

Decision-making in dynamic, continuously evolving environments: quantifying the flexibility of human choice

Maria Ruesseler¹, Lilian AE Weber^{1,2}, Tom R Marshall², Jill X O'Reilly², Laurence T Hunt^{1,2*}

- 1. Wellcome Centre for Integrative Neuroimaging, Department of Psychiatry, University of Oxford.** *Oxford Centre for Human Brain Activity (OHBA) University Department of Psychiatry Warneford Hospital, Oxford,, Oxford OX3 7JX*
- 2. Department of Experimental Psychology, University of Oxford.** *Anna Watts Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom*

* correspondence: laurence.hunt@psy.ox.ac.uk

Keywords: decision making; temporal response function; EEG; random dot kinetogram

Abstract

During perceptual decision-making tasks, centroparietal EEG potentials report an evidence accumulation-to-bound process that is time locked to trial onset. However, decisions in real-world environments are rarely confined to discrete trials; they instead unfold continuously, with accumulation of time-varying evidence being recency-weighted towards its immediate past. Confronted with time-varying stimuli, humans can appropriately adapt their weighting of recent evidence according to the statistics of the environment. The neural mechanisms supporting this adaptation currently remain unclear. Here, we show that humans' ability to adapt evidence weighting to different sensory environments is reflected in changes in centroparietal EEG potentials. We use a novel continuous task design to show that the Centroparietal Positivity (CPP) becomes more sensitive to fluctuations in sensory evidence when large shifts in evidence are less frequent, and is primarily sensitive to fluctuations in decision-relevant (not decision-irrelevant) sensory input. A complementary triphasic component over parietal cortex encodes the sum of recently accumulated sensory evidence, and its magnitude covaries with the duration over which different individuals integrate sensory evidence. Our findings reveal how adaptations in centroparietal responses reflect flexibility in evidence accumulation to the statistics of dynamic sensory environments.

Introduction

Unlike in most experiments, the choices that we make in daily life rarely occur in discrete trials. Naturalistic decisions instead arise organically and continuously in dynamic environments that evolve over time (Huk et al., 2018; Hunt et al., 2021), leading to uncertainty regarding the onset of decision-relevant changes in the environment (Orsolich et al., 2021; Shinn et al., 2022). When making decisions in such dynamic environments, more recently presented evidence should therefore usually be given greater weight and more historical evidence gradually discounted – a strategy known as ‘leaky’ evidence accumulation. Optimising leaky evidence accumulation involves adapting one’s behaviour to the overall statistics of the environment. This can be achieved by changing the rate at which previous evidence is leaked (the ‘decay’) and the amount of cumulative evidence required before a categorical decision is made (the ‘decision threshold’) (Glaze et al., 2015; Kilpatrick et al., 2019; Veliz-Cuba et al., 2016).

While it is known that humans (Glaze et al., 2018; Ossmy et al., 2013) and other animals (Levi et al., 2018; Piet et al., 2018) can adapt the decay and decision threshold of sensory evidence accumulation to different dynamic environments, the neural mechanisms that underlie this adaptation remain unclear. In conventional trial-based paradigms, neurophysiological correlates of perceptual decision-making have been well characterised, particularly signals that resemble an evidence accumulation-to-bound process (Gold and Shadlen, 2007; Hanks and Summerfield, 2017; O’Connell and Kelly, 2021). Two of the best-studied human electroencephalographic (EEG) correlates of decision formation are an effector-independent centroparietal positivity (CPP) that shows accumulator-like dynamics during decision formation (Kelly and O’Connell, 2013; O’Connell et al., 2012; Pisauro et al., 2017; Twomey et al., 2015), and motor preparation signals that emerge prior to a response (Donner et al., 2009; Steinemann et al., 2018; Wyart et al., 2012). Both signals can adapt their properties according to the overall statistics of a task. For example, pre-trial CPP amplitude is increased by emphasising speed over accuracy (Steinemann et al., 2018), while pre-trial motor lateralisation can reflect the prior expectation of a leftward or rightward action in the upcoming trial (de Lange et al., 2013; Kelly et al., 2021). Yet it remains unclear whether and how these signals reflect the ability to behaviourally adapt the decay of past information in continuous (dynamic) environments, which are more akin to many decisions faced in naturalistic settings.

Perhaps one reason why decision making in dynamic environments has been less studied than trial-based choice is the uncertainty concerning how best to analyse the time-varying neural data. For example, without clearly defined discrete trials, it appears unclear to which timepoint data should be epoched. As the stimulus is continuously changing, it is also ambiguous how to disentangle responses to previous versus current sensory evidence, which may be overlapping in time. However, recent innovations in trial-based task design (Brunton et al., 2013; Cheadle et al., 2014; Orsolich et al., 2021) and unmixing of overlapping EEG responses (Crosse et al., 2016; Ehinger and Dimigen, 2019; Hassall et al., 2021; Smith and Kutas, 2015) have suggested potential solutions to some of these challenges. By tightly controlling how sensory evidence fluctuates over time, it becomes possible to relate moment-to-moment stimulus fluctuations to subsequent behavioural and neural responses. In addition, by using data analysis techniques that explicitly target overlapping neural responses, it is also possible to establish the Temporal Response Function (TRF) to each new fluctuation in a continuous sensory evidence stream (Gonçalves et al., 2014). By combining these two approaches, we hypothesised that we would be able to characterise decision-related EEG responses in a continuous and dynamic setting, even in the absence of repeated experimental trials.

In this study, we examine how the EEG response to evidence fluctuations during a continuous perceptual decision task is affected by the overall statistics of the sensory environment. Participants were trained to attend to a continuously changing sensory evidence stream, in which brief “response periods” were embedded that were reported via button-press. We demonstrate a CPP-like potential that is sensitive to each fluctuation in the stream of continuous sensory evidence, and a triphasic centroparietal component that reflects the integrated sum of recently presented evidence. We then

show that subjects' behaviour adapts appropriately to different sensory environments, and that changes in centroparietal and motor preparation prior to button press signals reflect adaptation of leaky sensory evidence integration to different environments. This is not a simple feature of adaptation to sensory surprise, as the CPP-like potential largely responds to decision-relevant, not decision-irrelevant, evidence fluctuations. We also show substantial variability in the time constant of leaky evidence accumulation across subjects, which predicts the amplitude of the triphasic centroparietal TRF to absolute recent sensory evidence. Collectively these results provide a neural characterisation of human decision-making in a dynamic, continuously evolving perceptual environment and how this can adapt to the overall statistics of the environment.

Results

A novel task for exploring behavioural and EEG adaptations to the statistics of dynamic sensory environments

To study evidence accumulation in a continuous setting, we designed a novel variant of the classic random dot kinematogram (RDK) paradigm (Britten et al., 1992; Donner et al., 2009; Kelly and O’Connell, 2013; Newsome and Pare, 1988). Subjects continuously monitored a stream of time-varying sensory evidence for blocks of five minutes (fig 1a). During extended “baseline periods” (grey shaded area in fig 1b), the average level of motion coherence (black line) in the stimulus was zero, whereas during shorter intermittent “response periods” (green shaded area), the mean level of motion coherence became non-zero (either 30, 40, or 50% motion coherence). The participants’ task was to report whenever they detected such a period of coherent motion using a left or right button press. Importantly, the onset of response periods was not explicitly signalled to the participant. If they responded accurately (during a response period or within 500ms of it ending), they received a reward (+3 points); if they failed to report a response period (‘missed response period’), or they responded during a baseline period (‘false alarm’), they received a small punishment (-1.5 points). Participants also received a larger punishment (-3 points) if they reported the incorrect motion direction during a response period; in practice, such errors were very rare. Feedback was presented by changing the colour of the central fixation point (fig 1a). The accumulated total points were then converted into a monetary pay-out at the end of the task. Participants completed six runs, each consisting of four 5-minute blocks; they were given a short break between each block and a longer break between runs.

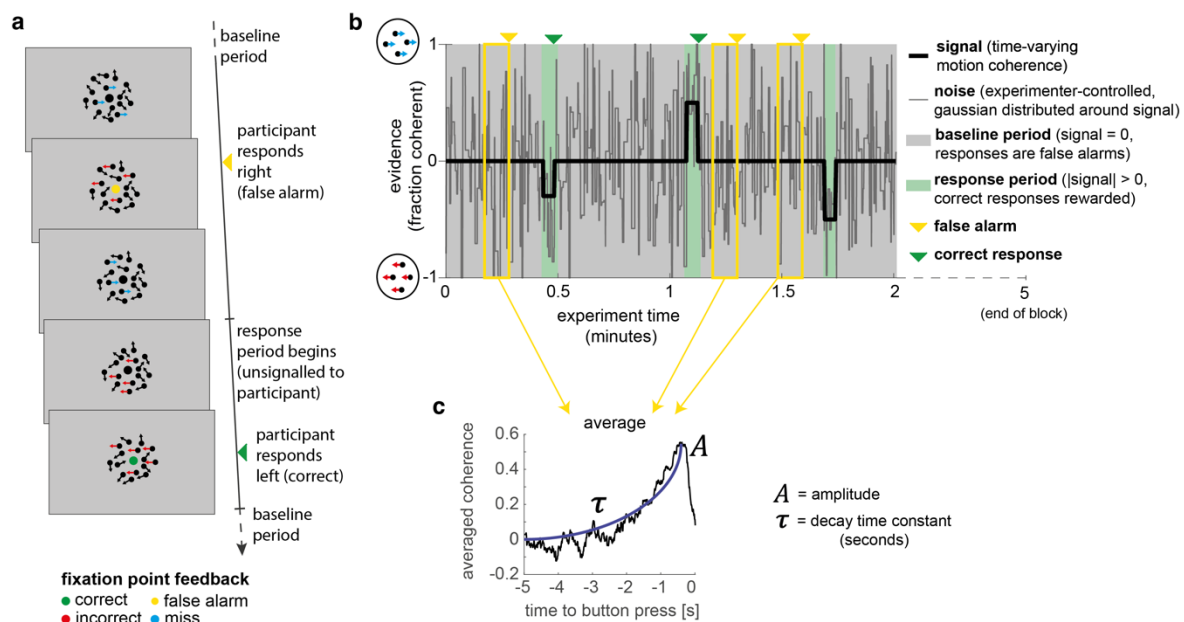


Figure 1 – A novel, continuous version of the random dot kinematogram (RDK) paradigm allows empirical measurement of participants’ leaky evidence integration kernels in dynamic environments.

(a) Task design. Participants continuously attend to a centrally presented RDK stimulus, for 5 minutes at a time. They aim to successfully report motion direction during ‘response periods’ (when coherent motion signal is on average non-zero) and withhold responding during ‘baseline periods’ (when signal is on average zero). (b) Task structure (example block; response periods are ‘rare’). During both baseline (grey) and response periods (green), the signal (black line) is corrupted with experimenter-controlled noise (grey line). The noise fluctuations that precede each response (arrows) can be averaged to obtain the evidence integration kernel. (c) The resulting evidence integration kernel is well described by an exponential decay function, whose decay time constant in seconds is controlled by the free parameter τ .

