

Research Articles | Behavioral/Cognitive

Contents of visual predictions oscillate at alpha frequencies

<https://doi.org/10.1523/JNEUROSCI.0474-25.2025>

Received: 6 March 2025

Revised: 28 August 2025

Accepted: 29 August 2025

Copyright © 2025 Hetenyi et al.

This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International license](#), which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

This Early Release article has been peer reviewed and accepted, but has not been through the composition and copyediting processes. The final version may differ slightly in style or formatting and will contain links to any extended data.

Alerts: Sign up at www.jneurosci.org/alerts to receive customized email alerts when the fully formatted version of this article is published.

Title

Contents of visual predictions oscillate at alpha frequencies

Short title: Visual predictions oscillate at alpha frequencies

Authors

Dorottya Hetenyi¹, Joost Haarsma^{1,2}, Peter Kok¹

¹Department of Imaging Neuroscience, UCL Queen Square Institute of Neurology, University College London, London, UK.

²Department of Cognitive Neuroscience, Faculty Psychology and Neuroscience, Maastricht University, Maastricht, 6229 ER

Corresponding author

Dorottya Hetenyi, dorottya.hetenyi.21@ucl.ac.uk

Pages

Manuscript: 35 pages

Figures

Manuscript: Figures: 5 + Supplemental figures: 5

Words

Abstract: 135

Introduction: 418

Discussion: 1500

References: 66

Acknowledgements

We would like to thank Gareth Barnes, Ole Jensen, Oscar Ferrante and Benjy Barnett for advice on MEG analyses, Benjy Barnett for feedback on an earlier version of the manuscript, and the Imaging Support Team at the Functional Imaging Laboratory for their help with data collection. This work was supported by a Wellcome/Royal Society Sir Henry Dale Fellowship [218535/Z/19/Z] and a European Research Council (ERC) Starting Grant [948548] to P.K. The Wellcome Centre for Human Neuroimaging was supported by core funding from the Wellcome Trust [203147/Z/16/Z].

Data and code availability

All original MEG data and all original code will be deposited at OSF (<https://osf.io>) and will be publicly available upon publication.

Conflict of interests

The authors declare no conflicts of interests.

Abstract

Predictions of future events have a major impact on how we process sensory signals. However, it remains unclear how the brain keeps predictions online in anticipation of future inputs. Here, we combined magnetoencephalography (MEG) and multivariate decoding techniques to investigate the content of perceptual predictions and their frequency characteristics. Thirty-two participants (23 female) were engaged in a shape discrimination task, while auditory cues predicted which specific shape would likely appear. Frequency analysis revealed significant oscillatory fluctuations of predicted shape representations in the pre-stimulus window in the alpha band (10 – 11Hz). Furthermore, we found that this stimulus-specific alpha power was linked to expectation effects on shape discrimination behaviour. Our findings demonstrate that sensory predictions are embedded in pre-stimulus alpha oscillations and modulate subsequent perceptual performance, providing a neural mechanism through which the brain deploys perceptual predictions.

Significance statement

Prior knowledge greatly influences how we perceive the world. However, it is unclear how the brain maintains and keeps prediction signals in the anticipation of future inputs. Our study reveals a neural mechanism by which the brain maintains sensory predictions. We demonstrated that perceptual predictions are encoded in pre-stimulus alpha oscillations (10–11 Hz). These oscillations not only reflect the content of predicted visual shapes but are also directly linked to enhanced perceptual performance. These findings provide key insights into how the brain prepares for and improves perception of future sensory inputs, contributing to our understanding of predictive processing in human cognition.

Keywords

– expectation; time-resolved multivariate analysis; perceptual inference; pre-stimulus oscillations; stimulus templates; predictive processing

Introduction

Predictions about how the world is structured play an integral role in perception (Friston, 2009; Bastos et al., 2012; Clark, 2013; de Lange et al., 2018). Our prior knowledge forms the basis for predictions of future sensory events, which are subsequently integrated with sensory input to form a perceptual experience. While there is a wealth of evidence supporting the idea that the brain deploys predictions to guide perception, the mechanisms through which the brain keeps these predictions online remain largely unknown. One likely candidate for conveying perceptual predictions are neural oscillations (Arnal and Giraud, 2012; Mayer et al., 2016; Auksztulewicz et al., 2017; Bastos et al., 2020).

Alpha rhythms (8 – 12Hz) are the predominant oscillations in the awake human brain (Berger, 1929), yet their functional role is controversial (Klimesch et al., 2007; Jensen and Mazaheri, 2010). The amplitude and phase of these ongoing oscillations is known to influence performance in visual tasks (Ergenoglu et al., 2004; van Dijk et al., 2008; Busch et al., 2009; Mathewson et al., 2009; Romei et al., 2010; Dugué et al., 2011; VanRullen et al., 2011; Hanslmayr et al., 2013; Iemi et al., 2019), and have been found to vary with experimental manipulations that target stimulus predictability (Sherman et al., 2016; Alilović et al., 2019). Specifically, pre-stimulus alpha oscillations have a similar topography to post-stimulus responses, implying a shared neural substrate in the processing of pre-existing information and external stimuli (Shen et al., 2023), alpha phase determines the influence of predictions on perception (Sherman et al., 2016), and alpha power has been shown to predictively encode the position of a moving stimulus (Turner et al., 2023). Furthermore, alpha-band activity modulates sensory processing of both low-level features and high-level object representations through selective attention, suggesting a hierarchical and predictive role in top-down control. (Noah et al., 2020; Meyyappan et al., 2024). However, whether these oscillations actually convey the contents of perceptual predictions remains unknown.

To test this hypothesis, we combined magnetoencephalography (MEG) with multivariate decoding analyses to resolve visual representations with millisecond resolution, and

characterise the temporal and frequency characteristics (Kerrén et al., 2018, 2022) of sensory predictions (Stecher and Kaiser, 2024; Stecher et al., 2025). Participants were engaged in a shape discrimination task where auditory cues predicted the identity of upcoming abstract shapes. We identified the neural representations of cued sensory predictions prior to stimulus onset and tested whether these sensory predictions had an oscillatory nature, as well as whether the power of such predictive oscillations modulated perceptual performance (Tarasi et al., 2022).

Materials and Methods

Participants

Sixty-two healthy right-handed participants (43 female) with normal or corrected-to-normal vision and no history of neurological disorders took part in the behavioural experiment. This experiment served as a pre-assessment process to familiarise the participants with the task and select only those whose average performance accuracy on the challenging shape discrimination task was above 70% across the four runs. Thirty-nine participants (28 female) met the performance inclusion criteria and participated in the MEG experiment. Seven participants were excluded from subsequent analyses due to excessive head movement (N=5) or not completing the full experiment (N=2), leaving thirty-two participants (23 female, age 26 ± 5 years, mean $\pm$ SD) for the MEG analysis.

Stimuli

The experiment employed the same design as (Kok and Turk-Browne, 2018), wherein participants discriminated between two consecutively presented shapes which were preceded by a predictive auditory cue. Each predictive cue was composed of three pure tones (440, 554, and 659Hz; 80ms per tones; 5ms intervals), played with rising or falling pitch, with a total duration of 250ms. Visual stimuli were generated using MATLAB (The MathWorks Inc., version 2021b) and Psychophysics Toolbox (Brainard, 1997). The visual stimuli consisted of complex abstract shapes defined by radial frequency components (RFCs), a method where

shape contours are created by modulating the radius of a circle as a function of angular position (Zahn and Roskies, 1972). This technique allows for precise control of shape complexity and variation by adjusting the frequency and amplitude of sinusoidal modulations around the shape's circumference. To generate the contours of the stimuli, seven RFCs (0.55Hz, 1.11Hz, 4.94Hz, 3.39Hz, 1.54Hz, 3.18Hz, 0.57Hz) were used which were based on a subset of stimuli from Op de Beeck et al.'s work (Op de Beeck et al., 2001); see their Fig. 1A). To construct a two-dimensional shape space, we systematically varied the amplitudes of two RFCs (specifically the 1.11Hz and 1.54Hz components) while keeping the amplitudes of the other components constant (indicated by the arrows on Fig.1D). These variations created the four distinct abstract shapes such as: shape A (baseline amplitudes, amplitude of $1 \times 1.11\text{Hz} = 8$; $1 \times 1.54\text{Hz} = 8$); shape B (amplitude of $1 \times 1.11\text{Hz} = 8$; $4 \times 1.54\text{Hz} = 26$); shape C (amplitude of $4 \times 1.11\text{Hz} = 26$; $1 \times 1.54\text{Hz} = 8$) and shape D (amplitude of $4 \times 1.11\text{Hz} = 26$; $4 \times 1.54\text{Hz} = 26$). These shapes were designed such that discrimination between shapes A and D (both amplitudes changed) was orthogonal to the discrimination between shapes B and C (only one amplitude varied), defining a clear two-dimensional shape space (Fig. 1D). As explained above, the shapes were designed in a very precise and controlled way, accounting for discrimination accuracy perceptually and neuronally. This allowed us to study signals reflecting complex visual shapes, rather than conventional gratings, yet in an experimentally controlled manner. Additionally, RFC-based warping was used to generate moderately distorted versions of the two main experiment shapes (shape A and D, Fig. 1D) for the benefit of the shape discrimination task. This warp to define the shape was achieved by modulating a different RFC's amplitude (3.18Hz) than the two used (1.11Hz and 1.54Hz) to define the shape space. This modulation could be either positive or negative (counterbalanced over conditions) and was orthogonal to the shape space used for the two main experiment shapes, and therefore to the cue predictions as well. The visual stimuli were displayed on a rear-projection screen using a projector (1024 x 768 resolution, 60 Hz refresh rate) against a uniform grey background.

