstrating that occipital alpha/beta oscillations carry stimulus-related information, in a temporally predictive fashion, we provide empirical evidence of these rhythms as a spectral 'fingerprint' of hierarchical predictive processing in the human visual system.

Significance Statement

'Hierarchical predictive coding' is a general model of sensory information processing in the brain. When *in silico* predictive coding models are constrained by neural transmission delays, their activity naturally oscillates in roughly the alpha range (~8-12 Hz). Using time-resolved EEG decoding, we show that neural rhythms in this approximate range (alpha/low-beta) over the human visual cortex predictively encode the position of a moving stimulus. From the amplitude of these oscillations we are able to reconstruct the stimulus' trajectory, revealing signatures of temporally-predictive processing. This provides direct neural evidence linking occipital alpha/beta rhythms to predictive visual processing, supporting the emerging view of such oscillations as a potential spectral 'fingerprint' of hierarchical predictive processing in the human visual system.

Introduction

'Predictive coding' is a general model of the hierarchical inference process underlying visual processing (Rao & Ballard, 1999). The functional architecture of the visual system implied by predictive coding is that of a hierarchical network of interconnected neural populations. The higher levels of this network attempt to predict the activity of lower levels, with the residuals of these predictions being passed back upwards.

In the predictive coding literature, the fact that neural signalling takes time has often been overlooked (but see Friston, 2008; Hogendoorn & Burkitt, 2019). Consideration of this fact places important constraints on predictive coding models, in that predictions and residuals can never be transmitted instantaneously, but must rather pass between levels with some delay. Recent theoretical work has suggested that when biologically plausible signalling delays are built into hierarchical predictive coding networks, the recursive network dynamics naturally generate oscillatory activity in approximately the alpha frequency range (~8-12 Hz, with the precise frequency depending on the signalling delay and neural time constant; Alamia & VanRullen, 2019). This is important because it suggests that oscillations in this general frequency range may (in some cases) be a signature of predictive sensory processing, arising from rhythmic 'message passing' between hierarchically-organised neural populations. If this is true, one might expect features of these rhythms, such as their power (squared amplitude), to carry information about the underlying stimulus being processed. However, this has yet to be directly tested. The primary aim of this study was therefore to examine whether the power of alpha oscillations over the occipital cortex carries stimulus-related information.

One complication which arises when incorporating neural delays into a hierarchical predictive coding framework is that for time-varying input, backwards predictions will always conflict with sensory input if neural delays are not accounted for. To effectively minimize prediction error, information processing must not only be hierarchically predictive, but also *temporally* predictive. That is, extrapolation mechanisms are needed that adjust forwards and backwards signals and correct for the lag incurred during signal transmission (Hogendoorn & Burkitt, 2019). Consequently, if alpha oscillations are a signature of predictive coding, the information they carry should display temporally predictive/anticipatory qualities. When prior expectations about the stimulus can be generated, these rhythms should carry information about expected input, even in the absence of feed-forward signals. While there is mounting evidence that neural activity patterns during visual processing do carry predictive information (e.g., Blom et al., 2020; Kok et al., 2014, 2017; Liu et al., 2021), the spectral locus of such information has typically not been investigated.

In the present study, we examined whether and how information about the position of a predictably moving stimulus manifests in oscillations over the occipital cortex. In two experiments (N = 32, 34), participants viewed an apparent motion stimulus (i.e. a series of spatially and temporally separated flashes that generate the percept of coherent motion) travelling along a circular path while EEG was recorded. In these sequences, the stimulus' trajectory was predictable, meaning its future position could be anticipated, although the end of each sequence was unexpected. Importantly, in a previously published analysis of the dataset from Experiment 1 we demonstrated that predictions about the upcoming stimulus position were evident in the EEG signal (Blom et al., 2020). Here, we use a complementary analysis strategy to investigate whether predictive representations manifest in specific oscillatory frequency bands. To do so, we develop a method for constructing probabilistic spatial maps from oscillatory power estimates. With this method, we demonstrate that the location of the stimulus can be decoded from occipital oscillations in the alpha/low-beta range (peak information at ~12 Hz). We also observe anticipatory activation of neighbouring unstimulated position representations at the end of motion sequences, suggesting that the processes underlying predictive spatial pre-activation manifest in alpha/beta-band oscillations. In a second experiment, we replicate and extend these findings, ruling out the possibility that the encoding of information in this range is driven by visual entrainment.

Materials and Methods

Experiment 1

This experiment includes data collected using two slightly different protocols. Note that separate investigation of this dataset has been previously reported (Blom et al., 2020, 2021).

Participants

Twelve observers (6 female, mean age 25 years) participated under the first protocol and twenty observers (12 female, mean age 23 years) participated under the second protocol. All had normal or corrected-to-normal vision. Both protocols were approved by the human research ethics committee of the University of Melbourne (Ethics ID 1954628), Australia. All observers gave written informed consent prior to participating and were reimbursed AUD15 per hour.

Procedure

The stimulus was a black, truncated wedge presented on a uniform 50% gray background. The stimulus could appear 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 (Figure 1). 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. Participants viewed the stimuli from a headrest at a distance of 60 cm.

Task

Participants viewed an apparent motion stimulus moving along a circular trajectory, while EEG was recorded. After moving for between 3 and 44 flash repetitions (300 ms and 4.4 s), the stimulus either disappeared or reversed its direction (Figure 1). Participants were tasked with making a button press whenever the stimulus was coloured red instead of black. This occurred 32 times per block under protocol 1 and 50 times per block under protocol 2. The task was designed to keep participants 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.

