Progressive multi-stage extrapolation of predictable motion in human visual cortex

William Turner[1,2*], Charlie Sexton[2], Philippa A. Johnson[3], Ella Wilson[2], & Hinze Hogendoorn[1,2]

[1]School of Psychology & Counselling, Queensland University of Technology, Brisbane.
[2]Melbourne School of Psychological Sciences, The University of Melbourne, Melbourne, Australia.
[3]Cognitive Psychology Unit, Institute of Psychology & Leiden Institute for Brain and Cognition, Leiden University, Leiden, Netherlands

Corresponding Author: williamfrancisturner@gmail.com

Acknowledgements: This work was supported by Australian Research Council Grants FT200100246, DP220101166, and DP180102268 awarded to HH and QUT ECRIS and Decision Science Hub Seed Funding Grants Awarded to WT.

Competing interests: The authors declare no competing interests.

Abstract

Neural processing of sensory information takes time. Consequently, to estimate the current state of the world, the brain must rely on predictive processes – for example, extrapolating the motion of a ball to determine its probable present position. Mounting evidence suggests that extrapolation occurs during early (retinal) processing, however it remains unclear whether extrapolation continues during later-stage (cortical) processing. Moreover, we currently lack a spatially precise characterisation of extrapolation effects in the human brain, with most studies relying on invasive neurophysiological techniques in animals. Here, we address these issues by demonstrating how precise probabilistic maps can be constructed from human EEG recordings. Participants (N = 18) viewed a stimulus moving along a circular trajectory while EEG was recorded. Using LDA classification, we extracted maps of stimulus location over time and found evidence of a widespread temporal shift occurring across distinct processing stages. This accelerated emergence of position representations indicates progressive extrapolation occurring at multiple stages of processing, with representations across the hierarchy shifted closer to real-time. We further show evidence of representational overshoot during early-stage processing following unexpected changes to an object's trajectory, and demonstrate that the observed dynamics can emerge spontaneously in a simulated neural network via spike-timing-dependent plasticity.

Introduction

Neural processing of visual information takes time. Retinal ganglion cells produce spikes at latencies of ~20-70 ms and additional delays accumulate as these signals pass on to downstream regions[1,2]. For time-sensitive interactions with dynamic environments, delays are problematic. For instance, imagine a hunter trying to take down bolting prey, or a tennis player trying to return a 200 km/h serve, while only having access to outdated visual information. For the tennis player, a delay of even just 50 ms will cause their ball-position estimates to be off by 2.8 metres. The fact that such behaviours are still possible raises an important question: if visual processing is delayed, how do we accurately localize moving objects in real time?

Neurophysiological recordings in non-human animals have revealed clear evidence of predictive motion extrapolation occurring during the earliest stages of visual processing (see[2] for a recent review). For example, retinal ganglion cells in salamanders, rabbits, mice and monkeys have been found to 'anticipate' the arrival of moving stimuli such that the peak of population-level activity approximately aligns with the leading edge of the stimulus despite phototransduction delays[3–7]. Similar anticipatory effects have also been observed in early cortical regions of cats[8,9] and monkeys[10–13].

Within a hierarchical predictive coding framework[14] motion extrapolation can help to minimize prediction error and thus the metabolic cost of sensory processing[15]. However, this relies upon extrapolation occurring at all stages of processing, not just the earliest, to avoid the re-emergence of delays and misalignment of sensory representations across hierarchical layers. Yet within the existing literature, extrapolatory effects have been observed predominantly in early visual regions (i.e. retina, LGN, and V1), and it is unclear to what degree neurons downstream from the retina simply inherit their extrapolated activity profiles, without actively driving further extrapolation. As such, an important open question is whether neural motion extrapolation is a multi-stage phenomenon, that occurs during later-stage cortical processing.

To address this question, one path forward lies in using more global measures of neural activity, to concurrently probe distinct stages of visual processing. To this end, anticipatory effects have recently been observed in human M/EEG [16–19] and fMRI recordings[20], with some recent evidence from our own lab suggesting that extrapolation may only occur during very early (i.e. pre-cortical) processing[17]. However, all of these past studies have been limited in their ability to examine stimulus representations with fine-grained spatial resolution, so have been unable to clearly resolve the underlying predictive dynamics.

Here, we develop a method for extracting high-resolution maps of a visual stimulus' position over time from EEG recordings. This allows us to precisely reconstruct the trajectories of moving stimuli, revealing evidence of overshoots when stimuli disappear or reverse direction (consistent with[18]). To determine whether widespread extrapolation occurs across multiple stages of visual processing, we train machine learning classifiers on the evolving cascade of neural responses which follow the onset of static stimuli. We find evidence of the same activity patterns occurring in response to smoothly moving stimuli, but with a cumulative, compensatory shift in their timing. Specifically, activity patterns associated with temporally distinct stages of processing, which are activated sequentially under unpredictable conditions, emerge earlier than expected when viewing smoothly moving (i.e. predictable) stimuli. This accelerated emergence of position representations leads to increased temporal-alignment of representations across processing stages, reducing the gap between the encoded and actual position of the stimulus. Finally, we provide a simple, biologically plausible

model that captures these dynamics, by demonstrating that they emerge spontaneously, without supervision, at all levels of a hierarchical neural network via spike-timing-dependent plasticity (STDP).

Results

Participants (N = 18) each completed two experimental sessions. In each session they viewed 2000 localiser stimuli (4000 total), consisting of a white wedge-shaped stimulus randomly flashed in 40 equally-spaced positions around an invisible circle (Figure 1A). Participants also viewed 960 smooth motion sequences in each session (1920 total), in which the same stimulus moved along a circular trajectory for 1.5-3.5s, before either disappearing or reversing its direction at an unpredictable location (Figure 1B). Following reversals (50% of trials) the stimulus continued moving for between 0.5-1s before disappearing.

Decoding stimulus-position information from EEG recordings

To characterise how stimulus-position information is encoded in neural activity, we first examined whether it was possible to predict the position of the static localiser stimuli from participants' neural response patterns. At each timepoint, we fit multivariate linear models to predict the position of a given stimulus. Because these positions are angular (i.e. circularly distributed), we trained models to predict the sine and cosine of the stimulus' angular position from the voltage at all electrodes. We scored the performance of these models by calculating the inverse absolute angular error ('decoding score') between the predicted and actual position of a stimulus (following[21], see Method).

Cross-validation revealed clear evidence of stimulus-position information in participants' neural activity emerging on average ~75 ms after stimulus onset and remaining sustained for ~500 ms ($p < .01$ cluster corrected, Figure 1C). A searchlight analysis (Figure 1C inset) indicated that a stimulus' position was best predicted from neural activity recorded over the occipital cortex. Examining average decoding performance (75-250 ms) at each localiser position, revealed that all positions could be accurately decoded, with a slight qualitative advantage for stimuli in the lower visual field (Figure 1D).

A temporal generalization analysis[22] revealed activity dynamics that were predominantly transient and evolving, with a strong diagonal response in the temporal generalization matrix and only brief, transient periods of peri-diagonal generalization (Figure 1E, $p < .01$ cluster corrected). Re-running the analysis on frequency-specific power estimates (i.e. the relative pattern of oscillatory power across electrodes at a given frequency), revealed that position information was predominantly encoded in the alpha/low-beta range (~10-20 Hz, see Figure 1F; consistent with[16]).

Taken together, these dynamics are consistent with the delayed propagation of position-specific activity patterns through a hierarchical network of brain regions following stimulus onset. The predominance of a diagonal pattern within the temporal generalization matrix suggests that stimuli trigger evolving sequences of neural activity, reflecting distinct stages of sensory processing[22–25]. Re-running the temporal generalization analysis with a focus on sequentially predicting the location of successive stimuli (Figure 1G) revealed that information about multiple stimuli is encoded across distinct stages of processing at any given timepoint (consistent with[23,24]). The fact that stimulus-position information was spectrally localised in the alpha/low-beta range is in line with the recent suggestion that such rhythms

may be an oscillatory 'fingerprint' of information processing within hierarchical predictive networks under neural delays[16,26].