Crucially, the net motion presented to the participant on each frame of the stimulus was not the *average* level of motion coherence (black line in fig 1b), but instead was a noisy sample from a Gaussian distribution about this mean (grey line). This noisy sample was resampled on average every 280ms (inter-sample interval drawn from an exponential distribution, truncated at 1000ms). This “experimenter-controlled sensory noise” confers several benefits.

First, the injection of sensory noise places a stronger demand on temporal evidence integration than a classical RDK task. This is because any individual period of strong motion evidence could be driven by a noisy sample during a baseline period, rather than necessarily signalling the onset of a response period (fig 1b). As such, continuous and temporally extended integration is essential to successfully disambiguate changes in the mean from noisy samples around the baseline. Indeed, participants would occasionally make ‘false alarms’ during baseline periods in which the structure of the preceding noise stream mistakenly convinced them they were in a response period (see figure 4, below).

Second, the noise fluctuations allow a “reverse correlation” approach to studying subjects’ evidence integration. Simply by averaging the noisy stimulus that was presented prior to each response, we could extract an ‘integration kernel’ that empirically reveals how far back in time evidence is being integrated – in other words, how quickly previous evidence is decaying in the participant’s mind – and how strong this evidence needed to be on average to support a choice. We performed this reverse correlation for both false alarm responses (example shown in fig 1c) and correct responses. The fact that evidence integration kernels naturally arise from false alarms, in the same manner as from correct responses, demonstrates that false alarms were not due to motor noise or other spurious causes. Instead, false alarms were driven by participants treating noise fluctuations during baseline periods as sensory evidence to be integrated across time.

Finally, and perhaps most importantly, the experimenter-controlled sensory noise allows us to characterise how continuous sensory evidence fluctuations cause changes in the simultaneously recorded continuous EEG signal. To study this, we used a deconvolutional general linear model (GLM) approach (Crosse et al., 2016; Ehinger and Dimigen, 2019; Gonçalves et al., 2014; Hassall et al., 2021) to estimate temporal response functions (TRFs) to various events relating to the time-varying sensory evidence. We describe this approach and the resulting TRFs in more detail below.

Behavioural adaptations to environments with different statistical properties

We used this paradigm to investigate whether and how participants adapted their evidence integration behaviour to the overall statistics of the sensory environment. To test this, we manipulated both the duration and frequency of ‘response periods’ in the task. We hypothesised that this would affect the decay of past sensory evidence and/or the decision threshold used to commit to a response. Importantly for our subsequent analyses, we kept the generative statistics of the gaussian noise during ‘baseline periods’ consistent across conditions. This allowed us to directly compare behavioural evidence integration kernels for false alarms and EEG TRFs across conditions without any potential confound from how the noise was structured.

Within each 20-minute run, participants completed four pseudorandomly ordered 5-minute blocks drawn from a 2*2 factorial design (fig 2a). Response periods were either LONG (5 seconds) or SHORT (3 seconds), and either FREQUENT (baseline periods between 3-8 seconds in duration) or RARE (baseline periods between 5-40 seconds). Participants were extensively trained on these trial statistics prior to completing the task, and were then explicitly cued which environment they were currently in. As a consequence, participants neither had to learn nor infer the higher-order statistics of the sensory environment during the task; instead, they had to adapt their decision behaviour according to the pre-learnt statistics of the cued environment.

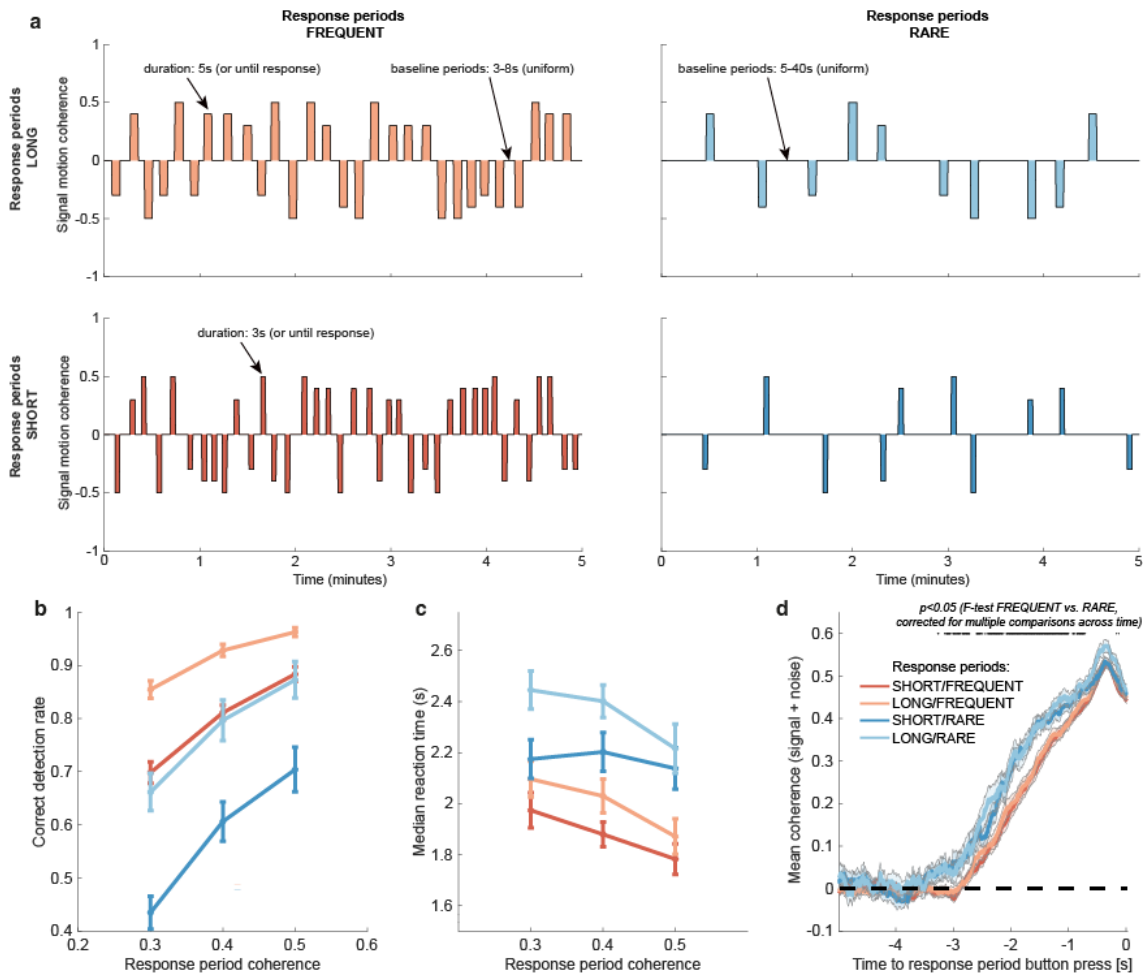


Figure 2 – Variations in response period structure across different environments elicit behavioural adaptations in decision making. (a) Structure of response periods (signal only, before noise was added to the stimulus stream) across the different environments. This was manipulated in a 2*2 design, where response periods were either FREQUENT or RARE, and LONG or SHORT. Participants were extensively trained on these statistics prior to the task, and the current environment was explicitly cued to the participant. **(b)** Correct detection rate for all response periods. Participants successfully detected more response periods when they were LONG than SHORT (as would be expected, because the response period is longer), but also detected more when they were FREQUENT than RARE. **(c)** Median reaction time for successfully reported response periods across the four conditions. Participants took longer in RARE vs FREQUENT conditions, and in LONG vs SHORT conditions. **(d)** Integration kernels for ‘response periods’ shows a main effect of FREQUENT vs. RARE response periods, but unexpectedly no effect of LONG vs. SHORT response periods. All plots show mean +/- s.e. across participants. Note that to make reaction times and integration kernels comparable between the four conditions, we only include those responses that were shorter than 3.5s in analyses for (c) and (d) (that is, the maximum response time in SHORT response periods).

Response periods. We first tested whether participants adjusted their behaviour across the four conditions by analysing detection behaviour for response periods. We used a three-way repeated measures ANOVA to test for effects of motion coherence, response period length, and response period frequency.

As expected, participants were faster ($F(2,46)=29.51$, $p=7.45*10^{-9}$) and more accurate ($F(2,46)= 17.6$, $p= 0.00035$) at detecting response periods with a higher level of average motion coherence (fig 2b/c).

Examining the effects of response period frequency, participants were less likely ($F(1,23)=41.99, p=1.30*10^{-6}$) and slower ($F(1,23)=83.24, p=6.23*10^{-9}$) to detect response periods when these were RARE (blue lines in fig 2b/c) than when they were FREQUENT (red lines in fig 2b/c). We hypothesised that this could be explained by participants requiring a stronger overall level of recently accumulated evidence before committing to a response when response periods were rare. Indeed, when we examined the ‘integration kernel’ of average sensory evidence for successful responses, significantly more cumulative evidence was required for RARE than FREQUENT response periods, stretching back up to 3.5s prior to the commitment to a response (fig 2d; F-test with permutation-based correction for multiple comparisons, $p<0.05$).

We also found that participants detected response periods more frequently when these were LONG than when they were SHORT (fig 2b; $F(1,23)=178.52, p=2.51*10^{-12}$). This is unsurprising, as there was simply more time to detect the change in motion during these response periods. In fact, participants were slightly more conservative in their responding when trials were LONG, as shown by a longer average reaction time for LONG response periods relative to SHORT (even when restricting this analysis to focus on LONG responses that were less than 3.5s, the maximum possible SHORT response duration including 500ms response tolerance period; fig 2c; $F(1,23)=70.00, p=2.78*10^{-8}$). This was also reflected in their false alarm frequency, as shown below. Surprisingly, the sensory evidence integration properties (i.e. the ‘integration kernels’) were *not* affected by length of response periods (fig 2d). This ran contrary to our initial hypothesis that participants would integrate evidence for longer when response periods were LONG. We suggest that this may result from the manipulation of response period duration being relatively small (3s vs. 5s) compared to the manipulation of response period frequency.

False alarms. We then examined whether false alarms differed across conditions, testing for effects of response period frequency and length on false alarm rate using a two-way repeated measures ANOVA. We found that despite the structure of the noise stream being identical across the four conditions, there was a lower overall frequency of false alarms in LONG versus SHORT conditions (Figure 3a; $F(1,23)=58.67, p=8.98*10^{-8}$). This provides further evidence that participants were overall more conservative in their responses in LONG conditions than SHORT. There was no effect of response period frequency on false alarm rate ($F(1,23)=0.37, p=0.55$).

We then examined what *caused* participants to false alarm during baseline periods. Were participants still integrating evidence continuously during these periods of the task, or might false alarms be driven by other spurious factors, such as motor noise? We tested this by calculating integration kernels derived from these responses. We found the recovered evidence integration kernels showed exponential decay weighting, implying that participants were indeed performing continuous evidence integration throughout baseline periods as well as response periods, and that evidence accumulation was more temporally extended when response periods were RARE rather than FREQUENT (figure 3b).