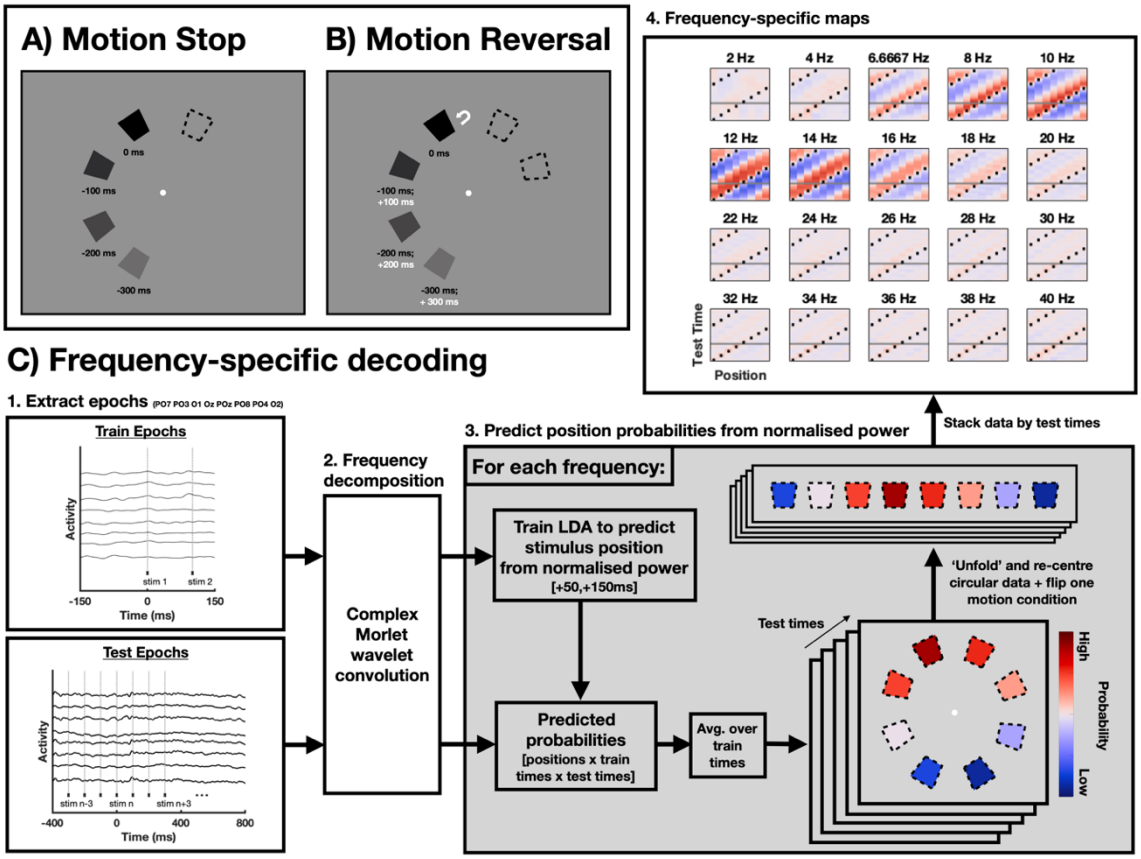

Figure 1. Stimulus display and analysis pipeline. Participants viewed an apparent motion stimulus moving through 8 positions around a circle. In Experiment 1, the position of the stimulus was updated every 100 ms (66 ms stimulus on screen, 33 ms ISI). In Experiment 2, the update rate was varied between 100 ms (10 Hz), 125 ms (8 Hz) and 150 ms (6.67 Hz) by adjusting the ISI. After moving for between 3 and 44 presentations (300 ms - 4.4 s) the stimulus would either **A)** disappear or, in a subset of trials, **B)** reverse its direction of motion. **C)** Recordings from 8 occipital electrodes were separated into a training set (epochs around the first flash in an apparent motion sequence) and a testing set (four specific epochs: Start, Middle, Stop and Reversal). Complex Morlet wavelet convolution was used to extract frequency-specific power estimates. For each frequency, LDA classifiers were then trained to predict the stimulus position from normalized power estimates at each training time point (+50 to +150 ms). Across testing timepoints, the average predicted posterior probabilities for the stimulus occupying each of the 8 possible positions was then taken. In other words, at each testing timepoint, predicted posterior probabilities were generated across all pre-trained temporally-specific classifiers (+50 to +150 ms) and the average across these was taken. The data was then re-centred on the presented stimulus position and one motion direction was 'flipped', yielding frequency-specific stimulus-position maps.

Experimental Design

Under protocol 1, participants completed six blocks of sequences across three testing sessions. Under protocol 2, participants completed two blocks across two testing sessions.

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).
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).
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).
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).

Because 32 target trials were appended to the trial list, each block encompassed 928 trials (in 16 sets of 58 trials). Each set was initiated with a button press. Each participant completed two blocks per session, with a block lasting ~30 min. In total, each participant 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).
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).
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).

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 set was initiated with a button press. In a session, participants completed one block, taking ~90 min. In total, each participants 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 re-referenced offline to the average of the two mastoid electrodes and resampled to 512 Hz. Eleven participants had one bad channel during one of the sessions. This channel was spherically interpolated using EEGlab (Delorme & Makeig, 2004).

All data were epoched relative to stimulus onset. For the decoding analysis, we make a distinction between training and test epochs. Training epochs (-150 to +150ms) were used to train temporally-specific LDA classifiers. Under both protocols, the training epochs were time-locked to the first presentation in a sequence. The initial stimulus was random and had no history, meaning its position could not be anticipated. The training data was initially epoched from -800 ms before stimulus onset to +800 ms after and baseline-corrected to the mean of the 200-ms period before stimulus onset. Reduced epochs (-150 to +150ms) were then extracted and concatenated prior to time-frequency decomposition.

Test epochs were extracted relative to the onset of four events of interest ('Start', 'Middle', 'Stop', and 'Reversal'). Initial epochs were again taken from -800 ms to +800 ms 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. Reduced epochs (-400 to +800ms) were then extracted and concatenated prior to time-frequency decomposition. Training and testing epochs in which the amplitudes across any of the 8 occipital electrodes exceeded 100 μV were rejected. Across all observers, 11.70% (SD = 6.98 %) of epochs were removed in this way.

Time-frequency decomposition and power-based decoding analysis

To focus on EEG activity recorded over the visual cortex, our analyses were restricted to the eight occipital electrodes (PO7 PO3 O1 POz Oz O2 PO4 PO8). To construct the training set, we extracted epochs between -150 and +150 ms after the onset of the first stimulus in the apparent motion sequences. To construct the testing set, we extracted epochs between -400 and +800 ms relative to the four events of interest ('Start', 'Middle', 'Stop', and 'Reversal').

We decoded from timepoint-specific normalized power estimates, to avoid potential issues with baselining (Hajonides et al., 2020). To extract these power estimates, time-frequency decomposition was performed using custom MATLAB code. The EEG time series was convolved with a set of complex Morlet wavelets, defined as Gaussian-windowed complex sine waves: $e^{i2\pi tf}e^{-t2/(2*\sigma 2)}$, where t is time, f is frequency (which increased from 2 to 40 Hz in 20 linearly spaced steps, although for consistency the third extracted frequency was set to 6.67 Hz to align with the slowest stimulus presentation rate in Experiment 2), and σ defines the width of each frequency band, defined as $n/2\pi f$, with n logarithmically increasing from 3 to 10. From the resulting analytic signal (z) we obtained power estimates defined as $p(t) = |z(t)|^2$.