Behavioural experiment

The study had two parts, a behavioural training and screening experiment, and an MEG experiment for those who passed the behavioural screening. In both parts, participants were engaged in a shape discrimination task. Each trial started with a fixation bullseye (diameter, 0.7°) for 100ms, followed by the presentation of two consecutive shape stimuli each for 250ms, and separated by a 500ms blank screen containing only a fixation bullseye (Fig. 1A). On each trial, the second shape was the same as the first or slightly warped. The modulation was either positive or negative, and the size of the modulation was determined by an adaptive staircasing procedure (Watson and Pelli, 1983), updated after each trial, in order to make the task challenging. Participants were instructed to report whether the two presented shapes were identical or different. After the response interval ended (750ms after disappearance of the second shape), the fixation bullseye was replaced by a single dot, signalling the end of the trial while still prompting participants to fixate. On each trial, one of the four shapes (A, B, C or D; Fig. 1D) was presented, in a counterbalanced (i.e., non-predictable) manner. Participants performed four runs (360 trials in total) of the shape discrimination task, maximum one week prior to the MEG session.

MEG experiment

The MEG experiment started with two localiser runs, containing the same four abstract shapes as in the behavioural task. To ensure participants were engaged, they performed a fixation dimming task (10% of total trials, ~24 of 248 trials per run). Each trial began with a fixation bullseye (visual angle: 0.7°) displayed for 100ms, followed by one of the four shapes presentation for 250ms. Following the stimulus presentation, the fixation bullseye reappeared and remained on the screen for a period between 1000 and 1200ms. In 10% of the trials, fixation bullseye dimmed for 150ms and participants had been instructed to press a button when this occurred. By using identical stimulus durations, these runs were designed to be as similar as possible in terms of stimulus presentation to the main experiment. During the localisers, participants correctly detected $95.3 \pm 0.7\%$ (mean $\pm$ SEM) of fixation dimming

events and incorrectly pressed the button on $4.9 \pm 2.2\%$ of trials, suggesting that participants were successfully engaged by the fixation task.

Following the localiser runs, participants performed 8 main task runs (2x training runs, 6x prediction runs), 64 trials per run, in total 512 trials. During the prediction runs, an auditory cue (falling vs. rising tones, 250ms) was presented 100ms after trial onset. Following a 1500ms interval, two consecutive shape stimuli were displayed (each for 250ms) and, separated by a 500ms blank screen (Fig. 1A). As in the behavioural session, participants' task was to indicate whether the two shapes were the same or different. The auditory cue predicted whether the first shape presented on that trial would be shape A or D. The cue was valid on 75% of trials, whereas on the other 25% of trials the unpredicted shape would be presented (Fig. 1B). For instance, if the cue was a falling auditory tone, it might lead to shape A in 75% of cases and shape D in the other 25% of cases. The prediction induced by the auditory cue (predicting the identity of the first shape) was thus orthogonal to participants' task (whether the two shapes were the same or different). This experimental design choice was made to prevent the cues from inducing response biases in the task (c.f. (Kok et al., 2012a, 2017; Kok and Turk-Browne, 2018). Despite this orthogonality, previous work using a similar design has revealed a significant benefit in behavioural performance of valid prediction cues, which was correlated with neural effects of the predictive cues (e.g., (Kok et al., 2012a, 2017; Kok and Turk-Browne, 2018). Additionally, ensuring the predictive cues were orthogonal to the task allowed us to manipulate expectation while keeping attention constant, since both validly and invalidly cues shapes were equally task-relevant.

Note that shapes B and C were never presented in the prediction runs. The contingencies between cues and shapes were flipped halfway through the experiment, and the order was counterbalanced over subjects. Prior to the first prediction run, and after the cue reversal halfway through, participants were trained on the cue–shape associations during training runs in the MEG and explicitly informed about the cue contingencies. In the training runs, the auditory cue was 100% predictive of the identity of the first shape.

Pre-processing

Whole-head neural recordings were obtained using a 273-channel MEG system with axial gradiometers (CTF Systems) at a rate of 600Hz located in a magnetically shielded room. Throughout the experiment, head position was monitored online and corrected if necessary using three fiducial coils that were placed on the nasion, right and left preauricular. If participants moved their head more than 5mm from the starting position, they were repositioned after each run. Eye movements were recorded using an EyeLink 1000 infrared tracker (SR Research Ltd.). The recorded eye-tracker data were used to identify eye-blink related artefacts in the MEG signal. Auditory tones were delivered using earplugs (Etymotic Research Inc.). A photodiode was placed at the bottom left corner of the screen to measure visual stimulus presentation latencies. The photodiode signal was used to realign the MEG signal with stimulus onset.

The data were pre-processed offline using FieldTrip (Oostenveld et al., 2011). The variance (collapsed over channels and time) was calculated for each trial in order to identify artefacts. Trials with large variances were subsequently selected for manual inspection and removed if they contained excessive and irregular artefacts. Number of trials removed per participants: valid condition 5.740 ± 4.495 ; invalid 5.761 ± 5.473 ; presented shape A 5.664 ± 4.4288 and presented shape D 5.826 ± 5.015 (mean $\pm$ SD). Next, independent component analysis was used to further remove cardiac and eye movements related artefacts. The independent components were correlated to the eye tracking signal to identify potentially contaminating components for each participant, and inspected manually before removal. Notch filters were also applied at 50, 100, and 150Hz to remove line noise and its harmonics. No detrending was applied for any analysis. Finally, main task data were baseline corrected on the interval of -200 to 0ms relative to auditory cue onset, and localiser data were baseline corrected on the interval of -200 to 0ms relative to shape onset.

Decoding analysis

To uncover the representational content of neural activity, we performed a decoding analysis using a customised linear discriminant analysis (LDA) decoder. LDA aims to find a linear transformation of the data, where the resulting signal is optimally discriminative between two labels. The data fed into the decoder was separated based on the number of MEG channels ($\mathbf{F}$) at each time point, therefore $\hat{\boldsymbol{\mu}}_1$ and $\hat{\boldsymbol{\mu}}_2$ were column vectors of length $\mathbf{F}$ that contain the neural responses in the training set for label 1 and 2. The weights vector ($\mathbf{w}$) that optimally discriminates between labels on the basis of the channels was calculated as: $\mathbf{w} = \tilde{\boldsymbol{\Sigma}}_c^{-1}(\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)$. $\tilde{\boldsymbol{\Sigma}}_c$ is the common regularised covariance matrix. The data to be decoded ($\mathbf{X}$) was set as a matrix of size channels X trials ($\mathbf{N}$), and the decoded data ($\mathbf{y}$) was then obtained, where $\mathbf{T}$ denotes the matrix transpose as: $\mathbf{y} = \mathbf{w}^T \mathbf{X}$. Rather than assigning discrete labels, the decoder outputs a continuous measure of label encoding strength in the signal for each single trial. Therefore, no binary cut-off was applied to the decoded data. Retaining single trial data allowed us to perform across-trial analysis (such as coherence and logistic regression against behavioural reports, see below). A normalisation factor was applied to allow comparing the data across time points (i.e., the mean difference in the decoded data between labels was equal to a value of one). The equation describing the decoder, including the normalisation factor, as described in (Mostert et al., 2015), was as follows:

$$\mathbf{w} = \frac{\hat{\boldsymbol{\Sigma}}_c^{-1} (\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)}{(\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)^T \hat{\boldsymbol{\Sigma}}_c^{-1} (\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)}$$

This decoder assessed how sensor-level activity varied based on a discriminability index, providing a continuous measure of shape decoding on a trial by trial basis. The decoding analysis was time-resolved. First, the data were down-sampled (from the original sampling rate of 600Hz to 200Hz) using a 28ms sliding time window centred on each time point, with the window advancing in 5ms steps. Then, the decoder was applied to the down-sampled data, resulting in a decoder output matrix of trials X training time X testing time.