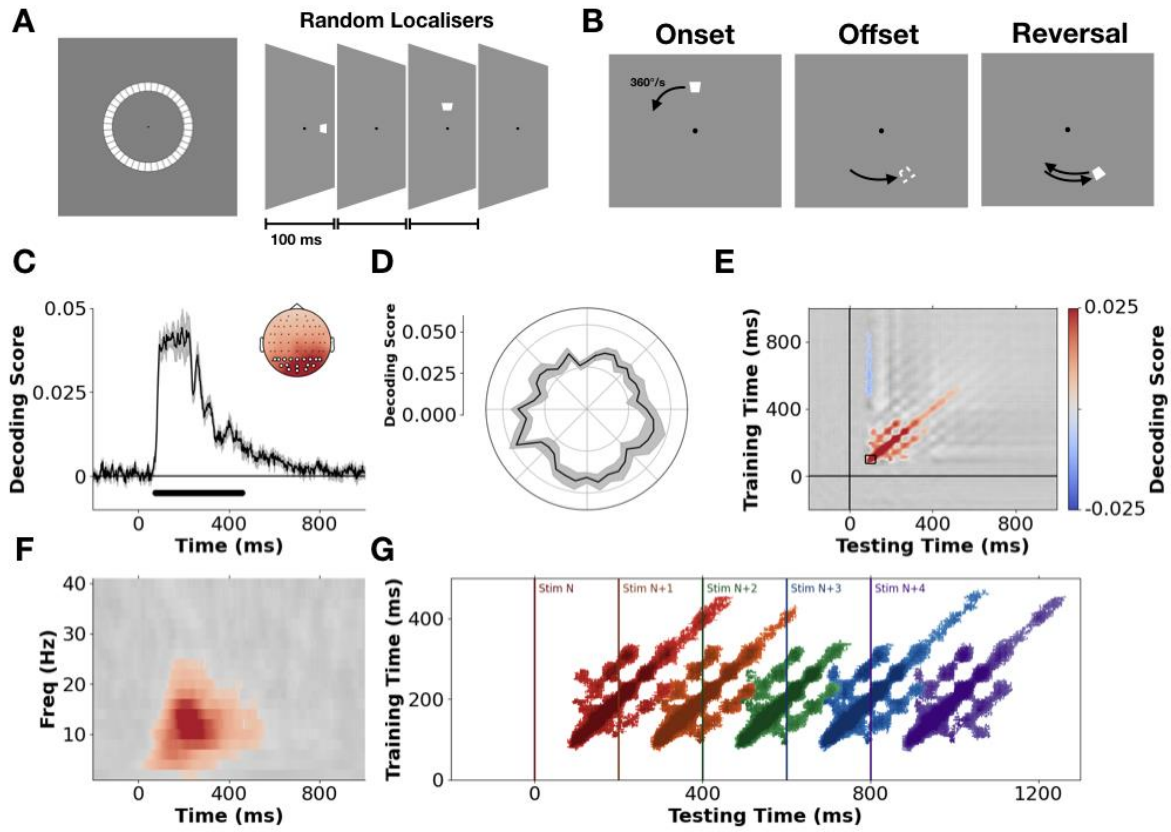

Figure 1. Stimulus design and characterisation of location-specific neural activity patterns. A) Localiser stimuli were randomly presented in 40 equally spaced positions tiling an invisible circle around fixation. **B)** In the smooth motion sequences, participants viewed the stimulus moving along a circular trajectory at 360°/s for 1.5-3.5s before disappearing. In 50% of trials, stimuli reversed their direction mid-way through the sequence. Following a reversal the stimulus continued moving for 0.5 to 1s. **C)** Model decoding score (fractional inverse absolute angular error, 0 = guessing, 1 = perfect prediction, see Method) is plotted over time. Predictions are based on brain activity recorded at varying time points relative to localiser onset (x-axis). Black dots indicate timepoints where prediction accuracy significantly differs from chance ($p < .01$ cluster corrected). Inset shows the results of a searchlight analysis over electrodes plus their immediate neighbours. Highlighted electrodes show the posterior/occipital sites used in the subsequent LDA-based analyses. **D)** Average decoding performance (75-250 ms) at each localiser position. **E)** Results of a temporal generalisation analysis in which the performance of timepoint-specific decoders is assessed across all testing time points. The full generalisation matrix is plotted in greyscale with cluster-corrected timepoints overlaid in colour ($p < .01$). The region marked with a black box indicates the training time-period (75-125 ms) for subsequent LDA-based mapping. **F)** Results of a frequency-specific decoding analysis in which normalized power estimates were used as the input features to the model. The full results are plotted in greyscale with cluster-corrected timepoints overlaid in colour. **G)** Temporally generalised decoding of successive localiser stimuli. Coloured regions show above chance generalization for each stimulus respectively ($p < .01$ cluster corrected).

Mapping the position of moving stimuli

Having confirmed that the localisers evoked position-specific activity patterns across successive hierarchical levels, we next examined whether we could leverage these patterns to continuously track the position of the stimuli during smooth motion. First, we trained multiclass linear discriminant analysis (LDA) classifiers on participants' neural responses to the localisers, treating each position as a distinct class. Then, from participant's neural responses to the smoothly moving stimuli, we extracted predicted posterior class probabilities via the pre-trained LDA models. In other words, we extracted a prediction as to the probability that the stimulus was in each localiser position at a given timepoint. Averaging the probabilities extracted from models trained 75-125 ms after localiser onset (see boxed region in Fig. 1E) we were left with a matrix of position probabilities over time (i.e. a probabilistic spatio-temporal map).

Figure 2 shows three such maps time-locked to either stimulus onset, stimulus offset or the reversal point in the motion sequences. Examining these, it is clear that the position of the moving stimuli can be tracked from participants' neural response patterns. This, in itself, is non-trivial as smoothly moving stimuli do not evoke the well-defined onset/offset responses that have previously been leveraged to decode the position of 'apparent motion' stimuli[16,18]. The fact that we can map the position of smoothly moving objects via a bank of pre-trained static position representations indicates that the position-specific activity patterns evoked by static and dynamic stimuli overlap, at least partially (consistent with[17]).

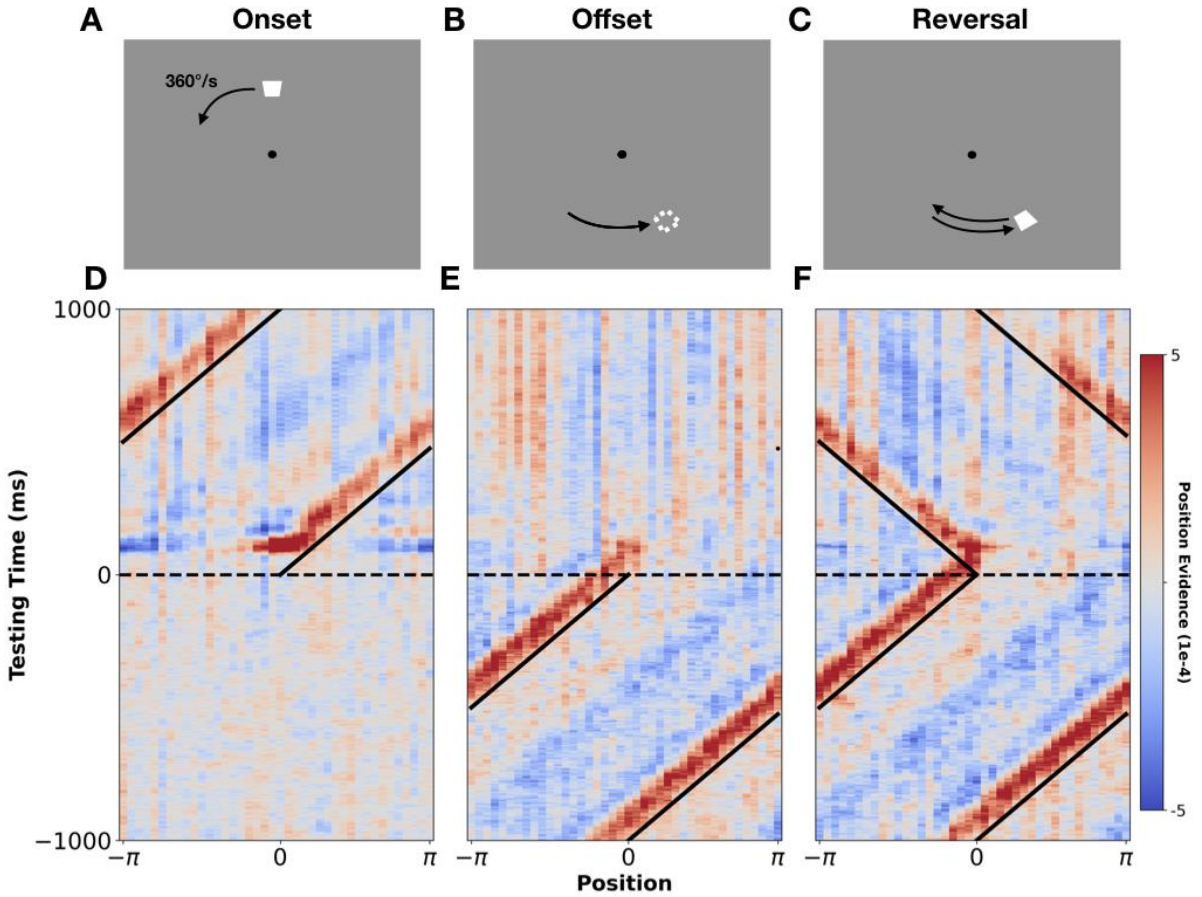

Figure 2. Mapping the position of moving stimuli. Panels **A) – C)** show the three events of interest: stimulus onset, stimulus offset, and stimulus reversal. Panels **D) – F)** show probabilistic spatio-temporal maps centred around these three events. Diagonal black lines mark the true position of the stimulus.

Horizontal dashed lines mark the time of the event of interest (stimulus onset, offset, or reversal). Red indicates high probability regions and blue indicates low probability regions relative to chance (0.025).

To examine the shape and precision of the decoded probability distributions, we took two averaged time slices: 1) directly following motion onset (75-125 ms; Figure 3A), and 2) in the lead-up to stimulus offset (-1000-0 ms; Figure 3B). Examining the first (Figure 3A), we could effectively get a snapshot of the decoded probability distribution, just after visually-evoked activity first reaches visual cortex. From this we see that participants' early neural responses already encoded a remarkably precise probability distribution over space (FWHM of above chance probabilities = 54° polar angle). Examining the second time slice (Figure 3B), we effectively get a snapshot of the decoded probability distribution, following extended exposure to motion – that is, the 'steady state' of decoded position information during motion. Notably, the distribution has narrowed (FWHM of 36° polar angle) and evolved from being more-or-less symmetric to become positively skewed in the direction of stimulus motion. The trailing edge of the distribution became suppressed and the leading edge became enhanced, mirroring the changes observed in directly imaged neural activity of non-human animals (see[2]). These changes serve to shift the high probability region closer towards the real-time position of the stimulus. Since these results are generated from models trained on early neural responses to the localisers (75-125 ms), this provides evidence of extrapolation occurring during early-stage human visual processing.

Considering Figure 3A&B, an important auxiliary question arises: how might a point-estimate of the stimulus' real-time location be read out from these probability distributions? One option is to take the point of maximum probability (green triangles in Figure 3A&B). However, even after considerable exposure to predictable motion (panel B) this estimate continues to lag the real-time position of the stimulus (marked with the vertical black line). An alternative option is to take the centroid (i.e. the vector average), to better leverage information contained within the entire distribution. Interestingly, this yields an estimate (orange triangles) which initially aligns with the distribution peak (panel A), but which then shifts, after ongoing exposure to predictable motion, to align with the real-time position of the stimulus (panel B). Plotting both peaks and centroids across time (Figure 3C&D), we can see that the peak estimate consistently lags the real-time position of the stimulus. In contrast, the centroid estimate approximately tracks the real-time stimulus position, yet overshoots when the stimulus disappears or reverses. This suggests that early visual processing may serve to encode a probability distribution over space in a manner which allows for different point estimates to be read out, depending on whether accurate instantaneous position readout is required. At any given time point, if a real-time estimate of an object's position is required, the centroid may be taken. However, if speed is not paramount (i.e. no action will be taken) the peak can instead be taken, as, after a brief delay, this will give a more reliable estimate of where the stimulus was (i.e. it will not suffer from overshoot when the stimulus changes direction). For consideration of how this speculative proposal may be further tested, see the Discussion.