To investigate the spectral locus of stimulus-position information, we trained LDA classifiers at each training time point (+50 and +150 ms) to predict the position of the initial stimulus from frequency-specific normalized power estimates. Across testing timepoints (-400 to +800ms) we then took individual trials, and computed the posterior probabilities associated with the stimulus being in each of the 8 possible positions. Averaging across testing trials this yielded 8 values indicating the probability the stimulus was in a given position, for a given training representation. To temporally-generalize this measure, we averaged the probabilities from each of the temporally-specific classifiers (+50 to +150 ms). This yielded a time-generalized measure of the relative probability that a stimulus was in each of the possible locations, at a given testing timepoint (i.e., a probabilistic map). Temporal generalization was necessary to allow for the fact that the timing of sensory processing likely changes when stimuli are predictable (Blom et al., 2020). Finally, we re-ordered the resulting probability values to centre the location of the presented stimulus at t=0, flipping one motion direction condition to align the probability estimates. Averaging across stimulus positions and motion directions this yielded frequency-specific maps of the stimulus position over time.

To examine the timecourses of position information encoding, and to test for evidence of temporal prediction, we extracted the position evidence timecourse for the location one step ahead of the position the stimulus occupied at t=0 (i.e. one position forwards along its original trajectory of motion). This allowed us to see whether future (expected) position representations

were activated when the stimulus unexpectedly reversed direction or disappeared (i.e. in the absence of direct visual input). To assess the frequency specificity of stimulus-position information encoding, we convolved a cosine function with each frequency-specific spatial tuning function (i.e. probabilistic maps constructed from the power of individual frequencies). Averaging across time (+50 to +150 ms), this yielded a single estimate of the strength of the stimulus-position information encoded at each frequency.

Statistical Analysis

We adopted a non-parametric approach to analysing the position evidence timecourses (Figure 2B and 3B). Specifically, we estimated a one-sided bias-corrected and adjusted bootstrapped confidence interval around the mean (10000 bootstrapped samples, alpha levels of 0.05 and 0.01). Timepoints where this interval exceeded 0 were taken as being significantly different from chance.

Experiment 2

To control for the underlying rhythmicity of the stimulus we ran a second experiment in which the update-rate of the stimulus was varied across three frequencies: 100 ms (10 Hz), 125 ms (8 Hz) and 150 ms (6.67 Hz). Unless otherwise stated the procedure employed in Experiment 2 was identical to Experiment 1.

Participants

Thirty four observers (17 female, mean age 25 years) participated in Experiment 2. All observers had normal or corrected-to-normal vision. The experimental protocol was approved by the human research ethics committee of the University of Melbourne (Ethics ID 1954628), Australia. All observers gave written informed consent prior to participating and were reimbursed AUD15 per hour.

Procedure

Participants completed three separate blocks of apparent motion sequences in a random order. In each block the interstimulus interval (ISI) was varied (Block Type 1: 33.33 ms ISI – 100 ms update rate; Block Type 2: 58.33 ms ISI – 125 ms update rate; Block Type 3: 83.33 ms ISI – 150 ms update rate), while the stimulus presentation time (66.66 ms) was held constant. Each block consisted of sequences of 4-12 consecutive presentations starting at each position and

moving in both directions, presented 6 times (9 sequence lengths × 8 starting positions × 2 directions × 6 repetitions = 864 trials).

EEG acquisition and preprocessing

All acquisition and screening procedures were identical to Experiment 1. Thirteen participants had one bad channel during one of the sessions. These channels were spherically interpolated using EEGlab (Delorme & Makeig, 2004). Across all participants, 14% (SD = 9.13%) of epochs were removed for exceeding the 100 uV limit. Identical time frequency decomposition, decoding, and statistical analysis methods to those used in Experiment 1 were used to analyse the data from Experiment 2.

Data and Code Availability

Code and data for recreating all analyses will be made available on the open science framework at the time of publication: https://osf.io/x8n9p/.

Results

Figure 2B shows probabilistic stimulus-position maps derived from the power of occipital oscillations in the alpha/low-beta range (10-16 Hz), split by motion direction and event of interest ('Start', 'Middle', 'Stop', 'Reversal', Figure 2A).