To test how effective the decoder was at revealing neural patterns, it was first trained and tested on shape A and D trials (between -100 and 600ms relative to stimulus onset) from the localiser runs, using a leave-one-block-out approach. Analogously, a shape B vs. C decoder was tested on shape B and C trials (Fig. S1D). To further validate the analysis, we tested the shape A vs. D decoder on the shape B and C trials, and the shape B vs. C decoder on shape A and D trials (Fig. S1BC). We expected significant decoding within shape categories (e.g. training and testing on shape A vs. D), but not across shape categories (i.e., training on shape A and D and testing on shape B and C, and vice versa).

Localiser decoding results were analysed using non-parametric cluster-based permutation tests. The data were represented as 2D matrices of decoding performance, with training time on one axis and testing time on the other. The statistical analysis focused on identifying significant 2D clusters in these matrices. To do so, univariate t-statistics were calculated for the entire matrix. Elements that were considered neighbours, i.e., directly adjacent in cardinal or diagonal directions, were collected into separate positive and negative clusters if they passed a threshold corresponding to a p-value of 0.001 (two-tailed). The significance of the clusters was assessed by summing the t-values within each cluster to obtain cluster-level test statistics. These test statistics were then compared to a null distribution, which was created by randomly shuffling the observed data 10,000 times. A cluster was considered significant if its resulting p-value was less than 0.05 (two-tailed).

In order to reveal predicted shape representations, the decoder was trained on shape A vs. D localiser trials (70 – 200ms) again in a time-resolved manner with the identical parameters as above, and subsequently tested on the pre-stimulus window (-1750 to 0ms relative to shape onset) during the prediction runs (Fig. S2). To address label imbalances resulting from trial rejections during pre-processing, random resampling was applied to the training sets, ensuring an equal number of each decoded classes (shapes) for every participant. Furthermore, we repeated the same procedure for each participant using a control decoder trained on shapes B vs. C localiser trials, i.e. shapes which were not presented during the prediction runs. This

results of applying this control decoder to the pre-stimulus prediction window served as a baseline (Baseline 2, Fig. 2C right) in further analyses. It is important to highlight that the shape B vs. C discrimination was orthogonal to shape A vs. D discrimination.

Frequency analysis of pre-stimulus decoding time series

Our primary aim was to test whether the decoded neural representations of predicted shapes had oscillatory dynamics. Therefore, we adapted the analysis approach of (Kerrén et al., 2018), investigating the frequency characteristics of decoder time series using FFT. This analysis examined whether the spatial shape-specific pattern identified by training an LDA decoder on shape-specific evoked signals (in the localiser runs) showed evidence of rhythmic fluctuations during the cue-stimulus prediction interval. The LDA shape decoders were trained on the trial-by-trial evoked responses during the shape localiser in a time-resolved manner (sliding window of 28ms, steps of 5ms). These decoders were applied to the cue-stimulus prediction intervals (-1250ms to 0ms relative to the onset of the first shape) in order to reveal shape-specific sensory prediction signals. Based on the results of cross-validated within localiser shape decoding (Fig. S1A), we averaged over decoding training timepoints in the 70-200ms post-stimulus window for subsequent analyses. The decoder output (as mentioned above) contained a continuous measure of shape discrimination, rather than binarized class information. Using decoding accuracy alone would not allow us to explore the oscillatory nature of the neural signal, as continuous decoder output reflects a non-graded shape information. Additionally, trial-by-trial phase misalignment may impose a challenge, as we did not expect the predictive signal to occur at the exact same latency across participants, averaging across trials to compute accuracy could harm meaningful temporal dynamics. Keeping the continuous output better preserves this information for subsequent frequency analysis.

For each participant, each trial of the pre-stimulus decoded time series was tapered with a Hann window covering the whole time period (-1250 to 0ms), and then subjected to an FFT.

In a control analysis, we used Fitting Oscillations and One-Over-F (FOOOF, as implemented in the Fieldtrip toolbox (Donoghue et al., 2020), which separates rhythmic activity from concurrent power-spectral 1/f modulations in electrophysiological data, to validate the oscillatory nature of the predictive representations (Fig. S3A). We also repeated the frequency analysis in a time-resolved manner (between -2000ms and 500ms, with a fixed-length 500ms Hann window running at every 50ms) to be able to visualise the extent of the alpha oscillations (Fig. S3C).

To assess the reliability of our results, we created an empirical baseline using decoders with randomly shuffled shape labels (Baseline 1, Fig. 2C left). The labels of the two shapes (shape A and D) were shuffled pseudo-randomly before training the decoder, 25 times per participant. Therefore, each participant yielded 25 permuted datasets. The analysis parameters for the baseline decoding were identical to the non-shuffled decoder, i.e. identical spectral analysis was performed for each of the 25 datasets per participant. We generated an empirical null distribution by bootstrapping the permuted datasets ($n = 10,000$) (Stelzer et al., 2013), and compared this to the frequency analysis results of the non-shuffled shape A vs. D decoder data (Kerrén et al., 2018). Frequency bins with higher power than the empirical null distribution (exceeding the 95th percentile) were considered significant. To further validate the findings, we also conducted the identical frequency analysis (same analytical parameters) using decoders trained to distinguish shape B vs. C rather than A vs. D as an additional baseline (Baseline 2, Fig. 2C right). Importantly, shapes B and C were highly similar to shapes A and D but lay on an orthogonal discrimination axis (Fig. 1D), yielding a very precise and stringent control analysis. To correct for multiple comparisons, we extracted the p-values for each frequency bin (1–99 Hz) relative to the Baseline 1 data and applied the False Discovery Rate (FDR) correction. The p-values were derived using a one-sided test vs. Baseline 1, to test specifically whether the power in the shape A vs. D decoder was higher than in the baseline. To test whether the alpha fluctuations observed in the decoder traces were indeed driven by alpha band specific MEG signals, we applied band-pass filtering using a one-pass Butterworth

filter with a filter order of four between 8 and 12Hz on the main task (prediction runs) sensor-level MEG data for all MEG channels. Then, the bandpass-filtered MEG time courses were subjected to the identical decoding and frequency analysis as described above (Fig. S3B).

In an exploratory analysis, we repeated the analysis for two shorter training time windows (90 – 120ms and 160 – 190ms), centred around the first (105ms) and second peak (175ms) of localiser shape decoding (Fig. 2A) to distinguish effects of earlier and later visual representations. These time windows were chosen since they appeared to form distinct clusters in the localiser decoding temporal generalisation matrix, with reduced cross-decoding between the two clusters suggesting qualitatively different representations (Fig. S1A).

Relating behavioural and neural effects of predictions

To investigate whether there was a relationship between shape-specific pre-stimulus alpha power and shape discrimination performance, we performed a logistic regression analysis separately for valid and invalid prediction trials. Based on the existing literature relating pre-stimulus oscillatory power and phase to behavioural performance (Hanslmayr et al., 2007; Mayer et al., 2016; Samaha et al., 2018), we limited the pre-stimulus decoding time series to -500ms to 0ms relative to stimulus onset. To be able to accurately estimate pre-stimulus alpha power, yet be as close as possible to stimulus onset, we used a fixed-length 500ms Hann window over the -500ms to 0ms time window, resulting in ~2Hz frequency resolution (alpha frequency bins: 7.812Hz, 9.375Hz, 12.500Hz). Separately for valid and invalid prediction trials, trial-based power estimates of the pre-stimulus (-500ms to 0ms) alpha activity were z-scored and averaged over for the three alpha frequency bins. Identical frequency analyses were applied to the raw MEG signal (-500ms to 0ms, averaged over occipital channels) to quantify trial-based raw alpha power. We balanced the trial numbers by randomly choosing (n = 1000 times) a subset of trials from the conditions with higher trial counts (i.e., valid). The dependent variable of the model was the behavioural outcome (correct or incorrect response), sorted separately again for valid and invalid predictions. Independent variables were shape-specific alpha power and raw alpha power. Raw alpha power was added to the model to control for

non-shape specific alpha effects on behaviour. The model parameter estimates (i.e., beta values) served as an indication of an underlying link between alpha power and behavioural performance. The valid and invalid condition beta values were statistically compared using a paired t-test.

