

Computational specialization within the cortical eye movement system

Tom R. Marshall¹, Maria Ruesseler², Laurence T. Hunt², & Jill X. O'Reilly^{1,3*}

1. Wellcome Centre for Integrative Neuroimaging, Dept Experimental Psychology, Oxford University. *Anna Watts Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom*
2. Wellcome Centre for Integrative Neuroimaging, Dept Psychiatry, Oxford University. *Centre for Human Brain Activity (OHBA) University Department of Psychiatry Warneford Hospital, Oxford, Oxford OX3 7JX*
3. Wellcome Centre for Integrative Neuroimaging, Nuffield Dept Clinical Neurosciences, Oxford University. *FMRIB Centre, John Radcliffe Hospital, Headington, Oxford OX3 9DU*

* jill.oreilly@psy.ox.ac.uk

Abstract

Animals actively sample their environment through actions such as whisking, sniffing, and saccadic eye movements. Computationally, sensorimotor control may be viewed as an interplay between two processes that place different demands on their neural circuits: a rapid/competitive process for promptly selecting each upcoming action, and a slow/integrative process weighing the outcomes of multiple prior actions to build a model of the environment. Using saccadic eye movements as a model system, we addressed the hypothesis that frontal and parietal cortex are computationally specialized for these two functions. Through biophysical modelling, we predicted neural signatures of the competitive and integrative processes. We localized these signals to the frontal eye fields and intra-parietal cortex, respectively, using whole-brain, high-temporal-resolution neuroimaging (MEG). This frontal/parietal specialization can be linked to the differential characteristics of cortical circuits, and thus may represent a more general organizing principle of sensorimotor function in the primate brain.

Introduction

Far from being passive recipients of sensory information, both humans and animals actively sample the environment using their sensory organs. They integrate information across many individual samples to build a model of the sensory environment. In rodents, active sampling processes include whisking and sniffing; in primates, the most important and best-studied process is the control of saccadic eye movements.

As an observer views a visual scene, several saccadic eye movements per second are generated in order to direct the eye's small focal window to points of potential interest. Information from multiple fixations is integrated to construct a 'model' of the full visual field^{1,2}. This model in turn influences the selection of targets for future saccades³. Therefore, the process of active sampling may be viewed as an interplay between two concurrent processes with distinct characteristics.

Firstly, the brain must generate the sampling action. Since only one saccade is made at once, each individual saccade must be selected by a process of *competition* between representations of alternative possible saccadic targets⁴; the process is by necessity *winner-take-all* in that the eyes can only fixate one location at a time⁵ and must operate on a fast *timescale*, driven by dynamics that ensure a new saccade is made every few hundred milliseconds⁶.

Secondly, the brain must *integrate* the limited information gained from many individual saccades – not only to construct the visual scene, but also to inform the selection of targets for future saccades. Behavioural modelling suggests that the likely information value of future saccades is represented in a *salience* map^{7,8} that can be regarded as a Bayesian prior distribution over potential saccadic targets^{3,9}. Far from having winner-take-all dynamics, this salience map must capture the *relative or probabilistic distribution* over all possible saccades. It therefore must integrate information across multiple previous saccades.

What are the neural substrates of these two processes, and are they anatomically separated? At least three regions in macaque neocortex are involved in the selection of saccadic targets:

the frontal eye field (FEF) and the posterior parietal cortex (specifically regions LIP and 7a, of which LIP is the more extensively studied). These regions, and their human homologues, have attracted extensive interest not only in terms of their role in selecting saccades per se, but also as a model system for decision-making¹⁰, and a possible substrate for selective attention. However, consensus on the differential roles of frontal and parietal regions remains elusive. This study addresses the hypothesis that the distinction between the rapid, action selection process and the slower integration process defines the division of labour between the frontal and parietal eye fields in the primate brain.

One reason for hypothesising this division of labour is theoretical. Neurons in LIP have longer time constants and stronger local connectivity than those in FEF¹¹ – properties that make LIP a suitable substrate for ‘bump attractor’ dynamics¹² that could integrate information across many saccades¹³. In contrast, models of FEF generally embody competitive processes in which neurons representing alternative saccades compete in a winner-take-all dynamic^{14,15}. Thus, structural differences in the network properties of cortex between the parietal and frontal regions, by determining neural time constants, could lead to an anatomical division of labour.

A second line of evidence favouring the hypothesis is neurological. Unilateral damage to the frontal eye fields results in an *active* motoric bias (gaze vergence) towards the ipsilateral field¹⁶, *as if a competitive process selecting between targets had become unbalanced*; conversely unilateral damage to the parietal cortex results in a paucity of voluntary saccades and attention towards the ipsilateral field (neglect syndrome¹⁷), *as if the part of the salience map or prior covering the contralateral field had been permanently set to zero*.

Despite these observations, extensive electrophysiological work has demonstrated the response functions of individual neurons in FEF and LIP, in saccade selection paradigms, to be strikingly similar^{10,18–20}. Using evidence accumulation paradigms cleverly designed to extend the saccade selection process over hundreds of milliseconds (most commonly based on random dot kinematograms), both FEF and LIP have been shown to contain spatially-selective neurons that respond more strongly when a saccade is planned to their preferred location in the visual field. Responses in both areas are consistent with a model in which each neuron’s

firing rate at a given moment represents the log odds that its preferred location will be the target of the next saccade, effectively ramping up to target selection or down to target rejection^{18,21}. When activity in neurons selective for one of the targets reaches a threshold, a saccade is released towards that target. The process can be described mathematically as a sequential probability ratio test¹⁰ which, when noise is also modelled, maps onto the well-studied drift-diffusion model of decision making²². The model, which was developed in the context of perceptual decision tasks (random dot kinematograms) holds for more abstract forms of evidence accumulation such as the weather prediction task²³ and value-based decision-making²⁴.

Despite this elegant model, a series of surprising studies in monkey LIP and the homologous rat posterior parietal cortex (PPC) have recently shown that inactivation of parietal cortex^{25,26}, unlike FEF²⁷ (or the homologous frontal orienting field or FOF in the rat), has no impact on performance in the classic RDK evidence accumulation tasks (or analogous auditory tasks in the rat). This calls into question the causal role of parietal cortex in evidence accumulation, instead suggesting that frontal and parietal contributions to eye movement selection may not be as similar as initial neurophysiological evidence suggested²⁸.

The paradox that LIP neurons track the log odds of each saccadic target with such fidelity, but apparently play no causal role in the process of evidence accumulation, could be resolved by hypothesising (as above) that whilst both LIP and FEF neurons follow the evidence accumulation process leading up to each individual saccade, the causal role of FEF is to resolve competition between alternative 'next' saccades, whilst the causal role of LIP is to accumulate evidence *across multiple saccades*. In the inactivation studies described^{25,26}, as in most evidence accumulation tasks, evidence accumulation is completed within the context of a single saccade selection episode. There is no need to integrate information across multiple eye movements. Each trial starts with a 'flat prior' – previous trials tell the observer nothing about the correct response on the current trial – and ends with the selection of a single action. In this context, a representation of action probabilities for the current action is identical to the optimal cross-saccade map, but there is no need to carry information over between saccades. Indirect support for this argument comes from a study²⁹ in which rats

with parietal inactivation showed a reduced perceptual attraction bias (evidence of the influence of a prior) in an interval judgement task.

We designed an experiment, using whole brain high-temporal-resolution imaging with magnetoencephalography (MEG) to directly test the hypothesis that LIP integrates the value of saccades across multiple previous saccades to build up a prior, in contrast to the specialization of FEF for the selection of the next action (saccade). We modified the classic random dot motion task³⁰ in two ways, which our modelling indicated should drive independent signals relating to competition between saccadic targets, evidence accumulation overall, and the influence of prior beliefs. Importantly, as prior beliefs may not be detectable in event-related MEG activity, but rather encoded via synaptic plasticity³¹, we used the method of rapid frequency tagging to detect changes in the gain on perceptual information as a function of prior beliefs. Using MEG, we were able to localize these processes both in space (to the frontal and parietal cortical regions) and in time (to the periods of target selection, and pre-trial preparation, respectively).

Results

To spatially and temporally localise the processes of (i) selecting individual saccades via competition and (ii) integrating across multiple saccades to form a prior, we modified the classic random-dot kinematogram (RDK) task³⁰ in two ways (fig 1A).

In the classic RDK task participants observe mixtures of *randomly* and *coherently*-moving dots and accumulate evidence over hundreds of milliseconds to determine the direction of coherent motion, responding with a congruent saccade. We introduced *within-trial competition* between left- and right options to drive the competitive neural process hypothesised to occur in FEF, and we introduced longer term, *cross-trial integration* favouring left- or right- response, to drive the computation of a prior hypothesised to occur in parietal cortex.

Within-trial competition. In classic RDK stimuli a single set of coherent dots move to the left or right. In our modified version, all trials included some level of evidence for *both* choice

options (left and right) concurrently, and participants reported the *dominant* motion direction (fig1B). This means that the level of *competition* (ratio of left to right dots, fig1D) was manipulated orthogonally to the signal-to-noise or *total coherence* (total number of left and right dots, compared to random dots, fig1C). This manipulation allowed us to test for the hallmark of a competitive process in frontal and parietal neural activity, namely that it should depend upon the strength of evidence for the losing option as well as the evidence for the winning option²². In particular, a neural mean-field instantiation of a drift diffusion-like competitive process³² makes precise predictions about the independent effects of competition and coherence on brain activity³³.

Cross-trial integration (prior). In classic RDK tasks, the dot direction on each trial is independent. Each option (left, right) is equally likely, and so subjects do not have to retain any information about the current trial after the trial ends. In our modified version, we introduced long-term correlations in the dot-motion direction. The probability that the next correct choice would be 'right' was not fixed at 50%, but took values of 20%, 50% or 80% for blocks of about 25 trials (changes in prior probability were un-signalled and occurred with a uniform hazard rate of 0.04 (fig1E, 'true probability' (solid grey line)). Since the dominant direction on previous trials could be useful for determining the direction on the current trial, participants could benefit from integrating information across trials to construct a *prior* over the dominant motion direction on the next trial (fig1E, 'belief strength' (purple trace)). We modelled this evidence integration process using a Bayesian ideal observer model [similar to ^{34,35}], which captured local variations in prior probability, and learning delays. We used the expected value from the Bayesian ideal observer (see 'Methods' eq5) of $p(\text{right} = \text{correct})$, and modelled belief strength as $|p(\text{right}) - 0.5|$, meaning that it became strong when there was a high prior probability of either leftward or rightward saccade. This allowed us to test whether these prior beliefs were reflected in brain activity, either as an influence on the decision process itself or in the 1s foreperiod prior to evidence presentation on each trial.