Between-participant variation in evidence integration. We then sought to characterise the time constant of leaky evidence integration within each individual participant. To do this, we fit an exponential decay model (fig 1c) to the empirical integration kernel from false alarm responses:

$$k(t) = Ae^{-\frac{t}{\tau}}$$

where $k(t)$ is the height of the integration kernel t seconds before its peak; A is the peak amplitude of the integration kernel (in units that denote the fraction of dots moving towards the chosen response direction), and τ is the decay time constant (in units of seconds). Our exponential decay model provided a good fit to data at a single subject level (see Extended Data figure 3.1 for example fits), as demonstrated by the strong reliability across conditions for both A and τ (Fig 3c and Extended Data Figure 3.2a; Pearson’s correlation between SHORT and LONG conditions: τ : $R(23) = 0.71$; A : $R(23) = 0.72$; Pearson’s correlation between RARE and FREQUENT conditions: τ : $R(23) = 0.87$; A : $R(23) = 0.81$; all $p<0.0001$). Indeed, a striking feature of these integration kernels was that variation across *individuals* exceeded the variation observed across *conditions* (see fig 3c, for example).

Consistent with our earlier analyses (figure 2d, figure 3b), we found that by fitting this single-subject model, τ was significantly longer when response periods were RARE than FREQUENT (paired $T(23)=3.62$, $p=0.0014$; fig 3c) but A did not differ between these conditions (paired $T(23)=0.03$, $p=0.97$; Extended Data fig 3.1a). Again consistent with our analyses of behaviour during response periods (figure 2d), there was no difference between these parameters for LONG versus SHORT response periods (τ : paired $T(23)=0.82$, $p=0.42$; A : paired $T(23)=-0.97$, $p=0.34$; Extended Data fig 3.1a).

In summary, these results indicate that participants adapted to response periods being rarer by accumulating sensory evidence with a longer time constant of integration, but that there was also substantial between-subject variability in evidence accumulation across participants.


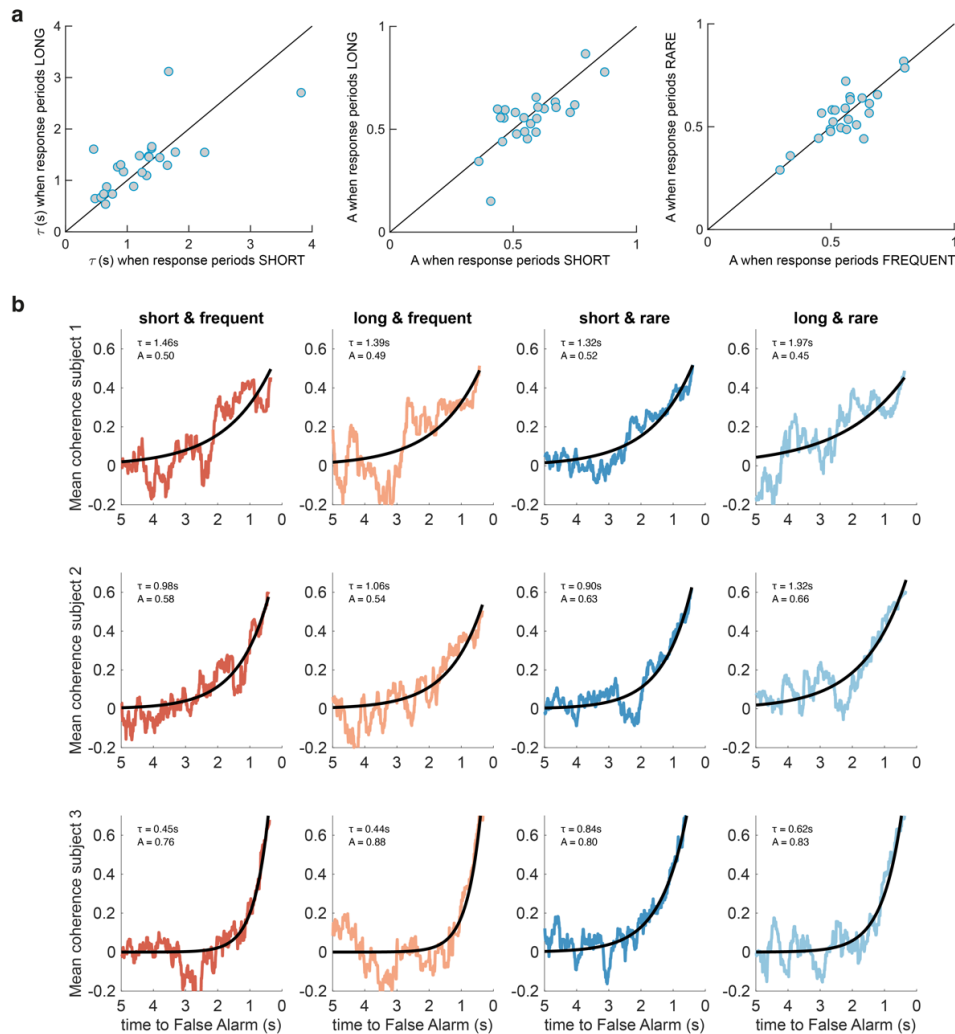


Figure 3 – Changes in false alarm response frequency and evidence integration kernels across environments with different statistical structure. (a) False alarm rates (responses during baseline periods) showed a main effect of response period duration – participants showed significantly lower false alarm rates when response periods were LONG versus SHORT, consistent with having a more cautious response threshold (also evidenced by longer reaction times during response periods, main figure 3c). (b) Integration kernels calculated for false alarms across the four conditions. (c) Exponential decay model fitted to individual participants’ kernels during false alarms shows a significantly longer decay time constant when response periods were RARE versus FREQUENT.



Extended Data Figure 3.1 – Between-subject variability in evidence integration kernels exceeds between-condition variability. (a) Left panel: Unlike for RARE versus FREQUENT (Extended Data Figure 3.1c), there was no significant difference in decay time constant for integration kernels for LONG versus SHORT. Similarly, there was no difference in amplitude parameter A for either of these comparisons (middle/right panels). (b) Variability across individuals, and comparative consistency across conditions, of integration kernels. This figure shows integration kernels from three example subjects in the experiment across the four conditions. Although our experiment primarily aimed to test whether integration kernels would be adapted to different environmental statistics (columns of figure), we found (unexpectedly) that different individuals had very different integration kernels (rows); some would integrate evidence over longer durations (e.g. top row), and others over far shorter durations (e.g. bottom row). All analyses are from false alarm responses only.

EEG correlates of continuous sensory evidence integration

Having established behavioural differences in evidence integration across individuals and across environments with different statistical structures, we then examined how participants' EEG responses reflected these differences. To test this, we examined the effects of the *noise fluctuations* on the EEG signal during baseline periods. We focussed on this time period for three reasons: (i) the generative statistics of the noise were identically matched across all four task conditions; (ii) behavioural evidence from 'false alarms' clearly indicated that participants were still integrating sensory evidence during baseline (fig 3, extended data fig 3.1); (iii) the large number of noise fluctuations embedded in the stimulus (>1000 per 5-minute block) meant that we had many events of interest to recover EEG temporal response functions (TRFs) with a high signal-to-noise ratio (Gonçalves et al., 2014; Lalor et al., 2006).

We therefore built a deconvolutional GLM to estimate TRFs to various events relating to the time-varying noise fluctuations during baseline periods. In particular, this GLM included regressors that described (fig 4): (a) "jump events" in the experimenter-controlled noise ('stick functions' that were 1 whenever the net motion changed, and 0 elsewhere); (b) the "change in evidence" associated with each jump event (stick functions with a parametric modulator of $|\Delta\text{evidence}|$, i.e. absolute difference between previous and current motion coherence); (c) the current $|\text{evidence}|$ (a continuous regressor, reflecting the absolute difference from 0 across time; note that this regressor is absolved to look for effector-independent signals processing current motion strength, as opposed to those signed towards leftward/rightward motion). We also included several further regressors to capture EEG correlates of the onset of response periods, the level of motion coherence, correct and false alarm button-presses (see methods for full details).

Using this approach, we found a set of consistent TRFs that reliably reflected the continuous updates in the time-evolving sensory evidence during baseline (fig 4). In particular, the two regressors that reflected changes in the sensory evidence ('jump events') and the magnitude of $|\Delta\text{evidence}|$ both elicited positive-going scalp topographies over centroparietal electrodes, peaking ~300ms after this change occurred (fig 4a/b). This scalp topography, timecourse and reporting of $|\Delta\text{evidence}|$ is consistent with the P300 component (Donchin, 1981; Duncan-Johnson and Donchin, 1977; Mars et al., 2008; Squires et al., 1976). The scalp topography is also consistent with the centroparietal positivity (CPP) (Kelly and O'Connell, 2013; O'Connell et al., 2012; O'Connell and Kelly, 2021), whose ramp-to-threshold dynamics have been proposed to account for many established effects in the P300 literature (Twomey et al., 2015). In addition, the "absolved sensory evidence" regressor elicited a triphasic potential over centroparietal electrodes (fig 4c). This triphasic potential is notably similar to EEG potentials reflecting 'decision update' signals during trial-based tasks that require integration of multiple, discrete pieces of evidence (Wyart et al., 2012).