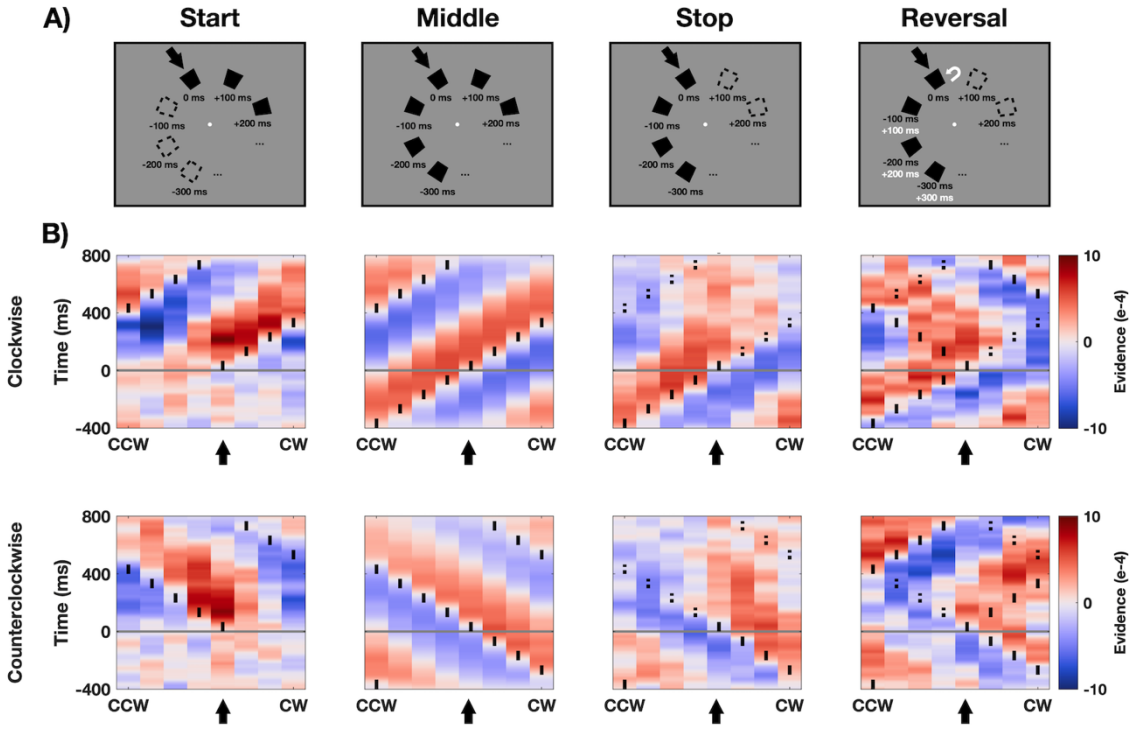

Figure 2. Tracking the position of moving stimuli from occipital alpha/low-beta power. A) Illustration of the four events of interest: 'Start': the initial stimulus in a motion sequence, 'Middle': presentation of a stimulus embedded within an ongoing sequence, 'Stop': presentation of the final stimulus in a motion sequence, and "Reversal": presentation of final stimulus before a motion reversal. **B)** Probabilistic stimulus-position maps derived from the power of occipital oscillations in the alpha/low-beta range. Red indicates high probability regions and blue indicates low probability regions. Time is shown on the y-axis and spatial position on the x-axis. The solid black lines in each map indicate the position of the stimulus as it moves along its motion trajectory. Dotted lines (in the Stop and Reversal maps) indicate positions the stimulus would have occupied if it had kept moving as expected, rather than stopped or reversed. Maps are split by the motion direction of the stimulus (clockwise, counterclockwise) and the event of interest. In panels A and B, solid black arrows mark the position of the relevant event of interest.

Figure 3A shows the same data, collapsed across motion direction. The probabilistic maps presented in Figures 2 and 3 reveal that occipital alpha/low-beta oscillations contain position information, with the high probability region (shown in red) consistently tracking the location of the stimulus, even when the stimulus unexpectedly reverses direction. This

demonstrates that it is possible to reconstruct the stimulus' trajectory from the power of these oscillations alone.

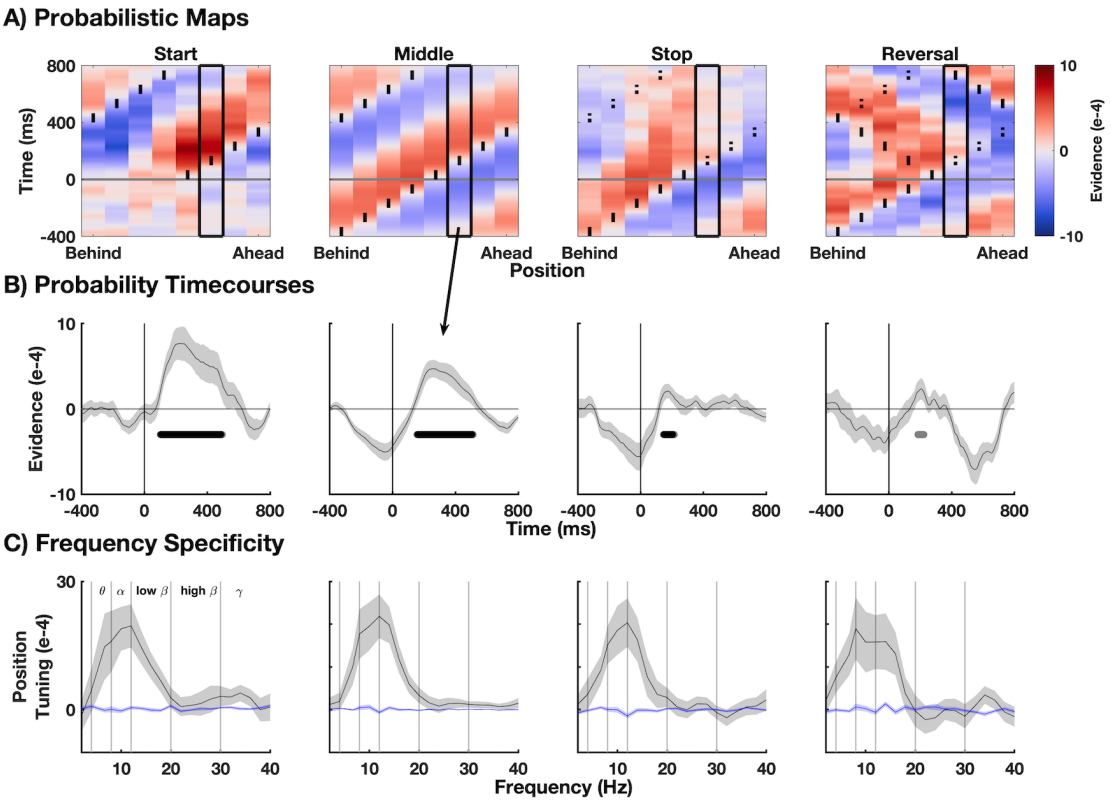

Figure 3. Experiment 1. A) Stimulus-position maps derived from the power of alpha/low-beta oscillations (10-16 Hz) over the occipital cortex. Red indicates high probability regions and blue indicates low probability regions. Time is shown on the y-axis and spatial position on the x-axis. The solid black lines in each map indicate the position of the stimulus as it moves along its motion trajectory. Dotted lines (in the Stop and Reversal maps) indicate positions the stimulus would have occupied if it had kept moving as expected, rather than stopped or reversed. **B)** Probability time-courses for the +1 ahead position. Timepoints where the lower bound of single-sided bias-corrected and accelerated (BCa) bootstrapped 95% or 99% confidence intervals exceeded zero are marked with grey and black dots respectively. Solid vertical lines mark the onset of the stimulus at t=0. **C)** Frequency specificity of positional information. Position tuning is quantified as the average cosine-convolved evidence between +50 and +150 ms. Grey and blue lines show position evidence calculated on actual and scrambled position maps. The vertical grey lines mark the approximate boundaries of the canonical EEG frequency bands (theta: 4-8 Hz, alpha: 8-12Hz, low beta: 12-20 Hz, high beta: 20-30 Hz, gamma: 30-40 Hz).

To examine whether position information was encoded in a temporally predictive fashion, we examined the evidence timecourse for the position one step ahead of the position the stimulus occupied at t=0 (i.e. one position forwards along its original trajectory of motion, Figure 3B). In line with our previous work using more conventional classification analysis applied to raw EEG amplitudes (Blom et al., 2020), these timecourses reveal that when the

stimulus stopped or reversed, there was anticipatory activation of the next expected position representation at the time of expected presentation.