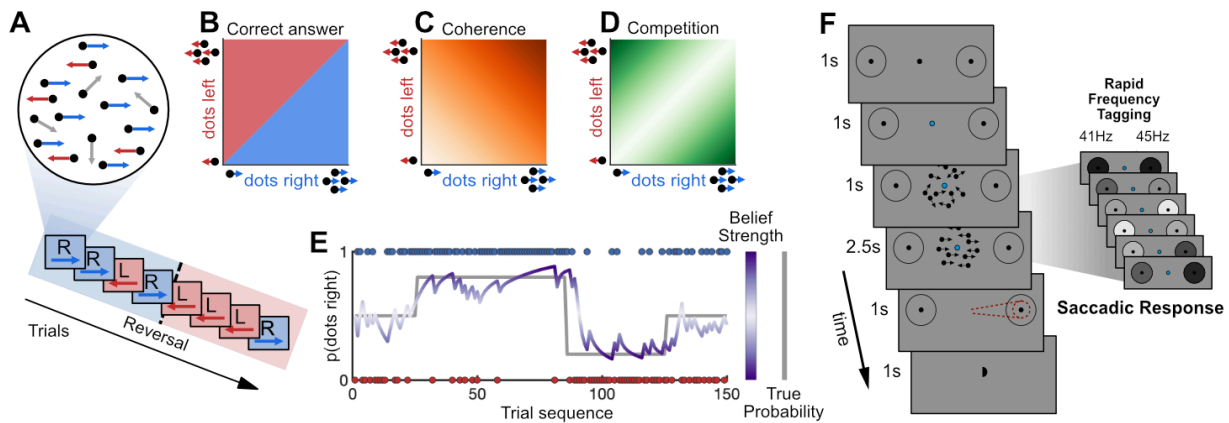


Figure 1: Adapted dot-motion task. A) Stimuli within a trial varied along two orthogonal dimensions; number of dots moving left, and right (all other dots moved randomly). B) The correct response is given by the dominant motion direction. C, D) 2-d stimulus space can also be parameterized as varying along two dimensions: **total coherence** (middle panel), or the total percentage of coherent motion to the left or right, and **competition** (right panel), the unsigned difference between proportion of coherent dots moving left and right. E) Additionally, the cross trial prior probability $p(\text{correct direction} = \text{right})$ varied across the experiment, with pseudo-random, unsignalled 'blocks' of trials in which the dominant direction was right 20%, 50% or 80% of the time (grey line). Bayesian learning models were used to estimate direction and strength of beliefs and observer should have about the current trial based on previous trials (purple line) F) Structure of a single trial: A get-ready cue indicated dots were about to appear. All trials began with 1 second of random motion, which temporally separated stimulus onset from evidence onset, meaning we were able to distinguish visual evoked neural responses from evidence accumulation processes. This also provided a temporally extended foreperiod in which neural activity reflecting prior beliefs could be measured. Random motion was followed by a 2.5s evidence accumulation period, in which coherent motion was present. When the dots disappeared, participants had a 1s response interval to make a saccade in the direction of perceived dominant motion. Participants were given unambiguous feedback on the correct answer (small centrally-presented hemisphere on the correct side), so they could learn the cross-trial prior independently of the quality of evidence on each trial. During the 1s foreperiod and 2.5s period of evidence accumulation, the two potential saccadic targets were 'tagged' with high frequency flicker to selectively entrain neural oscillations.

Choice behaviour is predicted by both within- and across-trial task features

Participants ($n = 29$, final analysis $n = 26$, for information on participant exclusions see 'Methods') performed 600 trials of a modified random-dot motion task with added within-trial competition and cross-trial integration, divided into 4 blocks with short breaks, while MEG was recorded. All participants also completed a practice session of 300 trials, outside the scanner, on a separate day.

We first confirmed that participants' choice behaviour was influenced by both the total coherence (signal to noise) and the competition between left- and rightward motion, by fitting logistic regressions to participants' saccade directions (left, right) as a function of percent coherent motion and proportion of coherent dots that moved right. As expected, there was an interaction such that participants were more likely to saccade right when a greater proportion of the coherent dots moved rightwards and this effect increased as the total amount of coherent motion increased ($t(25) = 9.74$, $p < 2 \cdot 10^{-10}$, fig2A).

Next, we tested whether participants learned and used the across-trial regularities in the stimulus sequence (the prior). We first confirmed that feedback from previous trials influenced participants' decision on the current trial using lagged logistic regression (figS1), indicating that participants were indeed retaining information across at least two previous trials. To quantify how this across-trial information should be integrated to inform behaviour we used a Bayesian learning model³⁴ to compute, for every trial in a stimulus sequence, the *prior belief* that an ideal observer should have, based on the feedback observed on all previous trials. For visualisation (fig2B), we divided trials using a tertile split according to whether the Bayesian prior strongly favoured rightward or leftward motion, or neither (neutral prior). On 'neutral prior' trials, the point of subjective equality (PSE) closely matched the point of *objective* equality (mean 49.7% right motion). Strong priors in either direction biased the PSE by 10-15% (strong left prior; PSE 55% right motion, strong right prior PSE 42.3% right motion, fig2B, inset).

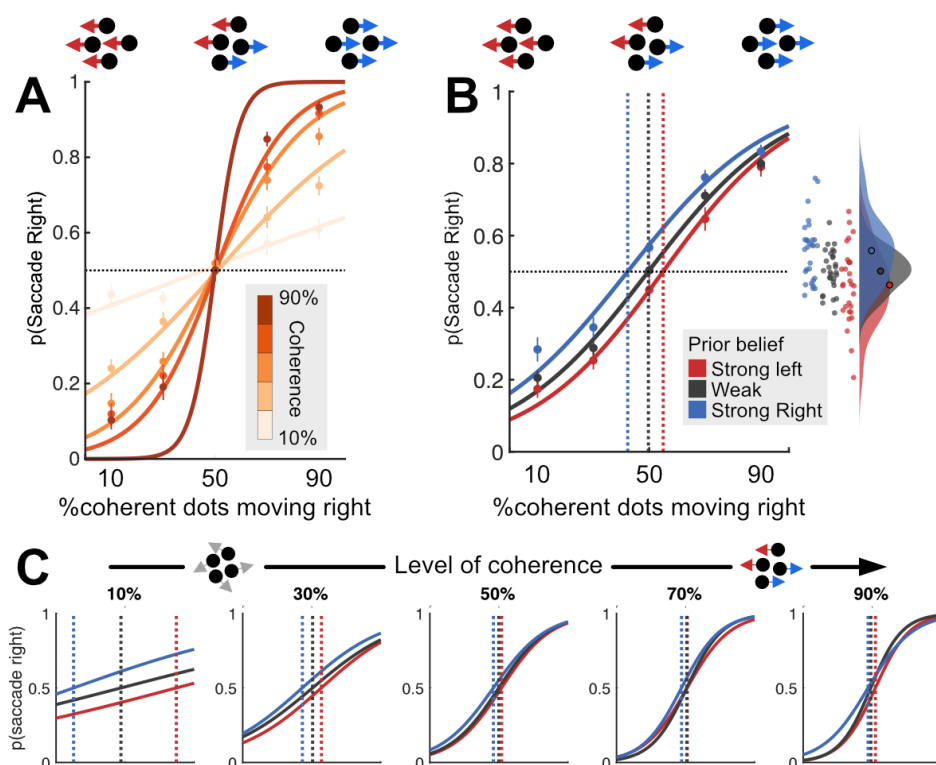


Figure 2: **Within- and across-trial influences on behaviour.** A) Increasing overall coherence (darker lines) produces a parametric performance improvement (steeper logistic curve). Dots indicate mean observed data, lines indicate logistic fits. B) Prior belief influences choice behaviour on current trial, leading to a shift in the point of subjective equality (dashed vertical lines). Inset: Raincloud plot showing values of $p(\text{Saccade right})$ for individual subjects on trials where 50% of dots moved right. C) Prior belief shifts point of subjective equality (dashed lines) most strongly when the least information is available on the current trial (coherence is lowest).

A multiple logistic regression confirmed that participants' choice behaviour was influenced both by the proportion of dots moving right (t-test on regression coefficients across the group: $t(25) = 14.79$, $p < 1 \cdot 10^{-14}$) and by the prior based on previous trials ($t(25) = 2.80$, $p = 0.0091$). These two effects did not interact ($t(25) = 0.45$, $p = 0.66$). This is evidence that decision-relevant properties of the currently-viewed stimulus (the proportion of coherent dots that moved right), and the beliefs participants had developed based on previously-viewed stimuli, made independent contributions to their choice behaviour.

We next determined whether properties of the current stimulus would determine the degree of influence prior belief had on choice behaviour; Bayesian theory suggests participants should rely more upon the prior when evidence in the current trial is weak³⁶ (low total

coherence) due to precision weighting. To visualise this effect, we repeated the Bayesian ideal observer analysis but divided the trials according to the level of total coherent motion on the current trial (fig1C). Based on the same tertile splits as above, prior belief biased current choice strongly when there was little decision-relevant motion, (at 10% coherence, PSE moved from $\approx 50\%$ to 10%/90%) but had little effect when a lot of motion was decision-relevant (at 90% coherence, PSE moved from $\approx 50\%$ to 48%/52%).

The effect shown in figure 2c was statistically confirmed using logistic regression. Percent total coherent motion had opposite effects on the influence of current evidence and prior belief on choice; when total coherent motion was high, the current evidence (proportion of coherent dots moving right) influenced behaviour more (Wilcoxon test on logistic regression coefficients for (coherence*prior) interaction across the group: $Z = 4.7$, $p < 3 \times 10^{-6}$, non-parametric test due to first-level outliers, see 'Methods') but prior belief influenced behaviour less ($Z = -2.52$, $p = 0.012$). This contrasts with the previous analysis where we found no interaction between the prior belief and the degree to which *competition* influenced choice behaviour. This confirms that participants selectively weighted the two sources of evidence available to them; up-weighting the impact of their prior belief – shaped by what they had seen *previously* - when little sensory evidence was *currently* available to guide their choice. This precision-weighting resembles the optimal Bayesian strategy for the task³⁶.

Neural models and neuroimaging

Next, we turned to our MEG data to assess whether the competitive process of selecting the current saccade and the integrative process of forming a prior across many saccades could be teased apart in space (frontal eye fields vs parietal cortex), time, and neural mechanism (impacts on event-related fields vs input gains respectively).

Biophysical model of the decision process

A key signature of a competitive selection process is that the decision variable depends on the strength of evidence for the *unselected* option, as well the selected option²². For this reason, in our modified random dots task, evidence for the chosen and unchosen dots direction was manipulated independently, allowing us to identify signals driven by

competition independently of signal to noise. To model neural mass signals detectable by MEG, we adapted a model of competitive selection between options, similar in algorithmic terms to the drift diffusion model and specified at the biophysical level, that has previously been described by Wang and colleagues^{32,33}. This model has been successfully used in a value-based choice task³³ to predict the independent effects of total value (analogous to total coherence in our task) and value difference (analogous to competition in our task) on the event related field (ERF) as measured in MEG (fig3A,B), and has previously been fit to brain activity in both parietal³² and frontal cortex^{33,37,38}. We adapted a neural mean-field version of this model to make predictions about how neural responses would vary with competition, coherence and prior for our task.