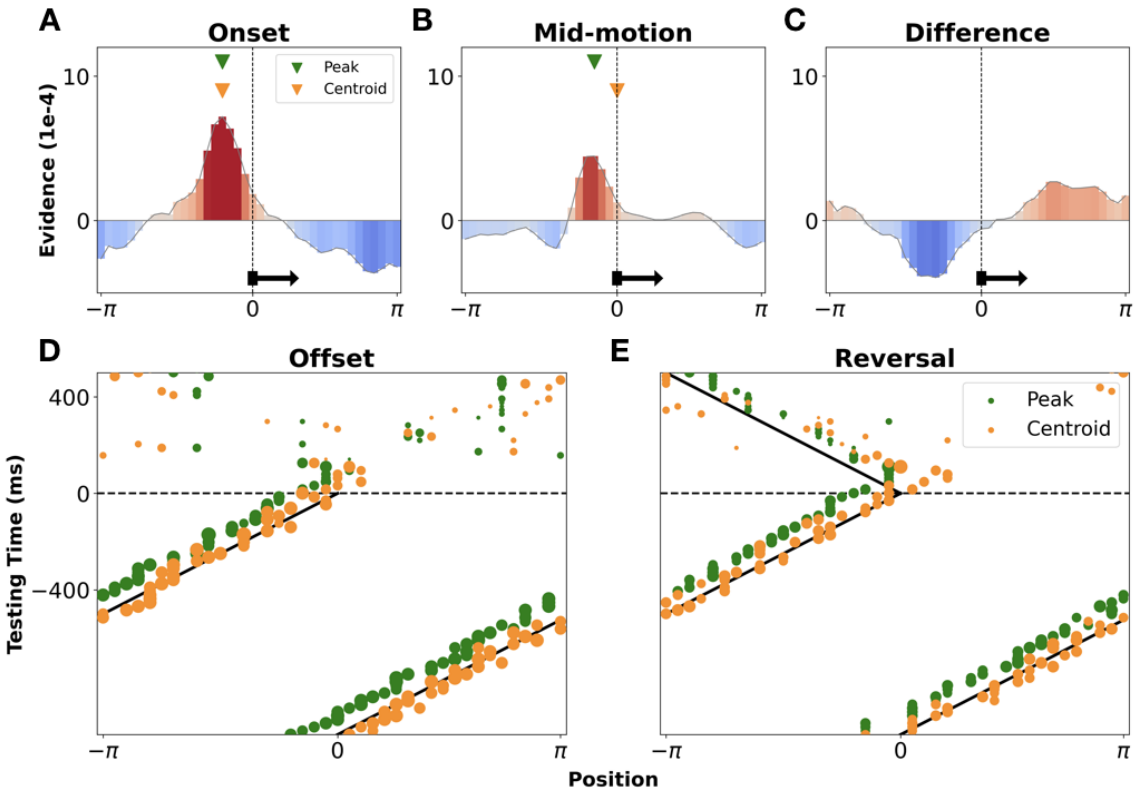

Figure 3. Examining the shape and potential discrete read-out of decoded probability distributions. Panels **A)** and **B)** show averaged time slices through the maps in Figure 2, at motion onset (75-125) and mid-motion (-1000-0 ms) respectively. Green inverted triangles show the point of peak probability and orange triangles show the centroid (vector average) of the distribution. Panel **C)** shows the difference between the probability distributions extracted from mid-motion and immediately following stimulus onset. In panels **A-C)** the true position, and direction, of the stimulus is marked by a vertical black dashed line and horizontal black arrow. Panels **D)** and **E)** show timepoint-specific markers of peak probability (green) and centroid (orange), overlaid on the real-time position of the stimulus around stimulus offset and reversal. Peaks and centroids are calculated every 15 ms, with the certainty of the estimate (taken as the peak height or vector average length, respectively) dictating the size of the plotted dot.

Accelerated emergence of representations across distinct processing stages

Turning to the primary question of interest, we examined how stimulus-position information was encoded across distinct stages of visual processing. Specifically, we extracted probability distributions over possible positions from decoders pre-trained on data recorded at different timepoints following localiser onset (Figure 4). In this way, training time effectively becomes a proxy for hierarchical level, with time-point-specific decoding performance reflecting the presence or absence of position information at specific levels of representation. Then, instead of averaging across decoders (as we did to generate Figure 2), we can simultaneously consider all extracted probability distributions at once. Importantly, the temporal generalisation matrix (Figure 1D) revealed a predominantly diagonal pattern, validating that decoders trained on different timepoints learn different activity patterns, thus indexing distinct stages of processing (stable activity would result in a constant, square pattern of generalisation, which we do not observe, see[22]). Using training time as a proxy for processing stage, we can

therefore examine how position-specific information is encoded across different stages of processing. (Note, this does not mean we can infer anything about the cortical locus of such processing, only that we are decoding from distinct activity patterns.)

To investigate the cascade of neural responses evoked by moving stimuli, after training the LDA models on localiser-evoked activity, we took 1s epochs of EEG data in the lead up to stimulus offset or reversal (i.e. during sustained periods of smooth motion), and extracted posterior probabilities over the localiser positions from the pre-trained, timepoint-specific models. We then realigned the probabilities to the true position of the moving object. This allowed us to generate an average snapshot of how location information is encoded, at any given time, across multiple stages of visual processing during ongoing exposure to motion. If moving stimuli evoke the same cascade of neural responses as static stimuli, without any delay compensation, then the decoded position of the stimulus would sit along the diagonal line in Figure 4A-B ('No Compensation' line). That is, position representations encoded at later stages of processing will progressively lag those encoded at earlier stages, due to a compounding of delays as information travels along the processing hierarchy. If, instead, perfect delay compensation is achieved, then the decoded position of the stimulus would sit on the vertical line in Figure 4A-B ('Full Compensation' line), with the real-time position of the stimulus encoded across all processing stages.

Examining Figure 4A, we see that the bulk off the high probability region (in red) is shifted away from the No-Compensation line towards the Full-Compensation line. This indicates a shift in the timing of evoked responses, with representations emerging earlier than would be expected without delay compensation. To generate point estimates of the encoded position of the stimulus across processing stages, we overlay the timepoint-specific peak probabilities (in white). We chose the peak, as opposed to the centroid, as the most conservative discrete estimator, since the centroid already showed extrapolative properties (see Figure 3). However, the results are unaffected when substituting the centroid (see Supplement 1).

Examining the peak position estimates in Figure 4A&B, we see a sustained forward shift away from the No-Compensation line, with the distance between the points and No-Compensation line increasing over time. This is indicative of sustained extrapolation occurring, leading the encoded position of the stimulus to increasingly deviate from what would be expected as a consequence of neural delays.

One reasonable concern here is that such a shift may be driven by autocorrelation of position signals at successive timepoints. That is, the position of a smoothly moving object is predictive of itself over short time windows. Hence, the position information within participants' neural responses will be autocorrelated. Since neighbouring positions are likely to generate more similar neural activity patterns, this could conceivably blur the extracted positional probabilities. While such smearing is unlikely to affect peak probability estimates, we nevertheless sought to empirically rule out the possibility that the observed shift is simply an artifact of autocorrelation.

To this end, we ran a control analysis on 'synthetic' EEG responses to simulated motion, which we constructed by averaging across successively lagged ERPs evoked by an ordered sequence of localiser stimuli. That is, we created a matrix in which each subsequent row contained the response to the neighbouring localiser, temporally offset by the time taken by the stimulus to move between these positions. Averaging across these, the resulting synthetic EEG response simulates the neural response to a localiser stimulus moving at 1 cycle per second around the display, without the presence of actual coherent motion. This allowed us

to create a control condition containing autocorrelated stimulus information but no predictive dynamics (since the constituent signals were evoked by individual localisers). Analysing these control data exactly as we did for the true motion trials (Figure 4B), we can see that the high probability region is centred along the 'no-compensation' line, demonstrating that autocorrelation cannot account for the observed spatio-temporal shift.

To statistically test for a difference in the slopes of the peak-position estimates from the actual and synthetic datasets, we fit a linear regression model to the group-level data. We predicted the peak position estimate from training timepoint and data type (actual vs synthetic) as well as their interaction. Crucially, the interaction term was significant (β = -0.018, std err = .002 t = -7.48, $p < .001$), indicating that the slopes of the fitted lines were different for the actual and synthetic datasets. Overall, these results suggest that the encoded position of the stimulus is shifted forwards after exposure to predictable motion, with this shift growing for later emerging neural representations. This is indicative of sustained, progressive extrapolation during cortical visual processing, resulting in the accelerated emergence of position representations and a gradual accumulation of position shifts along the processing hierarchy.