To assess the frequency specificity of stimulus-position information encoding, we convolved a cosine function with each frequency-specific spatial tuning function (i.e. probabilistic maps constructed from the power of individual frequencies) between +50 and +150 ms (Hajonides et al., 2021). Averaging across time, this yielded a single estimate of the strength of the stimulus-position information encoded at each frequency. Consistent with the theoretical work of Alamia and VanRullen (2019), this analysis revealed that stimulus-position information was strongly encoded in the alpha range (Figure 3C). Interestingly, peak encoding occurred at roughly the border of the canonical alpha and beta ranges (12 Hz), with clear information encoding extending into the low-beta range (~12-20 Hz). This potentially suggests that the relevant time-delay for visual processing is slightly shorter than Alamia & VanRullen (2019) originally assumed (see Discussion). We note that this entire pattern of results also holds after first subtracting the condition-specific ERPs from the data, suggesting non-phase-locked power effects, and not simply VEP amplitude differences, are driving decoding (results not shown, see online data).

One interpretation of the results from Experiment 1 is that occipital alpha/low-beta oscillations are a spectral signature of ongoing recursive signaling between hierarchically-organised regions of the visual system, with their power carrying (spatial) information about the underlying stimulus being processed. However, because the apparent motion stimulus in this experiment was updated every 100 ms (i.e. at 10 Hz), it is possible that stimulus-related information was entrained in the alpha range by the underlying rhythmicity of stimulus-evoked activity.

To examine this possibility, we ran a second experiment in which we varied the stimulus update rate across three frequencies: 10 Hz (100 ms), 8 Hz (125 ms), and 6.67 Hz (150 ms). A new group of participants (N = 34) viewed an otherwise identical apparent motion stimulus moving along a circular path, while EEG was recorded. Figure 4A shows stimulus-position maps derived from the power of occipital alpha oscillations in Experiment 2. Different rows illustrate different stimulus update rates (10 Hz, 8 Hz, or 6.67 Hz), for the Start (left column), Middle (middle column) and End of the sequence (right column). Note that there were no motion reversals in this experiment.

The results of Experiment 2 replicated those of Experiment 1. The location of the stimulus could again be tracked from the power of occipital alpha/low-beta oscillations (10-16

Hz), across variations in stimulus update rate. Similarly, we again saw an increase in probability for the next expected position when the stimulus disappeared (Figure 4B).

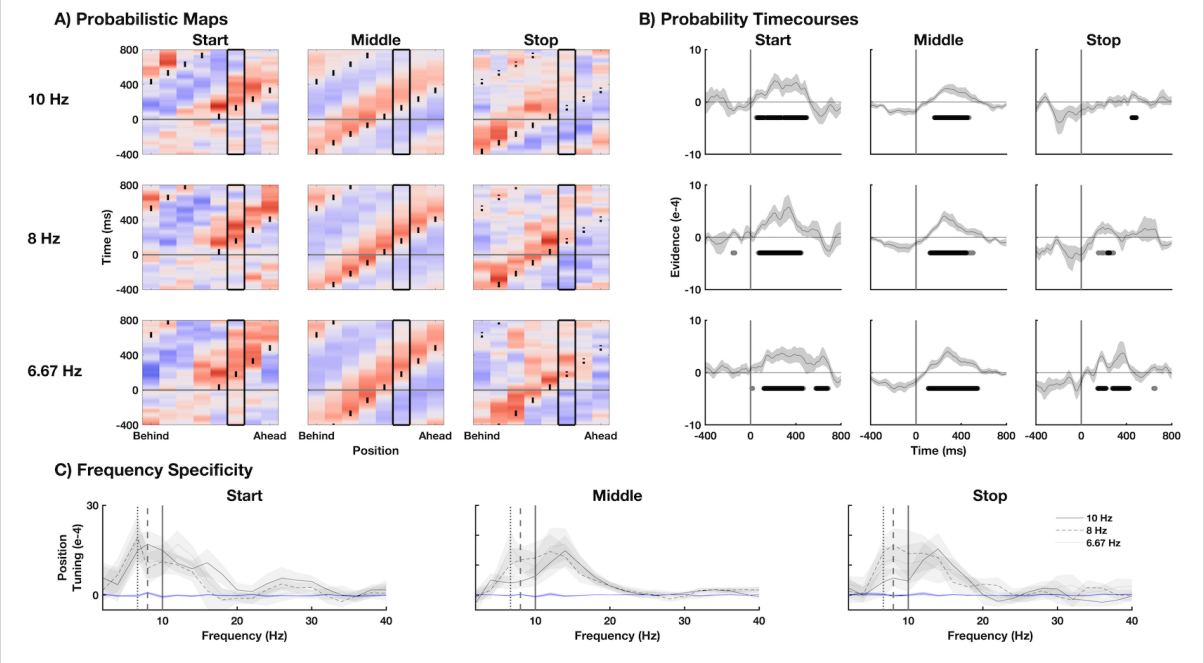

Figure 4. Experiment 2. A) Stimulus-position maps derived from the power of 10-16 Hz oscillations over the occipital cortex in Experiment 2. Plotting conventions are the same as in Figure 2, except the data has been split by stimulus update rate (rows) and trial type (columns). Each map shows data at the start (left column), middle (middle column) and end (right column) of a motion sequence. **B)** Probability time-courses for the +1 ahead position. Timepoints where the lower bound of single-sided bias-corrected and accelerated (BCa) bootstrapped 95% or 99% confidence intervals exceeded zero are marked with grey and black dots respectively. **C)** Frequency specificity of position information and average peak frequency of position tuning, calculated from cosine-convolved evidence between 50-150 ms. The vertical lines mark the stimulation frequency across the 6.67 Hz, 8 Hz, and 10 Hz conditions.

Figure 4C shows that stimulus-position information was again strongly encoded in the alpha/low-beta range, across variations in stimulus update-rate. Even after extended exposure to driven input at lower frequencies (i.e. in the Middle and Stop epochs), there is minimal effect on information encoding within the alpha/beta range. However, at lower frequencies there is slightly more (albeit inconsistent) variability across update rates. For the 6.67 and 8 Hz conditions, there is qualitatively stronger information encoding, although this does not perfectly scale with the stimulation frequency (see Discussion).