Briefly, the model comprises two neuronal pools coding for different choice options, with strong recurrent excitation within a pool and strong inhibition between pools. The between-pool inhibition mediates a winner-take-all competition between options resulting in one pool reaching a high-firing attractor state and the other pool a quiescent attractor state. The inputs to the two pools were proportional to the number of dots moving left, and the number of dots moving right.

To model the effect of the prior as a driving input, the input to each pool occurred in three phases; the first representing a participant's prior belief about the upcoming stimulus, the second reflecting undifferentiated activity due to random dot motion during the 1-second 'incoherent motion' epoch, and the third reflecting properties of the stimulus itself – i.e., the number of dots moving left and right - during the 2.5-second 'coherent motion' epoch (fig3A). The timing of the driving inputs reflects electrophysiological findings that the initial response of neurons in both FEF and LIP^{39,40} to target presentation, prior to evidence accumulation, weakly reflect the influence of prior beliefs on saccade selection^{31,41}.

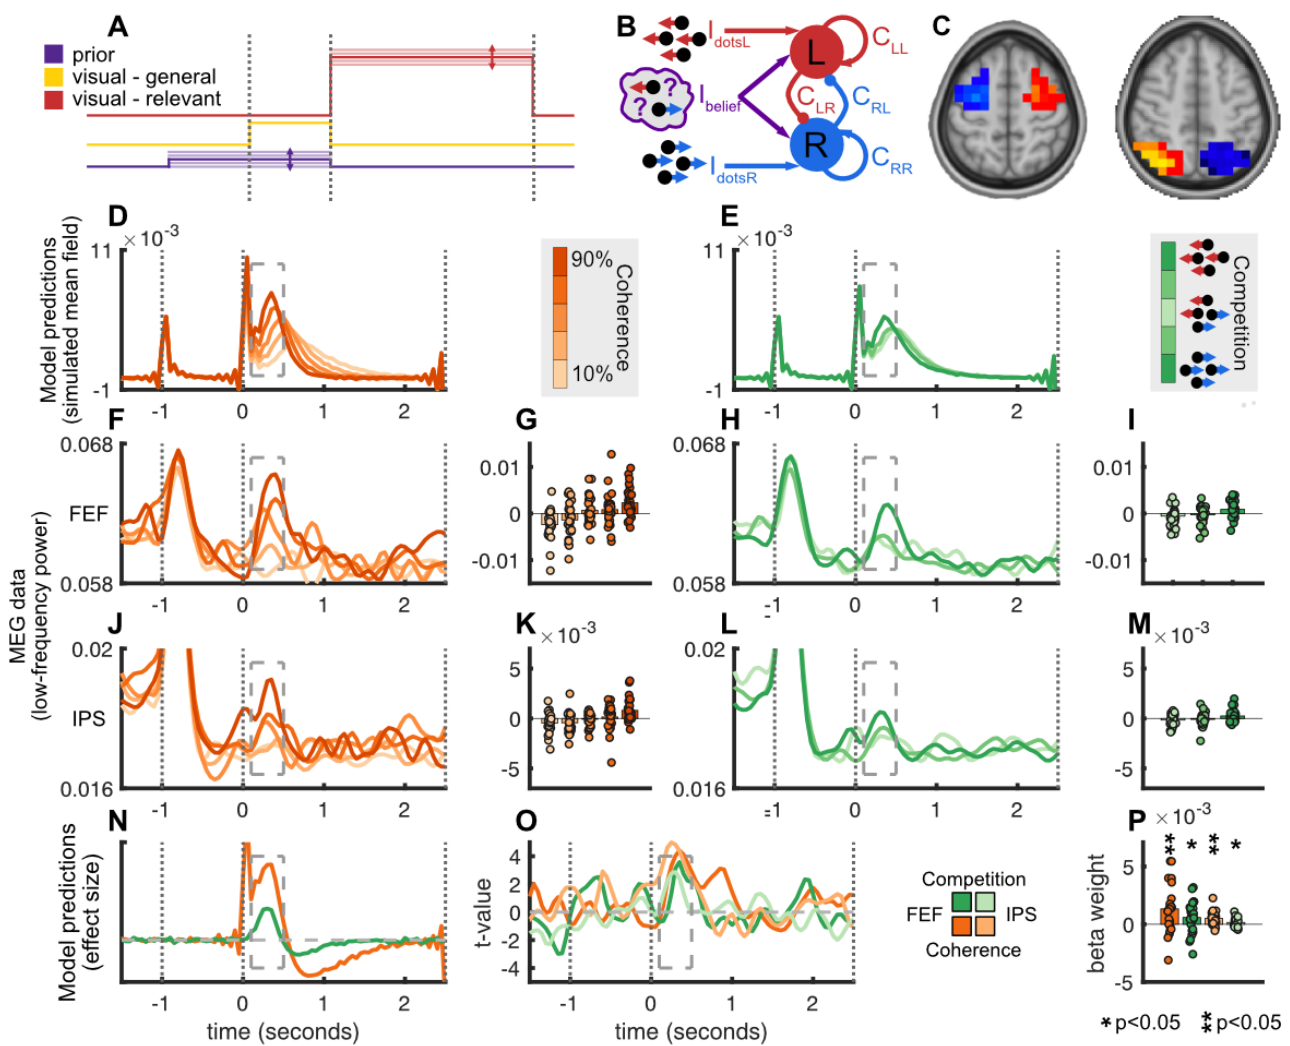


Figure 3: A) Input to neural mean-field model. Each node received weak variable input corresponding to prior knowledge (purple), weak fixed input corresponding to motion in all directions (yellow) and strong variable input corresponding to coherent motion (red). B) Properties of neural mean-field model. Each node contained recurrent excitatory connections and inhibitory connections to the other node. C) Regions-of-interest used to reconstruct activity in FEF (based on Wang & Kastner atlas⁴²) and IPS (based on a combination of data and anatomy – see ‘Methods’). Top row: Neural mean-field model predicts a parametric increase in activity as a function of stimulus coherence (D), and of stimulus competition (E). Second row: Low-frequency FEF activity displays parametric modulations as a function of coherence (F,G) and competition (H,I), consistent with predictions of neural mean-field model. Third row: As middle row, but for IPS low-frequency activity (J-M). Bottom row: N) General linear model reveals effects of both coherence and competition in neural mean-field model in post-stimulus period. O,P) Significant effects of coherence and competition are observed in both FEF and IPS in the same post-stimulus window.

The neural mean-field model generated three distinct predictions about the activity that should be observed in a competitive system. Firstly: Increasing total *coherence* (sum of dotsL and dotsR) should produce a parametric increase in neural activity following target onset (fig3D) in the time window 100-500ms following onset of coherent motion. Secondly: As *competition* decreases (i.e. greater absolute difference between dotsL and dotsR), there should be a parametric increase in neural activity following target onset (fig3E) in the same time window. Thirdly: Weakly increasing input due to prior knowledge in the prestimulus period (fig 3A, figS2A) should not significantly alter prestimulus activity, but rather bias the initial conditions of the competitive accumulation process such that activity evoked by the much stronger stimulus-driven input parametrically increased with prior strength, even though the prior input was no longer active, presumably due to the weak prior input biasing the state of the network before the stronger inputs began.

Event-related field in both Frontal Eye Field and Parietal cortex reflects the competitive process of selecting the current saccade

We next identified whether the pattern of results predicted by the neural mean-field model were observed in Frontal Eye Field and parietal cortex. We transformed subjects' MEG data to source space using LCMV beamforming⁴³, and extracted time series from regions-of-interest (ROIs) in the Frontal Eye Field⁴² and parietal cortex (fig3C, see 'Methods' for ROI definitions). We extracted the time-varying power in the low frequency range (2.8 – 8.4 Hz) as a proxy for the event related field or ERF, the magnetic field arising from local field potentials in cortex⁴⁴. We used low frequency power as a proxy because the ERF itself becomes sign-ambiguous after beamforming, but data in the time-frequency domain does not. We focussed on low frequencies rather than higher frequencies because the network model is dominated by low-frequency responses and does not exhibit higher frequency oscillations³³. This analysis revealed two transient responses in both regions (fig3F,H,J,L); the first shortly after onset of incoherent motion, the second 100-500ms after onset of coherent (i.e., choice-relevant) motion.

Concordant with ramping activity observed in electrophysiological studies^{10,18}, in both FEF and parietal cortex there was a clear parametric modulation of the low-frequency MEG signal 100-500ms after the onset of coherent motion as a function both of total coherence

(fig3F,G,J,K) and stimulus competition (fig3H,I,L,M). As predicted, stronger evoked activity was observed with higher levels of total coherence; this is consistent with the observation that ramping evidence-accumulation signals in single-unit studies rise faster for higher coherence levels in standard dot motion tasks¹⁰. Stronger evoked activity was observed for *lower* levels of stimulus competition, indicating that the presence and strength of evidence conflicting with a decision affects the decision process – a hallmark of a competitive system. Statistically, a linear multiple regression with parameters coherence, competition and prior strength confirmed that the amplitude of the evoked response varied as a function of both coherence and competition; specifically, and as predicted by the neural mean-field model (fig3N) activity increased as a function of stimulus coherence (FEF; $t(25) = 3.42$, $p = 0.002$, Parietal; $t(25) = 4.67$, $p = 9e-5$), and decreased as a function of competition (FEF; $t(25) = 2.13$, $p = 0.043$, Parietal; $t(25) = 2.18$, $p = 0.039$, fig3O,P). Both effects were tested in a time window 100-500ms after the onset of coherent motion, defined based on the predictions of the biophysical model. To test whether there was any difference in timing of effects between FEF and parietal cortex, we used permutation testing (permuting the waveforms between ROIs within subjects), and found no significant difference in timing between regions (Coherence; $p = 0.12$, Competition; $p = 0.92$).

The qualitative correspondence of the MEG activity to a neural mean-field model of a competitive decision process suggests that both Frontal Eye Fields and parietal cortex could be engaged in resolving saccadic choices via competition by mutual inhibition, compatible with previous observations of evidence accumulation signals in both regions. However, evidence from inactivation studies^{25,26} suggests that the activity in parietal cortex is not necessary for saccade selection whereas activity in FEF is. The signal observed in parietal cortex may, we hypothesised, play a parallel role in integrating this information in a cross-saccade prior.

Contrary to our predictions, prior strength did not significantly modulate the low-frequency MEG signal in either FEF or parietal cortex (FEF; $t(25) = -1.22$, $p = 0.23$, Parietal; $t(25) = 0.80$, $p = 0.43$, figS2). However, the absence of an effect of prior should perhaps be interpreted with caution, as for any null result; perhaps the effect was too subtle to be detected in the present paradigm.

A prior integrating across multiple saccades is represented in parietal cortex via gain modulation

Although participants' behaviour was influenced by the prior, we were unable to detect activity corresponding to the prior in low frequency MEG signal (approximating the event related field, mainly driven by synchronised post-synaptic potentials) from FEF or parietal cortex. However, theory suggests that the model of the prior as a driving input is a poor match for the physiological mechanisms by which prior beliefs are represented and influence choice: representations sustained over inactive periods or delays may be more efficiently represented by non-spiking mechanisms⁴⁵.

