Simple Human Response Times are Governed by Dual Anticipatory Processes with Distinct and Distributed Neural Signatures

Ashwin G. Ramayya[1]*, Vivek Buch[1], Andrew Richardson[2], Timothy Lucas[3], and Joshua I. Gold[4]

Affiliations:

[1]Department of Neurosurgery, Stanford University, Stanford, CA 94305, USA

[2]Department of Neurosurgery, Hospital of University of Pennsylvania, Philadelphia, PA 19104, USA

[3]NeuroTech Institute, Columbus, OH 43210, USA

[4]Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA

*Corresponding author. Email: aramayya@stanford.edu

Abstract

Human behavior can be highly sensitive to anticipation, but the mechanisms underlying this sensitivity are poorly understood. We obtained intracranial electrocephalography (iEEG) measurements in neurosurgical patients as they performed a simple sensory-motor task with variable (short or long) foreperiod delays that affected anticipation of the cue to respond. Participants showed two forms of anticipatory response biases, distinguished by more premature false alarms (FAs) or faster response times (RTs) on long-delay trials. These biases had distinct neural signatures in prestimulus neural activity modulations that were distributed and intermixed across the brain: the FA bias was most evident in preparatory motor activity immediately prior to response-cue presentation, whereas the RT bias was most evident in visuospatial activity at the beginning of the foreperiod. These results suggest that human anticipatory behavior emerges from a combination of motor-preparatory and attention-like modulations of neural activity, implemented by anatomically widespread and intermixed, but functionally identifiable, brain networks.