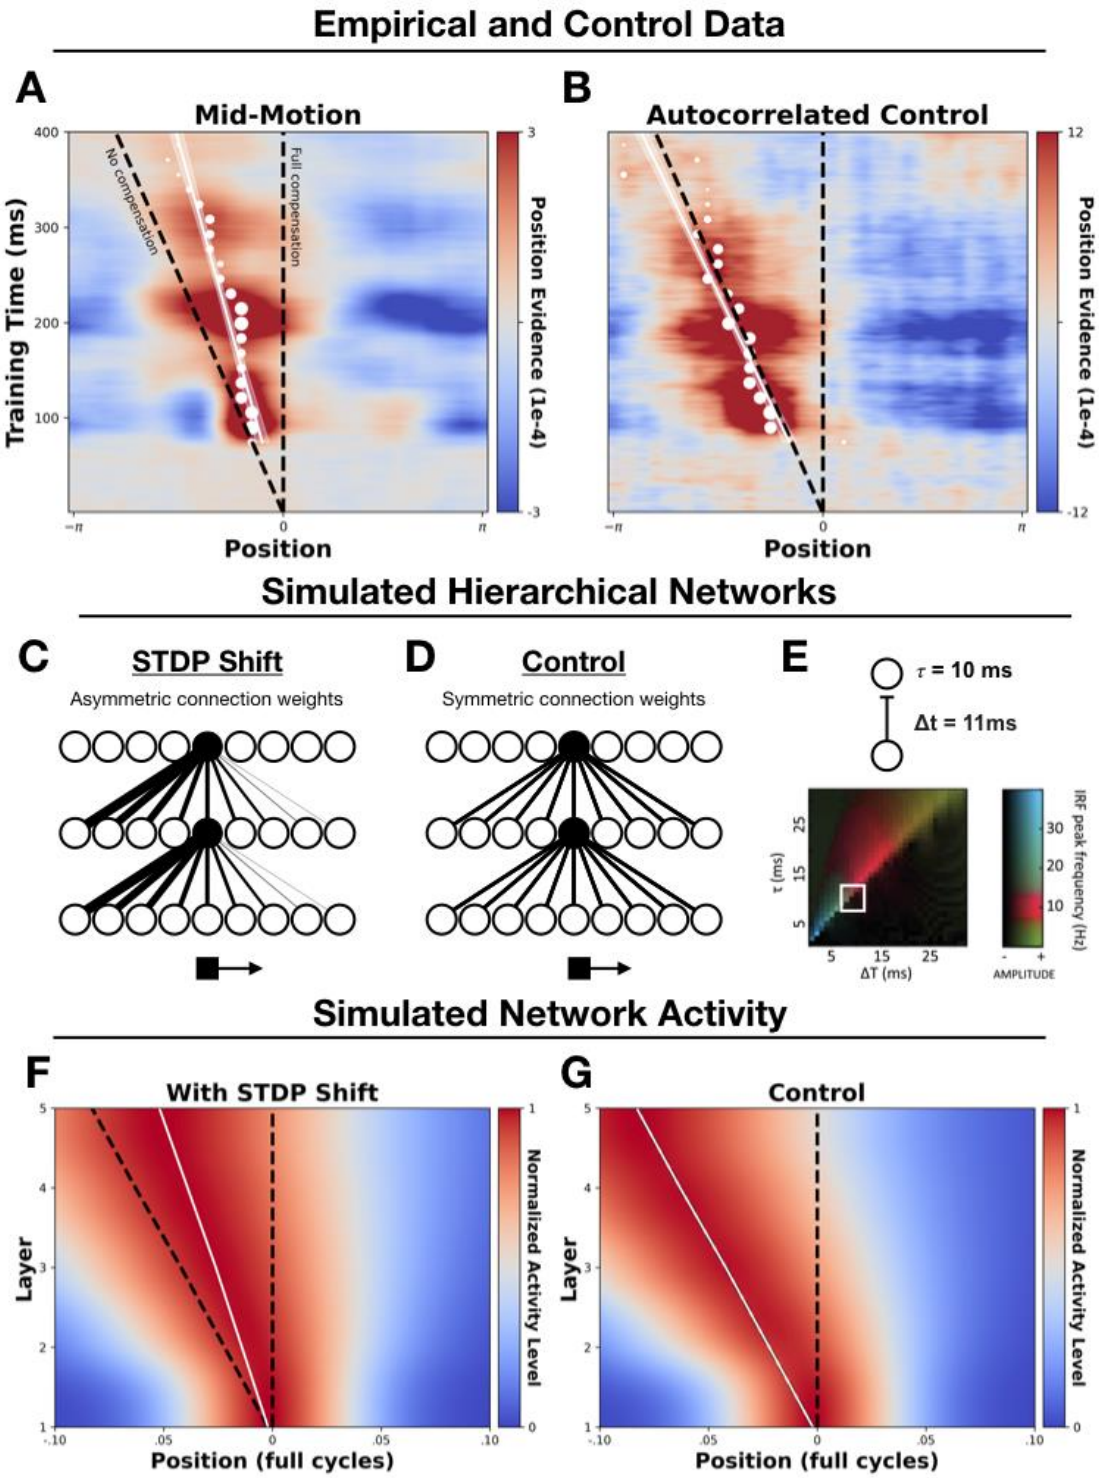

Figure 4. Decoding position information across distinct processing stages. Panels **A) – B)** show probability distributions over spatial positions (x-axis) extracted from timepoint specific decoders (y-axis). White dots show the points of peak probability (in 15 ms steps), with their size being proportional to the size of the peak estimate, with overlaid regression lines and standard errors. In all panels the vertical black dashed line marks the real-time position of the stimulus ('Full Compensation'), and the diagonal dashed line marks the delayed position of the stimulus ('No Compensation'). Panels **C) – D)** show schematic depictions of the simulated STDP and control networks. In the STDP network, the receptive fields of neurons (example neurons are highlighted in black) shift in the direction opposite

to motion after STDP-driven learning, allowing neurons to effectively 'anticipate' the arrival of a moving stimulus. In the control model, no shift occurs and receptive fields are symmetrical. Panel **E)** shows the between-layer signalling delay (Δt) and synaptic time constant (τ) used in the simulations. Also shown is Figure 1D from[26] which demonstrates that hierarchical predictive networks with delays in this approximate range (boxed white region) have oscillatory impulse response functions (IRFs) in the alpha/low-beta range (the range in which we found position-specific information could be best decoded). Panels **F-G)** show simulated activity from the STDP and control networks, smoothed for visualisation purposes. Plotting conventions are the same as in panels **A-B)**.

Accounting for the observed dynamics in a hierarchical network via STDP

What might drive progressive extrapolation during cortical visual processing? Below, we provide a minimal biologically plausible model that captures the observed dynamics.

Given that we observed shifts across multiple processing stages (as indexed by timepoint-specific decoders), we implemented a general learning mechanism that is known to exist across cortical regions: spike-timing dependent plasticity (STDP). This is a form of synaptic plasticity whereby synapses are strengthened when presynaptic cells fire shortly before a postsynaptic action potential, and weakened when they fire shortly after[27,28]. Recent research has shown how this simple associative learning mechanism can drive motion extrapolation as activity passes along the visual hierarchy[29–31]. Applying STDP to feedforward connections spontaneously produces an asymmetrical connectivity pattern, whereby the receptive fields of downstream neurons shift in the opposite direction to motion (see Figure 4C). This allows these neurons to 'anticipate' the arrival of a stimulus that is about to enter their (original) receptive field, driving a forwards shift in the population-level activity distribution (see[2] for review).

We simulated a 5-layer network with transmission delays, with feedforward connectivity profiles subject to the STDP-driven receptive field shifts reported by Sexton and colleagues[31]. The network comprised of 21 subpopulations of neurons tuned to velocities between -2 and 2 cycles/s (where negative velocities indicate counter-clockwise motion). Each level of the network comprised 1000 neurons, with 11 ms inter-layer signalling delays and a synaptic time constant of 10 ms. Crucially, hierarchical predictive networks with delays in this general range have been found to produce activity which oscillates in the alpha/low-beta range (see[26] Figure 1D, reproduced here in Figure 4E); the same range in which we found stimulus-position information could be best decoded (see Figure 1F). We simulated firing rates across the network in response to a stimulus traversing a circular trajectory at 1 cycle/s (see the Method). We also included an initial period of velocity estimation, starting at the onset of motion, in which information about the stimulus velocity is integrated. During the earliest timepoints, activity across the velocity subpopulations is widespread, but quickly becomes centred on populations with tuning at or close to the actual stimulus velocity (see Method and Supplement 1).

Figure 4F-G shows maps of the activity across two simulated networks: one in which STDP-driven learning has occurred (Figure 4F) and an otherwise identical control network in which STDP-driven learning has not occurred (Figure 4G). To compare the activity of these networks to the EEG results, we computed the average activity across all neural subpopulations per level and timepoint – in effect extracting a macroscopic 'neural image' of the stimulus at each level of the network. Re-centering these activity profiles and averaging over time (as in Figure 4A&B) allowed us to then compare the position of the peak population-level response, relative to the No-Compensation and the Full-Compensation lines, in the same

way as for the EEG data. Following STDP-driven learning, the activity evoked by a moving stimulus is shifted forwards off the no-compensation line (Figure 4F). Conversely, in the control model (Figure 4G) no shift occurs, with activity centred on this line.

This simulation serves to demonstrate that STDP-driven learning, and the resultant asymmetries in the receptive fields of hierarchically organised neurons, is a biologically plausible and minimally sufficient mechanism that can generate the accelerated emergence and forwards shifting of stimulus-position representations which we observed. Strikingly, these latency shifts occur entirely unsupervised. Through a known organisation principle (velocity tuned sub-populations; see[15]) and local synaptic learning rule (STDP[27,32]), these dynamics emerge at all levels of the processing hierarchy, with a gradual accumulation of position shifts across levels (i.e. progressive extrapolation), mirroring what we observed in the EEG recordings of the human observers. (For an additional exploratory analysis of the timecourse/emergence of this shift see Supplement 2).

Discussion

We have shown that probabilistic maps of the position of a moving stimulus can be generated from EEG recordings. We first reconstructed the trajectory of moving stimuli from early visual responses and found evidence of predictive overshoots following unexpected trajectory changes. Then, after training classifiers on the evolving cascade of neural activity patterns which follow the onset of static localiser stimuli, we found that the same activity patterns are triggered by moving stimuli (evidenced by successful cross-generalization), but with a clear shift in their timing. Specifically, we observed an accelerated emergence of object-position representations, corresponding to a forwards shift in the neurally encoded position of the moving stimulus at higher levels of processing. As a simple, biologically plausible model of this progressive shift, we demonstrated that these dynamics emerge at all levels of a simulated hierarchical neural network via spike-timing-dependent plasticity (STDP).

To our knowledge, this study is the first to provide evidence of progressive, cumulative motion extrapolation during later-stage (cortical) visual processing. Evidence of neural motion extrapolation in early visual regions has previously been reported in both human and non-human animals (e.g., [3,9,10,17,18]). However, an important outstanding question has been whether motion extrapolation is a widespread, multi-level phenomenon, which continues beyond the earliest stages of processing. Indeed, while pre-cortical extrapolation mechanisms have been well-characterised[2], it has remained unclear whether cortical regions simply inherit their extrapolated activity profiles from these upstream pre-cortical regions. Our observation of a progressive shift in the encoded position of smoothly moving objects, across distinct processing stages, provides a clear answer to this question, indicating that continued extrapolation of object-position information in the cortex does occur.

A recent paper from our own lab[17] also reported evidence of extrapolation in the early visual response, but did not observe any later-stage extrapolation. However, there are several key differences in the experimental and analysis approaches across these studies that may explain this difference. Firstly, whereas the stimuli in Johnson et al. (2023) moved linearly across a hexagonal grid, our stimuli moved along a circular trajectory at constant eccentricity, providing greater trial numbers and an improved signal-to-noise ratio. Secondly, stimuli in the present study moved much more quickly (~48 dva/s at its inner edge) than in Johnson et al. (2023; ~10 dva/s), and so potentially activated a largely distinct population of neurons. Finally, whereas Johnson and colleagues inspected the evolution of neural codes over time for