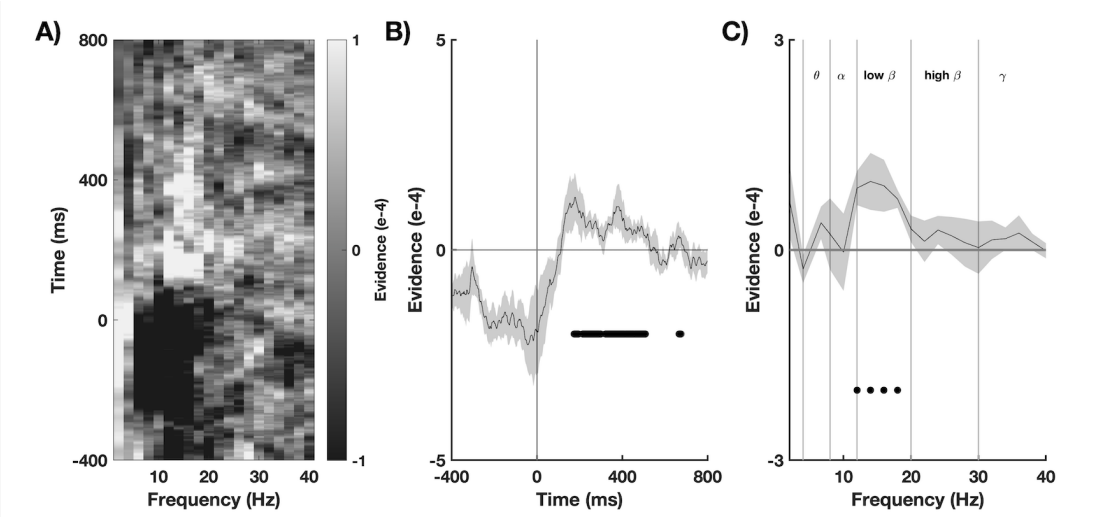

373 **Figure 5. Examining the timecourse and spectral locus of predictive position information**
374 **(collapsing across both experiments). A)** A frequency by time image of probability values
375 for the +1 ahead position, collapsing across all Stop epochs. This reveals a transition from
376 below-chance probabilities for the +1 ahead position (in black) to above-chance probabilities
377 (in white), centred on the alpha/low-beta range. **B)** The timecourse of probability values when
378 averaging within the alpha/low-beta range (10-16 Hz). **C)** Frequency specificity of predictive
379 position information (averaged between 100 – 400 ms). In panels B and C, timepoints where
380 the lower bound of single-sided bias-corrected and accelerated (BCa) bootstrapped 95% or
381 99% confidence intervals exceeded zero are marked with grey and black dots.

382 Finally, to examine the timecourse and spectral locus of predictive position information
383 in greater detail, we collapsed the data from all Stop epochs across both experiments. Focussing
384 on the +1 ahead position (i.e. the boxed region in Figure 3A), we calculated the median
385 probability timecourse across participants for all stop conditions (Experiment 1 has one
386 condition, Experiment 2 has 3 conditions for the three speeds, respectively). Averaging across
387 these yields a frequency by time image of probability values. Examining this, we can see that
388 up until the presentation of the last stimulus (at 0 ms) there is a clear suppression of positional
389 probability for the +1 ahead position, occurring in the alpha/low-beta range. This is
390 unsurprising given the fact that the stimulus is presented in other locations during this time
391 period. Around the expected time of stimulus presentation, however, we see evidence of a
392 switch to above chance decoding, even though no stimulus is actually presented. Figure 5B
393 shows the timecourse of this effect within the alpha/low-beta range (10-16 Hz). Importantly,
394 Figure 5C shows that this predictive effect is specific to the alpha/low-beta frequency range.

Discussion

Across two experiments we investigated whether stimulus-related (spatial) information is encoded in the power (squared amplitude) of neural oscillations over the occipital cortex. We also examined whether information is encoded in a *temporally* predictive fashion, as is required for predictive coding networks to effectively minimise prediction error under neural delays (Hogendoorn & Burkitt, 2019).

In Experiment 1, we demonstrated that the location of a moving stimulus could be decoded from the power of occipital oscillations in the alpha/low-beta frequency range (with peak encoding at ~12 Hz). Strikingly, we found it was possible to track the position of the moving stimulus and reconstruct its trajectory from the power of these rhythms alone. We also observed anticipatory activation of the expected but unstimulated stimulus position following the end of a motion sequence. This demonstrates that the previously-reported pre-activation revealed by analysis of raw EEG amplitudes (Blom et al., 2020) is likely encoded in alpha/low-beta band activity. In Experiment 2, by varying the update-rate of the stimulus we demonstrated that the encoding of information in this frequency range is not driven by visual entrainment.

This study contributes to an emerging line of research examining potential links between the *in silico* oscillatory dynamics of hierarchical predictive coding networks and rhythmic activity patterns in human EEG recordings (Alamia et al., 2020; Alamia & VanRullen, 2019). To our knowledge, this study is the first to demonstrate that the power of occipital oscillations in the alpha/low-beta range carries predictive stimulus-related information. This finding is broadly consistent with the theoretical predictions of Alamia and VanRullen (2019). Interestingly, we found that peak information occurred at the border of the canonical alpha/low-beta frequency ranges (12 Hz in Experiment 1). This potentially suggests that the relevant inter-regional delay for visual processing may be shorter than originally assumed (i.e., < 12 ms), leading to a higher frequency macroscopic spectral signature (Alamia & VanRullen, 2019). From the results of Experiment 2, there was some evidence that, qualitatively speaking, the strength of information encoding at lower frequencies depended on the stimulus-update rate. One interpretation of this overall pattern of results is that there is both a stimulus-independent oscillatory signature in the alpha/low-beta range, which emerges due to the inherently rhythmic dynamics of hierarchical predictive coding under signaling delays (Alamia & VanRullen, 2019; Hogendoorn & Burkitt, 2019), and an additional stimulus-dependent component, which emerges when the stimulation frequency sufficiently deviates from this range. Ultimately however, further studies in which the stimulus update rate is varied across a wider range are needed to fully address this possibility.

To briefly account for the observed effects, we think that stimulus onset may have generated waves of activity originating from retinotopically-specific locations in visual cortex. Indeed, previous theoretical and empirical work has established the existence of such waves, and has shown that they manifest as oscillations in the alpha frequency range (Alamia et al., 2020; Alamia & VanRullen, 2019; Lozano-Soldevilla & VanRullen, 2019). Crucially, given their spatially-specific nature, these waves of activity will have been registered on the electrode-level as relative power differences. By examining normalized oscillatory power across electrodes, it was therefore possible to determine the on-screen position of the stimulus.