For example, in an extension of the biophysical model used above, prior beliefs were modelled as modifications to the synaptic weights from evidence-tuned inputs to the decision pools³¹. These gain modulations could manifest through short term synaptic plasticity on the timescale of seconds. It has been proposed that short term synaptic plasticity may act as an activity-silent store for working memories (or in this case, a prior) that is later reactivated as a bump attractor^{12,46}, a proposed mechanism for working memory³¹.

To probe for possible gain changes between visual inputs and the parietal cortex, we exploited the method of rapid frequency tagging⁴⁷. During the time that moving dot stimuli were present on the screen the two saccade targets were rhythmically flickering at two different high frequencies (41 and 45Hz, fig1F) that were indistinguishable to observers. Flicker at such high frequencies is typically not perceived (they are for example close to the refresh rate of a standard 60Hz computer monitor – note that here we used a 1kHz projector to present stimuli). Indeed, no participant reported awareness that the stimuli were flickering. Flickering stimuli have previously been shown to produce detectable rhythmic activity in M/EEG signals⁴⁸ including at higher frequencies above 40Hz^{49,50}, presumably by producing synchronised neural activity in visual cortex that propagates through the visual streams^{51,52}. Parietal neurons are known to increase^{51,52} their neuronal gain when an attentional or saccadic target is present in their receptive field^{19,53–55}. Therefore, when one or other target is expected to be the saccadic target (for example due to a prior belief) we would expect the tag frequency for that target to propagate more effectively into parietal cortex;

additionally we would expect the effect to be seen mainly in the contralateral hemisphere due to the lateralized representation of the visual field in occipital and parietal cortex.




Figure 4: A) Illustration of frequency tagging in a single trial. Onset of moving dots and visual flicker at saccade targets produces detectable high-frequency activity in the MEG signal (grey trace). Sliding-window fourier analysis tracks the power envelope at this frequency (green trace). B) Parietal region of interest, identical to fig 3C. C) Time course of frequency tagging activity as a function of prior belief (contralateral, weak, or ipsilateral to region). D) Direction of prior belief modulates parietal activity before onset of coherent motion. E) Time course of frequency tagging activity as a function of congruence of prior belief and dominant direction of moving dots (congruent, incongruent). Dashed box indicates time-window where cluster values were maximal. F) Individual datapoints averaged across time-window denoted by cluster test. G) Time course of tag activity as a function of dominant dot-motion direction. No significant dot-motion-related activity modulation was observed.

In the MEG data we recovered clear spectral peaks at the 41Hz and 45Hz tag frequencies (figS3A) that were present during the foreperiod, detectable from about 500ms after flicker onset (see 'Methods') and persisted during the entire stimulus period (fig4A). The frequency

tagging effect was strongest in parietal and occipital regions and notably absent in frontal regions (figS3B), suggesting the signal propagated forward from occipital cortex but only through a few synapses. There was a clear lateralized effect such that the tag frequency presented at the right target was strongest in the left hemisphere of the brain, and vice versa. For each region of interest (left parietal, right parietal) we defined ‘frequency tag activity’ as the time-resolved power at the specific flicker frequency of the contralateral flickering saccade target, corrected for the main effects of flicker frequency and hemisphere (for full description see ‘Methods’).

We defined regions of interest (ROIs) in the parietal cortex in each hemisphere based on the conjunction of parietal cortex defined by the Harvard-Oxford cortical structural atlas and regions showing a strong response at the tag frequencies (fig 4B, for full details see ‘Methods’). This conjunction ROI overlapped substantially with two regions of interest thought to be homologous to eye-movement regions LIP and 7a in the macaque⁵⁶; however the spatial resolution of MEG is not sufficient to say with certainty which intra-parietal sub-regions were the source.

The cross-saccade prior is represented prior to evidence accumulation

If the parietal cortex represents a prior expectation based on the integration of previous saccades, this should be in evidence in the foreperiod, when only incoherent motion was present. We defined a time window from 500ms after flicker onset (the point at which tag activity was first evident in parietal cortex – fig3A) until the onset of coherent motion. We divided trials with a tertile-split into those on which the prior strongly favoured the contralateral target, strongly favoured the ipsilateral target, or was close to neutral.

To model participants’ prior beliefs, we used a Bayesian ideal observer model (see ‘Methods’ for full description). The ‘ground truth’ or generative probability of rightwards motion was either high (70 or 90%), low (30 or 10%), or neutral 50%; these probabilities changed (un-signalled) about every 25 trials (see fig1 and ‘Methods’). The Bayesian ideal observer model estimated the probability that the dominant motion direction on the upcoming trial would be right or left, based on the feedback from previous trials. The advantage of this approach over

using ‘ground truth’ prior probabilities is to capture local fluctuations in probabilities, and learning delays. A simple lagged regression also showed that participants’ judgements were affected by the outcome of at least two prior trials (figS1).

Concordant with our hypothesis, we found that during the foreperiod, frequency tagging activity reflected the direction of the prior – activity was strongest on trials when the prior favoured the contralateral target (linear contrast comparing tertiles strong-contra, weak, and strong-ipsi: $t(25) = 2.27$, $p = 0.028$, fig4C,D). To test for a parametric effect of prior strength on frequency tag strength, we conducted a linear regression of tag strength on prior strength (defined as the prior probability of contralateral motion on the upcoming trial, $p(\text{contra})$) – this failed to reach significance (group t-test on regression coefficients, $t(25) = 1.50$, $p=0.073$) suggesting that either the representation of the prior is largely categorical, or that we had insufficient sensitivity to detect a parametric effect.

Gain modulation in parietal cortex is sensitive to dot motion, in the context of the prior

Whilst the frequency-tag data provided evidence of gain modulation relating to prior beliefs, it did not track evidence accumulation (absolute or relative coherence on a single trial basis) in the same way as the ERF (figS6). This is compatible with a model in which the prior is represented as activity-silent synaptic plasticity³¹, or by tonic changes in baseline firing rate as in a bump attractor¹² since tonic activity would likely not be detected in the band-passed MEG signal.

The gain modulation signal relating to prior beliefs was observed in the foreperiod of the task, during which the only lateralized effect is prior belief. If parietal cortex integrates incoming evidence with the prior belief to form a posterior, that will become the prior for the next trial (belief updating), we might expect to see a further gain-modulation signal representing the evidence, or decision, *relative to the prior*, reflecting this update process.

We coded trials as ‘congruent’ – i.e., the target favoured by the prior was also favoured by within-trial evidence – or ‘incongruent’. Indeed, stronger activity at the contralateral tagging frequency was observed on congruent than incongruent trials, during the evidence

accumulation phase of the trial (fig4E; linear regression followed by cluster-based permutation test on regression coefficients, $t_{\text{maxsum}} = 21.0023$, cluster $p = 0.0238$ see 'Methods'). Interestingly, this effect was observed to be strongest in a later time window (800-1150ms after stimulus onset) than the decision-related activity in FEF. This suggests that, in the context of a paradigm in which information can be integrated across many saccades, evidence coding in parietal cortex represents the combination of this evidence with the prior, akin to the formation of a posterior distribution that could be used to guide selection of future saccades³⁴.

At no point did we observe a main effect of the ultimate saccade direction in parietal cortex (fig4G); this is perhaps unsurprising as the frequency tagging stimulus ended at the end of the coherent motion period, several hundred ms before the saccade was made.

Time course of neural effects suggests a sequential interplay between frontal and parietal cortex

An advantage of MEG over other human neuroimaging methods is its high temporal resolution. Combined with the fact that MEG is a whole-brain method, this allowed us to examine the sequence of effects unfolding between FEF and LIP (fig5).

To directly compare the sequence of events we ran a series of ANOVAs plotting the effect sizes for each of the factors affecting activity in FEF and IPS as a function of time. A clear sequence of events emerged: initially the prior is represented via gain modulation in parietal cortex, then the evolution of activity in both FEF and parietal cortex reflects the evidence accumulation process as captured in a competitive mean field model, and finally the interaction of evidence and prior is again reflected in gain modulation in the parietal cortex. Before the onset of coherent motion, differences in activity across trials must relate to the prior. In this 'foreperiod' time window, we indeed observed a lateralized effect in parietal cortex as a function of the participants' prior belief. Then, shortly after stimulus onset we observed parametric effects of stimulus coherence and stimulus competition in frontal eye field, predicted by the neural mean-field model and likely the signature of the choice process itself. Finally, having resolved competition between choice options, we observed a later

effect in parietal cortex which related to congruence between the participants' initial prior (i.e., a prediction about the upcoming stimulus direction) and their eventual decision; namely, a strong suppression of activity when the stimulus is incongruent with the prior belief. This likely reflects the integration of evidence from the current trial with the prior belief – perhaps reflecting a *belief update* process whereby information about the current trial is used to modify the state of parietal cortex to guide future saccades.




Figure 5: Summary of neural effects in frontal and parietal cortex. A) Time courses of significant statistical effects; repeated measures ANOVAs comparing contralateral, weak, and ipsilateral prior belief in parietal cortex (top row), all levels of stimulus coherence (middle row, orange) and stimulus competition (middle row, green) in FEF, and comparing ‘evidence (dots) congruent with belief’ trials with ‘evidence (dots) incongruent with belief’ trials in parietal cortex (bottom row). Shaded boxes indicate preselected analysis windows (top and middle rows, see fig3F,H, J,L and 4C) and window delineated by cluster permutation test (bottom row, see fig4E). Coloured dashed lines indicate maxima of the respective F-statistics. B) Putative schematic of computations in frontal and parietal cortex and how they influence one another.

The above is consistent with a model in which activity in both FEF and parietal cortex reflect the evidence accumulation process, but a prior integrating over multiple previous saccades and their outcomes is also represented, in parietal cortex, by the modulation on input gains³¹. This prior is represented even outside the time period in which evidence for different saccades is being weighed. The representation of the prior that is present before each

saccade (in the foreperiod in our experimental task) serves to bias the saccade selection process, influencing choice behaviour (fig2B,C). During evidence accumulation, competitive dynamics lead to the selection of one or other saccadic target. As in previous electrophysiological studies, we see the neural correlates of this process in both FEF and parietal cortex, but inactivation studies^{25,26} suggest that only the FEF is causally involved in the perceptual decision. Once the saccade selection process is resolved (but before the saccade is physically executed), an ‘update’ signal is observed in parietal cortex, perhaps reflecting the integration of the new evidence into the prior belief to form a posterior.

Discussion

Optimal exploration of the visual environment through eye movements requires the brain both to select among currently competing saccadic targets, and to integrate information across saccades to construct a Bayesian prior in order to plan future saccades. Using whole-brain imaging at high temporal resolution with MEG, combined with computational modelling, we demonstrated a temporal cascade of computations, consistent with an ongoing temporal interplay of activity across frontal and parietal cortex; a *network* account of optimal visual exploration.

Selecting the next saccade by resolving competition

The MEG mean field signals in both the Frontal Eye Field and parietal cortex reflect competition between currently-presented choice options. In our modified random dots task, evidence for the unchosen option was manipulated independently of the chosen option, allowing us to identify a competition-driven signal as well as the classic signal relating to the total evidence (overall coherence level) in these regions. These signals matched the predictions of a biophysical model of choice.