411 evidence of extrapolation, here we were able to extract probability maps across space (i.e.
412 akin to the 'neural image' of the stimulus, see[7]), allowing us to directly estimate spatial
413 extrapolation without additional curve-fitting steps. Altogether, by fixing the eccentricity of
414 the stimulus and collecting more repetitions of the same trajectory, we significantly improved
415 the signal-to-noise ratio of the decoding, enabling us to adopt a more direct decoding
416 approach and generate clear position maps (with a 5-fold increase in precision relative to our
417 own recent attempts[16]). Given that the temporal generalization analysis provided evidence of
418 distinct processing stages unfolding over a longer time scale, we were able to consider shifts
419 in position representations up until 400 ms post stimulus onset (as compared to 150 ms[17]).
420 Overall, this revealed a clear, progressive shift in the timing of positional representations
421 across temporally distinct processing stages.
422 One salient feature of the current results is that, while we find evidence of an
423 extrapolative temporal shift, this shift is not complete. That is, peak positional probabilities
424 never align with the real-time position of the stimulus (although centroid estimates do align
425 with real-time in the early visual response, see Figure 3 and Supplement 1). As we have
426 discussed elsewhere[2], it is difficult to tell whether partial shifts such as these are due to
427 incomplete extrapolation, or rather to the fact that EEG recordings necessarily tap into a
428 mixture of signals from different cortical regions (some of which may be fully extrapolated,
429 and some may not be extrapolated at all). In the latter case, only signals involved in time
430 sensitive localisation may be fully extrapolated. Ultimately, distinguishing between these
431 possibilities will likely require the use of invasive recording procedures, where activity can be
432 isolated to precise neural populations. What the current pattern of results does tell us is that
433 not all position-related activity is fully extrapolated. This is sensible, as many visually-evoked
434 signals will carry auxiliary position information in addition to the actual feature they encode
435 (since they arise from a retinotopically specific location). Because extrapolation comes with a
436 cost (i.e. extrapolated activity must be ignored/corrected when expectations are violated), a
437 more efficient strategy may simply be to bind featural and position information
438 retrospectively.
439 Central to the question of whether partial or full extrapolation is achieved is an
440 understanding of how the position of an object is actually read out from visual regions by
441 downstream/effector areas. Addressing this is beyond the scope of the current study, however
442 the present findings raise an interesting hypothesis. We have shown that a probability
443 distribution over possible location can be linearly decoded from early brain activity patterns,
444 and that the shape of this distribution changes after exposure to motion – shifting towards
445 the real-time stimulus position. Considering the observed changes, a hypothesis which
446 emerges is that the shape of this distribution is modulated in such a way as to allow different
447 point estimates of a moving object's location to be read out, depending on current task
448 demands. Specifically, if accurate real-time interaction is paramount, then an estimate of the
449 stimulus' real-time position may be derived from the centroid of the distribution. However,
450 when the stimulus disappears or changes trajectory this can overshoot and be unreliable. As
451 such, when real-time interaction is less important, the peak of the distribution may be taken,
452 yielding a more stable, but delayed, estimate of the stimulus' position (i.e. it reliably tells you
453 where the stimulus just was). In other words, at any given moment, different estimators may
454 be used, either to determine where a stimulus just was with high-reliability (peak), or where
455 a stimulus probably now is, with lower reliability (centroid). To properly test this idea, future
456 studies will need to vary the speed of the moving stimulus. The central test will be whether
457 the centroid shifts to align with the real-time position of the stimulus regardless of its speed,

serving as an effective real-time estimate. If this occurs, then future studies would also need to determine whether/how downstream regions can extract summary measures (such as a centroid) of population-level activity in early visual regions, and ultimately how these relate to behaviour (such as reaching or related targeting actions).

By simulating the activity of a hierarchically-organised network of neurons, we showed how STDP-driven learning can drive the accelerated emergence of object-position representations, in an unsupervised fashion. Our main aim in presenting this result is to show how a simple cortical mechanism can in principle drive the sustained, progressive extrapolation we observed. Indeed, given the existence of velocity-tuned neural subpopulations, and the ubiquity of STDP-based learning throughout the cortex, the simulations provide compelling support for our empirical findings, as they suggest that something additional would be needed to *prevent* such temporal shifts from occurring. Crucially, these shifts will only occur for sequences of input to which we are frequently exposed (like smooth motion), and will also depend on the specific plasticity of the underlying sensory region.

Importantly, in presenting the network simulations, our intention is not to claim that STDP is the sole drive of neural motion extrapolation. In prior work a variety of other extrapolation mechanisms have been well catalogued, with many operating during pre-cortical visual processing (see[2]). These can drive forwards shifts in the evoked population-level distribution of neural activity, and may well have contributed to the motion-induced shape changes we observed in the probability distributions decoded from the early visual response (i.e. < 150 ms). As such, we caution against taking the STDP network model as a complete model of neural extrapolation. Nevertheless, given its generality, we feel that STDP is an important candidate mechanism to consider – especially when accounting for widespread, progressive effects such as those we observed. Indeed, given the ubiquity of STDP, an interesting avenue of future research may involve examining whether sustained temporal shifts in the pattern of evoked neural responses can be found after ongoing exposure to 'motion' through more abstract feature space, such as colour, luminance, or numerosity. Finally, a remarkable feature of the current model (and[30,31]) is that spontaneous motion extrapolation is achieved during purely feed-forward processing. However, future studies should consider further developing these models, building in recurrent and/or horizontal connections. These may act as putative mechanisms for fine-tuning the magnitude of STDP-driven extrapolation in a delay-dependent fashion (i.e. calibrating the degree of shift to a given delay).

In sum, we have shown how precise probabilistic maps of the position of a moving object can be generated from EEG recordings. Using this approach, we have provided clear evidence of progressive neural motion extrapolation occurring during visual processing in the human brain. Most strikingly, we have shown that after ongoing exposure to smooth motion, there is an accelerated emergence of position-specific activity patterns across distinct processing stages, corresponding to a forwards shift in the neurally encoded position of the moving stimulus. This provides the first clear evidence of cortical neural motion extrapolation in the human visual system. Finally, we have shown how these dynamics would be expected to emerge spontaneously, without supervision, at all levels of a hierarchical neural network via spike-timing-dependent plasticity (STDP) – providing a mechanism for widespread neural extrapolation/delay compensation.

Method

Participants

Eighteen observers (15 female, 18-35 years old with mean age of 23 years) participated in the experiment. Each observer completed 2 sessions across separate days. All had normal or corrected-to-normal vision, gave written informed consent at the beginning of each session, and were reimbursed AUD15 per hour. The experimental protocol was approved by the human research ethics committee of the University of Melbourne (Reference Number: 2021-12985-16726-4).

Stimuli

Stimuli were generated using the Psychophysics Toolbox (Brainard, 1997) in MATLAB 2016a (Mathworks). Stimuli were presented on an ASUS ROG PG258 monitor with a resolution of 1,920 × 1,080 running at 120 Hz. The stimulus was a white, truncated wedge presented on a uniform grey background (Figure 1). The inner and outer edges of the wedge were 7.7 degrees of visual angle (dva) and 9.4 dva away from fixation. The wedge covered 9° of polar angle with 1.19 dva at the inner and 1.46 dva at the outer edge. During localiser trials the stimulus could appear in one of 40 locations tiling an invisible circle centred on the fixation point (see Figure 1A). Localiser stimuli were presented for 100 ms, with an interstimulus interval of 100 ms (i.e. onset rate of 5 Hz). Smooth motion sequences began and ended in randomly determined localiser positions, with sequences separated by an interval of 500 ms. The smoothly moving stimulus had a velocity of 360° of polar angle per second (i.e. a 3° offset per frame).

Task

Participants viewed the stimuli while EEG was recorded. In the localiser block, participants viewed the stimulus being randomly presented in 40 equally spaced positions around fixation (50 repetitions per position, per session). In the smooth motion block, participants reviewed the same stimulus moving for at least 1.5 seconds before either disappearing or reversing direction at a randomly determined localiser position. Following reversals the stimulus continued to move for 0.5-1 seconds. Participants viewed 960 motion sequences per session (12 repetitions per position and motion direction). Participants were given a self-paced break halfway through the localisers and 5 self-paced breaks during the smooth motion sequences. To maintain participants' attention, they were tasked with pressing a button whenever the stimulus changed from white to purple. This occurred 20 times during the localiser block and 50 times during the smooth motion block. Neural responses to these 'catch' stimuli were not analysed. The order in which participants viewed the localiser and smooth motion blocks was randomised in each session.

EEG acquisition and pre-processing

64-channel EEG data, as well as data from six EOG electrodes (placed above, below, and next to the outer canthi of each eye) and two mastoid electrodes, 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. Bad channels noted during

data collection (mean of 1 per session, max of 3) were spherically interpolated using the MNE 'interpolate_bads' function[33].

For the support vector regression (SVR)-based decoding analyses (used to initially characterise the neural encoding of position-specific information) we extracted epochs of data (-200 to 1000 ms) relative to localiser onset. These were baseline corrected to the mean of the 200 ms period prior to stimulus onset. For the LDA-based mapping analyses, we draw a distinction between training and testing epochs. Training epochs were extracted from -200 to 400 ms relative to localiser onset, and were baseline corrected (-200 to 0 ms prior to stimulus onset). Testing epochs were extract from the smooth motion sequences (-1000 to 1000 ms) relative to the events of interest (onset, offset, reversal), and were baseline corrected to the mean of the preceding 1000 ms period (i.e. one full cycle of motion). For the time-frequency based decoding analyses (see below), power estimates were extracted at 20 linearly spaced frequencies between 2 and 40 Hz using the tfr_morlet function in MNE, with the number of cycles, which are used to define the width of the wavelet's Gaussian window (n_cycles/2*pi*f), logarithmically increasing from 3 to 10 across frequencies.

Decoding analyses

Principal component analysis was applied before decoding to capture 99% of the variance (transformation computed on training data and applied to testing data), to help de-noise the data[34].

SVR-based analyses

To initially characterise the neural encoding of stimulus-position information we trained multivariate linear models (support vector regression with L2 loss) to predict the position of the localiser stimuli from participants' neural activity patterns (following[21]). Specifically, we trained models to predict the sine and cosine of the angular position of the stimulus. A custom scoring function was used to calculate the fractional inverse absolute angular error ('decoding score') between the predicted and actual position of the stimulus:

$$\text{Decoding score} = ((\pi/2) - (\frac{1}{n}\sum_{j=1}^{n} abs\left(\arg\left(\frac{e^{i*\hat{y}_j}}{e^{i*y_j}}\right)\right)))/\pi \qquad (1)$$

where $\hat{y}_j$ is the predicted angular position on trial j, and y_j is the actual angular position of the stimulus on trial j. This is designed such that a score of 0 indicates chance performance and a score of 1 indicates perfect accuracy. Custom five-fold cross-validation was used to evaluate out-of-sample prediction accuracy ensuring no leakage between test and training sets.

Temporal generalization analysis was conducted by examining how well models trained on neural activity patterns at one specific timepoint could predict the position of stimuli based on data from other time points (see[22]). A searchlight analysis across the scalp was conducted by running the decoding analysis using data from a single electrode plus its immediate neighbours (following[35]). The spectral locus of position information was examined by re-running the decoding analysis on frequency-specific normalized power estimates (i.e. the relative pattern of oscillatory power across electrodes in a given frequency band, see [16]). Finally, to examine how successive stimuli were encoded (following[23,24]) we re-ran the temporal generalization analysis after splitting the localiser data in half. From one half we

extracted training epochs (0 to 500 ms) relative to the onset of each localiser. From the other half, we extracted testing epochs (-200 to 1300 ms) relative to every 5th localiser stimulus. We then iteratively predicted the spatial location of these stimuli followed by the locations of the four subsequent stimuli. Finally, we re-ran the analysis after switching training and testing sets.