Coherence between shape-specific pre-stimulus fluctuations and raw alpha oscillations

To further assess the relationship between shape-specific alpha fluctuations (i.e., fluctuations in the decoder time courses) and neural alpha oscillations, we estimated the coherence between the two signals on a trial-by-trial basis at the sensor-level and in source-space. The coherence value is a number between 0 and 1, which reflect the consistency of the phase difference between the two signals at a given frequency. Cross-spectral density between z-scored MEG signals of all MEG-channel combinations, and between the MEG-channels and the z-scored decoder traces were calculated across trials at our frequency of interest (10Hz with 2Hz smoothing) between -1250 to 0ms. In order to localise the neuronal sources which were coherent with the decoder traces, we applied the DICS (Dynamic Imaging of Coherent Sources (Gross et al., 2001) beamformer, which is specifically designed to localise sources coherent with another time-series signal. We applied a regularisation parameter of $\lambda = 5\%$ to decrease the effect of noise on the source estimates. To additionally visualise the source of raw alpha oscillations, we repeated the beamformer analysis on the frequency decomposed MEG data at 10Hz with 2Hz smoothing including all MEG channels. The spatial filter was calculated (i.e., 'common filter') on the MEG data containing both pre (-1250 to 0ms) and post-stimulus (0 to 1250ms), then applied separately both time periods. The main advantage of using a common filter is that more data is used to construct the spatial filters (resulting in more reliable filters), and any difference in source activity can be attributed to power differences between the conditions, rather than to discrepancies between the filters. Based on previous studies demonstrating that the anatomical specificity gain of using subject-specific anatomical images is negligible (Holliday et al., 2003), we did not collect individual

anatomical MRI scans for our subjects. We followed a group-based template approach using a template MRI (in MNI space) in combination with a single shell head model and a standard volumetric grid (8mm resolution), as present in the Fieldtrip toolbox. Participants' individual fiducials were used to generate a participant-specific forward model in MNI space by wrapping the template head and source model to the participants' fiducials. Then, these participant-specific head- and source-models were included in the beamforming analysis.

Source localisation of shape decoding

To visualise the underlying neural sources during decoding, we applied source localisation analyses using an LCMV beamformer (Van Veen et al., 1997). The spatial distribution of the underlying signal during classification in LDA is primarily influenced by the magnetic field difference between the two experimental conditions. Therefore, one can visualise the source of a decoder by estimating the sources of the two different conditions, and compute the difference (Haufe et al., 2014). The spatial filter was computed for the time window of interest (70 – 200ms, i.e. decoder training window) in the averaged data, which was subsequently applied separately to the two conditions of interest (shape A and D trials). For shape A vs. D decoding, a percentage absolute signal change was computed in source space, to determine which source signals were involved in discriminating between shape A and D without making assumptions about the sign of the dipole (Fig. S1E).

Results

Predictive cues lead to improved shape discrimination accuracy

We tested whether the predictive cues affected behavioural performance. As a reminder, participants were required to indicate whether two shapes presented in succession were the same or different. It should be noted that any effects of the predictive cues on performance are not trivial, given that the shape discrimination task was orthogonal to the prediction manipulation (i.e., the cues predicted the identity of the first shape, but did not inform participants whether the two shapes would be identical or different). Still, valid predictive cues might improve performance indirectly by enhancing processing of the initial shape, facilitating discrimination of the subsequent shape (Kok et al., 2012a, 2017). Vice versa, invalid cues might perturb performance by impeding the processing of the initial shape. In line with this, shape discrimination accuracy was significantly influenced by whether the auditory cue correctly predicted the identity of the first shape (accuracy valid = $70\% \pm 1.2\%$ and accuracy invalid = $67\% \pm 1.3\%$, mean $\pm$ SEM; $t(31) = 3.215$, $p = 0.003$; Fig. 4A). There was no difference in reaction times (valid = $614\text{ms} \pm 1.3\%$ and invalid = $615\text{ms} \pm 1.3\%$, mean $\pm$ SEM; $p = 0.626$, $t(31) = -0.492$). Together, this suggests that valid predictions facilitated shape processing, leading to improved discrimination performance.

Predictions oscillate at alpha frequencies

To test whether perceptual predictions are conveyed by oscillations, thirty-two participants performed a challenging visual shape discrimination task, reporting whether two consecutive shapes were the same or different (Fig. 1A). On each trial, an auditory cue predicted the identity of the first shape (shape A or D, 75% valid; Fig. 1B). The shape discrimination task was orthogonal to the prediction manipulation, i.e., the auditory cue did not convey any information about whether the two shapes would be identical or different.

To reveal the representational content of neural activity, a decoding analysis was applied. We used an LDA decoder, which described how activity at the sensor-level varied as a function of

a discriminability index. Unlike conventional LDA which separates data into discrete categories, our customised decoder calculated the distances of each test sample to the hyperplane, treating these distances as discriminant evidence (Mostert et al., 2015; Kok et al., 2017). Thereby, we obtained a continuous measure of which shape was encoded in the neural signals, providing finer resolution in analysing the neural representations. For details on the implementation of the LDA, see methods (Mostert et al., 2015). The decoding analysis was performed in a time-resolved manner by applying it sequentially at every single time point on data prior down-sampled (with a sliding time window of 28ms in steps of 5ms). First, we identified shape-specific neural signals by applying the LDA to MEG responses evoked by task-irrelevant shapes (70 to 200ms post-stimulus) during separate shape localiser runs (Fig. 1C; Fig. S1A). Localiser runs consisted of the presentation of four abstract shapes (Fig. 1D), which were designed to lie on two orthogonal axes of perceptual and neural discriminability, allowing us to train a decoder that was sensitive to predicted shape information (shape A vs. D) as well as a baseline decoder that was sensitive to highly similar but unpredicted shapes (shape B vs. C; see Methods for further details). First, we tested whether the LDA was able to uncover neural representations of the presented shapes, we trained and tested a shape A vs. D decoder within the localiser runs in a cross-validated manner (-100 to 600ms, relative to stimulus onset, Fig. 2A). These decoders were then used to test whether the predictive auditory cues induced oscillatory representations of the predicted shapes (Fig. 1E). Specifically, we applied the shape decoders (i.e., spatial filters) to the MEG data recorded in the interval between the predictive cues and stimulus onset (-1250 to 0ms relative to the onset of the first shape) in a time-resolved manner.

{Figure 1}

That is, the time-resolved data (sliding window of 28ms, steps of 5ms, Fig. 2B, Fig. S2) from the prediction interval were put through a static shape-specific spatial filter obtained on the basis of the shape localiser data, yielding a shape-specific signal for each trial every 5ms (Fig. 2B). Averaging these shape-specific MEG signals over trials revealed prediction-like

representations around the auditory cue offset (at ~200ms training and ~1550ms testing time) and closer to the onset of visual stimuli (at ~500ms testing time) (Fig. S2). However, these clusters did not survive correction for multiple comparisons, unlike in a previous study from our lab (Kok et al., 2017). There are several possible explanations for this discrepancy. First, in the current study, the auditory cues predicted complex shapes rather than low-level features (grating orientation), and fMRI work has shown that predicted-but-omitted gratings can be decoded from the early visual cortex, whereas prediction-but-omitted shapes cannot (Aitken et al., 2020). Second, here we used a longer cue-stimulus onset asynchrony (1750ms here, 750ms in Kok et al., 2017). A longer prediction interval can 1) lead to more jitter in the onset of pre-stimulus prediction signal due to increased temporal uncertainty or 2) incentivise the brain to use oscillations rather than sustained above baseline activity to keep predictions online, for metabolic efficiency.

To test this, the shape-specific time courses were subjected to a Fast Fourier Transform (FFT) in order to reveal their frequency spectra (Fig. 2B). To establish the specificity of the neural signals induced by predictions, we created two separate baseline measurements. First, we shuffled the shape labels before training the decoder (N=25 permutations per participant) to create a bootstrapped baseline (Baseline 1, Fig. 2C left). Second, we trained a decoder to distinguish two shapes which were only presented in the localiser runs, but not in the prediction runs (shapes B and C). This discrimination was designed to be orthogonal to shape A vs. D discrimination, which was confirmed by an absence of cross-decoding between the two decoders (Fig. S1BC). The shape B vs. C decoder thus provides a highly stringent baseline (Baseline 2, Fig. 2C right), since it was trained to pick up neural representations of highly similar but orthogonal shapes to those that were predicted by the auditory cues. This analysis revealed that the decoded predictions oscillated at low frequencies, predominantly in the alpha frequency band (9 – 11Hz) (Fig. 3A top). We identified significant power differences between the shape A vs. D decoding data and Baseline 1, specifically at 10Hz and 11Hz, exceeding the 95th percentile of the empirical null distribution (both $p = 0.003$, corrected for multiple

comparisons with a False Discovery Rate of 0.01). It is important to note that Baseline 1 was based on the exact same pre-stimulus data; the only difference lies in the shuffling of the shape labels for the training of the decoder. We observed increased alpha power in Baseline 1 at 10 and 11 Hz, which aligns with the generally high alpha power observed in the pre-stimulus raw MEG signal (Fig. 3B bottom, 3C bottom). The decoder used in this analysis functions as a spatial filter, applying a linear combination of weights to all MEG sensors.

{Figure 2}

As a result, it is understandable that increased alpha power also appears in the Baseline 1 measurement. Importantly, alpha power was higher for the true shape-specific signal than for Baseline 1. This suggests that the topography of pre-stimulus alpha oscillations was more similar to the topography of the predicted shape-specific signal than expected by chance. There was also a negative difference in the power of very low frequencies (2 – 7Hz) when comparing the shape A vs. D decoder data to Baseline 1 data. The nature of this power difference currently is not fully understood and requires further investigation. However, these patterns have been observed in previous research using this methodology (Kerrén et al., 2018, 2022).