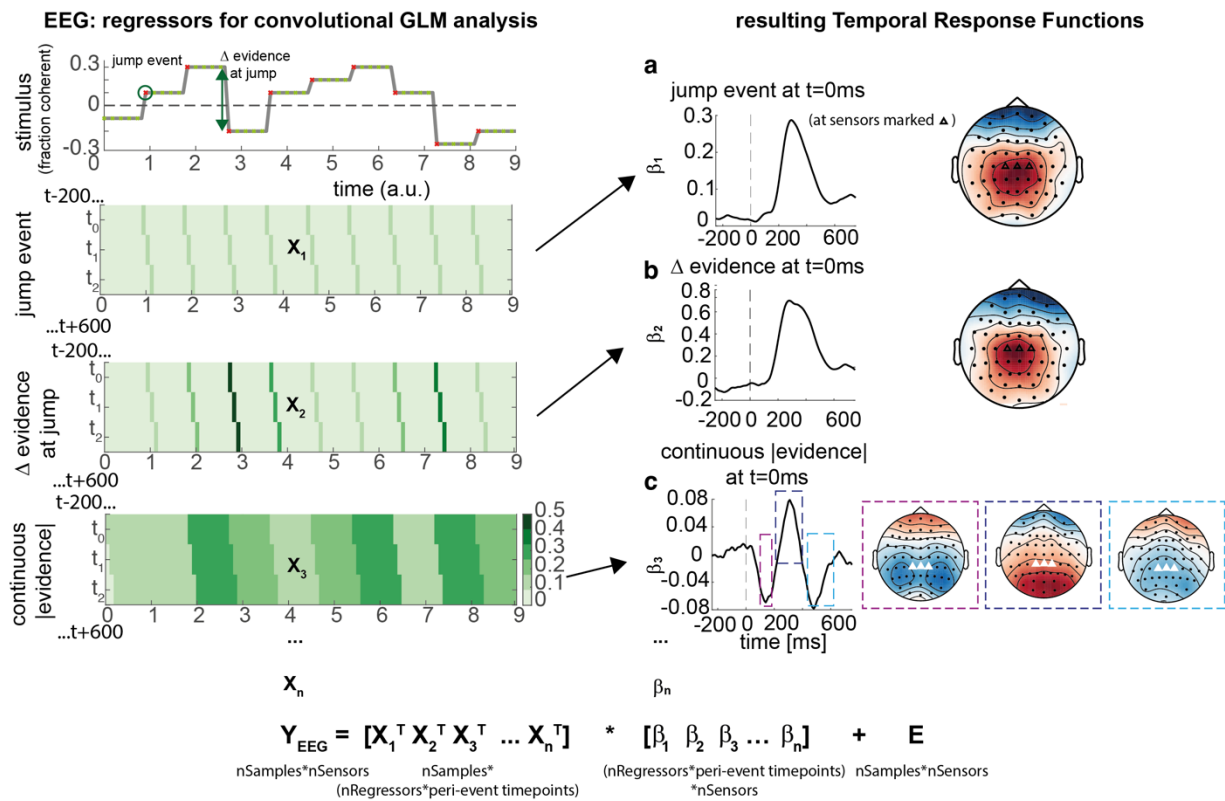


Figure 4 – Deconvolutional general linear model to estimate EEG temporal response functions to continuous, time-varying decision regressors. The left-hand side of the figure shows an example evidence stream during the baseline period (note that inter-sample intervals are shown as fixed duration for clarity, rather than Poisson distributed as in the real experiment). Three example regressors are shown: (a) ‘jump event’, when there was a change in the noise coherence level; (b) ‘ $|\Delta$ evidence|’, reflecting the magnitude of the jump update at each jump event; (c) continuous |evidence|, reflecting the continuous absolute motion strength. For each of these regressors, a lagged version of the regressor timeseries is created to estimate the TRF at each peri-event timepoint. This is then included in a large design matrix X , which is regressed onto continuous data Y at each sensor. This leads to a set of temporal response functions for each regressor at each sensor, shown on the right-hand side of the figure.

Increased CPP responses to Δ evidence and response thresholds when response periods are rare

We then examined whether these TRFs to noise fluctuations were adapting across the different sensory environments. Given our behavioural findings concerning integration kernels (figures 2d and 3b), we reasoned that we would most likely identify differences as a function of response period frequency rather than length. Indeed, we found that the centroparietal response to the same change in sensory evidence was larger when response periods were RARE than when they were FREQUENT (figure 5a; $p=0.017$, cluster-based permutation test). As large changes in mean evidence are less frequent in the RARE condition, the increased neural response to $|\Delta$ evidence| may reflect the increased statistical surprise associated with the same magnitude of change in evidence in this condition. In addition, when making a correct response, preparatory motor activity over central electrodes reached a larger decision threshold for RARE vs. FREQUENT response periods (figure 5b; $p=0.041$, cluster-based permutation test). This is again consistent with the behavioural finding of lower detection rates and slower reaction times in this condition (figure 2 b/c). By contrast, we found no statistically significant difference for either of these regressors between SHORT vs. LONG response periods (Extended Data Figure 5.1).

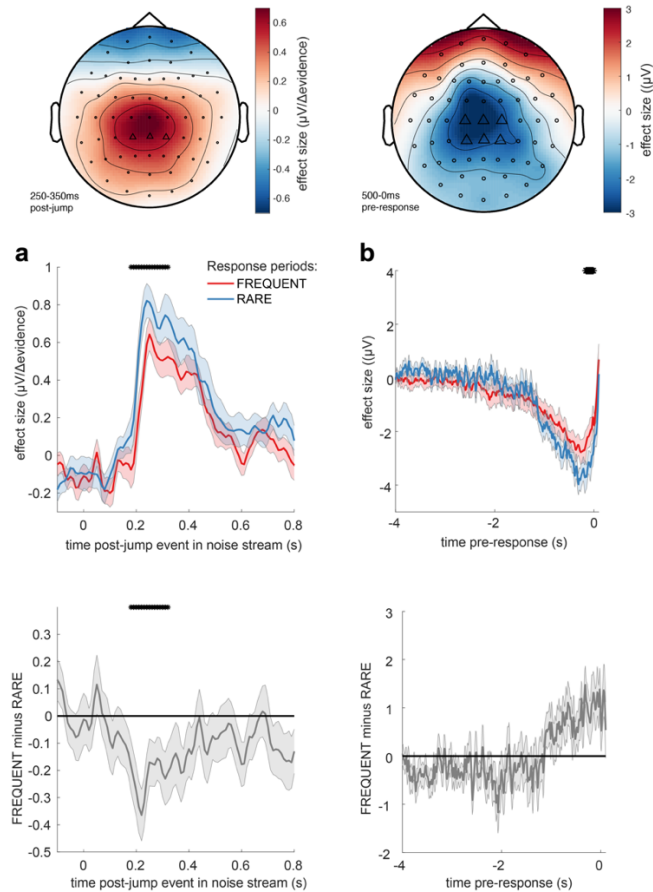
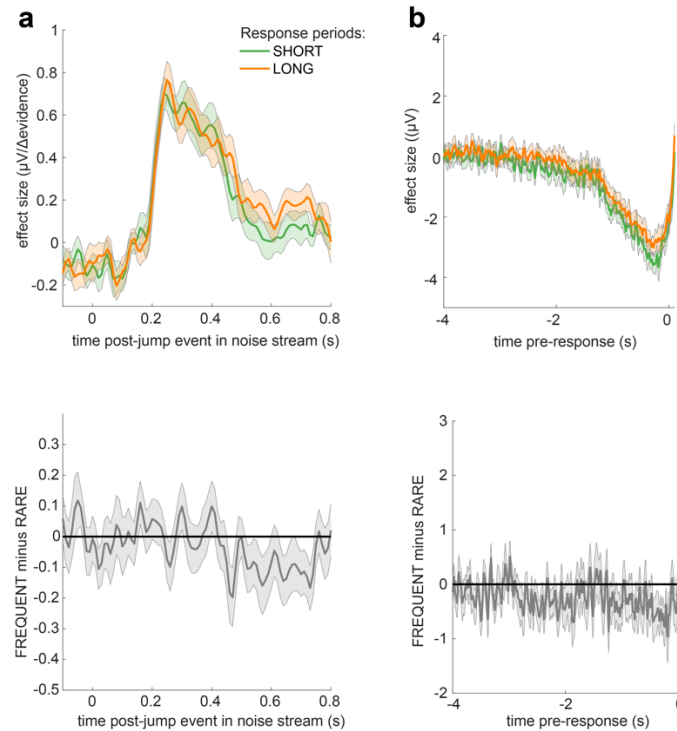


Figure 5 – Adaptations of EEG responses to sensory environments where response periods are RARE versus FREQUENT. (a) Centroparietal electrodes (see triangles in scalp topography) showed a significantly greater response to Δ evidence during ‘jump events’ in the noise stream when response periods were RARE than when they were FREQUENT. (b) Central and centroparietal electrodes showed a significantly greater negative-going potential immediately prior to a button press during response periods. * denotes significant difference between FREQUENT and RARE ($p < 0.05$, cluster corrected for multiple comparisons across time).



Extended Data Figure 5.1 – No significant differences in EEG responses between conditions where response periods were SHORT versus LONG. Panels arranged as in main figure 4.

Responses to $|\Delta\text{evidence}|$ reflect decision-relevant, not decision-irrelevant, statistics of stimulus

Given the potential role of $|\Delta\text{evidence}|$ in surprise detection, we next asked whether the centroparietal response to $|\Delta\text{evidence}|$ reflected low-level sensory properties of changes in the motion stimulus, or higher-level signals relevant to decision making. To test this, we collected an additional control dataset where the stimulus contained both horizontal motion (decision-relevant) that subjects had to integrate, as in the main experiment, but also vertical motion (decision-irrelevant) that had the same low-level sensory statistics. As in the main experiment, we found that centroparietal responses reflected both ‘jump events’ and their associated $|\Delta\text{evidence}|$ for decision-relevant motion, but these were substantially reduced for regressors that reflected changes in decision-irrelevant evidence (figure 5). This implies that low-level sensory surprise alone does not account for the centroparietal responses to $|\Delta\text{evidence}|$ in our continuous paradigm. Instead, the neural response is better described as reporting change detection that is relevant to signal detection and discrimination. It is possible that such change detection would be useful to indicate when a response period is more likely to arise in the task (Shinn et al., 2022).

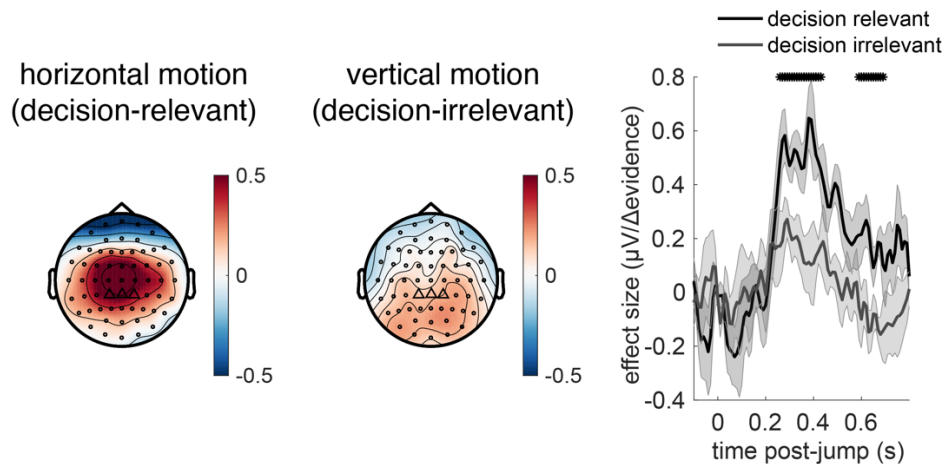


Figure 6 – Control experiment demonstrates that response to $|\Delta\text{evidence}|$ is primarily found to decision-relevant horizontal motion, but not decision-irrelevant vertical motion (with identical generative statistics).

Behavioural-neural correlation between evidence integration kernels and TRFs to continuous sensory evidence

Finally, given the consistency and between-subject variability in integration time constants shown in fig 3c and Extended Data Figure 3.1, we explored whether any components relating to processing of sensory evidence might reflect *cross-subject* variation in evidence integration. We therefore performed a behavioural-neural correlation between participants' integration time constants τ and their TRFs to sensory noise fluctuations. We found such a correlation for the triphasic potential elicited by the continuous “absoluted sensory evidence” regressor (see figure 4c). From approximately 420ms onwards, the amplitude of the final, negative component of this component showed a negative correlation with τ across participants (figure 7). In other words, this negative-going component was *larger* in amplitude (i.e. *more* negative) in participants who would integrate sensory evidence over *longer* durations (i.e. had a *higher* value of τ).