LDA-based mapping

To map the location of the smoothly moving stimuli, we trained multiclass linear discriminant analysis (LDA) classifiers on individual participants' neural responses to the localiser stimuli, treating each position as a distinct class. We then extracted predicted posterior class probabilities from these pre-trained LDA models, based off participants neural responses to the smoothly moving stimuli. Averaging the posterior probabilities across models trained 75-125 ms after localiser onset yielded a single matrix of probabilities (i.e. a spatio-temporal probabilistic map of the stimulus' position over time). To calculate the centroid at a given time point we took a weighted vector average:

$$\bar{a} = \frac{1}{n} * \sum_{j=1}^{n} w_j * \exp(i, a_j) \quad (2)$$

where j indexes position, w_j is the posterior probability at position j, i is the imaginary operator, and a_j is the polar angle of position j (in radians). The centroid was then taken as the argument of $\bar{a}$.

Statistical analyses

For the SVR-based decoding, we performed cluster-corrected one sample t-tests against zero using a critical alpha level of .01 and a cluster forming threshold of .01 (via the MNE function 'spatio_temporal_cluster_1samp_test', with 2^{12} permutations).

Neural network modelling

To simulate representations of stimulus position at different stages of cortical processing, we used a hierarchical network similar to that described by Sexton and colleagues[31] (see[30] for a two-layer implementation). The network consists of N_l layers, each comprising N_v velocity-tuned subpopulations. Within each layer and velocity subset, there are N_n neurons with spatial tunings distributed across a circular interval [0, 1], each receiving feedforward activity from neurons in the layer below (excepting the first layer, which encodes the stimulus position). The feedforward weights *W* for each neuron follow a Gaussian distribution with width σ_w and mean μ_w:

$$\mu_w = x_i^l + \beta STDP_{vl} \quad (3)$$

where x_i^l is the spatial position of the neuron *i* at layer *l*, β is the STDP shift scaling parameter and $STDP_{vl}$ is the magnitude of the receptive field shift for the velocity *v* and layer *l*, as taken from[31] (see *Estimating STDP-driven receptive field shift magnitudes)*. The addition of the $STDP_{vl}$ term in Eq. (3) means that the distribution of feedforward activity does not remain symmetric as it progresses through each layer, but is shifted in line with the receptive field

shift magnitude. We included the scaling parameter β in order to find STDP shift magnitudes that best fit the EEG data (see *Fitting procedure).* Firing rates are recorded during a simulation carried out across N_t timepoints, during which a point stimulus traverses a circular trajectory at 1 cycle/s. The stimulus position is encoded at each timestep t by the firing rate of neurons in the first (input) layer, $r_v^1(t)$, described by a Gaussian distribution centred on the stimulus position and with width σ_p. The input layer neurons have a baseline firing rate r_b, with stimulus magnitude equal to Sr_b. Unlike in Sexton et al. (2023), no spikes are generated: only firing rates are transferred between layers, subject to a transmission delay t_{delay}.

Firing rates at each higher layer are based upon (delayed) input from the layer below:

$$r_v^l(t) = r_v^l(t-\Delta t)e^{-\frac{\Delta t}{\tau_m}} + \frac{\Delta t}{\tau_m} I_v^l(t) \quad (4)$$

where τ_m is the passive membrane time constant, Δt is the length of the timestep used, and $I_v^l(t)$ is the delayed input to the velocity subpopulation at this layer:

$$I_j^l(t) = W_v^{l-1} \cdot r_v^{l-1}\left(t - t_{\text{delay}}\right) \quad (5)$$

Following the simulation, a global estimate of the represented stimulus position is generated for each timepoint and layer, by taking a weighted average of firing rate distributions across all velocity subpopulations. The weights for each velocity are time-dependent: at motion onset activity is widespread across all velocity subpopulations, before becoming primarily dominated by the neural activity tuned to the stimulus velocity. (Note: this assumption is for the purposes of the analyses presented in Supplement 2, and the results of the main analysis do not hinge upon it.)

Specifically, we defined a range of timepoints [0, T_i] in which information about the stimulus velocity is integrated. Starting at t=0, weights across velocity subpopulations are generated according to a Gaussian distribution centred on v=0, with width σ_g(t=0). At each subsequent timestep, the mean of the Gaussian, μ_g(t), is shifted in the direction of the true stimulus velocity by an interval such that μ_g(t=T_i) is centred on the true stimulus velocity. Likewise, σ_g(t) is decreased incrementally across each of the integration timesteps (see Figure 5). Because information about stimulus velocity is also subject to transmission delays, the change in weights thus described is applied to each layer in a delayed manner. The weighted average of firing rates across all velocity subpopulations is calculated for each layer and timepoint, then normalized, to generate a final estimate of the global position representation at each layer and timepoint.

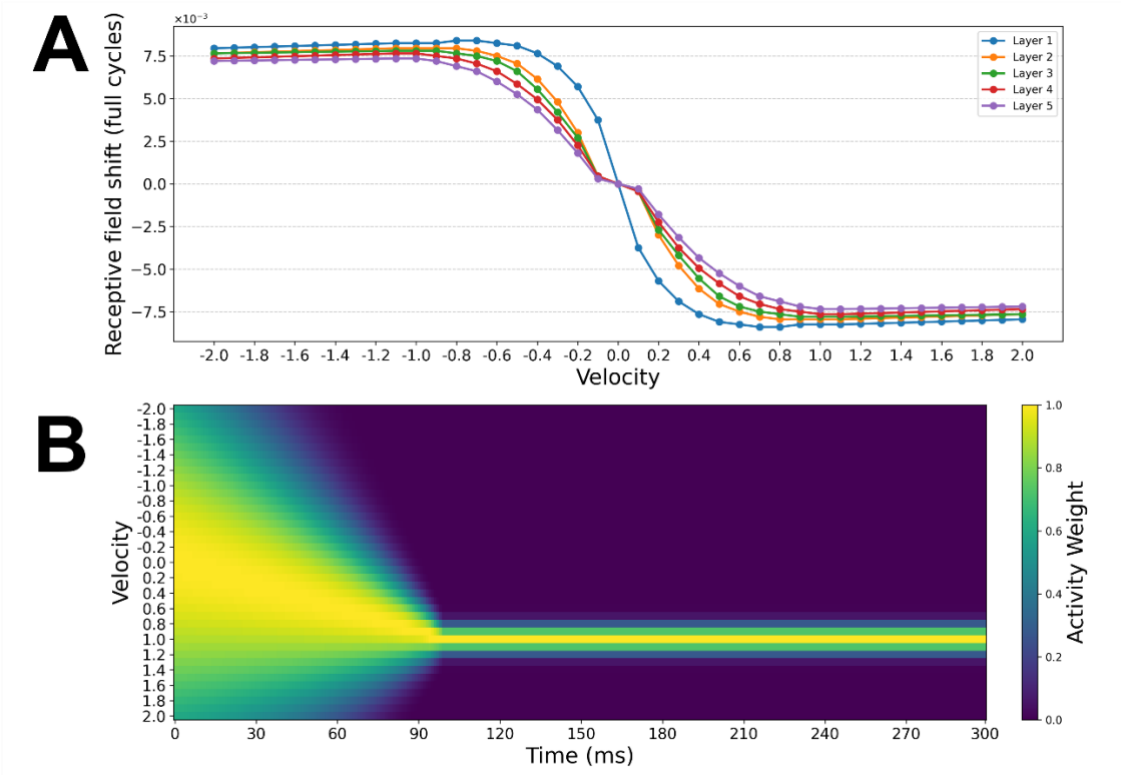

Figure 5. Contribution of velocity-specific receptive field shifts to global position representations. A) Magnitude of STDP-driven receptive field shifts across the velocity range -2 to 2 cycles/s, estimated based on the shifts reported by Sexton and colleagues[31] (see *Estimating STDP-driven receptive field shift magnitudes)*. Individual lines show STDP shift magnitudes for each layer of the network. **B)** Temporal evolution of activity weights for each velocity subpopulation during the simulation. At onset the activity is broadly tuned around 0 cycles/s, then shifts toward the true stimulus velocity (1 cycle/s) during the velocity integration period. Following the initial integration period, activity weights remain stable and centred on the true stimulus velocity.

The precise parameter values used in the simulation are shown in Table 1, and were taken from Sexton and colleagues[31] where applicable.

Table 1

Parameter Values Used in Numerical Simulation

Name	Value	Description
N_l	5	Number of layers
N_n	1000	Number of neurons per layer
N_v	21	Number of velocities
σ_w	1/32	Standard deviation of anatomical connections
σ_p	1/32	Standard deviation of stimulus input
r_b	5 Hz	Baseline firing rate at input layer
S	20	Stimulus intensity (a.u.)
Δt	1 ms	Timestep length
N_t	300	Total simulation timesteps
T_i	100	Velocity integration timesteps

τ_m	10 ms	Membrane time constant
t_{delay}	11 ms	Transmission delay
β	.6	STDP shift scaling parameter
$\sigma_g(t=0)$	2	Standard deviation of activity weights at onset
$\sigma_g(t=Ti)$	1/8	Standard deviation of activity weights at T_i

Estimating STDP-driven receptive field shift magnitudes

The STDP shift magnitudes estimated by Sexton and colleagues[31] are reported in the range 0.1 to 1 cycles/s, as well as 2, 3, 4 and 5 cycles/s. We wished to include a range of velocities that was symmetric around zero, and for which the EEG experiment velocity used (1 cycles/s) was an intermediate value, to avoid any edge effects during averaging. Therefore, the values reported by Sexton and colleagues[31] were extended by linearly interpolating between the values given for 1 and 2 cycles/s, and then inverting all positive values to generate a mirrored set of velocities tuned to the opposite directions. This allowed us to generate an estimate of STDP magnitudes for velocities in the range from -2 to 2 cycle/s, in 0.1 cycle/s increments (see Figure 5).

Fitting procedure

To compare the simulated network to the position representations decoded from EEG, we performed the recentering of the global position representations such that the real-time position of the stimulus is always centred on the midpoint, with lagged representations indicated by activity at any locations counter-clockwise of this central position. The resulting map plots position relative to stimulus on the horizontal axis and network depth (layer number) on the vertical axis.

As with the EEG data, we can define two lines within this rotated map of positional representations: 1) a diagonal 'No Compensation' line, connecting the points at each layer/training timepoint where the stimulus would be represented in the absence of any extrapolation, given the stimulus velocity, temporal lag between layers and the compounding effect of the synaptic time constant on peak activity, and 2) a vertical 'Full Compensation' line representing the real-time stimulus position (perfect delay compensation/extrapolation). The degree of extrapolation can therefore be quantified by measuring where peak activity sits as a ratio of the difference between the No-Compensation and Full-Compensation lines (Figure 6A). This ratio measurement provides us with a simple, straightforward way of comparing the degree to which positional representations are shifted in the EEG and simulated networks.