While stimulus-position information was encoded most strongly in the alpha/low-beta range, it should not be concluded that spatial information is exclusively encoded at this (relatively slow) temporal scale. Indeed, it is likely that spatial information is also processed on a more fine-grained timescale. In the present study, we may not have been able to observe this due to the relatively poor resolution of scalp-based EEG recordings for high frequency oscillations. Overall, the fact that stimulus-position information was predominantly encoded in the alpha/low-beta range, aligns with the view that oscillations in this general frequency range may be a macroscopic signature of predictive message passing between hierarchically organised regions of the visual system (Alamia & VanRullen, 2019). While such signalling almost certainly operates over many temporal and spatial scales, neural delays potentially cause these oscillations to be the most prominent macroscopic rhythmic 'fingerprint' of network activity.

In further support of this view, the probabilistic spatial maps we constructed displayed temporally predictive activation patterns. Specifically, when the stimulus disappeared or unexpectedly reversed its direction, we observed an increase in the probability the stimulus was occupying the next position along its original trajectory of motion, at the expected time of presentation. One could argue that these anticipatory activation patterns may be due to spatial smearing or variability in decoding, however we have observed similar dynamics in our previous work using more conventional classification analysis applied to raw EEG amplitudes (Blom et al., 2020). Moreover, similar anticipatory patterns have also been directly observed in numerous animal neurophysiology studies (Benvenuti et al., 2020; Berry et al., 1999; Chemla et al., 2019; Jancke et al., 2004; Liu et al., 2021; Trenholm et al., 2013) as well as in more recent human fMRI experiments (Ekman et al., 2017, 2022). The novelty of our current work therefore lies in the demonstration that these anticipatory spatial representations manifest in alpha/low-beta oscillations, consistent with recent computational predictions (Alamia & VanRullen, 2019).

Considering potential neural mechanisms underlying these dynamics, there are two main possibilities. First, it is possible that anticipatory activation is facilitated by an omni-directional spreading of activity between retinotopically organised neural populations. This could be facilitated by within-region lateral connectivity (Benvenuti et al., 2020; Liu et al., 2021) or between-region divergent connectivity (Baldo & Caticha, 2005). A second possibility, is that more complex sequence learning mechanisms are involved. For example, it has recently been shown that after repeated exposure to visual sequences, activity in the visual cortex associated with these sequences can be predictively activated ('pre-played') by the presentation of just a single stimulus (Ekman et al., 2017, 2022). It is possible that in our experiments, similar predictive associations between neighboring stimulus position were generated, leading to anticipatory activation (although why pre-play of an ongoing sequence did not occur would need to be accounted for). To arbitrate between these possibilities future studies could examine the dynamics that arise when participants are exposed to arbitrary, non-contiguous sequences of flashes (as in Ekman et al., 2022), using the decoding approach developed in the current study.

While considering the question of temporal prediction, one point which should be made is that despite showing temporally predictive qualities (activation of likely future positions in the absence of direct input), the bulk of activity in the spatial probability maps still lagged behind the stimulus (although activity onset did align with stimulus onset). This raises the question of whether there was sufficient temporal prediction to fully compensate for neural signaling delays (Hogendoorn & Burkitt, 2019). Ultimately, given the temporal smearing inherent to time-frequency based analyses, it is beyond the scope of this paper to fully address that question. This may be better tackled using alternative methods in which more fine-grained temporal resolution can be achieved (see Blom et al., 2020; Johnson et al., 2023).

In appraising the current findings, it is important to consider two potential non-stimulus-driven sources of information that may have influenced our decoding analyses: 1) eye movements, and 2) spatial attention differences. Eye movements are an insidious potential artifact in neuroimaging experiments, that must be considered when employing classification analyses (Quax et al., 2019). However, three factors greatly limit the possibility that eye movements confounded the current analyses. Firstly, in earlier analyses of the data from Experiment 1 (Blom et al., 2020), we demonstrated in a control sample of participants that the position of the stimulus could not be decoded from eye-movement traces. Secondly, the current analyses were restricted to only occipital electrodes. Since eye-movement-related muscle activity manifests predominantly at frontal electrodes, the likelihood that we are picking up on

eye movements is further reduced. Finally, training epochs were limited to the first 150 ms after the (unpredictable) onset of a motion sequence. Since saccade onsets and corresponding eye-movement-related biases in decoding performance typically occur >200 ms after stimulus onset (Quax et al., 2019), this further reduces the likelihood of eye-movement confounds. By restricting our analyses to an early time window, and only analysing activity recorded directly over the visual cortex, it is more likely that the current analyses are tapping into the initial feed-forwards sweep of visual information processing, rather than eye-movement-related information.

It is also important to consider whether our decoding analyses were confounded by position-related differences in spatial attention. This is because it has been demonstrated that position-specific differences in covert spatial attention can be decoded from the power of alpha oscillations (Foster et al., 2017). However, the theoretical work of Alamia and VanRullen (2019) potentially prompts a subtle but significant re-interpretation of this earlier finding. While Foster et al. (2017) clearly showed that shifts in attention co-occur with changes in alpha power, this does not mean that alpha oscillations necessarily directly reflect the deployment of spatial attention. Rather, top-down shifts in attention (occurring at ~300 ms in Foster et al., 2017) likely alter neural activity patterns in the visual system. Under Alamia and VanRullen's account (2019), this would, in turn, change the pattern/amplitude of occipital alpha oscillations. In that sense, alpha oscillations would reflect the knock-on effect that spatial attention differences have on macroscopic network dynamics, rather than the deployment of spatial attention directly. Considering the present results, the fact that our training epochs were restricted to 50-150 ms after initial stimulus onset, again makes it more likely that we are tapping into the first sweep of visual information processing rather than spatial-attention differences, which one might expect to manifest over a slower timescale.

In conclusion, consistent with recent *in silico* simulations (Alamia & VanRullen, 2019) we have shown that occipital alpha/low-beta oscillations carry predictive stimulus-related information. By examining the power of these rhythms, we could reconstruct the trajectory of a moving stimulus, tracking its position even across unexpected motion reversals. Moreover, we found that future position representations were anticipatorily activated in the absence of direct visual input, indicative of temporally predictive processing. Collectively, these results support the view of alpha/low-beta oscillations as a potential spectral 'fingerprint' of hierarchical predictive processing in the human visual system.