Although this across-subject correlation was discovered via exploratory analyses, it replicated across all four independent conditions, substantially increasing the likelihood of it being a true positive result (Response periods FREQUENT and SHORT: Spearman's $\rho = -0.45$, $p = 0.027$; FREQUENT and LONG: $\rho = -0.52$, $p = 0.0099$; RARE and SHORT: $\rho = -0.53$, $p = 0.0080$; RARE and LONG: $\rho = -0.46$, $p = 0.025$). By contrast, we found no evidence for an across-subject correlation between fitted integration decay time constants and EEG regressors encoding evidence updates at jump events ($\Delta\text{evidence}$), nor between fitted amplitude parameters (A) and any EEG regressors.

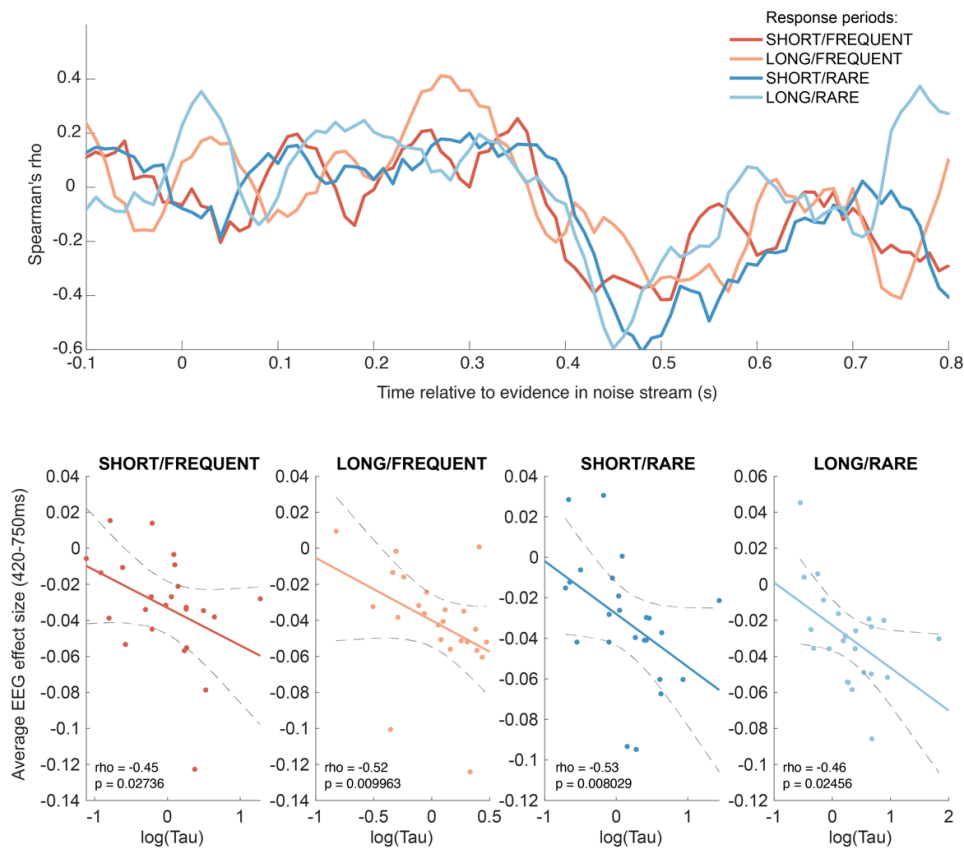


Figure 7 – Behavioural-neural correlation (across subjects) of integration decay time constant and response to absolute sensory evidence in stimulus (see figure 2c). Top panel shows Spearman's rank correlation between the time-varying EEG beta for absolute sensory evidence and individual subjects' τ parameter, separately for each of the four conditions. The negative-going correlation found in all four conditions from ~420ms onwards coincides with the third, negative-going limb of the triphasic response to absolute sensory evidence shown in figure 2c. Bottom panels shows the correlation plotted separately for each of the four conditions.

Discussion

Real-world decisions typically demand continuous integration of incoming evidence in dynamic environments, without external cues as to when evidence integration should be initiated or responses made. This contrasts with a longstanding tradition in decision making and psychophysics research to confine decisions to discrete trials, which are typically externally cued to the participant. In the current work, we developed an approach to measure how participants weighed their recent history of sensory evidence when performing a continuous perceptual decision task. We found that participants' behaviour was well described as a process of leaky (exponentially decaying) evidence accumulation, and that participants adapted the properties of this process to the overall statistics of the sensory environment (Ossmy et al., 2013) across four different experimental conditions. We also found that there was substantive inter-individual variability in leaky evidence integration, with some participants rapidly discounting past evidence and others integrating over several seconds. We then demonstrated that both sources of variability (between-condition and between-subject) were reflected by changes in centroparietal EEG responses to time-varying fluctuations in sensory evidence.

Understanding the neural mechanisms supporting simple perceptual decision making remains a key goal for the neurosciences, and the scalp topographies and signals that we have identified in the current study match well with known findings in the literature on changepoint detection and sensory evidence accumulation. In particular, the well-known P300 component has long been argued to be associated with detection of statistical surprise (Donchin, 1981; Duncan-Johnson and Donchin, 1977; Mars et al., 2008; Squires et al., 1976) or, more recently, a correlate of a continuous time-evolving decision variable (Twomey et al., 2015) that is equivalent to the centroparietal positivity (Kelly and O'Connell, 2013; O'Connell et al., 2012). Our centroparietal responses to $|\Delta\text{evidence}|$ (fig 4b) are consistent with a changepoint detection account of continuous decision making, in which decision-relevant input (fig 6) is evaluated for a change in latent state from a baseline period to a response period (Nassar et al., 2019). This account would also explain why these signals are enhanced when response periods are rarer, as a large $|\Delta\text{evidence}|$ is more statistically surprising when response periods are rare than when they are common. Such changepoint detection may be useful to transiently suppress neuronal activity and attend to salient incoming sensory evidence to guide choices (Shinn et al., 2022). Alongside this, we also identify a triphasic potential that is similar in timecourse and scalp topography to previous studies of evidence accumulation (Wyart et al., 2012), sensitive to the continuous incoming sensory evidence (fig 4c). Although we did not find variation in this potential across conditions, we did find that its amplitude reliably predicted the time-constant of leaky evidence accumulation across participants (fig 7). This suggests a key role for this component in translating incoming sensory evidence into a continuous representation of the decision variable across time, but further work is needed to understand how this variable then supports the sensorimotor transformation into a final commitment to making a choice (Steinemann et al., 2018).

Our work sits within a broader trend in recent research of moving away from trial-based designs towards continuous decision paradigms (Huk et al., 2018). In addition to being more naturalistic, a key advantage of such paradigms is that it is possible to relate time-varying properties of the continuous input to continuous output variables. In 'tracking' paradigms, for example, not only is the sensory input continuous and time-varying but also the behavioural responses made by participants. This approach has been used to dramatically reduce the length of time needed from individual participants, meaning that experiments that previously required many thousands of trials over several hours to recover psychophysical functions can now be completed in a matter of minutes (Bonnen et al., 2015; Knöll et al., 2018; Straub and Rothkopf, 2021). In our paradigm, the behavioural responses remained discrete and sparse (as participants were completing a signal detection/discrimination task, rather than a tracking task), but the EEG data is continuous and time-varying, and our analysis approach similarly benefits from being able to relate continuous variations in sensory input to this continuous neural signal. This has been shown in previous work to yield considerable improvements in signal-to-noise ratio compared to traditional trial-based event related

potentials (Dimigen and Ehinger, 2021; Gonçalves et al., 2014; Lalor et al., 2006), meaning that our approach should allow us to characterise the neural response to sensory evidence integration in less recording time than in previous work. Indeed, when we examined neural responses from individual participants in the present study, we found a high degree of individual-subject reliability. This efficiency is due to the high density of events in the continuous task design (>1000 jump events in each 5-minute block), limiting the amount of 'dead time' present in the experimental design (Henson, 2007). In sum, this leads to a more efficient experimental design than in conventional trial-based experiments, and may potentially allow for more rapid and reliable estimation of single-subject responses.

An intriguing finding in the present work is the substantive variability in integration decay time constants across individuals. Indeed, such inter-individual variability exceeded the between-condition variability that was observed due to our experimental manipulations. We consider two possible explanations of this inter-individual variability. The first is that it is a stable, trait-like feature of sensory evidence integration that is not unique to our task, but instead reflects true variability in perceptual evidence integration across individuals. Such a hypothesis would imply that it would predict variability in integration time-constants in other domains (for example, auditory evidence integration (Brunton et al., 2013; Keung et al., 2019; McWalter and McDermott, 2018) or more broadly cognitive tasks that involve continuous maintenance and manipulation of information across time in working memory). If so, it may also be possible to relate variability in behavioural time constants to underlying neurobiological causes by measuring the resting autocorrelation structure of neural activity, for example in MEG or fMRI data (Cavanagh et al., 2020; Manea et al., 2022; Raut et al., 2020). On the other hand, the individual variability we observe may be a consequence of the prior expectations that our participants have about the overall task structure, combined with learning over the course of training. One argument in support of this hypothesis comes from modelling of the task to optimise overall reward rates. In initial simulations, we found that the optimal combination of decay and threshold for an Ohnstein-Uhlenbeck process (Bogacz et al., 2006) occupied a 'ridge' in parameter space - in which reward could be maximised when either the decay was slow and the threshold was higher, or decay was fast and threshold lower (Ruesseler, 2021). The exact location of this ridge in parameter space varied across conditions, supportive of the need to change integration parameters across conditions. This existence of a ridge rather than a global maximum would imply that participants could have optimised reward rate during training via gradient descent in this parameter space, but ended up at different locations in the ridge of optimal parameters, and consequently displayed a range of integration time constants. Future work might test this hypothesis more directly by introducing additional experimental constraints on reward maximisation within the task, and/or explicitly modelling the learning/optimisation process.

Our findings of behavioural adaptations according to the overall statistics of the sensory environment are consistent with findings from previous research that have examined the same question in the absence of neural measures (Ossmy et al., 2013; Piet et al., 2018). Ossmy et al. (2013) used a trial-based paradigm with similarities to ours, involving unpredictable signal detection in the context of time-varying background noise in combination with a manipulation of signal-period duration. They used this to show that participants adapted the time constants of evidence integration in different environments. In our paradigm, we benefitted from being able to directly recover empirical integration kernels as opposed to estimating them in a model, and we also found behavioural differences that resulted from varying response period duration. However, we surprisingly found little effect on integration kernels from this manipulation; we instead found that our manipulation of response period frequency had a greater effect on the weighting of past sensory evidence. We hypothesise that this difference between our results and those of Ossmy et al. may simply result from our manipulation of response period duration not being sufficiently large (3s vs. 5s) to require a substantial change in evidence weighting to optimise rewards. This is also consistent with our neural results, where the between-condition variation in responses to time-varying evidence was primarily found as a function of response period frequency, rather than duration.