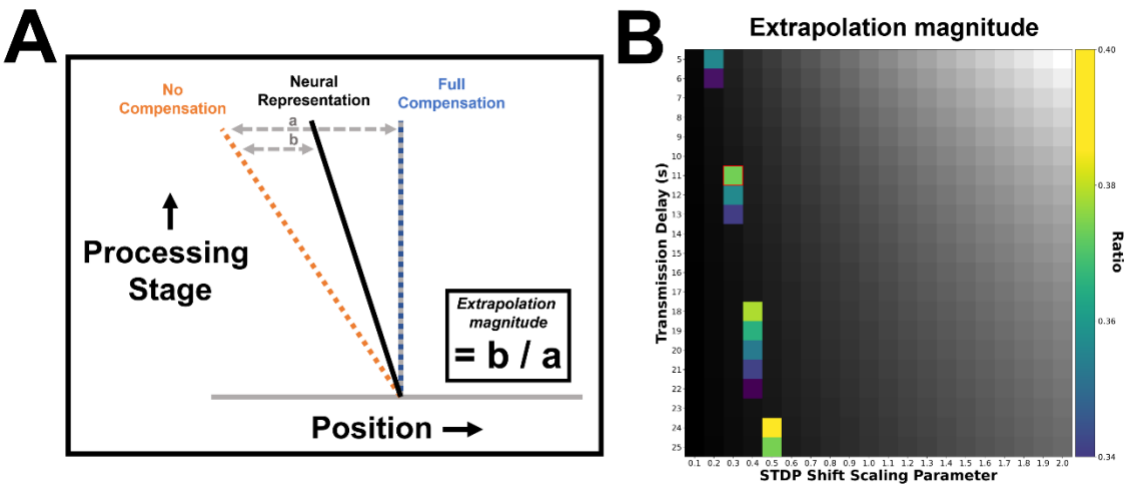

Figure 6. Calculating extrapolation magnitude. A) Neural activity in the rotated position representation maps can be characterised with reference to the 'No Compensation' line, indicating where the stimulus would be represented in the absence of any extrapolatory mechanisms, and the 'Full Compensation' line, indicating the current stimulus position and therefore where the stimulus would be represented in the absence of delays. The extrapolation magnitude is quantified by the actual neural representation (determined by the peak position of stimulus-driven activity), as a ratio of the distance between the No-Compensation and Full-Compensation lines. A single measurement of extrapolation magnitude was made at the final timepoint of the stimulation, in order to allow enough time for activity to propagate fully throughout the network and for the peak of neural activity (relative to the stimulus) to stabilise. **B)** Extrapolation magnitudes measured while varying the transmission delay and the STDP shift magnitude scaling parameter (β). Values within the range measured in the EEG analysis (.34 to .40) are colour coded. Values outside this range are shown in grayscale. The point in this parameter space used in the main analyses (Figure 4) is highlighted in red.

In the simulated network, the extrapolation magnitude for a given velocity is largely determined by the transmission delay and time constant (affecting primarily the position of the No-Compensation line, relative to the Full-Compensation line), and the magnitude of the STDP shifts (affecting primarily the position of the neural representation line, relative to the No-Compensation line). We ran the simulation while varying both transmission delay ($t_{delay;}$ 5 to 25ms) and STDP shift magnitude (via the scaling factor β; .1 to 2 times the values reported in Sexton and colleagues[31]), to find parameter values which led to an extrapolation magnitude that best fit the ratio derived from the EEG data (.36). The extrapolation magnitude measurement was made based on activity recorded at the final timepoint of the simulation, and at 400 ms for the EEG data. Figure 6B shows the values of delay time and STDP shift scaling parameter which produced extrapolation magnitudes most closely matching the EEG data. In line with Alamia and VanRullen[26], we ultimately constrained the network to have an 11ms transmission delay (such that with recurrent connectivity it would, in principle, generate oscillatory activity in the alpha/low-beta range) and took the best fitting value of β within this row, for use in the subsequent analyses (t_{delay} = 11 ms, β = .3).

Data and Code Availability

EEG analysis code is available at: https://github.com/bootstrapbill/neural-location-decoding. The pre-processed EEG data files and neural network simulation code will be publicly available by the time of publication at: https://osf.io/sn4a7/.

References

1. Lamme, V. A. F. & Roelfsema, P. R. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci.* **23**, 571–579 (2000).
2. Turner, W., Sexton, C. & Hogendoorn, H. Neural mechanisms of visual motion extrapolation. *Neurosci. Biobehav. Rev.* **156**, 105484 (2024).
3. Berry, M. J., Brivanlou, I. H., Jordan, T. A. & Meister, M. Anticipation of moving stimuli by the retina. *Nature* **398**, (1999).
4. Trenholm, S., Schwab, D. J., Balasubramanian, V. & Awatramani, G. B. Lag normalization in an electrically coupled neural network. *Nat. Neurosci.* **16**, (2013).
5. Johnston, J. & Lagnado, L. General features of the retinal connectome determine the computation of motion anticipation. *eLife* 1–19 (2015) doi:10.7554/eLife.06250.
6. Liu, B., Hong, A., Rieke, F. & Manookin, M. B. Predictive encoding of motion begins in the primate retina. *Nat. Neurosci.* **24**, 1280–1290 (2021).
7. Schwartz, G., Taylor, S., Fisher, C., Harris, R. & Berry, M. J. Synchronized Firing among Retinal Ganglion Cells Signals Motion Reversal. *Neuron* **55**, 958–969 (2007).
8. Orban, G. A., Hoffmann, K. P. & Duysens, J. Velocity selectivity in the cat visual system. I. Responses of LGN cells to moving bar stimuli: a comparison with cortical areas 17 and 18. *J. Neurophysiol.* **54**, 1026–1049 (1985).
9. Jancke, D., Erlhagen, W., Schöner, G. & Dinse, H. R. Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex. *J. Physiol.* **556**, 971–982 (2004).
10. Benvenuti, G. *et al.* Anticipatory responses along motion trajectories in awake monkey area V1. *bioRxiv* 1–42 (2020).
11. Guo, K. *et al.* Spatio-temporal prediction and inference by V1 neurons. *Eur. J. Neurosci.* **26**, 1045–1054 (2007).
12. Subramaniyan, M. *et al.* Faster processing of moving compared with flashed bars in awake macaque V1 provides a neural correlate of the flash lag illusion. *J. Neurophysiol.* **120**, 2430–2452 (2018).
13. Sundberg, K. A., Fallah, M. & Reynolds, J. H. A Motion-Dependent Distortion of Retinotopy in Area V4. *Neuron* **49**, 447–457 (2006).
14. Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex : a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
15. Hogendoorn, H. & Burkitt, A. N. Predictive Coding with Neural Transmission Delays: A Real-Time Temporal Alignment Hypothesis. *eNeuro* **6**, 1–12 (2019).
16. Turner, W., Blom, T. & Hogendoorn, H. Visual Information Is Predictively Encoded in Occipital Alpha/Low-Beta Oscillations. *J. Neurosci.* **43**, 5537–5545 (2023).
17. Johnson, P. A., Blom, T. & van Gaal, S. Position representations of moving objects align with real-time position in the early visual response. *eLife* (2023).
18. Blom, T., Feuerriegel, D., Johnson, P., Bode, S. & Hogendoorn, H. Predictions drive neural representations of visual events ahead of incoming sensory information. *Proc. Natl. Acad. Sci.* **117**, 7510–7515 (2020).
19. de Vries, I. E. J. & Wurm, M. F. Predictive neural representations of naturalistic dynamic input. *Nat. Commun.* **14**, 3858 (2023).
20. Lee, C. S., Aly, M. & Baldassano, C. Anticipation of temporally structured events in the brain. *eLife* **10**, e64972 (2021).

21. King, J.-R., Pescetelli, N. & Dehaene, S. Brain Mechanisms Underlying the Brief Maintenance of Seen and Unseen Sensory Information. *Neuron* **92**, 1122–1134 (2016).
22. King, J.-R. & Dehaene, S. Characterizing the dynamics of mental representations: The temporal generalization method. *Trends Cogn. Sci.* (2014) doi:10.1016/j.tics.2014.01.002.
23. Gwilliams, L., King, J.-R., Marantz, A. & Poeppel, D. Neural dynamics of phoneme sequences reveal position-invariant code for content and order. *Nat. Commun.* **13**, 6606 (2022).
24. King, J.-R. & Wyart, V. The Human Brain Encodes a Chronicle of Visual Events at Each Instant of Time Through the Multiplexing of Traveling Waves. *J. Neurosci.* **41**, 7224–7233 (2021).
25. Gwilliams, L. & King, J.-R. Recurrent processes support a cascade of hierarchical decisions. *eLife* **9**, e56603 (2020).
26. Alamia, A. & VanRullen, R. Alpha oscillations and traveling waves: Signatures of predictive coding? *PLOS Biol.* **17**, e3000487 (2019).
27. Bi, G. & Poo, M. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *J. Neurosci.* **18**, 10464–10472 (1998).
28. Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).
29. Fu, Y.-X., Shen, Y., Gao, H. & Dan, Y. Asymmetry in Visual Cortical Circuits Underlying Motion-Induced Perceptual Mislocalization. *J. Neurosci.* **24**, 2165–2171 (2004).
30. Burkitt, A. N. & Hogendoorn, H. Predictive Visual Motion Extrapolation Emerges Spontaneously and without Supervision at Each Layer of a Hierarchical Neural Network with Spike-Timing-Dependent Plasticity. *J. Neurosci.* **41**, 4428–4438 (2021).
31. Sexton, C. M., Burkitt, A. N. & Hogendoorn, H. Spike-timing dependent plasticity partially compensates for neural delays in a multi-layered network of motion-sensitive neurons. *PLOS Comput. Biol.* **19**, e1011457 (2023).
32. Feldman, D. E. The Spike-Timing Dependence of Plasticity. *Neuron* **75**, 556–571 (2012).
33. Gramfort, A. *et al.* MNE software for processing MEG and EEG data. *NeuroImage* **86**, 446–460 (2014).
34. Grootswagers, T., Wardle, S. G. & Carlson, T. A. Decoding Dynamic Brain Patterns from Evoked Responses : A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *J. Cogn. Neurosci.* 677–697 (2017) doi:10.1162/jocn.
35. Hajonides, J. E., Nobre, A. C., van Ede, F. & Stokes, M. G. Decoding visual colour from scalp electroencephalography measurements. *NeuroImage* **237**, 118030 (2021).