For further validation, we also compared the shape A vs. D decoding power spectrum to the spectrum of shape B vs. C decoding (Baseline 2). Based on our initial findings of significant differences in the 10Hz and 11Hz frequency bins between Shape A vs. D decoder and Baseline 1, here we averaged over these two frequency bins (10 and 11Hz) of the two spectra. This analysis revealed significantly higher alpha power for shape A vs. D decoding than for shape B vs. C decoding in the pre-stimulus window (paired one-sided t-test, $p = 0.0304$, $t(31) = 1.9452$) (Fig. 3A bottom). As before, it is important to note that the two spectra were based on the exact same pre-stimulus MEG data, the only difference lies in which shapes the decoders were trained to discriminate. If the pre-stimulus alpha fluctuations reflected more generic shape representations, this comparison would yield no significant differences. Therefore, the difference between these two decoding spectra demonstrates that these

signals were highly specific to the shapes predicted by the auditory cues. It is important to note that the shape decoders were trained on sensory signals evoked by task-irrelevant shapes during the localiser (i.e., participants performed a fixation dimming task), ruling out contributions of explicit decision-making signals. These alpha power effects were also present in a control analysis designed to remove non-rhythmic signals, confirming the oscillatory nature of the decoded predictions (Fig. S3A). In sum, both analyses revealed that visual predictions induced by auditory cues led to neural representations of the predicted shapes fluctuating at an alpha rhythm prior to stimulus onset.

In order to establish a link between these shape-specific prediction signals and neural alpha oscillations, we estimated the coherence (cross-spectral density) between the two signals at the sensor level and in source space at 10Hz (with 2Hz spectral smoothing) between -1250 to 0ms.

{Figure 3}

Coherence values indicate how consistent the phase synchrony is between two time-series signals. Sensor-level coherence analysis revealed high oscillatory synchrony between the shape-specific prediction signals and raw alpha oscillations, mainly over the occipital lobe (Fig. 3B top), where alpha power was also the highest (Fig. 3B bottom). Coherence analysis in source space using Dynamic Imaging of Coherent Sources (DICS) (Gross et al., 2001) confirmed that the coherence between the two signals was strongest in the occipital lobe (Fig. 3C top), as was raw for alpha power (Fig. 3C bottom). No such coherence was present between raw alpha signals and signals obtained by training decoders using shuffled labels (i.e., Baseline 1, Fig. S4). Taken together, these results suggest that the shape-specific prediction signals revealed in the current study originate from alpha oscillations in the occipital lobe.

Note that this demonstration of predictions fluctuating with alpha oscillations crucially relied on training a decoder on a static, phase-aligned signal (obtained from the localiser). In a situation such as the current one where signals are not phase-aligned between trials (because

they are modulated by endogenous oscillations), a more typical cross-validated within-experiment decoding analysis would not have succeeded, since there would not have been a consistent signal to train and test the decoder on.

Oscillatory power of predicted shape representations modulates behavioural expectation effects

If the strength of sensory predictions indeed modulates perceptual discrimination, there should be an opposite relationship between shape-specific pre-stimulus oscillations and behavioural performance on valid and invalid trials. To test this hypothesis, we performed a logistic regression analysis predicting behavioural accuracy from shape-specific oscillatory power in the alpha frequency range (8 – 12Hz), separately on valid and invalid trials. In line with previous literature relating oscillatory power to behavioural outcome (Hanslmayr et al., 2007; Mayer et al., 2016; Samaha et al., 2018), we limited the time window of interest to -500ms to 0 pre-stimulus, since prediction signals immediately preceding stimulus onset are most likely to impact perceptual performance (Kok et al., 2017). To assess the relationship between shape – specific decoder and raw occipital alpha, also control for non-specific alpha effects, raw alpha power (i.e., sensor-level alpha (8 – 12Hz) averaged over occipital MEG channels, for the identical time window of -500 to 0ms) was included in the model as an additional predictor. Calculating the across-trial correlation between shape-specific alpha and raw alpha power revealed that while the two signals were correlated (mean $r = 0.2725$, $p = 2.640 \times 10^{-14}$, $t(31) = 12.90$), they were not collinear (mean $r^2 < 0.1$), allowing both to be included as predictors in the same regression model. This indicates a weak but consistent positive relationship between the magnitude of shape-specific and raw occipital alpha activity. The analysis revealed a significant difference between valid and invalid prediction trials, with a numerically positive relationship between pre-stimulus shape-specific alpha power and performance on valid trials, and a numerically negative relationship on invalid trials ($p = 0.014$, $t(31) = 2.593$, Fig. 4B, left). On the other hand, raw occipital alpha power did not predict behavioural outcome (Fig. 4B, right), reflected in a non-significant difference between betas

of valid and invalid prediction trials ($p = 0.934$, $t(31) = 0.083$). To rule out spurious effects caused by correlated regressors, we also ran the logistic regression with shape-specific alpha power as the only predictor. This control analysis reproduced the significant difference between valid and invalid trials ($p = 0.021$, $t(31) = 2.4364$).

{Figure 4}

Note that the individual parameter estimates for valid and invalid trials were not significantly different from zero, while the difference between the two was. This likely reflects the fact that the individual conditions also contain unexplained, non-specific trial-by-trial variance in behavioural performance (e.g. due to slow fluctuations in alertness) that are subtracted out in the valid vs. invalid comparison. Importantly, the differential relationship between alpha power and behaviour dependent on prediction validity rules out any non-specific explanations of our results, and demonstrates a strong link between neural and behavioural effects of prediction. In short, pre-stimulus content-specific alpha oscillations modulated subsequent shape discrimination accuracy, such that the difference in accuracy between validly and invalidly predicted shapes was greater when shape-specific pre-stimulus alpha power was higher.

Stimulus predictions are driven by relatively late sensory representations

In an exploratory analysis, we investigated whether perceptual predictions reflected early or late visual representations, by dividing the training time period (70 – 200ms) into two separate windows, centred around the first (105ms) and second (175ms) peaks of the localiser decoding results, respectively (Fig. S1A). We then repeated the same frequency analysis performed on the original 70–200ms time window for each of these separate windows. Additionally, we created distinct baseline measurements for the early and late time windows. Note that these two distinct time windows appeared to form two distinct clusters (90 – 120ms and 160 – 190ms) in the temporal generalisation matrix of within-localiser decoding, with reduced cross-decoding between the two clusters suggesting qualitatively different representations (Fig. S1A). For the early training window (90 – 120ms), frequency analysis of

the pre-stimulus decoder time series revealed no power differences in the alpha band (10Hz: $p = 0.223$; 11Hz: $p = 0.090$) between the shape A vs. D decoding data and an empirical null distribution (Fig. 5A). Furthermore, training the decoder on the later time window (160 – 190ms) revealed significantly higher pre-stimulus alpha power (10Hz: $p = 0.016$; 11Hz: $p = 0.002$) in the shape A vs. D decoding data compared to a null distribution (Baseline 1, Fig. 5B).

{Figure 5}

Lastly, there was a significant difference in the average power in the 10 and 11Hz frequency bins of pre-stimulus shape A vs. D decoding between the early and late training time windows ($p = 0.005$, $t(31) = -3.022$). This is striking since these power spectra were calculated on the exact same MEG pre-stimulus data, the only difference was the localiser time window (90 – 120ms vs. 160 – 190ms) on which the decoder was trained. In sum, oscillating predictions seem to reflect relatively late sensory representations (160 – 190ms), rather than early feedforward-sweep-like signals.

Discussion

The present study examined the mechanisms through which predictions exert their influence on perception. Specifically, we tested whether content of perceptual predictions was represented in oscillations, and whether the power of this representation modulated performance on a visual discrimination task. To this end, we used multivariate decoding of MEG data to obtain the frequency spectrum of predicted shape representations. We revealed that predicted shape representations were strongest in the alpha frequency band (10 – 11Hz) (Fig. 3A). Source localisation analyses suggest that these signals originated from the occipital lobe (Fig. 3CD). Furthermore, we found that this stimulus-specific alpha power modulated task performance, such that higher alpha power resulted in stronger expectation effects on shape discrimination (Fig. 4B left). Together, these findings suggest that sensory templates of predicted visual stimuli are represented in pre-stimulus alpha oscillations, which subsequently modulate performance on a perceptual discrimination task.