A limitation of our results in frontal eye field was the inability to resolve lateralized effects. Activity in the frontal eye fields *as a whole* resembled the neural mean-field model *as a whole*. Since neurons in FEF predominantly code for saccades to the contralateral hemisphere^{57,58} we also expected *lateralised* activity in FEF to resemble the mean-field of a single evidence-accumulator in our model, i.e., a parametric modulation of activity as a

function of overall evidence for a single choice option, but we did not observe this (figS2,S3). This may be due to several factors, most notably MEG-specific signal processing limitations. Left and right FEF are relatively proximate and their mean-field signals are likely to be temporally correlated, conditions which make neural sources difficult to separate by beamforming⁴³. Furthermore, although coding for contralateral hemisphere is the norm, some FEF neurons also possess ipsilateral response fields⁵⁹. It may be that contralateral dominance is not strong enough to be detectable in the mean field.

Prior beliefs as modulations to input gain

We were unable to detect changes in the mean field driven by prior beliefs, although these beliefs did affect behaviour. However, using the method of rapid frequency tagging, we did find evidence that parietal cortex represented a cross-saccade prior through modification of input gains, perhaps via synaptic plasticity³¹. Driving oscillations at 41 or 45Hz propagated more effectively into the parietal cortex from the target favoured by the prior belief, than the unfavoured target. This effect was evident during the foreperiod in which no evidence was presented for either option (all dots moving incoherently). This effect was hemispherically lateralised; when participants had a belief the upcoming dots would move right, frequency tagging activity was strongest in left IPS, and vice versa.

The fact that inputs to parietal cortex reflect a cross-saccade prior is consistent with rodent work showing that parietal cortex tracks stimulus history²⁹. It may be that the pre-stimulus representation of belief in parietal cortex serves to bias the decision process later observed in FEF.

Later in the trial, after decision-related activity in the neural mean-field model and the FEF had quiesced, so – we infer – the decision process had presumably been resolved, we then observed modulation of IPS activity by *belief-decision congruence*. On trials where participants had a strong prior belief, late IPS activity was strongest in the hemisphere contralateral to that belief when their eventual decision was congruent with their pre-trial belief, and suppressed when the eventual decision violated their pre-trial belief. This congruence effect may reflect an updating of the prior or salience map to be used on future trials.

A network account of frontal and parietal eye fields

Overall, our findings imply that evidence accumulation is indeed occurring in both regions. However, the apparent paradox that parietal inactivation does not affect evidence accumulation^{25,26} may be resolved by the insight that evidence accumulation in the two regions is for different computational purposes; in FEF, to resolve competition between options, and in parietal cortex to learn the across-trial statistics. As the outputs of these computations need to be integrated to produce optimal behaviour, this also explains the co-activation of the dorsal attention network commonly observed in fMRI. The slow dynamics of the BOLD response may cause the network to *appear* to co-activate, where in fact a temporally-precise cascade of events – from parietal cortex, to frontal, back to parietal – occurs in a short window (fig5).

Network structure may be constrained by intrinsic dynamics

Neurally representing each type of task-relevant information – currently available and previously learned – requires computations that operate on different intrinsic timescales, and display different neural dynamics. Resolution of competition between two currently-available choice options requires a fast, winner-take-all neural system that rapidly converges to one of two stable attractor states. In contrast, representation of previously-learned knowledge must maintain a broadly similar state over a long timescale and incrementally change as new datapoints are incorporated. It is attractive to think that the division of labour between FEF and parietal cortex is partly a result of the different intrinsic properties of cortical networks in these different brain regions, thus linking structure to function. Recent work in which multiple neurons were recorded in parallel in FEF and LIP has shown that the intrinsic timescale of activity in LIP is slower than FEF, lending LIP to the kind of ‘bump attractor’ dynamics that can integrate information over longer timescales¹².

Inactivation studies have suggested a double dissociation of function between FEF and parietal cortex, with only FEF essential for within-trial evidence accumulation²⁵ and only parietal cortex important for cross trial integration²⁹. However, in the neurophysiological literature, signals in both LIP and FEF reflect evidence accumulation within each trial^{18,19}. The present results suggest that these findings can be reconciled as follows: both parietal and frontal eye fields showed neural signals that were consistent with the predictions of the

winner-take-all biophysical competition model (fig3), whereas only parietal cortex showed neural signals consistent the predictions of the across-trial Bayesian learning model (fig4). It is possible that the apparent ‘fast-timescale’ signature of winner-take-all decision making in parietal cortex (fig3J-M) is in fact a consequence of the decision being primarily resolved in FEF and then rapidly relayed back to parietal cortex via top-down projections.

Our findings may therefore be relevant for understanding why the cortical network for a single behaviour (saccade selection, or attention allocation) is distributed across frontal and parietal cortex rather than localised to one region; an organising principal that seems inefficient at first glance. It is known that brain regions have different intrinsic time constants⁶⁰ and these time constants follow principles of cortical hierarchy⁶¹; sensory cortical activity changes on a fast timescale and ‘higher-order’ areas are slower. Recent work has also shown that monkey LIP neurons tend to have slower intrinsic timescales than FEF neurons¹¹. Since optimal allocation of attention requires the integration of fast and slow computations, it may be efficient to distribute these computations across regions with maximally different intrinsic timescales, then integrate the outputs via long-range connections. The need to integrate computations on different timescales may be *why* the cortical eye movement network has the topology it does; indeed, it may be why distributed topologies are a feature of many brain networks.

Methods

Behavioural task

Participants (N=34) performed a variant of the classic dot-motion task. 1 second after a preparatory cue a field of 100 dots covering 4.2 degrees of visual angle (on-screen width 8.8cm, screen distance 120cm) was presented at fixation for 3.5 seconds. For the first second all dots moved randomly (uniform distribution over motion directions). After one second some dots changed direction and moved either to the right or left. There were five levels of overall coherence (sum of left-moving and right-moving dots): 10%, 30%, 50%, 70% and 90%, and three levels of competition: Coherently-moving dots either moved 90% or 70% in the direction of dominant motion, or 50% in each direction (i.e., no correct answer). Two saccade targets subtending 4 degrees of visual angle were concurrently presented at 6.8 degrees eccentricity, 0.8 degrees below the horizontal midline. Participants were instructed to maintain fixation until the dots disappeared, then to make a saccade to the saccade target on the side of dominant motion as quickly as possible. After 1s feedback (a semicircle at fixation) was given indicating the correct side.

Across trials the correct side was drawn from a generative distribution with $p(\text{right})$ either 20%, 50%, or 80%. The generative distribution changed randomly across a block with a fixed probability of 4%, i.e., a switch approximately every 25 trials. On trials with no correct answer (i.e., where equal numbers of dots moved left and right) feedback was drawn from the same generative distribution as the dots, however a programming error led to the feedback in this condition being the reverse of what was intended. This error affected approximately 8% of all trials.

All participants were trained on the task in a separate session outside the MEG for 300 trials (approximately 40 minutes). 5 participants were excluded due to low accuracy in the practice session and 29 participants completed the MEG session. During MEG acquisition data were recorded in four runs of max 20 minutes, each consisting of 150 trials. Total task time was approximately 80 minutes.

Power calculation

Sample size was determined by simulation. Based on the z-scores and sample size from the MEG analysis in³³, we calculated the effect size d as the z-score divided by the square root of the sample size. We then simulated a population of 10,000 virtual participants with this effect size, randomly sampled subsets of 10-60 participants from this population, calculated the sub-sample z-score, and computed statistical power for each sub-sample size as the proportion of z-scores for that sample size that exceeded 1.65 (equivalent to $p < 0.05$ in a one-tailed test). Simulation results indicated statistical power of 80.5% at our final sample size of $N=26$.

Neural Models

Neural mean-field model

To model decision dynamics in Frontal Eye Field and parietal cortex we implemented a mean-field reduction of a spiking network model⁶². Full details of the biophysical parameters of the model are given in³³. Briefly, the model consists of two connected neuronal pools each coding for one motion option (left, right), with within-pool excitatory connections and across-pool inhibition. Each unit receives noisy background inputs simulating endogenous cortical noise, plus three task-related inputs: Firstly; a weak (0 to 1.6Hz) input to one pool in the pre-stimulus period, simulating an initial bias in the decision process due to parietal input. Secondly; a weak (5Hz) input to both pools during the incoherent motion period, capturing the presence of low levels of motion in both decision-relevant directions. Thirdly; a strong input (ranging from 10.1Hz to 18.1Hz) to each pool in the coherent motion period proportional to the number of dots in each motion direction.

We focused on the synaptic inputs (I_1, I_2) to facilitate comparability with the MEG data, since MEG is known to be primarily sensitive to postsynaptic potentials (Hamalainen 1993). For optimal comparison with the MEG data which was bandpass-filtered and transformed to the frequency-domain, removing the DC component, we used the temporal derivative of the signal from the mean-field model.

Bayesian learning model

Because the modified dots task had temporal structure (dominant motion direction on trial j could be predicted, but not perfectly, from trials $1 \dots t-1$), performance could be facilitated by tracking the true generative distribution of dominant motion directions, including when it changed. To model participants' learning strategies we used a Bayesian ideal observer model; a 'virtual participant' that was fed the sequences of feedback given to the human participants and constructed, via Bayesian inference, a belief about the generative distribution on the upcoming trial.

Model details

On each trial, the prior probability that the dominant motion direction would be 'right' followed a Bernoulli distribution with parameter q_t , i.e. the prior probability that the correct answer would be 'right' on trial t was q_t .

The prior distribution over q_t was initiated as a uniform on the range (0,1) on the first trial, and thereafter obtained from the posterior over q_{t-1} , based on the outcomes $x_{1:t-1}$, combined with a uniform 'leak'; the posterior and 'leak' distributions were weighted by a factor H representing the true hazard rate:

$$p(Q_t = q | X_{1:t-1}) = (1 - H) \cdot p(Q_{t-1} = q | X_{1:t-1}) + H \cdot U(0,1) \quad \text{Eq. 1}$$

Where:

$$H = p(q_t \neq q_{t-1}) = 1/25 \quad \text{Eq. 2}$$

i.e. it was assumed that participants know approximately the true value of H following extensive pre-training.

The posterior $p(Q_t = q|x_{1:t})$ was obtained iteratively from the prior $p(Q_t = q|x_{1:t-1})$ and the likelihood:

$$p(Q_t = q|X_t) = p(X_t|x \sim B(q))$$

Eq. 3

... using Bayes' theorem:

$$p(Q_t = q|X_{1:t-1}) \propto p(X_t|x \sim B(q)) \cdot p(Q_t = q|X_{1:t-1})$$

Eq. 4

where the posterior over q_t was normalized to integrate to 1.

Where a scalar value for the 'prior' is used in data analysis, this is the expected value of q_t based on the prior distribution

$$E(q_t) = \int p(Q_t = q|X_{1:t-1}) \cdot q dq$$

Eq. 5

Analysis of saccade data / behavioural data

Custom matlab code extracted saccade direction (left, right, or no saccade) based on the eyetracker data. Due to a technical error one participant's eyetracker data were over-written and saccade information was reconstructed from the horizontal EOG.