In conclusion, our work demonstrates that it is possible to accurately measure the time-course and neural correlates of sensory evidence integration in continuous tasks, and how this adapts to the overall properties of the environment. This work provides a framework for future work to investigate how evidence integration is adapted to other features of decision tasks, and how this may vary across individual participants. Our approach will also be useful for future work to investigate how the properties of evidence integration change in clinical populations, and how they are affected by various interventions (e.g. pharmacological, electrical/magnetic stimulation, cognitive training).

Methods

Task Design

In the continuous task, participants observed a stream of randomly moving dots in a circular aperture (Figure 1). A fraction of these dots move coherently to the left or to the right; the motion coherence is the proportion of dots moving in the same direction, whereas the other dots move randomly. In this study, the coherence varied between -1 (all dots move to the left) and 1 (all dots move to the right). At 0 coherence, all dots move randomly.

Unlike in trial-based versions of the task, and even previous RDM tasks where motion is continuous (e.g., (Kelly and O'Connell, 2013)), during the present task the coherence *changed* constantly. During 'baseline' periods, the *average* of these constantly changing values remains 0. Within this stream of constantly changing coherence, there were response periods in which the *average* coherence was either to the left or to the right (see Figure 1b, 2a). The aim of the subject performing this task was to detect these stable periods of predominantly leftward or rightward motion (response periods). For participants, this means that they should not respond as soon as they think they know in which direction the dots are moving coherently, as is the case for a discrete trial version of the RDM task. Instead, participants must weigh the recent history of motion directions to detect periods where the average motion direction of the dots was consistently leftwards or rightwards.

Participants indicated their decision about the average motion direction by pressing keyboard button 'L' to report a response period with average rightward motion and 'A' to report a response period with average leftward motion. Every time a button was pressed, a change in colour of the central fixation point provided feedback: correct responses were indicated by a green fixation point, a red fixation point followed an incorrect response during the response period, and false alarms (i.e., button presses during baseline periods) were indicated by a yellow fixation point. Whenever a response period was missed (no response made), the fixation point turned blue after 500ms (note that a button press within these 500ms was still counted as a correct response, to account for non-decision time and allow participants to integrate over the entire length of the response period). Following correct or incorrect responses made during a 'response period', the response period was terminated immediately, and the stimulus returned to baseline.

Participants were rewarded for correct responses but lost points for any other response. They received 3 points for correct responses, punished with -3 points for incorrect responses, and missed response periods or false alarms were both punished with -1.5 points. A reward bar was shown at the end of each 5-minute block to indicate how many points participants have won in total (the reward bar was shown continuously onscreen during training). As participants won more points, their reward increased to the right until they hit the right border of the reward bar (equivalent to a net gain of 15 points), the bar was reset to the middle of the screen and they received £0.50 bonus to take at the end of the experiment. In rare cases where participants were performing poorly and losing points on average, they hit the left border of the reward bar (-15 points) and had £0.50 deducted from their take-home bonus.

Structure of the noise

An essential feature of this task paradigm was the noise structure, which leads to continuously varying coherence levels. Notably, the noise was placed under experimental control rather than randomly generated, meaning that we could examine how fluctuations in the noise impact participants' behavioural and neural data. More generally, the noise can be described as a series of short intervals that vary in duration and coherence (Figure 1B). The interval duration was sampled from an exponential distribution with a mean duration of 270ms. This distribution was then truncated, with a minimum duration of 10ms and a maximum duration of 1000ms for each step. Interval coherence was sampled randomly from a normal distribution whose standard deviation of 0.3 for response periods and 0.5 for baseline periods. The mean of this normal distribution depends on whether the step occurred during baseline or a response period. During a baseline period the mean

of the normal distribution is 0. That means it is equally likely that negative or positive coherences are drawn. During response periods, the mean of the normal distribution was sampled uniformly from the set [-0.5, -0.4, -0.3, 0.3, 0.4, 0.5].

Design of the random dot motion stimulus

The task was coded in Psychtoolbox (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) and parameters were chosen similar to (Shadlen and Newsome, 2001). Participants were seated 87 cm in front of the screen. Moving dots were displayed in a circular aperture subtending with a radius of 5 visual degrees on a Dell monitor with a refresh rate of 100Hz. Dots had a size of 0.1 visual degrees and were displayed with a density of 2.5 dots per squared visual degree. The fixation point in the centre of the screen had a size of 0.3 visual degrees. All dots were black and displayed on a mid-grey background (rgb: 0.5, 0.5, 0.5).

Dots were equally divided in 3 sets. These sets were shown sequentially, meaning only one set per frame was shown. Each time a set reappeared on the screen the coherence on that frame dictated, the likelihood of that dot either being displaced randomly or in the direction of the coherence. Randomly displayed dots moved like Brownian motion particles, with no particular speed. Dots that moved coherently were displaced according to a speed of 7 degrees per second (Shadlen and Newsome, 2001). This approach means that subjects were forced to integrate across the entire field of moving dots to establish the motion direction; tracking a single dot is not reliable because it only reappears on every third frame and does not necessarily move coherently for more than 2 frames.

Conditions

To understand whether participants can flexibly adapt their integration kernel, we tested the continuous evidence integration task under different conditions following a within-subject 2x2 design. In different 5-minute blocks, participants were told that they would either have long (5 seconds) or short (3 seconds) response periods, and either frequent (baseline period range 3s-8s) or rare (baseline period range 5s-40s) response periods. Subjects received extensive training (see below) so that these environmental statistics were well learnt prior to the experimental session. During the experiment, they were cued as to which condition they were currently in, by: (i) displaying in text at the beginning of each block (e.g. “response periods are LONG and FREQUENT”); (ii) having a different shape of fixation point (triangle, square, circle, star) for each of the four blocks. This meant that there was no inference nor memory required from the subjects to know which condition they were currently in.

Training

Our training protocol was designed to overtrain participants to reach a high level of performance on the task and to minimize learning effects during the main testing session. Our training taught participants about the structure of the long/short and rare/frequent trial periods, how to discriminate such trial periods from background noise fluctuations, and crucially incentivised participants to maximise their overall reward rate.

Training consisted of a sequence of different tasks that incrementally trained participants on the random dot motion stimulus and the continuous nature of the task. First, participants were introduced to the conventional RDM task based on discrete trials (Shadlen and Newsome, 2001), initially with very strong motion coherences (-0.9, 0.9), and then progressively with motion coherences resembling those found in the main experiment [-0.5, -0.4, -0.3, 0.3, 0.4, 0.5].

Next, participants completed intermediate task versions that still consisted of discrete trials, but had features of the continuous task: in particular, noise fluctuations was superimposed on the mean motion coherence, as in the final continuous task version coherences would fluctuate throughout the trial. This meant that participants had to estimate the average motion direction across the entire trial, but respond before the trial ended. Trial lengths were 5s (length of long response periods) or 3s (length of short response periods) and the fixation point shrank over the course of the trial so that participants would have an idea of how much time they had to respond. Trials with a mean

coherence of 0% were also included, to simulate baseline periods of the continuous task that participants had to contrast to the other trials; participants had to suppress a response in these trials.

Then, participants moved on to the continuous version of the task, but they were first trained on a paradigm with higher mean coherences during response periods; in addition, the fixation dot would change its colour to white to indicate the onset of response periods. Gradually, all conditions of the experiment and the final mean coherence levels of signal periods were introduced, and a change of the fixation dot colour to white was disabled. At this point, when the paradigm was the same as in the full task, they had an extended period of practice on the task across all four conditions.

We progressed participants through the different versions of training by checking psychometric functions to establish that they had fully learnt each stage before progressing onto the next stage of the task. Participants had to perform at 80% correct or higher on discrete trials to move on to the continuous tasks, and performance in the continuous task was checked qualitatively by plotting the stimulus stream and responses after each block. Verbal feedback was given to participants based on their performance, which also helped participants to improve their behaviour during training. If, in the latter parts of training, participants still missed more than half of the response periods, they were excluded from the subsequent EEG session. Participants ($n=3$) who failed to progress from the training session were paid for that session and did not progress on to the EEG session.

After participants completed training successfully, they participated in the EEG testing session not more than one week later. In this session, participants first performed a 'reminder' where they practised one run of the full task for 20 minutes (all 4 conditions, presented for 5 minutes each, in randomised order); this was performed while the experimenter put on the EEG cap. Then, while EEG data was collected, they completed 5-6 task 'runs', each lasting 20 minutes. Each run consisted of all 4 conditions in randomised order.

Data collection

We tested 33 participants (13 male). Of those, three were unable to learn the continuous task and did not progress beyond training. One participant was excluded for falling asleep during the EEG session. Another 5 were excluded from the analysis due to technical issues matching the continuous EEG with the stimulus stream, and/or issues with EEG data quality after pre-processing. This means 24 subjects were included in the analysis. Each subject completed 6 runs, except for one subject who completed only 5 runs. For the control experiment with superimposed vertical motion (figure 6), a further 6 participants were tested (4 male). All participants were aged 18-40, had normal or corrected-to-normal vision and gave written consent prior to taking part in the study. The study was approved by the local ethics committee (CUREC R60298).

Behavioural analysis

Detection rate/reaction times. We calculated correct detection rate (fig 2b) as the proportion of response periods in which correct responses were made. We calculated this separately for each run within each level of motion coherence (collapsing across leftward/rightward correct responses), and then averaged across the six runs, to obtain three values (0.3, 0.4 and 0.5 motion coherence) for each of the four conditions per subject. We performed a similar analysis on reaction times for these correct responses (fig 2c), but here we excluded responses in LONG conditions that exceeded 3.5s, such that the average response time could be directly compared between LONG and SHORT conditions. (We note that due to noise and the associated uncertainty concerning response period onset, participants are incentivised to respond as quickly as possible whenever they thought they were in a response period, as delaying responses would lead to 'missed trials'). We analysed the effects of coherence, response period length and response period frequency on detection rate and reaction time across the 24 participants using a three-way repeated measures ANOVA.

Integration kernels. We calculated integration kernels by averaging the presented motion coherence for 5s preceding every buttonpress (having first multiplied this by -1 for leftward button

presses, so that left and right responses can be averaged together). We did this separately for false alarms (fig 3b) and for correct responses (fig 2d). We excluded correct responses in LONG conditions that exceeded 3.5s, for similar reasons as outlined above. We also note that the integration kernel in this period includes a mixture of ‘signal’ (shift in mean coherence) plus noise, whereas the integration kernel from false alarms is driven by noise alone. This explains why the segment of the integration kernel that reflects non-decision time (i.e. immediately prior to button press) returns close to 0 in fig 3b but is closer to 0.5 in fig 2c, and also why the false alarm integration kernel is more clearly an exponential decay function. To fit the exponential decay function, we first found the peak of the empirical integration function (using *max* in MATLAB), and set this timepoint to $t=0$ in equation 1 (see main text). We then fit A and τ to the empirical integration kernel for all timepoints up to and including $t=0$ using *fminsearch* in MATLAB, using a least squares cost function between the fitted model and data with an L2 regularisation term that penalised large values of either A or τ ($\lambda = 0.01$).