Previous studies have hypothesised that oscillations play a critical role in conveying perceptual predictions (Ergenoglu et al., 2004; VanRullen et al., 2011; Mayer et al., 2016; Auksztulewicz et al., 2017; Samaha et al., 2018). This is largely based on indirect evidence, consisting of a range of studies finding that pre-stimulus alpha oscillations modulate performance on perceptual discrimination tasks (Ergenoglu et al., 2004; Hanslmayr et al., 2007; van Dijk et al., 2008; Mathewson et al., 2009; Sherman et al., 2016). Further, there is a second body of evidence that links experimental manipulations regarding stimulus predictability to the power and phase of low frequency oscillatory activities (Mayer et al., 2016; Sherman et al., 2016; Shen et al., 2023; Turner et al., 2023). For instance, expectations were found to rhythmically influence perceptual decision-making, where pre-stimulus occipital alpha phase modulated the strength of expectation effects on behaviour (Sherman et al., 2016). Furthermore, greater representational strength was observed in alpha-band activity in cue-induced imagery (relative to self-generated). This may reflect a more precise and temporally structured encoding of predicted features, consistent with the idea that prediction signals fluctuate in the alpha

rhythm, especially when externally cued and aligned to task-relevant content (Hu and Yu, 2023). Studies of selective attention have also shown that pre-stimulus alpha-band topographies contain category-specific information of to-be-attended objects, indicating that alpha activity not only reflects attentional engagement but also carries content-specific signals (Noah et al., 2020; Meyyappan et al., 2024). This also raises the important question whether expectation and attention act through the same or different neural mechanisms. Conceptually, expectation may be defined as the anticipation of likely future sensory events based on prior probabilities, and attention as the selective prioritisation of task-relevant stimuli or features. While some previous work suggests that they may operate independently (Kok et al., 2012a, 2016), many studies have revealed complex interactions between the two (e.g. (Kok et al., 2012b; Richter and de Lange, 2019); and for a review, see (Schröger et al., 2015), and disentangling the two is beyond the scope of the current study.

Lastly, a recent study has demonstrated a link between pre-stimulus high alpha/low beta power and the occurrence of high confidence false percepts (Haarsma et al., 2025). However, the key hypothesis that content of perceptual predictions is embedded in pre-stimulus alpha oscillations has remained untested. In the current study, we present evidence that the representation of predicted shapes fluctuated with pre-stimulus alpha oscillations. Therefore, we speculate that pre-stimulus alpha oscillations mediate content-specific feedback signalling, meaning that stimulus content oscillates with an alpha rhythm. Our findings align with a recent study on scene imagery, which successfully decoded individual imagined scenes from alpha band activity, demonstrating that complex visual image content can be encoded in alpha rhythms (Stecher and Kaiser, 2024). Furthermore, our findings are also in line with recent simultaneous EEG-fMRI recordings demonstrating alpha oscillations relating to feature-specific BOLD in superficial and deep layers (Clausner et al., 2024), suggesting an active involvement of alpha oscillations in stimulus processing. Together, this supports the more general proposal that oscillations can represent visual contents (Stecher et al., 2025).

While the role of pre-stimulus alpha oscillations has been extensively studied, it remains controversial. Several studies have reported that alpha-band activity is typically stronger in the absence of visual stimuli or when stimuli are unattended (Jensen and Mazaheri, 2010; Mathewson et al., 2011; Bonnefond and Jensen, 2012), leading to the influential hypothesis that alpha is predominantly an inhibitory rhythm. For example, during spatial attention tasks, alpha power shows hemispheric lateralisation, with decreased alpha power contralateral to the attended location and increased power ipsilaterally (Worden et al., 2000). This pattern has been interpreted as a mechanism of top-down distractor suppression (Ferrante et al., 2023). Together, these findings have shaped the dominant view of alpha as a rhythm that plays a key role in gating information flow via inhibition.

However, our results suggest that alpha oscillations are not solely inhibitory, but have an active role in cortical communication by representing the contents of feedback signals, and shaping perceptual priors and stimulus templates under conditions of expectation (Bonnefond and Jensen, 2012; Samaha et al., 2018). The link between stimulus-specific pre-stimulus alpha power and expectation effects on perception revealed in the current study suggests that whether alpha facilitates or inhibits sensory processing may depend on whether inputs match or mismatch current predictions. These findings do not necessarily challenge the prevailing view of alpha as an inhibitory rhythm. Rather, they suggest that fluctuations in the predicted-shape decoding time courses may reflect a distinct functional role of alpha in this context. Specifically, the magnitude of stimulus-specific alpha-band activity may represent the strength of predictive signals, with its amplitude modulating how effectively incoming sensory information is processed in the visual cortex. We believe that the prediction signals embedded in alpha oscillations in the current study reflect such stimulus-specific predictions. It is striking that these predictions modulated behavioural performance, despite the fact that the predictions were orthogonal to the task; there was a behavioural accuracy benefit when the first shape matched the prediction, and this benefit was stronger when shape-specific alpha power was higher.

While there is convergent evidence that the brain contains predictive signals (Kok et al., 2012a; Mayer et al., 2016; Aitken et al., 2020; Haarsma et al., 2023), the mechanisms through which the brain deploys these predictions remain largely unknown. Predictive coding has been suggested to involve rhythmic interactions between different frequency band activities (Arnal and Giraud, 2012; Bastos et al., 2012), where high frequency gamma is responsible for feedforward signalling (originating predominantly from superficial layers) and alpha/beta oscillations exert top-down control (feedback predictions), emerging from deep cortical layers. Indeed, animal work investigating the frequency characteristics and cortical layer specificity of predictable information processing (Bastos et al., 2015, 2020; Chen et al., 2023) revealed that pre-stimulus alpha power is an indicator of stimulus predictability, originating from cortical layers involved in feedback signalling (Bastos et al., 2020). Our results extend these intracranial electrophysiological observations by relating pre-stimulus alpha oscillations to the *contents* of feedback signalling.

Our findings reveal that prediction signals manifest in a spatially specific activity pattern that fluctuates in an alpha rhythm, predominantly originating from the occipital lobe (Fig.3BC, Fig.S7). This may be explained by neurons in the visual cortex tuned to different shapes receiving increased feedback modulation in the alpha band when their preferred shape is predicted. Alternatively, neurons tuned to the *unpredicted* shape may be suppressed in a rhythmic pattern, aligning with the pulse inhibition alpha theory (Klimesch et al., 2007; Mathewson et al., 2011). Future work using single cell recordings is required to distinguish between these alternatives.

Rather than predictions being actively conveyed in an alpha rhythm, an alternative explanation of our results may be that prediction signals passively ride on ongoing alpha oscillations. Alpha oscillations are the most prominent frequency band in the awake human brain, especially in the visual cortex, and even a non-oscillatory top-down signal arriving in visual cortex may inherit these alpha rhythms. Given our finding that shape-specific alpha power has opposite effects on behaviour dependent on the validity of the predictions, such a more passive

explanation seems less likely. However, future research is indeed needed to properly distinguish between these hypotheses.

Exploratory analyses revealed that the oscillating prediction signals reflected relatively late sensory representations (160 – 190ms localiser training window, Fig. 5B). We speculate that during this time period, the sensory representations captured by the decoder reflected an integration of bottom-up inputs and top-down recurrence, rather than solely the first feedforward sweep. Like the current study, previous studies have also revealed top-down modulations that reflected relatively late post-stimulus representations (i.e., 120 – 200ms) (Kok et al., 2017; Dijkstra et al., 2018). This may explain why predictions have been shown to modulate later sensory processing, while leaving the early feedforward sweep (< 80ms post-stimulus) mostly untouched (Alilović et al., 2019; Aitken et al., 2020).

Many prominent and influential theoretical frameworks have long speculated on the role of neural oscillations in perception (Klimesch et al., 2007; Jensen and Mazaheri, 2010; VanRullen et al., 2011). Here, we shed light on this by showing that the content of visual predictions fluctuated at alpha rhythms, and these rhythmic predictions modulate subsequent perceptual performance. These findings enrich current models of perceptual inference in the human brain by revealing a possible neural mechanism through which predictions are kept online in order to guide perception.