To analyse saccade data we fit logistic regression models. Firstly, we asked whether the probability of saccading to the right on trial t depended on the proportion of coherent dots moving right on trial t (%R = 10,30,50,70 or 90%), the total coherence on trial t (coh = 10,30,50,70 or 90%), and the interaction of these (%R-mean(%R) x coh) (fig1A).

Secondly, we asked whether the probability of saccading to the right on trial t depended on the proportion of dots moving right on trial t (%R = 10,30,50,70 or 90%), the participants'

prior belief about the dots on trial t formed from observing feedback on trials $1, 2, \dots, t-1$ $E(q_t)$ as defined in Eq 5 above), and the interaction of these (fig1B).

Thirdly, we asked whether the observed effects of proportion dots moving right and prior belief on saccade direction were altered as a function of the level of stimulus coherence. To test this we fit a first-level logistic regression with evidence (%R = 10,30,50,70 or 90%), and prior belief $E(q_t)$ at each level of total coherence (coh = 10,30,50,70 or 90%). We then computed linear contrasts over the effects for evidence and prior belief across the five levels of total coherence. Due to the presence of outliers at the first level (generalized extreme studentized deviate many-outlier procedure⁶³) we used nonparametric Wilcoxon signed rank tests at the second level to test against the null hypothesis of zero median.

MRI acquisition

To enable localisation of cortical source generators of the MEG signal, a high-resolution structural MRI was acquired for each participant using a Siemens 3T PRISMA MRI scanner with voxel resolution of $1 \times 1 \times 1 \text{ mm}^3$ on a $232 \times 256 \times 192$ grid. The anatomical MRI scan included the face and nose to improve co-registration with the MEG data (see 'MEG processing and analysis'). MRIs could not be acquired for 3 participants due to drop-out and screening contraindications. All imaging analysis was therefore conducted on the remaining 26 datasets.

MEG acquisition

Data were recorded at 1.2KHz with an Elekta Neuromag VectorView 306 MEG system with 102 magnetometers and 102 pairs of orthogonal planar gradiometers. Head position indicator (HPI) coils were placed at four locations on the head to record head position relative to the MEG sensors at the start of each run. Head landmarks (pre-auricular points and nasion) and 200 points on the scalp, face and nose were digitized using a Polhemus Isotrack II system. EEG Electrodes were placed above and below the left eye and on the temples to record horizontal and vertical EOG, and on each wrist to record ECG.

Rapid Frequency Tagging

During presentation of the dot-motion stimulus, the two saccade targets flickered rapidly at 41Hz and 45Hz (counter-balanced across subjects). The targets flickered sinusoidally between black and white (greyscale) at a refresh rate of 1.44KHz. To achieve this high rate of presentation we used a PROPixx DLP LED projector (VPIxx Technologies Inc., Saint-Bruno-de-Montarville, Canada). In post-experiment debriefing no participant reported awareness of the flicker.

MEG processing and analysis

MEG analysis was performed using fieldtrip⁶⁴, OSL (<https://github.com/OHBA-analysis/osl-core>), and custom MATLAB scripts.

Data were first Maxwell filtered using the MaxFilter program (Elekta Instrument AB, Stockholm, Sweden); Maxfilter is a method for separating parts of the recorded MEG signal that arise from external noise and neuronal activity respectively. MRI and MEG data were co-registered using RHINO (Registration of Headshapes Including Nose in OSL). Data were downsampled to 200Hz, bandpass filtered between 1 and 80Hz and bandstop filtered around the line-noise frequency of 50Hz. Trials containing outlier values were automatically detected and removed using function `osl_detect_artefacts` with default settings, and Independent Component Analysis was used to automatically remove additional artifacts associated with eye-blinks, ECG, and line noise.

A single-shell forward model was constructed from each participant's anatomical MRI. Sensor data were projected onto an 8mm grid using an LCMV vector beamformer^{43,65} carried out on each 20-minute MEG run separately. The grid was constructed using a template (MNI152) brain, then each participants' anatomical MRI was warped to the template brain and the inverse warp applied to the grid. This ensures comparability of source reconstructions across participants. Because MaxFilter considerably reduces the dimensionality of the data – to approximately 64 – the data covariance matrix was reduced to 50 dimensions using PCA. Eigenvalue decomposition of magnetometer and gradiometer channels was performed in

order to normalise each sensor type to ensure that both sensor types contributed equally to the covariance matrix calculation⁶⁶.

Frontal Eye Field and parietal cortex low-frequency ROI analysis

Frontal Eye Field was defined with reference to a probabilistic atlas of human visual and oculomotor cortical regions⁴². This atlas was used to construct ‘virtual channels’ for left and right FEF based on a symmetric orthogonalization method⁶⁷.

The parietal ROI was defined with respect to both anatomy (the Harvard Oxford cortical atlas) and function (band-limited power at the specific frequency-tag frequencies) – see ‘Frequency tagging analysis’.

Because source analysis creates an inherent sign ambiguity (i.e., a given current in a given direction can be equally well expressed as an equal, opposite current in the opposite direction) all MEG analysis was performed in the frequency domain. Oscillatory power was then computed via time-frequency analysis using a 500ms sliding window multiplied with a Hanning taper, at 0.8Hz frequency resolution. We then averaged across the low-frequency 2.8-8.4Hz band as a proxy for evoked activity. This resulted in FEF-specific single-trial time series.

Inspection of the MEG data revealed two evoked responses with comparable durations; one 100-500ms following the onset of the incoherent motion, and one 100-500ms following the onset of coherent motion. To visualise the effects of task variables on these evoked responses we took the average low-frequency power across each time window for each level of coherence (10%, 30%, 50%, 70%, 90% of dots moved either left or right, fig1E,F), competition (50%, 70% or 90% of coherently-moving dots moved in the same direction, fig1H,I), and prior strength (unsigned value output from Bayesian model $abs(E(q_t) - 0.5)$, figS2). Because the latter is a continuous variable we performed a tertile split into ‘weak’ (p(right) close to 0.5), ‘medium’, and ‘strong’ (p(right) close to 0 or 1) values for visualisation purposes.

To determine the effect of stimulus variables (coherence competition) and prior belief strength on low-frequency power in the FEF and parietal cortex we applied linear regressions with these three predictor variables in two ways: Firstly, for visualisation, at every time point (fig3F,H,J,L,O). Secondly, for statistical purposes, over the average power values in a preselected box from 100-500ms after coherent motion onset (fig3G,I,K,M,P), based on the results from the neural mean field model (fig3D,E,N).

Frequency tagging analysis

To analyse the effect of the rapidly flickering saccade targets on posterior brain regions we used a combination of anatomical and data-driven selection criteria to focus on the brain regions that produced the strongest tagging response. We selected a 3s time-window from 800ms before the onset of coherent motion to 2.2s seconds afterwards - i.e., almost the entire period the targets were flickering - and calculated the fourier transform at all voxels for all participants and averaged across all artifact-free trials. We then compared power at the left-target frequency (41Hz for half the participants and 45Hz for the other half) with power at the right-target frequency, which revealed strong effects of tagging in voxels in parietal and occipital regions. We then used an anatomical atlas to create weighted maps defined by the conjunction of statistically significant differences between tag-frequency power and the parietal cortex anatomical label. The left parietal ROI consisted of all voxels in left parietal cortex that showed significantly greater power at the left-hemisphere tagging frequency than the right, and the reverse was true for the right parietal ROI. We multiplied these weight maps with the source-space data to create 'virtual channels' in left and right parietal cortex. We then performed time-frequency analysis on each virtual channel at the relevant tag frequency (41Hz or 45Hz, depending on the flicker rate of the contralateral frequency tag), using a longer 1000ms sliding window to increase frequency resolution. This produced time-resolved estimates of power at the relevant tag frequency in the left and right parietal ROIs.

For the analysis of prior belief (fig4C,D) we expected effects prior to coherent motion onset and therefore focused on the 1-second epoch of incoherent motion, focusing on a time

window 500-100ms before coherent motion onset due to low SNR in the first part of the incoherent motion epoch as frequency tagging activity ramped up (fig4A).

We performed a tertile split on the prior, but – as we expected *hemispherically lateralised* effects due to the retinotopic organisation of parietal cortex – we used the signed prior, splitting into ‘strong left’, ‘weak’, and ‘strong right’. We then averaged tag power values in the selected time window, at each level of prior and compared these using a 1 x 3 ANOVA with linear contrast. We also conducted a linear regression using the parametric prediction of the prior probability of dots moving right, $E(q_t)$ as per Eq. 5, as explanatory variable, in the same time window.

Because we did not have a strong a priori prediction about the time window in which parietal activity might be driven by belief-decision congruence we used a nonparametric cluster-based permutation test⁶⁸ over all time points to determine whether there was a difference between the congruent and incongruent conditions. We again performed a tertile split on the prior; strong left, weak, strong right, and a three-way split on the dot directions; mostly left motion (10% or 30% right), equal motion (50% right), or mostly right motion (70% or 90% right). We then defined trials where the dominant motion direction and prior agreed as ‘congruent’, and trials where they were opposite as ‘incongruent’, and performed the cluster-based permutation test on the participant conditional means.

Time course analysis

To illustrate how the computational components of active sampling evolve in time we calculated, for every time point in the task epoch (from the onset of incoherent motion until the cessation of coherent motion) repeated measures ANOVAs comparing relevant task variables. We compared, respectively; frequency tagging activity in parietal cortex as a function of prior strength (strong contralateral, weak, strong ipsilateral), low-frequency activity in FEF and parietal cortex as a function of stimulus properties (coherence; 10%, 30%, 50%, 70%, 90%, competition; low, medium, high), and parietal frequency tagging activity as a function of prior/stimulus congruence (congruent, incongruent). To illustrate the temporal cascade of events we found the time point at which each F-statistic was maximal.

Acknowledgements

The authors would like to thank Sven Braeutigam for technical assistance with data collection, and Eelke Spaak for helpful comments on an earlier version of this manuscript.

JOR is supported by a Career Development Fellowship from the Medical Research Council (MR/L019639/1). MR is supported by a PhD studentship from the Wellcome Trust (109064/Z/15/Z). LTH is supported by a Sir Henry Dale Fellowship from the Royal Society and the Wellcome Trust (208789/Z/17/Z). The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust (203139/Z/16/Z).