Supplementary Information

Supplement 1: Replication of main latency shift using the distribution centroid as a discrete position estimate.

To ensure that the main latency shift effect is robust across specific analysis choices, we re-plotted Figure 4A&B overlaying the centroid (i.e. vector average), instead of the point of peak probability, as a discrete position estimate. Examining Figure S1, the same effect can be observed as we report in the main text, building confidence that this does not depend on the specific read-out method we choose to employ. In the main text we use the peak probability estimate as a more conservative read-out method, since this does not display the same extrapolative properties as the centroid during early processing (see Figure 3).

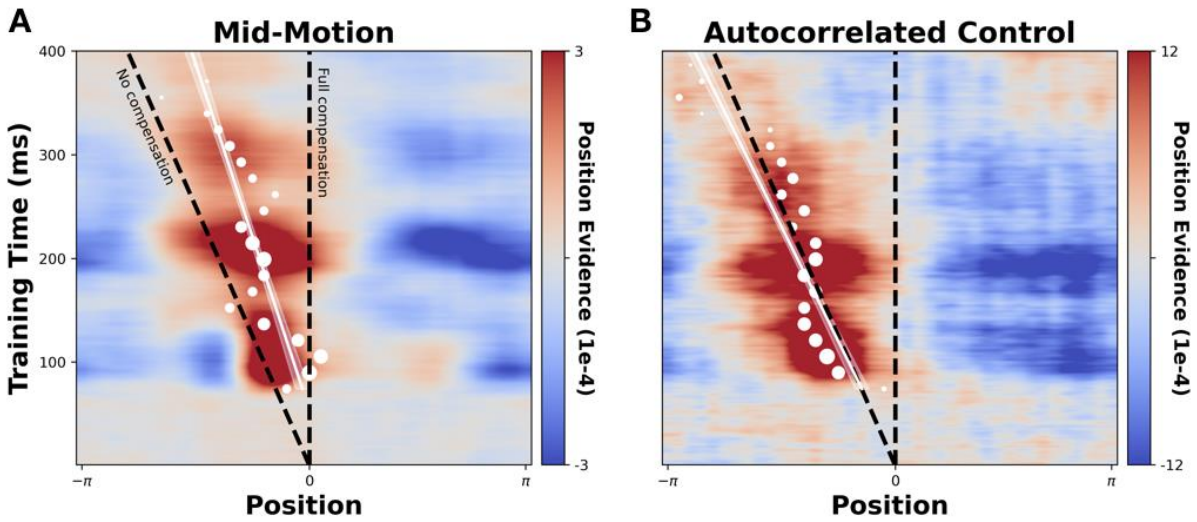

Figure S1: Recreation of Figure 4A&B using the centroid, rather than peak, probability estimate. All plotting conventions are the same as in Figure 4.

Supplement 2: Examining prediction formation in empirical and model position representations

In the main analysis we consider the 'steady state' of position representations formed after sustained exposure to motion along a predictable trajectory, and found evidence of progressive extrapolation. However, after the initial appearance of a moving object, position representations must necessarily lag, since motion extrapolation is only possible once the object's velocity has been established. Here, we conduct an additional exploratory analysis to examine how rapidly motion extrapolation arises when a moving object first appears, and how the temporal evolution of this effect may be accounted for in the STDP network model.

In Figure S2 (below) we compare how both decoded and simulated positional representations evolve over timepoints immediately following stimulus onset. Initially, the decoded maps generated from the raw and synthetic EEG data (top panels) are similar. However, from ~150 ms the bulk of the high probability region in the raw map begins to shift forwards, with only a small portion of activity left travelling diagonally along the No-Compensation line. No such shift occurs in the synthetic map, with activity remaining centred on the No-Compensation line. This indicates that it takes ~150 ms for the 'steady state' temporal shift which we observed after sustained exposure to smooth motion to emerge.

For the simulated maps, we can see that the same forwards shift occurs in the population-level activity of the trained STDP model, but not the untrained (control) model. In the trained model peak activity initially following the No-Compensation line, but then gradually shifts forwards across later timepoints. This occurs because of the velocity estimation process, in which the individual weights for each velocity subpopulation change as a function of time. At the earliest timepoints, all subpopulations are assumed to be active. Taking the average of each of these yields a global representation in which the constituent STDP shifts are effectively cancelled out (given that the range of velocities is symmetric around zero). As the information about the stimulus velocity is integrated, the global response gradually becomes dominated by the subpopulation tuned to 1 cycle/s, and the represented position shifts according to the STDP shift magnitude associated with this velocity. This shift in the global position representation occurs at each layer in a time delayed fashion, causing the angle of the line tracking activity peaks to change across layers during the intermediate timepoints. At the later timepoints, all layers have received full information about the stimulus velocity, resulting in a straight line connecting the activity peaks.

While this qualitatively captures the dynamics in the EEG-derived maps, it is important to emphasis the speculative nature of this additional analysis and assumed velocity estimation mechanism. These hinge on the assumption that all velocity-tuned subpopulations are activated by the onset of the stimulus (i.e. the onset transient), with the population tuned to the true velocity eventually winning out (either due to direct competition, or a passive fading of onset evoked activity). However, it is also possible that no velocity-tuned populations are initially activated by stimulus onset, and that activity simply gradually builds up in the population tuned to the true velocity. Arbitrating between these two possibilities remains an interesting avenue for future research that will likely require the use of direct intracranial recording techniques. Crucially, the assumed velocity estimation process only alters the early temporal dynamics of the network, and has no influence on its ultimate 'steady state' behaviour. As such the main simulations are independent of this specific component/assumption, and only rely on the reasonable assumption that neurons tuned to a specific velocity are active after sustained exposure to motion.

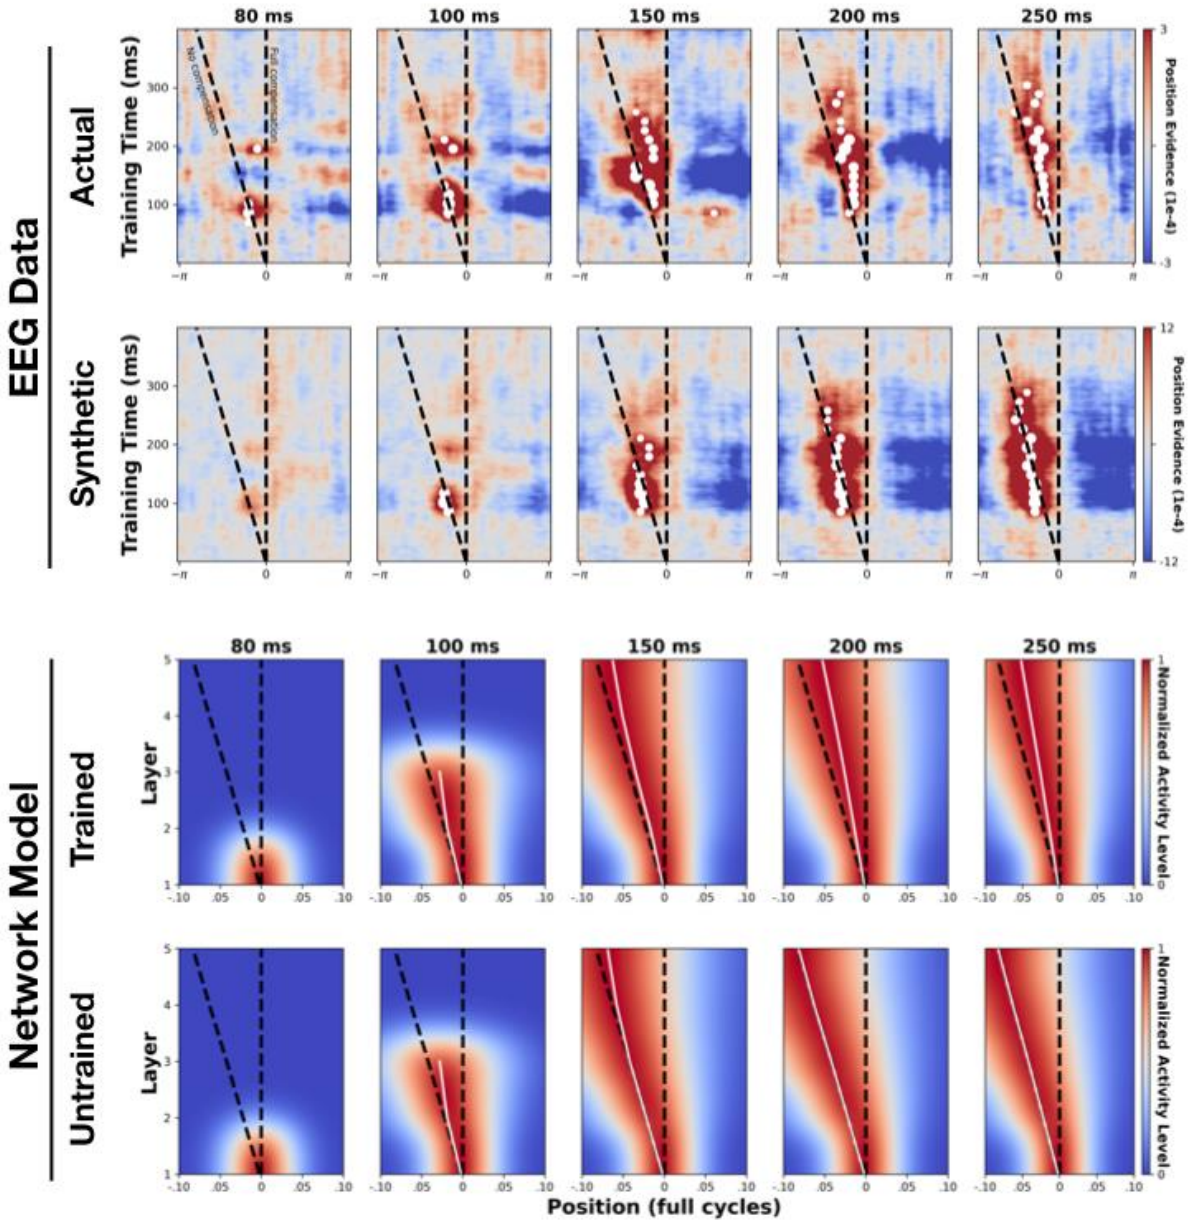

Figure S2. Tracking the formation of predictive time shifts. The top two rows show the temporal evolution of position information in the 250 ms following stimulus onset in the EEG data (top row: actual data, bottom row: synthetic control data). Bottom two panels show the equivalent for simulated network activity (top row: trained STDP model, bottom row: untrained control model).