References

- Aitken F, Turner G, Kok P (2020) Prior Expectations of Motion Direction Modulate Early Sensory Processing. *J Neurosci* 40:6389–6397.
- Alilović J, Timmermans B, Reteig LC, van Gaal S, Slagter HA (2019) No Evidence that Predictions and Attention Modulate the First Feedforward Sweep of Cortical Information Processing. *Cerebral Cortex* 29:2261–2278.
- Arnal LH, Giraud A-L (2012) Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences* 16:390–398.
- Auksztulewicz R, Friston KJ, Nobre AC (2017) Task relevance modulates the behavioural and neural effects of sensory predictions Jensen O, ed. *PLoS Biol* 15:e2003143.
- Bastos AM, Lundqvist M, Waite AS, Kopell N, Miller EK (2020) Layer and rhythm specificity for predictive routing. *Proc Natl Acad Sci USA* 117:31459–31469.
- Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ (2012) Canonical Microcircuits for Predictive Coding. *Neuron* 76:695–711.
- Bastos AM, Vezoli J, Bosman CA, Schoffelen J-M, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P (2015) Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron* 85:390–401.
- Berger H (1929) Über das Elektrenkephalogramm des Menschen. *Archiv f Psychiatrie* 87:527–570.
- Bonnefond M, Jensen O (2012) Alpha Oscillations Serve to Protect Working Memory Maintenance against Anticipated Distracters. *Current Biology* 22:1969–1974.
- Brainard DH (1997) The psychophysics toolbox. *Spatial vision* 10:433–436.
- Busch NA, Dubois J, VanRullen R (2009) The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience* 29:7869–7876.
- Chen L, Cichy RM, Kaiser D (2023) Alpha-frequency feedback to early visual cortex orchestrates coherent naturalistic vision. *Sci Adv* 9:eadi2321.
- Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36:181–204.
- Clausner T, Marques JP, Scheeringa R, Bonnefond M (2024) Frequency and Laminar Profile of Feature Specific Visual Activity Revealed by Interleaved EEG-fMRI. Available at: <http://biorxiv.org/lookup/doi/10.1101/2024.07.31.605816> [Accessed August 28, 2025].
- Cousineau D (2005) Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in quantitative methods for psychology* 1:42–45.
- de Lange FP, Heilbron M, Kok P (2018) How Do Expectations Shape Perception? *Trends in Cognitive Sciences* 22:764–779.
- Dijkstra N, Mostert P, Lange FP de, Bosch S, van Gerven MA (2018) Differential temporal dynamics during visual imagery and perception. *eLife* 7:e33904.

Donoghue T, Haller M, Peterson EJ, Varma P, Sebastian P, Gao R, Noto T, Lara AH, Wallis JD, Knight RT, Shestyuk A, Voytek B (2020) Parameterizing neural power spectra into periodic and aperiodic components. *Nat Neurosci* 23:1655–1665.

Dugué L, Marque P, VanRullen R (2011) The Phase of Ongoing Oscillations Mediates the Causal Relation between Brain Excitation and Visual Perception. *J Neurosci* 31:11889–11893.

Ergenoglu T, Demiralp T, Bayraktaroglu Z, Ergen M, Beydagi H, Uresin Y (2004) Alpha rhythm of the EEG modulates visual detection performance in humans. *Cognitive Brain Research* 20:376–383.

Ferrante O, Zhigalov A, Hickey C, Jensen O (2023) Statistical Learning of Distractor Suppression Downregulates Prestimulus Neural Excitability in Early Visual Cortex. *J Neurosci* 43:2190.

Friston K (2009) The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences* 13:293–301.

Gross J, Kujala J, Hämäläinen M, Timmermann L, Schnitzler A, Salmelin R (2001) Dynamic imaging of coherent sources: Studying neural interactions in the human brain. *Proc Natl Acad Sci USA* 98:694–699.

Haarsma J, Deveci N, Corbin N, Callaghan MF, Kok P (2023) Expectation Cues and False Percepts Generate Stimulus-Specific Activity in Distinct Layers of the Early Visual Cortex. *J Neurosci* 43:7946–7957.

Haarsma J, Hetenyi D, Kok P (2025) Shared and diverging neural dynamics underlying false and veridical perception. *J Neurosci*:e1479242025.

Hanslmayr S, Aslan A, Staudigl T, Klimesch W, Herrmann CS, Bäuml K-H (2007) Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage* 37:1465–1473.

Hanslmayr S, Volberg G, Wimber M, Dalal SS, Greenlee MW (2013) Prestimulus Oscillatory Phase at 7 Hz Gates Cortical Information Flow and Visual Perception. *Current Biology* 23:2273–2278.

Haufe S, Meinecke F, Görgen K, Dähne S, Haynes J-D, Blankertz B, Bießmann F (2014) On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage* 87:96–110.

Holliday IE, Barnes GR, Hillebrand A, Singh KD (2003) Accuracy and applications of group MEG studies using cortical source locations estimated from participants' scalp surfaces. *Human Brain Mapping* 20:142–147.

Hu Y, Yu Q (2023) Spatiotemporal dynamics of self-generated imagery reveal a reverse cortical hierarchy from cue-induced imagery. *Cell Reports* 42:113242.

Iemi L, Busch NA, Laudini A, Haegens S, Samaha J, Villringer A, Nikulin VV (2019) Multiple mechanisms link prestimulus neural oscillations to sensory responses. *eLife* 8:e43620.

798 Jensen O, Mazaheri A (2010) Shaping Functional Architecture by Oscillatory Alpha Activity:
799 Gating by Inhibition. *Frontiers in Human Neuroscience* 4 Available at:
800 <https://www.frontiersin.org/article/10.3389/fnhum.2010.00186>.

801 Kerrén C, Linde-Domingo J, Hanslmayr S, Wimber M (2018) An optimal oscillatory phase for
802 pattern reactivation during memory retrieval. *Current Biology* 28:3383–3392.

803 Kerrén C, van Bree S, Griffiths BJ, Wimber M (2022) Phase separation of competing
804 memories along the human hippocampal theta rhythm. *eLife* 11:e80633.

805 Klimesch W, Sauseng P, Hanslmayr S (2007) EEG alpha oscillations: The inhibition–timing
806 hypothesis. *Brain Research Reviews* 53:63–88.

807 Kok P, Jehee JFM, de Lange FP (2012a) Less Is More: Expectation Sharpens
808 Representations in the Primary Visual Cortex. *Neuron* 75:265–270.

809 Kok P, Mostert P, Lange FP de (2017) Prior expectations induce prestimulus sensory
810 templates. *Proceedings of the National Academy of Sciences* 114:10473–10478.

811 Kok P, Rahnev D, Jehee JFM, Lau HC, de Lange FP (2012b) Attention Reverses the Effect
812 of Prediction in Silencing Sensory Signals. *Cerebral Cortex* 22:2197–2206.

813 Kok P, Turk-Browne NB (2018) Associative Prediction of Visual Shape in the Hippocampus.
814 *The Journal of Neuroscience*.

815 Kok P, van Lieshout LLF, de Lange FP (2016) Local expectation violations result in global
816 activity gain in primary visual cortex. *Sci Rep* 6:37706.

817 Mathewson KE, Gratton G, Fabiani M, Beck DM, Ro T (2009) To See or Not to See:
818 Prestimulus α Phase Predicts Visual Awareness. *J Neurosci* 29:2725–2732.

819 Mathewson KE, Lleras A, Beck DM, Fabiani M, Ro T, Gratton G (2011) Pulsed Out of
820 Awareness: EEG Alpha Oscillations Represent a Pulsed-Inhibition of Ongoing
821 Cortical Processing. *Front Psychology* 2 Available at:
822 <http://journal.frontiersin.org/article/10.3389/fpsyg.2011.00099/abstract> [Accessed
823 September 5, 2023].

824 Mayer A, Schwiedrzik CM, Wibral M, Singer W, Melloni L (2016) Expecting to See a Letter:
825 Alpha Oscillations as Carriers of Top-Down Sensory Predictions. *Cereb Cortex*
826 26:3146–3160.

827 Meyyappan S, Ding M, Mangun GR (2024) Hierarchical Organization of Visual Feature
828 Attention Control. *bioRxiv*:2024.10.02.615879.

829 Morey RD (2008) Confidence intervals from normalized data: A correction to Cousineau
830 (2005). *Tutorials in Quantitative Methods for Psychology* 4:61–64.

831 Mostert P, Kok P, de Lange FP (2015) Dissociating sensory from decision processes in
832 human perceptual decision making. *Sci Rep* 5:18253.

833 Noah S, Powell T, Khodayari N, Olivan D, Ding M, Mangun GR (2020) Neural Mechanisms
834 of Attentional Control for Objects: Decoding EEG Alpha When Anticipating Faces,
835 Scenes, and Tools. *J Neurosci* 40:4913–4924.

836 Oostenveld R, Fries P, Maris E, Schoffelen J-M (2011) FieldTrip: Open Source Software for
837 Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data.
838 Computational Intelligence and Neuroscience 2011:1–9.

839 Op de Beeck H, Wagemans J, Vogels R (2001) Inferotemporal neurons represent low-
840 dimensional configurations of parameterized shapes. *Nat Neurosci* 4:1244–1252.

841 Richter D, de Lange FP (2019) Statistical learning attenuates visual activity only for attended
842 stimuli. *eLife* 8:e47869.

843 Romei V, Gross J, Thut G (2010) On the Role of Prestimulus Alpha Rhythms over Occipito-
844 Parietal Areas in Visual Input Regulation: Correlation or Causation? *Journal of*
845 *Neuroscience* 30:8692–8697.

846 Samaha J, Boutonnet B, Postle BR, Lupyan G (2018) Effects of meaningfulness on
847 perception: Alpha-band oscillations carry perceptual expectations and influence early
848 visual responses. *Sci Rep* 8:6606.

849 Schröger E, Marzecová A, SanMiguel I (2015) Attention and prediction in human audition: a
850 lesson from cognitive psychophysiology. *Eur J of Neuroscience* 41:641–664.