This research was funded in whole, or in part, by the Wellcome Trust. For the purpose of Open Access, the author has applied a CC-BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

1. Yarbus, A. L. *Eye movements and vision*. (1967). doi:10.1016/0028-3932(68)90012-2
2. Land, M., Mennie, N. & Rusted, J. The roles of vision and eye movements in the control of activities of daily living. *Perception* **28**, 1311–1328 (1999).
3. Itti, L. & Baldi, P. Bayesian surprise attracts human attention. *Vision Res.* **49**, 1295–1306 (2009).
4. Itti, L. & Koch, C. Computational modelling of visual attention. *Nat. Rev. Neurosci.* **2**, 194–203 (2001).
5. Usher, M. & McClelland, J. L. The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
6. Amit, R., Abeles, D., Bar-Gad, I. & Yuval-Greenberg, S. Temporal dynamics of saccades explained by a self-paced process. *Sci. Rep.* (2017). doi:10.1038/s41598-017-00881-7
7. Gottlieb, J. From Thought to Action: The Parietal Cortex as a Bridge between Perception, Action, and Cognition. *Neuron* **53**, 9–16 (2007).
8. Koch, C. & Ullman, S. Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology* **4**, 219–227 (1985).
9. Friston, K., Thornton, C. & Clark, A. Free-energy minimization and the dark-room problem. *Front. Psychol.* **3**, 1–7 (2012).
10. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
11. Hart, E. & Huk, A. C. Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network. *Elife* **9**, 1–22 (2020).
12. Wimmer, K., Nykamp, D. Q., Constantinidis, C. & Compte, A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.* **17**, 431–439 (2014).
13. Esnaola-Acebes, J. M., Roxin, A. & Wimmer, K. Bump attractor dynamics underlying stimulus integration in perceptual estimation tasks. *bioRxiv* 1–31 (2021).
14. Schall, J. D., Purcell, B. A., Heitz, R. P., Logan, G. D. & Palmeri, T. J. Neural mechanisms of saccade target selection: Gated accumulator model of the visual-motor cascade. *Eur. J. Neurosci.* (2011). doi:10.1111/j.1460-9568.2011.07715.x
15. Cohen, J. Y. *et al.* Cooperation and competition among frontal eye field neurons during

- visual target selection. *J. Neurosci.* (2010). doi:10.1523/JNEUROSCI.4600-09.2010
16. Kennard, M. A. Alterations in response to visual stimuli following lesions of frontal lobe in Monkeys. *Arch. Neurol. Psychiatry* **41**, 1153–1164 (1939).
 17. Heilman, K. M. & Watson, R. T. Mechanisms underlying the unilateral neglect Syndrome. *Adv. Neurol.* **18**, 93–106 (1977).
 18. Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* **2**, 176–185 (1999).
 19. Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
 20. Ding, L. & Gold, J. I. Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* **22**, 1052–1067 (2012).
 21. Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* **86**, 1916–1936 (2001).
 22. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
 23. Knowlton, B. J., Squire, L. R. & Gluck, M. A. Probabilistic classification learning in amnesia. *Learn. Mem.* **1**, 106–120 (1994).
 24. Dorris, M. C. & Glimcher, P. W. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* (2004). doi:10.1016/j.neuron.2004.09.009
 25. Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D. & Brody, C. D. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* **4**, 1–28 (2015).
 26. Katz, L. N., Yates, J. L., Pillow, J. W. & Huk, A. C. Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* **535**, 285–288 (2016).
 27. Piet, A. T., Erlich, J. C., Kopec, C. D. & Brody, C. D. Rat Prefrontal Cortex Inactivations during Decision Making Are Explained by Bistable Attractor Dynamics. *Neural Comput.* **29**, 2861–2886 (2017).
 28. Hanks, T. D. *et al.* Distinct relationships of parietal and prefrontal cortices to evidence

- accumulation. *Nature* **520**, 220–3 (2015).
29. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).
 30. Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**, 4745–4765 (1992).
 31. Soltani, A. & Wang, X. J. Synaptic computation underlying probabilistic inference. *Nat. Neurosci.* **13**, 112–119 (2010).
 32. Wong, K.-F., Huk, A. C., Shadlen, M. N. & Wang, X.-J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* **1**, 6 (2007).
 33. Hunt, L. T. *et al.* Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* **15**, 470–476 (2012).
 34. O'Reilly, J. X., Jbabdi, S., Rushworth, M. F. S. & Behrens, T. E. J. Brain Systems for Probabilistic and Dynamic Prediction: Computational Specificity and Integration. *PLoS Biol.* **11**, (2013).
 35. Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–21 (2007).
 36. Cox, R. T. Probability, Frequency and Reasonable Expectation. *Am. J. Phys.* (1946). doi:10.1119/1.1990764
 37. Bonaiuto, J. J., De Berker, A. & Bestmann, S. Response repetition biases in human perceptual decisions are explained by activity decay in competitive attractor models. *Elife* (2016). doi:10.7554/eLife.20047
 38. Hunt, L. T., Behrens, T. E. J., Hosokawa, T., Wallis, J. D. & Kennerley, S. W. Capturing the temporal evolution of choice across prefrontal cortex. *Elife* **4**, e11945 (2015).
 39. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352 (2011).
 40. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
 41. Hauser, C. K., Zhu, D., Stanford, T. R. & Salinas, E. Motor selection dynamics in FEF

- explain the reaction time variance of saccades to single targets. *Elife* (2018).
doi:10.7554/eLife.33456
42. Wang, L., Mruczek, R. E. ., Arcaro, M. J. & Kastner, S. Probabilistic maps of visual topography in human cortex. *Cereb. Cortex* **25**, 3911–3931 (2015).
 43. Van Veen, B. D., van Drongelen, W., Yuchtman, M. & Suzuki, A. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Trans. Biomed. Eng.* **44**, 867–80 (1997).
 44. Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J. & Lounasmaa, O. V. Magnetoencephalography theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev. Mod. Phys.* **65**, 413–497 (1993).
 45. Stokes, M. G. ‘Activity-silent’ working memory in prefrontal cortex: A dynamic coding framework. *Trends Cogn. Sci.* **19**, 394–405 (2015).
 46. Barbosa, J. *et al.* Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. *Nat. Neurosci.* **23**, 1016–1024 (2020).
 47. Adrian, E. . D. & Matthews, B. H. C. The Berger Rhythm: Potential Changes from the Occipital Lobes in Man. *Brain* **57**, 355–385 (1934).
 48. Müller, M. M., Malinowski, P., Gruber, T. & Hillyard, S. A. Sustained division of the attentional spotlight. *Nature* **424**, 309–312 (2003).
 49. Zhigalov, A., Herring, J. D., Herpers, J., Bergmann, T. O. & Jensen, O. Probing cortical excitability using rapid frequency tagging. *Neuroimage* **195**, 59–66 (2019).
 50. Drijvers, L., Jensen, O. & Spaak, E. Rapid invisible frequency tagging reveals nonlinear integration of auditory and visual information. *Hum. Brain Mapp.* 1–15 (2020).
doi:10.1002/hbm.25282
 51. Moratti, S., Clementz, B. A., Gao, Y., Ortiz, T. & Keil, A. Neural mechanisms of evoked oscillations: Stability and interaction with transient events. *Hum. Brain Mapp.* **28**, 1318–1333 (2007).
 52. Thomas, N. W. D. & Paré, M. Temporal processing of saccade targets in parietal cortex area LIP during visual search. *J. Neurophysiol.* **97**, 942–947 (2007).
 53. Colby, C. L., Duhamel, J. R. & Goldberg, M. E. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J. Neurophysiol.* **76**, 2841–2852 (1996).

54. Bisley, J. W. & Goldberg, M. E. Neuronal Activity in the Lateral Intraparietal Area and Spatial Attention. *Science* (80-.). **299**, 81–86 (2003).
55. Ipata, A. E., Gee, A. L., Gottlieb, J. P., Bisley, J. W. & Goldberg, M. E. LIP responses to a popout stimulus are reduced if it is overtly ignored. *Nat. Neurosci.* **9**, 1071–1076 (2006).
56. Mars, R. B. *et al.* Diffusion-Weighted Imaging Tractography-Based Parcellation of the Human Parietal Cortex and Comparison with Human and Macaque Resting-State Functional Connectivity. *J. Neurosci.* **31**, 4087–4100 (2011).
57. Bruce, C. J. & Goldberg, M. E. Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* **53**, 603–635 (1985).
58. Bruce, C. J., Goldberg, M. E., Bushnell, M. C. & Stanton, G. B. Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J. Neurophysiol.* **54**, 714–734 (1985).
59. Crapse, T. B. & Sommer, M. A. Frontal eye field neurons with spatial representations predicted by their subcortical input. *J. Neurosci.* **29**, 5308–18 (2009).
60. Murray, J. D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).
61. Wang, X. Macroscopic gradients of synaptic excitation and inhibition in the neocortex. *Nat. Rev. Neurosci.* **21**, 169–178 (2020).
62. Wong, K.-F. & Wang, X.-J. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
63. Rosner, B. Percentage points for a generalized esd many-outlier procedure. *Technometrics* **25**, 165–172 (1983).
64. Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J.-M. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* **2011**, 156869 (2011).
65. Woolrich, M. W., Hunt, L. T., Groves, A. & Barnes, G. R. MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization. *Neuroimage* **57**, 1466–1479 (2011).
66. Quinn, A. J. *et al.* Task-evoked dynamic network analysis through Hidden Markov Modeling. *Front. Neurosci.* **12**, 1–17 (2018).
67. Colclough, G. L., Brookes, M. J., Smith, S. M. & Woolrich, M. W. A symmetric

- multivariate leakage correction for MEG connectomes. *Neuroimage* **117**, 439–448 (2015).
68. Maris, E. & Oostenveld, R. Nonparametric statistical testing of EEG- and MEG-data. *J. Neurosci. Methods* **164**, 177–90 (2007).
69. Mesulam, M.-M. Spatial attention and neglect: parietal, frontal and cingulate contributions to the mental representation and attentional targeting of salient extrapersonal events. *Philos. Trans. R. Soc. London* **354**, 1325–1346 (1999).

Supplementary Information

Lagged logistic regression reveals integration kernels over previous trials

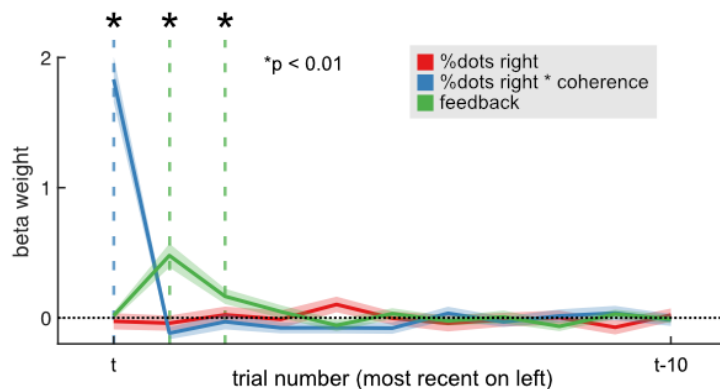


Figure S1: Lagged logistic regression of current and previous task variables on current choice. Error bars indicate standard error of the mean across participants.

To determine whether participants indeed integrated information from previous trials we performed a logistic regression with dependent variable; saccade direction (left, right), and independent variables; proportion of coherent dots moving right, interaction of proportion coherent dots moving right with coherence (effectively the absolute number of dots moving right) and post-trial feedback. To determine whether task parameters on *previous* trials influenced choice on the *current* trial we included the above regressors at ‘lags’ ranging from zero (i.e., the task parameters on the current trial t) to 10 (parameters on trial $t-10$).