EEG acquisition

EEG data was collected at a sampling rate of 1000 Hz with Synamps amplifiers and Neuroscan data acquisition software (Compumedics) and 61 scalp electrodes following the 10-20 layout. Additionally, bipolar electrodes were placed below and above the right eye and on the temples to measure eyeblinks as well as horizontal and vertical eye movements (HEOG and VEOG channels). A ground electrode was attached to the left elbow bone. The EEG signal was referenced to the left mastoid but later re-referenced to the average of left and right mastoids. Impedances of electrodes were kept below 15 k Ω .

EEG pre-processing

Data were pre-processed using spm12 (<http://www.fil.ion.ucl.ac.uk/spm/>) (Litvak et al., 2011), the FieldTrip toolbox for EEG/MEG-analysis (<http://fieldtriptoolbox.org>) (Oostenveld et al., 2010) and MATLAB (Version R2018b, The MathWorks, Inc., Natick, Massachusetts, United States). Each session for each participant was pre-processed as continuous data. First, each session was downsampled to 100Hz. Then, the data was rereferenced to the average of left and right mastoid electrodes and bandpass filtered the data between 0.1Hz and 30Hz using the function *spm_eeg_filter* with default settings (fifth-order Butterworth filter, passed in both directions). In a next step, we used signal space projection methods in SPM to perform eyeblink correction. The bipolarized VEOG channel was used to build a spatial confound topography of eye blinks to delineate ocular source components, Segments of 1000ms around eye blink events in the VEOG channel were generated and averaged. Principal component analysis was then used to define the noise subspace of eyeblinks across all channels, and the first principal component was regressed out of the continuous EEG data (Berg and Scherg, 1994; Hunt et al., 2012). For each participant and session the spatial confound map of the first component was visually checked to ensure it showed a typical eye blink topography before the regression was applied. The EEG data was further thresholded to remove artefacts that were $\geq 100\mu\text{V}$ in a single channel by labelling a 500ms window around the peak of the artefact, and removing these timewindows when estimating the deconvolutional GLM.

Deconvolutional GLM analysis

We used triggers sent to each jump in the noise stream to align the continuous EEG data with the continuous stream of sensory evidence (and other experimental events, such as button presses). As the downsampled EEG was at the same sampling rate as the refresh rate as the display (100 Hz), we simply used the continuous stream of evidence presented on each frame of the experiment from then onwards. In addition to the 5 participants excluded due to technical issues with trigger recording and alignment (see ‘Data collection’ above), there was one further participant in our main EEG sample ($n=24$) who had 3 out of 24 blocks missing due to technical issues; this participant was nevertheless taken forward into the main analysis with the remaining 21 recorded blocks.

We then constructed a design matrix X for the continuous EEG data, with 11 regressors in total:

$$\begin{aligned}
 EEG \sim & jump_{event} + jump_{level} + jump_{|Devidence|} + continuous_{|evidence|} \\
 & + continuous_{evidence(signed)} + response\ period\ onset_{event} \\
 & + response\ period\ onset_{coherence} + buttonpress_{correct} \\
 & + buttonpress_{false\ alarm} + (left - right\ buttonpress)_{correct} \\
 & + (left - right\ buttonpress)_{false\ alarm}
 \end{aligned}$$

The ‘buttonpress’ regressors and the regressors with subscript ‘event’ are ‘stick functions’ (1 at the timepoint that they occurred, and 0 at all other timepoints). Other regressors are parametric modulators of these, except for the two continuous regressors which were valued at all timepoints of the experiment (reflecting the current motion onscreen, either absolved (reported in the main paper) or signed (not discussed)). In this paper we focus on responses to $jump_{event}$ (fig 4a), $jump_{|Devidence|}$ (fig 4b, 5a, Extended Data figure 5.1a, 6), $continuous_{|evidence|}$ (figs 4c, 7), and $buttonpress_{correct}$ (fig 5b, Extended Data figure 5.1b). For all of these except for the $buttonpress_{correct}$, we only estimate the EEG response during the baseline periods, when participants are still integrating evidence (as shown empirically in figure 3) but the statistics of the stimulus stream across all four experimental conditions are matched. We calculated the correlation between the key regressors of interest (table 1) to ensure that they were sufficiently decorrelated from one another to reliably obtain parameter estimates in the general linear model.

To obtain the deconvolved response to each of these regressors, we time-expanded the design matrix into a large design matrix X_{dc} (see (Ehinger and Dimigen, 2019) for recent review). We used simple ‘staircasing’ of the regressors to create this design matrix (as illustrated in figure 4), rather than a time-Fourier basis set (Litvak et al., 2013) or time-Spline basis set (Ehinger and Dimigen, 2019); any of these approaches might be suitable for future studies. The number of timepoints for each of the regressors varied slightly between different regressors (for example, we were primarily interested in activity *after* stimulus changes but *before* buttonpresses; the number of pre- and post-event lags reflected this). We then estimated parameter estimates for the deconvolved regressor at each sensor for each subject with ordinary least squares (using the method of (Courrieu, 2008) to facilitate fast computation of the pseudoinverse of the design matrix).

Table 1: (Average) explained variance between regressors of the convolutional GLM for baseline periods. Columns and rows are the different regressors used to investigate baseline periods. Between each pair of regressors for the key continuous variables, the explained variance (squared correlation coefficient) was calculated to ensure that these regressors were not correlated with each other prior to estimating the GLM.

R^2	Jump	Jump level	Prediction error	Absoluted stimulus
Jump	1	0	0.01	0
Jump level	0	1	0.18	0.03
Prediction error	0.01	0.18	1	0.01
Absoluted stimulus	0	0.03	0.01	1

Table 2: Design matrix. This table describes the size of the design matrix assuming a sampling frequency of the EEG signals of 100Hz. For each regressor the number of lags pre- and post event and the total number of rows this regressor covers in the design matrix are described. The same number of lags were applied to vertical motion regressors for the control study.

Regressor	Pre-event time in time-expanded design matrix (ms)	Post-event time in time-expanded design matrix (ms)	Total rows in the time-expanded design matrix
Jump event (stick function)	1000	1500	251
Jump level (evidence at each jump event)	1000	1500	251
Jump $ \Delta$ evidence (at each jump event)	1500	1500	301
Continuous evidence	1500	1500	301
Continuous (signed) evidence	1500	1500	301
Correct buttonpresses (stick function)	5000	3500	851
Correct buttonpresses (+1 for right, -1 for left)	5000	3500	851
False alarm buttonpresses (stick function)	5000	3500	851
False alarm buttonpresses (+1 for right, -1 for left)	5000	3500	851
Onset of response period (stick function)	500	8000	851
Response period coherence of response period (stick function)	500	8000	851

Permutation test for convolutional GLM analysis

To test for significant differences between deconvolved EEG responses for LONG versus SHORT response periods, and for RARE versus FREQUENT response periods, we performed a non-parametric paired T test controlling for multiple comparisons across time, using the FieldTrip function *ft_timelockstatistics* (Maris and Oostenveld, 2007). We first selected electrodes and time-windows of interest based upon the average response to the key regressors across all four conditions (see figure 4); we note that because this selection vector is orthogonal to the difference between conditions (and the number of observations are matched between conditions), then it provides an unbiased method for selecting a window of interest (Kriegeskorte et al., 2009). In practice, this meant that the cluster-based permutation test was performed on an average of three centro-parietal electrodes (CP1,CP2,CPz) and a time window from 0 to 800ms post-event for jump-locked events (figure 5a, Extended Data 5.1a, 6a), and an average of six centro-parietal and central electrodes (C1,C2,Cz,CP1,CP2,CPz) and a time window from 2000ms to 0ms pre-event for buttonpress events (figure 5b, Extended Data 5.1b). 1000 permutations were generated with the Monte Carlo method, and clusters were selected based on a T-statistic threshold of 2.07 for initial cluster formation (except for $|jump|_{|evidence|}$, where a slightly lower threshold of $T > 1.80$ was used), and an alpha of 0.05 (two-

tailed) was then used for significance detection of clusters, corrected for multiple comparisons across time.

For the behavioural-neural correlations in Figure 7, we first temporally smoothed single subject betas with a gaussian kernel with 75ms FWHM (to further improve single subject SNR), and then calculated the Spearman's correlation at each timepoint between the estimated betas for the *continuous*_{|evidence|} regressor and the τ parameters fit to the empirical evidence integration kernels for false alarms. We did this separately for the four conditions, providing four separate tests of the same behavioural-neural correlation (we note that these tests are independent in the sense that they consist of separate data for each correlation, but not in the sense that different participants were used to generate the data). In figure 7b, we report the behavioural-neural correlation for the time window 420-750ms after the evidence.

Acknowledgements and licensing

We thank Neb Jojanovic with help in collecting the control experiment with simultaneous vertical and horizontal motion. MR is supported by a PhD studentship from the Wellcome Trust (109064/Z/15/Z). JOR is supported by a Career Development Fellowship from the Medical Research Council (MR/L019639/1) and an MRC Transition Support award (MR/T031344/1). 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 authors have applied a CC-BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