851 Shen L, Wu Z, Yue Z, Li B, Chen Q, Han B (2023) Prior knowledge employs pre-stimulus
852 alpha-band oscillations and persistent post-stimulus neural templates for conscious
853 perception. *J Neurosci*:JN-RM-0263-23.

854 Sherman MT, Kanai R, Seth AK, VanRullen R (2016) Rhythmic Influence of Top–Down
855 Perceptual Priors in the Phase of Prestimulus Occipital Alpha Oscillations. *Journal of*
856 *Cognitive Neuroscience* 28:1318–1330.

857 Stecher R, Cichy RM, Kaiser D (2025) Decoding the rhythmic representation and
858 communication of visual contents. *Trends in Neurosciences*:S0166223624002480.

859 Stecher R, Kaiser D (2024) Representations of imaginary scenes and their properties in
860 cortical alpha activity. *Sci Rep* 14:12796.

861 Stelzer J, Chen Y, Turner R (2013) Statistical inference and multiple testing correction in
862 classification-based multi-voxel pattern analysis (MVPA): Random permutations and
863 cluster size control. *NeuroImage* 65:69–82.

864 Tarasi L, di Pellegrino G, Romei V (2022) Are you an empiricist or a believer? Neural
865 signatures of predictive strategies in humans. *Progress in Neurobiology* 219:102367.

866 Turner W, Blom T, Hogendoorn H (2023) Visual Information Is Predictively Encoded in
867 Occipital Alpha/Low-Beta Oscillations. *J Neurosci* 43:5537–5545.

868 van Dijk H, Schoffelen J-M, Oostenveld R, Jensen O (2008) Prestimulus Oscillatory Activity
869 in the Alpha Band Predicts Visual Discrimination Ability. *J Neurosci* 28:1816–1823.

870 Van Veen BD, Van Drongelen W, Yuchtman M, Suzuki A (1997) Localization of brain
871 electrical activity via linearly constrained minimum variance spatial filtering. *IEEE*
872 *Trans Biomed Eng* 44:867–880.

873 VanRullen R, Busch NA, Drewes J, Dubois J (2011) Ongoing EEG Phase as a Trial-by-Trial
874 Predictor of Perceptual and Attentional Variability. *Front Psychology* 2 Available at:
875 <http://journal.frontiersin.org/article/10.3389/fpsyg.2011.00060/abstract> [Accessed
876 September 4, 2023].

877 Watson AB, Pelli DG (1983) Quest: A Bayesian adaptive psychometric method. Perception
878 & Psychophysics 33:113–120.

879 Worden MS, Foxe JJ, Wang N, Simpson GV (2000) Anticipatory Biasing of Visuospatial
880 Attention Indexed by Retinotopically Specific α -Band Electroencephalography
881 Increases over Occipital Cortex. J Neurosci 20:RC63–RC63.

882 Zahn CT, Roskies RZ (1972) Fourier Descriptors for Plane Closed Curves. IEEE
883 Transactions on Computers C-21:269–281.

884

JNeurosci Accepted Manuscript







Figure 1: Experimental paradigm. **A:** During prediction runs, an auditory cue preceded the presentation of two consecutive shape stimuli. On each trial, the second shape was either identical to the first or slightly warped with respect to the first, and participants' task was to report whether the two shapes were the same or different. **B:** The auditory cue (rising vs. falling tones) predicted whether the first shape on that trial would be shape A or shape D. The cue was valid on 75% of trials, whereas in the other 25% of (invalid) trials the unpredicted shape was presented. **C:** During shape localiser runs no predictive auditory cues were presented and participants performed a fixation dimming task. **D:** Four different shapes were presented in the localiser runs, appearing with equal (25%) likelihood. Only shape A and D were presented in the prediction runs. The amplitudes of two RFCs (1.11, and 1.54Hz components) were varied in order to create a two-dimensional shape space (multiplication indicated shape A (1:1); shape B (4:1); shape C (1:4); shape D (4:4), such that shape A vs. D discrimination was orthogonal to shape B vs. C discrimination. **E:** Schematic of the hypothesis: cue-induced predictions oscillate in the alpha frequency band (~10Hz) in the interval between predictive cue and stimulus onset. **F:** Participants were able to discriminate the two presented shapes more accurately when the auditory cue validly predicted the identity of the first shape (* $p < 0.05$). Dots represent individual participants, error bars indicate within-participant SEM (Cousineau, 2005; Morey, 2008).

Figure 2: Shape prediction frequency analysis pipeline. **A:** Neural representations of the presented shapes within the localiser runs in a cross-validated manner (-100 to 600ms, relative to stimulus onset). The decoder was highly accurate at discriminating the shapes based on the MEG signal. The presented shapes were successfully decoded from 65ms to 450ms ($p < 0.001$). For all subsequent analyses, decoding traces were averaged over a training window of 70 to 200ms, during which shape decoding peaked. **B:** A decoder was trained to discriminate between shapes A and D in the localiser runs. This decoder was applied to the pre-stimulus time window in prediction runs (-1250 to 0ms). Trial-based pre-stimulus decoding time series were subjected to FFT. **C:** The resulting power spectrum was compared to the 95th percentile of an empirical null distribution generated by bootstrapping decoders trained with pseudo-randomised labels (Baseline 1, left), as well as to a decoder trained on shapes only presented in the localiser (shapes B and C) (Baseline 2, right).


Figure 3: Auditory cue-induced prediction templates fluctuate at alpha frequencies. **A - top:** The power spectrum of pre-stimulus (-1250 to 0ms) shape decoding shows significant deviations from an empirical null distribution at 10Hz and 11Hz (** $p = 0.003$). The baseline power spectrum (dark blue line) was obtained by bootstrapping ($n = 10,000$) shuffled label decoding data ($n = 25$ per participant). Mean and shaded regions indicate SD. Solid orange line indicates the 95th percentile of the null distribution. Error bars indicate SEM. **A - bottom:** Pre-stimulus (-1250 to 0ms) MEG data revealed significantly higher 10 – 11Hz power for shape A vs. D decoding than for shape B vs. C decoding (* $p < 0.05$). Bars indicate power of shape A vs. D decoding; dark purple line indicates power of B vs. C decoding (applied to identical pre-stimulus prediction data). Shaded regions and error bars indicate SEM. **B - top:** Topography of sensor-level alpha (10Hz) coherence between shape-specific signals (i.e., decoder output) and pre-stimulus sensor-level raw MEG signals (-1250 to 0ms). **B - bottom:** Topography of raw pre-stimulus sensor-level alpha power (10Hz, -1250 to 0ms). **C - top:** Source-localised alpha coherence between shape-specific and pre-stimulus raw MEG signals (-1250 to 0ms). **C - bottom:** Source-localised raw alpha power (10Hz with 2Hz spectral smoothing, -1250 to 0ms).

Figure 4: Oscillatory power of predicted shape representations modulates behavioural accuracy. **A:** A representative participant's scatter plot of shape-specific alpha power as a function of valid vs incorrect prediction trials. We expected to observe an opposite trend between valid and invalid prediction trials, meaning for correct responses higher overall shape-specific alpha power for valid trials vs. incorrect valid trials, whereas higher power was expected for incorrect invalid trials vs. correct invalid trials. For all participants' scatter plots, see Fig. S5. **B (left):** Parameter estimates (i.e., betas) of the logistic regression between the power of pre-stimulus decoding alpha (-500 to 0ms), averaged over 8 – 12Hz frequency bins, and discrimination performance, separately for valid and invalid prediction trials (* $p < 0.05$). **B (right):** Parameter estimates (i.e., betas) of the logistic regression between the power of pre-stimulus raw alpha (-500 to 0ms), averaged over 8 – 12Hz frequency bins, and discrimination performance, separately for valid and invalid prediction. No significant difference between. Dots represent individual participants, error bars indicate within-participant SEM (Cousineau, 2005; Morey, 2008).


Figure 5: Late sensory representations drive stimulus predictions. **A:** Power spectrum of the -1250 to 0ms prediction time window shape A vs. D decoding, trained on the 90 to 120ms post-stimulus localiser window. No significant distinctions between the shape A vs. D decoder data and an empirical null distribution at 10Hz and 11Hz. Mean and shaded regions indicate SD. Dark solid orange line indicates the 97.5th percentile of the null distribution, implementing a one-sided test at $p < 0.05$ while correcting for the two time windows tested here. **B:** Power spectrum of the -1250 to 0ms prediction time window shape A vs. D decoding, trained on the 160 to 190ms post-stimulus localiser window. Statistically significant difference from an empirical null distribution at 10Hz and 11 Hz (* $p < 0.002$). The baseline power spectrum (dark blue line) was calculated as before. Mean and shaded regions indicate SD. Solid orange line indicates the 97.5th percentile of the baseline distribution.



A






B





C

Baseline 1: Shuffled label Shape A vs D decoder


Baseline 2: Shape B vs C decoder



FFT



A



B