Lagged multiple regression revealed that absolute number of right-moving dots on the current trial strongly predicted choice on the current trial ($t(25) = 10.51$, $p = 3.2e-11$), but on previous trials (lags > 0) did not predict current choice. In contrast, feedback on trial t did not predict choice on trial t ($t(25) = 0.56$, $p = 0.6$); this was entirely expected since feedback *followed* choice. However, feedback on trial $t-1$ ($t(25) = 5.2$, $p = 2e-5$) and trial $t-2$ ($t(25) = 2.8$, $p = 0.009$) did predict choice on trial t . Proportion of dots right on current or previous trials did not predict current choice beyond the other regressors (all $t < 1.8$).

The above strongly indicates that participants’ choice on trial t was influenced by the stimulus properties on trial t , and the history of post-trial feedback on previous trials $t-1$ and $t-2$.

No evidence that prior belief strength modulates low-frequency activity in frontal or parietal cortex

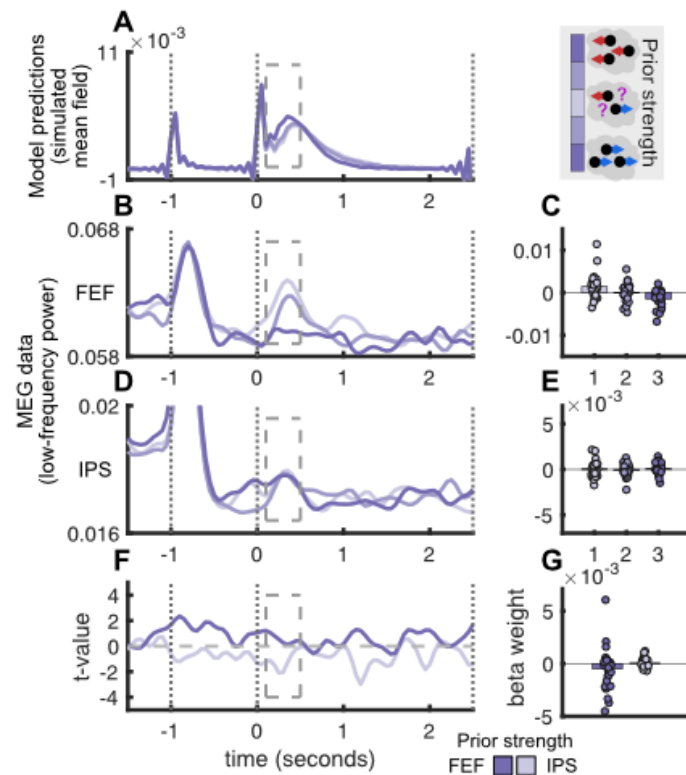


Fig S2: Signals related to strength of prior belief. All plotting conventions as fig3, main text. A) Neural mean-field model predicts a parametric increase in activity as a function of strength of prior belief. B,C: FEF low-frequency power. D,E: Parietal cortex low-frequency power. F,G: Results from general linear model.

Our neural mean field model (fig3A, main text) operationalised prior belief as a weak input in the prestimulus period. However, rather than showing strong changes in *prestimulus* signal, the model rather exhibited changes in *stimulus-evoked* activity, presumably by biasing the prestimulus state. A strong prestimulus input produced the strongest stimulus-evoked activity (figS1A). However, in contrast to this prediction we did not observe significant changes in stimulus-evoked activity in the MEG data. In FEF we observed a non-significant effect in the opposite direction; stimulus-evoked activity was strongest on trials where participants had a *weak* prior belief close to 50/50 (general linear model, $t(25) = -1.22$, $p = 0.23$, figS2B,C). In parietal cortex no trend was observed in either direction ($t(25) = 0.8$, $p = 0.43$, figS2D,E).

Spatial and spectral resolution of band-limited power at the two tag frequencies

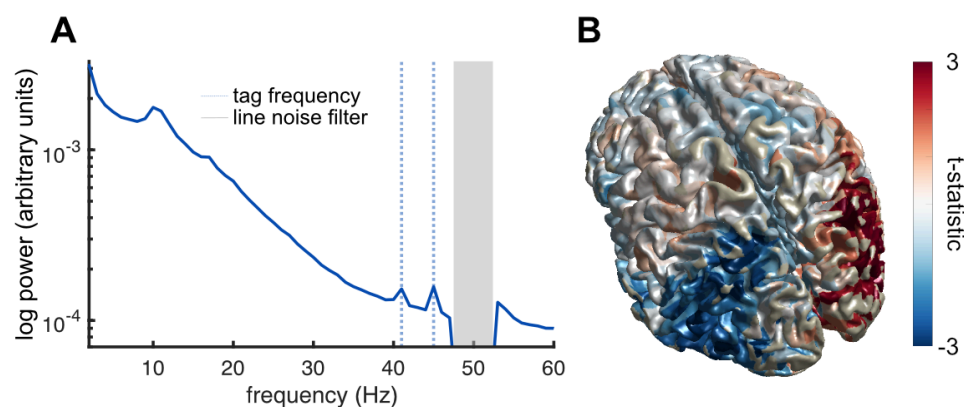


Figure S3: A) Average Fourier spectrum of left and right parietal cortex ROIs displays clear peaks at the 41Hz and 45Hz tagging frequencies (dashed blue vertical lines). Note grey box indicates line noise filter. B) Unthresholded statistical map of differences between 'power at left tag frequency' and 'power at right tag frequency'. Only parietal and lateral occipital regions showed lateralised differences in tag frequency power, presumably due to weaker SNR in frontal sources.

Effects of stimulus variables (coherence, competition) on frequency-tagging activity in parietal cortex

For comparison with the analysis of stimulus-driven effects (fig3) we applied the same model to the frequency tagging data from parietal cortex, pooled across left and right hemisphere ROIs. As we were able to observe lateralized effects of the prior in parietal cortex (fig4), we also tested for a lateralized effect of evidence, i.e., whether ‘more dots moving towards the contralateral side’ produced a stronger tag response than ‘more dots moving towards the ipsilateral side’.

Visually inspection suggested no strong effects of either coherence (figS4A,B) or competition (figS4C,D) on or signed/lateralized competition (figS4E,F) in the frequency tagging signal. To confirm this we fit a general linear model with these three effects to the frequency tagging data on single trials, then comparing the beta-weights from the first level model to zero. No significant effects were observed (coherence, $t(25) = -0.37$, $p = 0.72$; competition, $t(25) = -0.93$, $p = 0.36$; lateralized competition, $t(25) = 1.02$, $p = 0.32$).




Figure S4: Analysis of parietal cortex frequency tagging data analogous to low-frequencies (fig3)
A) Parietal frequency tagging time series as a function of stimulus coherence. B) Averaged across 100-500ms time window, dashed box in A. C,D) As A,B, but as a function of stimulus competition. E,F) As A,B, but as a function of ‘signed competition’, i.e., proportion of coherent motion moving contralateral to region of interest.

Ability to detect lateralised task processing in Frontal Eye Fields and Parietal Cortex

Frontal Eye Field and LIP neurons predominantly code for saccades to the contralateral hemispace^{57,58}. Therefore we might have expected the left- and right- hemisphere regions of interest to show lateralized effects in addition to the overall evidence accumulation signals reported in figure 3.

To obtain lateralized predictions, we looked at the inputs to individual pools in the neural mean-field model, rather than the mean field produced by summing all inputs (see main text). In the model our prediction was confirmed; as contralateral coherent motion increased, the simulated evoked response grew larger (figS5A,B). However, we did not observe a comparable effect in the FEF MEG data (fig S5C-F). Observation of the evoked-response window suggested a U-shaped pattern, where strong evidence for *either* motion direction (10%, 90% conditions) produced a large evoked response, whereas high competition (50% condition) produced the weakest evoked response.

Fitting a general linear model with % coherent dots moving right as a regressor confirmed an absence of evidence for left and right FEF acting as individual evidence accumulators. No significant linear effect was observed in either the left FEF ($t(25) = -0.61$, $p = 0.55$) or right FEF ($t(25) = -0.96$, $p = 0.35$). In fact qualitatively the left FEF data appear to show a quadratic effect entirely consistent with the main effect of competition (fig3); trials with the largest number of dots moving in one direction produced the strongest evoked responses, and 50/50 trials with maximum competition produced the weakest.

In parietal cortex (figS5G-J) a different pattern was observed: Here a linear effect of proportion coherent dots moving right was observed in the left ($t(25) = 3.32$, $p = 0.003$) but not the right hemisphere ($t(25) = -0.86$, $p = 0.40$).

The absence of evidence for our two predictions in FEF could be due to a lack of hemispherically-lateralised decision processing in FEF. It may be the case that neurons accumulating evidence for both choice options are present in both hemispheres. However, an alternative possibility is that we are unable to fully disentangle signals from left and right

FEF via beamforming⁴³. Similarly, in parietal cortex neurons responding to left- and right visual fields are not completely segregated, with the left hemisphere showing a more lateralized response than the right⁶⁹, perhaps consistent with the fact that the lateralized signal in parietal cortex was strong in the left hemisphere. However, again the limitations of beamforming may partly explain the difficulty in detecting lateralized responses.




Figure S5: A) Simulated synaptic inputs for right-motion accumulator of neural mean-field model, as a function of percentage coherent dots moving right. B) As A, but left-motion accumulator. C) Source-reconstructed activity from left FEF, same colour convention. D) Individual participant data in stimulus-evoked window. E-F) As C-D, but right FEF. G-J) As C-F but parietal cortex.

Additionally, although we did not detect an overall effect of the prior on evidence accumulation signals in the ERF, for completeness we tested whether a lateralized effect of the prior could be found. To generate predictions we re-ran our neural mean-field model with five levels of pre-stimulus input, ranging from ‘strong left’ (strong input to left motion pool, no input to right motion pool), to ‘strong right’ (reversed). This produced a linear modulation of the simulated evoked response that was proportional to the strength of the pre-stimulus input (figS6A,B).

To compare with our MEG data we fit a general linear model, with prior belief predicted by the Bayesian model as a regressor. However, our prediction was not confirmed (figS6C-F). In contrast to the neural mean-field model, no evidence was observed of linear trends in the MEG data as a function of prior belief (left, $t(25) = 1.46$, $p = 0.16$, right, $t(25) = -0.52$, $p = 0.61$, figS6C,F). Qualitatively, a quadratic effect was observed whereby weak prior belief produced a strong evoked response and strong belief in either direction produced a weak response. This could reflect an overall process by which a competitive system in FEF is engaged more strongly when the prior is weak; however, this prediction was outside the scope of our model.




Figure S6: A) Simulated synaptic inputs for right-motion accumulator of neural mean-field model as a function of prior belief. B) Source reconstructed activity in left FEF. C) Individual participant data in stimulus-evoked window. D-F) As A-C, but left-motion accumulator and right FEF.

In parietal cortex we observed a different pattern: Here, no significant linear effect was observed in left parietal cortex ($t(25) = 0.88$, $p = 0.39$, figS6G,H). A significant trend was observed in right parietal cortex ($t(25) = 2.18$, $p = 0.039$, figS6I,J), but this was in the opposite direction to the prediction from the neural mean field model, with strong *ipsilateral* (rightward) belief showing the largest response.