- Berg P, Scherg M. 1994. A multiple source approach to the correction of eye artifacts. *Electroencephalography and Clinical Neurophysiology* **90**:229–241. doi:10.1016/0013-4694(94)90094-9
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* **113**:700–765. doi:10.1037/0033-295X.113.4.700
- Bonnen K, Burge J, Yates J, Pillow J, Cormack LK. 2015. Continuous psychophysics: Target-tracking to measure visual sensitivity. *J Vis* **15**:14. doi:10.1167/15.3.14
- Brainard DH. 1997. The Psychophysics Toolbox. *Spatial Vision* **10**:433–436. doi:10.1163/156856897X00357
- Britten KH, Shadlen MN, Newsome WT, Movshon JA. 1992. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* **12**:4745–4765. doi:10.1523/JNEUROSCI.12-12-04745.1992
- Brunton BW, Botvinick MM, Brody CD. 2013. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**:95–98. doi:10.1126/science.1233912
- Cavanagh SE, Hunt LT, Kennerley SW. 2020. A Diversity of Intrinsic Timescales Underlie Neural Computations. *Front Neural Circuits* **14**:615626. doi:10.3389/fncir.2020.615626
- Cheadle S, Wyart V, Tsetsos K, Myers N, de Gardelle V, Hecce Castañón S, Summerfield C. 2014. Adaptive gain control during human perceptual choice. *Neuron* **81**:1429–1441. doi:10.1016/j.neuron.2014.01.020
- Courrieu P. 2008. Fast Computation of Moore-Penrose Inverse Matrices (No. arXiv:0804.4809). arXiv. doi:10.48550/arXiv.0804.4809
- Crosse MJ, Di Liberto GM, Bednar A, Lalor EC. 2016. The Multivariate Temporal Response Function (mTRF) Toolbox: A MATLAB Toolbox for Relating Neural Signals to Continuous Stimuli. *Front Hum Neurosci* **10**. doi:10.3389/fnhum.2016.00604
- de Lange FP, Rahnev DA, Donner TH, Lau H. 2013. Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *J Neurosci* **33**:1400–1410. doi:10.1523/JNEUROSCI.1094-12.2013
- Dimigen O, Ehinger BV. 2021. Regression-based analysis of combined EEG and eye-tracking data: Theory and applications. *Journal of Vision* **21**:3. doi:10.1167/jov.21.1.3
- Donchin E. 1981. Surprise!...Surprise? *Psychophysiology* **18**:493–513. doi:10.1111/j.1469-8986.1981.tb01815.x
- Donner TH, Siegel M, Fries P, Engel AK. 2009. Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Curr Biol* **19**:1581–1585. doi:10.1016/j.cub.2009.07.066
- Duncan-Johnson CC, Donchin E. 1977. On quantifying surprise: the variation of event-related potentials with subjective probability. *Psychophysiology* **14**:456–467. doi:10.1111/j.1469-8986.1977.tb01312.x
- Ehinger BV, Dimigen O. 2019. Unfold: an integrated toolbox for overlap correction, non-linear modeling, and regression-based EEG analysis. *PeerJ* **7**:e7838. doi:10.7717/peerj.7838
- Glaze CM, Filipowicz ALS, Kable JW, Balasubramanian V, Gold JI. 2018. A bias–variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nat Hum Behav* **2**:213–224. doi:10.1038/s41562-018-0297-4
- Glaze CM, Kable JW, Gold JI. 2015. Normative evidence accumulation in unpredictable environments. *Elife* **4**. doi:10.7554/eLife.08825
- Gold JI, Shadlen MN. 2007. The neural basis of decision making. *Annu Rev Neurosci* **30**:535–574. doi:10.1146/annurev.neuro.29.051605.113038
- Gonçalves NR, Whelan R, Foxe JJ, Lalor EC. 2014. Towards obtaining spatiotemporally precise responses to continuous sensory stimuli in humans: A general linear modeling approach to EEG. *NeuroImage* **97**:196–205. doi:10.1016/j.neuroimage.2014.04.012
- Hanks TD, Summerfield C. 2017. Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron* **93**:15–31. doi:10.1016/j.neuron.2016.12.003
- Hassall CD, Harley J, Kolling N, Hunt LT. 2021. Temporal Scaling of Human Scalp-Recorded Potentials. *bioRxiv* 2020.12.11.421180. doi:10.1101/2020.12.11.421180

- Henson R. 2007. Efficient Experimental Design for fMRI Statistical Parametric Mapping. Elsevier. pp. 193–210. doi:10.1016/B978-012372560-8/50015-2
- Huk A, Bonnen K, He BJ. 2018. Beyond Trial-Based Paradigms: Continuous Behavior, Ongoing Neural Activity, and Natural Stimuli. *J Neurosci* **38**:7551–7558. doi:10.1523/JNEUROSCI.1920-17.2018
- Hunt LT, Daw ND, Kaanders P, MacIver MA, Mugan U, Procyk E, Redish AD, Russo E, Scholl J, Stachenfeld K, Wilson CRE, Kolling N. 2021. Formalizing planning and information search in naturalistic decision-making. *Nat Neurosci* **24**:1051–1064. doi:10.1038/s41593-021-00866-w
- Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MFS, Behrens TEJ. 2012. Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience* **15**:470–476. doi:10.1038/nn.3017
- Kelly SP, Corbett EA, O’Connell RG. 2021. Neurocomputational mechanisms of prior-informed perceptual decision-making in humans. *Nat Hum Behav* **5**:467–481. doi:10.1038/s41562-020-00967-9
- Kelly SP, O’Connell RG. 2013. Internal and external influences on the rate of sensory evidence accumulation in the human brain. *J Neurosci* **33**:19434–19441. doi:10.1523/JNEUROSCI.3355-13.2013
- Keung W, Hagen TA, Wilson RC. 2019. Regulation of evidence accumulation by pupil-linked arousal processes. *Nat Hum Behav* **3**:636–645. doi:10.1038/s41562-019-0551-4
- Kilpatrick ZP, Holmes WR, Eissa TL, Josić K. 2019. Optimal models of decision-making in dynamic environments. *Current Opinion in Neurobiology, Computational Neuroscience* **58**:54–60. doi:10.1016/j.conb.2019.06.006
- Kleiner M, Brainard DH, Pelli D, Ingling A, Murray R, Broussard C. 2007. What’s new in Psychtoolbox-3. *Perception* **36**:1–16. doi:10.1068/v070821
- Knöll J, Pillow JW, Huk AC. 2018. Lawful tracking of visual motion in humans, macaques, and marmosets in a naturalistic, continuous, and untrained behavioral context. *Proceedings of the National Academy of Sciences* **115**:E10486–E10494. doi:10.1073/pnas.1807192115
- Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI. 2009. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* **12**:535–540. doi:10.1038/nn.2303
- Lalor EC, Pearlmutter BA, Reilly RB, McDarby G, Foxe JJ. 2006. The VESPA: A method for the rapid estimation of a visual evoked potential. *NeuroImage* **32**:1549–1561. doi:10.1016/j.neuroimage.2006.05.054
- Levi AJ, Yates JL, Huk AC, Katz LN. 2018. Strategic and Dynamic Temporal Weighting for Perceptual Decisions in Humans and Macaques. *eNeuro* **5**. doi:10.1523/ENEURO.0169-18.2018
- Litvak V, Jha A, Flandin G, Friston K. 2013. Convolution models for induced electromagnetic responses. *NeuroImage* **64**:388–398. doi:10.1016/j.neuroimage.2012.09.014
- Litvak V, Mattout J, Kiebel S, Phillips C, Henson R, Kilner J, Barnes G, Oostenveld R, Daunizeau J, Flandin G, Penny W, Friston K. 2011. EEG and MEG Data Analysis in SPM8. *Computational Intelligence and Neuroscience* **2011**:e852961. doi:10.1155/2011/852961
- Manea AM, Zilverstand A, Ugurbil K, Heilbronner SR, Zimmermann J. 2022. Intrinsic timescales as an organizational principle of neural processing across the whole rhesus macaque brain. *eLife* **11**:e75540. doi:10.7554/eLife.75540
- Maris E, Oostenveld R. 2007. Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* **164**:177–190. doi:10.1016/j.jneumeth.2007.03.024
- Mars RB, Debener S, Gladwin TE, Harrison LM, Haggard P, Rothwell JC, Bestmann S. 2008. Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. *J Neurosci* **28**:12539–12545. doi:10.1523/JNEUROSCI.2925-08.2008
- McWalter R, McDermott JH. 2018. Adaptive and Selective Time Averaging of Auditory Scenes. *Current Biology* **28**:1405–1418.e10. doi:10.1016/j.cub.2018.03.049
- Nassar MR, Bruckner R, Frank MJ. 2019. Statistical context dictates the relationship between feedback-related EEG signals and learning. *eLife* **8**:e46975. doi:10.7554/eLife.46975
- Newsome WT, Pare EB. 1988. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* **8**:2201–2211. doi:10.1523/JNEUROSCI.08-06-02201.1988
- O’Connell RG, Dockree PM, Kelly SP. 2012. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* **15**:1729–1735. doi:10.1038/nn.3248
- O’Connell RG, Kelly SP. 2021. Neurophysiology of Human Perceptual Decision-Making. *Annu Rev Neurosci* **44**:annurev-neuro-092019-100200. doi:10.1146/annurev-neuro-092019-100200

- Oostenveld R, Fries P, Maris E, Schoffelen J-M. 2010. FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience* **2011**:e156869. doi:10.1155/2011/156869
- Orsolich I, Rio M, Mrcic-Flogel TD, Znamenskiy P. 2021. Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making. *Neuron* **109**:1861-1875.e10. doi:10.1016/j.neuron.2021.03.031
- Ossmy O, Moran R, Pfeffer T, Tsetsos K, Usher M, Donner TH. 2013. The timescale of perceptual evidence integration can be adapted to the environment. *Curr Biol* **23**:981-986. doi:10.1016/j.cub.2013.04.039
- Pelli DG. 1997. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision* **10**:437-442. doi:10.1163/156856897X00366
- Piet AT, El Hady A, Brody CD. 2018. Rats adopt the optimal timescale for evidence integration in a dynamic environment. *Nat Commun* **9**:4265. doi:10.1038/s41467-018-06561-y
- Pisauro MA, Fouragnan E, Retzler C, Philiastides MG. 2017. Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nat Commun* **8**:15808. doi:10.1038/ncomms15808
- Raut RV, Snyder AZ, Raichle ME. 2020. Hierarchical dynamics as a macroscopic organizing principle of the human brain. *Proceedings of the National Academy of Sciences* **117**:20890-20897. doi:10.1073/pnas.2003383117
- Ruesseler M. 2021. Perceptual decision-making under uncertainty in humans (<http://purl.org/dc/dcmitype/Text>). University of Oxford.
- Shadlen MN, Newsome WT. 2001. Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *Journal of Neurophysiology* **86**:1916-1936. doi:10.1152/jn.2001.86.4.1916
- Shinn M, Lee D, Murray JD, Seo H. 2022. Transient neuronal suppression for exploitation of new sensory evidence. *Nat Commun* **13**:23. doi:10.1038/s41467-021-27697-4
- Smith NJ, Kutas M. 2015. Regression-based estimation of ERP waveforms: I. The rERP framework. *Psychophysiology* **52**:157-168. doi:10.1111/psyp.12317
- Squires KC, Wickens C, Squires NK, Donchin E. 1976. The effect of stimulus sequence on the waveform of the cortical event-related potential. *Science* **193**:1142-1146. doi:10.1126/science.959831
- Steinemann NA, O'Connell RG, Kelly SP. 2018. Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nat Commun* **9**:3627. doi:10.1038/s41467-018-06117-0
- Straub D, Rothkopf CA. 2021. Putting perception into action: Inverse optimal control for continuous psychophysics. doi:10.1101/2021.12.23.473976
- Twomey DM, Murphy PR, Kelly SP, O'Connell RG. 2015. The classic P300 encodes a build-to-threshold decision variable. *Eur J Neurosci* **42**:1636-1643. doi:10.1111/ejn.12936
- Veliz-Cuba A, Kilpatrick ZP, Josić K. 2016. Stochastic Models of Evidence Accumulation in Changing Environments. *SIAM Rev* **58**:264-289. doi:10.1137/15M1028443
- Wyart V, de Gardelle V, Scholl J, Summerfield C. 2012. Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* **76**:847-858. doi:10.1016/j.neuron.2012.09.015