References

Alamia, A., Timmermann, C., Nutt, D. J., VanRullen, R., & Carhart-Harris, R. L. (2020). DMT alters cortical travelling waves. *ELife*, *9*, e59784. https://doi.org/10.7554/eLife.59784

Alamia, A., & VanRullen, R. (2019). Alpha oscillations and traveling waves: Signatures of predictive coding? *PLOS Biology*, *17*(10), e3000487. https://doi.org/10.1371/journal.pbio.3000487

Baldo, M. V. C., & Caticha, N. (2005). Computational neurobiology of the flash-lag effect. *Vision Research*, *45*(20), 2620–2630. https://doi.org/10.1016/j.visres.2005.04.014

Benvenuti, G., Chemla, S., Boonman, A., Perrinet, L., Masson, G. S., & Chavane, F. (2020). Anticipatory responses along motion trajectories in awake monkey area V1. *BioRxiv*, 1–42.

Berens, P. (2009). CircStat: A MATLAB Toolbox for Circular Statistics. *Journal of Statistical Software*, *31*. https://www.jstatsoft.org/article/view/v031i10

Berry, M. J., Brivanlou, I. H., Jordan, T. A., & Meister, M. (1999). Anticipation of moving stimuli by the retina. *Nature*, *398*(March).

Blom, T., Bode, S., & Hogendoorn, H. (2021). The time-course of prediction formation and revision in human visual motion processing. *Cortex*, *138*, 191–202. https://doi.org/10.1016/j.cortex.2021.02.008

Blom, T., Feuerriegel, D., Johnson, P., Bode, S., & Hogendoorn, H. (2020). Predictions drive neural representations of visual events ahead of incoming sensory information. *Proceedings of the National Academy of Sciences*, *117*(13), 7510–7515. https://doi.org/10.1073/pnas.1917777117

Chemla, S., Reynaud, A., Volo, M. di, Zerlaut, Y., Perrinet, L., Destexhe, A., & Chavane, F. (2019). Suppressive Traveling Waves Shape Representations of Illusory Motion in Primary Visual Cortex of Awake Primate. *Journal of Neuroscience*, *39*(22), 4282–4298. https://doi.org/10.1523/JNEUROSCI.2792-18.2019

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

Ekman, M., Kok, P., & de Lange, F. P. (2017). Time-compressed preplay of anticipated events in human primary visual cortex. *Nature Communications*, *8*(1), Article 1. https://doi.org/10.1038/ncomms15276

562 Ekman, M., Kusch, S., & de Lange, F. P. (2022). *Successor-like representation guides the*
563 *prediction of future events in human visual cortex and hippocampus* [Preprint].
564 Neuroscience. https://doi.org/10.1101/2022.03.23.485480
565 Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-Band
566 Oscillations Enable Spatially and Temporally Resolved Tracking of Covert Spatial
567 Attention. *Psychological Science*, *28*(7), 929–941.
568 https://doi.org/10.1177/0956797617699167
569 Friston, K. (2008). Hierarchical Models in the Brain. *PLoS Computational Biology*, *4*(11),
570 e1000211. https://doi.org/10.1371/journal.pcbi.1000211
571 Hajonides, J. E., Nobre, A. C., van Ede, F., & Stokes, M. G. (2021). Decoding visual colour
572 from scalp electroencephalography measurements. *NeuroImage*, *237*, 118030.
573 https://doi.org/10.1016/j.neuroimage.2021.118030
574 Hogendoorn, H., & Burkitt, A. N. (2019). Predictive Coding with Neural Transmission
575 Delays: A Real-Time Temporal Alignment Hypothesis. *ENeuro*, *6*, 1–12.
576 Jancke, D., Erlhagen, W., Schöner, G., & Dinse, H. R. (2004). Shorter latencies for motion
577 trajectories than for flashes in population responses of cat primary visual cortex. *The*
578 *Journal of Physiology*, *556*(3), 971–982.
579 https://doi.org/10.1113/jphysiol.2003.058941
580 Johnson, P. A., Blom, T., & van Gaal, S. (2023). Position representations of moving objects
581 align with real-time position in the early visual response. *ELife*.
582 Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior expectations evoke stimulus
583 templates in the primary visual cortex. *Journal of Cognitive Neuroscience*, *26*(7),
584 1546–1554.
585 Kok, P., Mostert, P., & Lange, F. P. D. (2017). Prior expectations induce prestimulus sensory
586 templates. *Proceedings of the National Academy of Sciences*, *114*(39).
587 https://doi.org/10.1073/pnas.1705652114
588 Liu, B., Hong, A., Rieke, F., & Manookin, M. B. (2021). Predictive encoding of motion
589 begins in the primate retina. *Nature Neuroscience*, *24*, 1280–1290.
590 https://doi.org/10.1038/s41593-021-00899-1
591 Lozano-Soldevilla, D., & VanRullen, R. (2019). The Hidden Spatial Dimension of Alpha:
592 10-Hz Perceptual Echoes Propagate as Periodic Traveling Waves in the Human Brain.
593 *Cell Reports*, *26*(2), 374-380.e4. https://doi.org/10.1016/j.celrep.2018.12.058

Luo, H., & Poeppel, D. (2007). Phase Patterns of Neuronal Responses Reliably Discriminate Speech in Human Auditory Cortex. *Neuron*, *54*(6), 1001–1010. https://doi.org/10.1016/j.neuron.2007.06.004

Quax, S. C., Dijkstra, N., van Staveren, M. J., Bosch, S. E., & van Gerven, M. A. J. (2019). Eye movements explain decodability during perception and cued attention in MEG. *NeuroImage*, *195*, 444–453. https://doi.org/10.1016/j.neuroimage.2019.03.069

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87.

Ronconi, L., Oosterhof, N. N., Bonmassar, C., & Melcher, D. (2017). Multiple oscillatory rhythms determine the temporal organization of perception. *Proceedings of the National Academy of Sciences*, *114*(51), 13435–13440. https://doi.org/10.1073/pnas.1714522114

Trenholm, S., Schwab, D. J., Balasubramanian, V., & Awatramani, G. B. (2013). Lag normalization in an electrically coupled neural network. *Nature Neuroscience*, *16*(2). https://doi.org/10.1038/nn.3308

Author Contributions

W.T., T.B., and H.H. contributed to conception and design. T.B. programmed both experiments and oversaw data collection. W.T. analysed the data and drafted the article. All authors reviewed and revised the manuscript. H.H. funded and supervised the project.