Graphical Abstract

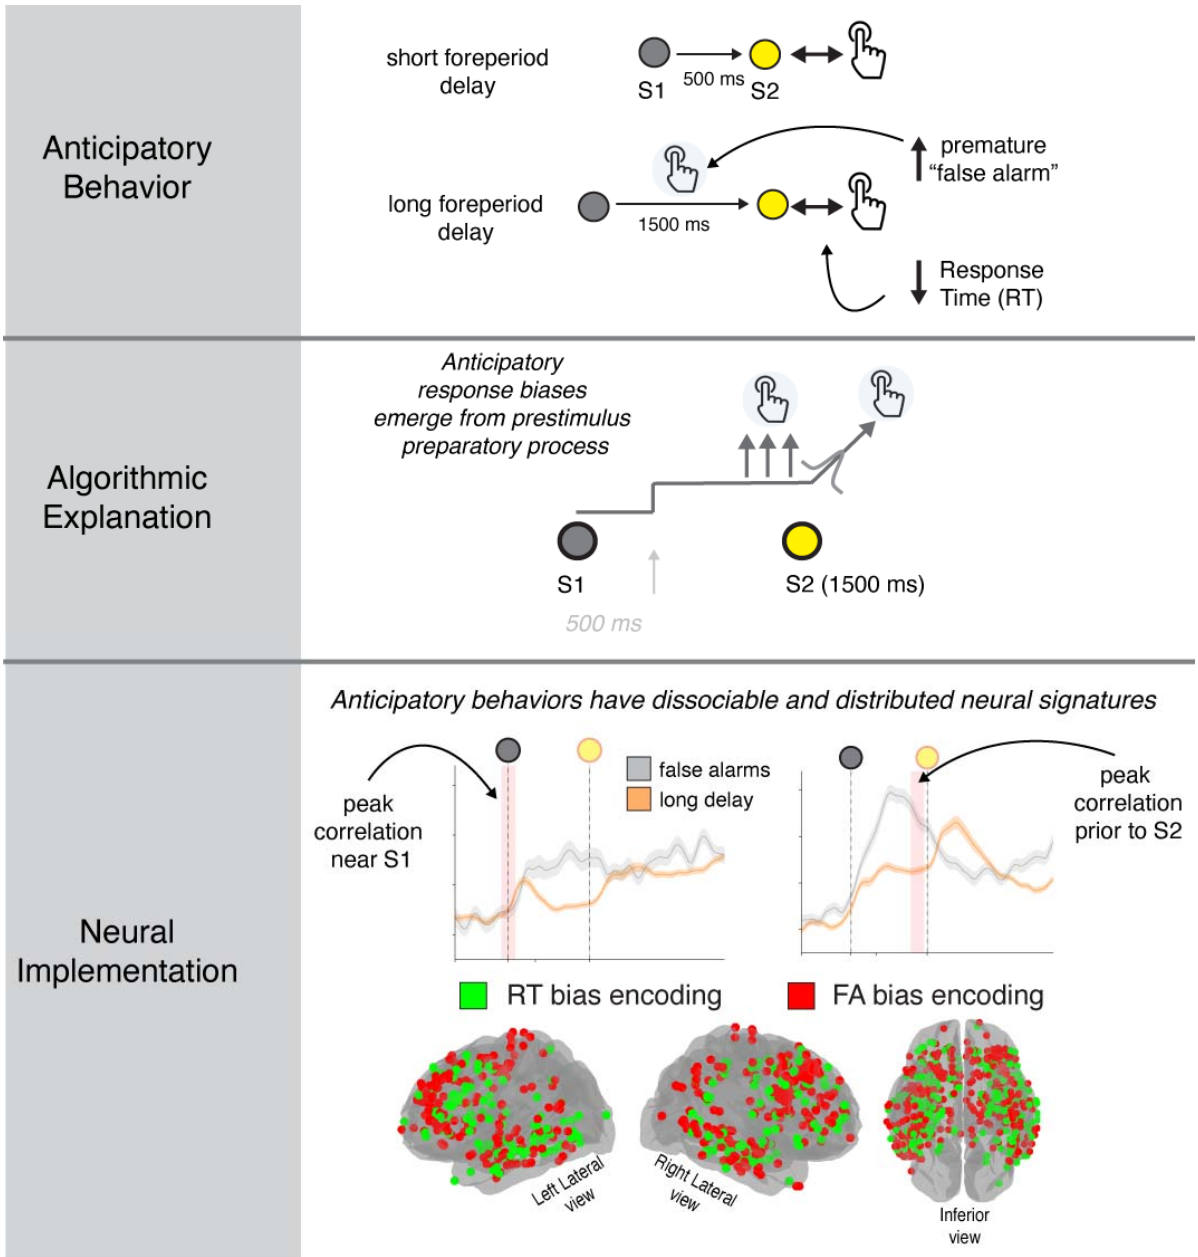

Introduction

Expectations of future events shape human behavior (Friston, 2010; A. C. Nobre & Van Ede, 2018). Even simple stimulus detection is not a passive, purely stimulus-driven process but involves predictive inference that combines incoming sensory information with expectations learned from prior experience (Helmholtz, 1866; Rao & Ballard, 1999). However, the mechanisms in the human brain that allow expectations to influence impending sensory-motor processing (a set of phenomena that we refer to as "anticipation") remains unclear.

Anticipatory influences on human behavior have been quantified extensively using simple sensory-motor detection tasks (Cattell, 1886; Luce, 1986). For example, the variable foreperiod-delay paradigm has been used to operationalize anticipatory behavior in terms of changes in response times (RTs) and premature responses ("false alarms", or FAs) as a function of changes in the length and predictability of the foreperiod delay (Luce, 1986; Niemi & Näätänen, 1981; Ollman & Billington, 1972). These effects are thought to reflect modulations of preparatory motor processes during the foreperiod delay (Los et al., 2001, 2014; A. C. Nobre & Van Ede, 2018; Salet et al., 2022). These preparatory motor processes are often modeled via "rise-to-bound" dynamics that account for endogenous RT variability as arising from a stochastically varying processes that triggers a motor response (Brown & Heathcote, 2005; Carpenter & Williams, 1995; Noorani & Carpenter, 2016; Ratcliff, 1978; Ratcliff et al., 2016; Stone, 1960; Usher & McClelland, 2001). In these models, anticipatory response biases are often assumed to arise from prestimulus elevations in the baseline, or "starting point," of the stochastic rising process, resulting in faster RTs (Carpenter & Williams, 1995; Noorani & Carpenter, 2016). A compelling feature of these models is that, in addition to providing parsimonious accounts of

behavior, they have algorithmic components that are thought to map directly onto the activity patterns of localized sets of neurons that contribute to motor preparation and execution (Hanes & Schall, 1996a).

However, exactly how these models relate to sensory-motor processing in the human brain is not well understood, reflecting a lack of brain measurements with appropriate combinations of high spatiotemporal resolution and broad anatomical scale. Scalp electroencephalography (EEG) studies have provided support for anticipatory processing during the foreperiod delay (Miniussi et al., 1999; Pfeuty et al., 2005; Rohenkohl & Nobre, 2011; Walter et al., 1964), but it is difficult to interpret these signals in terms of specific neural circuits because they aggregate activity across large brain regions. Functional MRI studies have shown regionally distributed hemodynamic correlates of anticipatory processing, but these findings have been inconsistent and are difficult to relate to RT variability because of limited temporal resolution (J. T. Coull & Nobre, 1998; Cui et al., 2009; Vallesi, 2010).

To overcome these limitations, we obtained high-resolution intracranial electroencephalography (iEEG recordings) from 23 patients with medically refractory epilepsy with indwelling intraparenchymal electrodes in widespread brain regions as they performed a stimulus-detection task with a variable foreperiod delay (Fig. 1A, Table S1). We focused on high-frequency iEEG activity, which reflects local spiking activity (70–200 Hz power; Dubey & Ray, 2019; Leonard et al., n.d.; Manning et al., 2009) sampled broadly across many parts of cortex and certain subcortical structures (Parvizi & Kastner, 2018). We tested the hypothesis that anticipatory processes modulate prestimulus activity of preparatory motor neural populations in the human brain that encode endogenous RT variability. We focused on a broad set anatomical

regions because neural correlates of RT variability have been identified in several motor-preparatory brain regions, including activity patterns that map directly (Hanes & Schall, 1996b; O'Connell et al., 2012a) or indirectly (Hauser et al., 2018; Heitz & Schall, 2012; Paraskevopoulou et al., 2021) onto rise-to-bound model dynamics. As detailed below, our results identify two behaviorally and neurally distinct processes that govern anticipatory effects on sensory-motor behavior and highlight the complex but identifiable mappings between algorithm- and implementation-level explanations of human behavior.

Results

Twenty-three participants performed a variant of a commonly used "foreperiod-delay task" that has been used extensively to investigate anticipatory influences on sensory-motor behaviors (Klemmer, 1957; Luce, 1986; Niemi & Näätänen, 1981). Briefly, each trial began with the presentation of a visual target ("warning signal," S1) on a computer screen that changed color after a randomly selected foreperiod delay of 500 ("short") or 1500 ("long") ms. Participants were instructed to respond via button press as soon as they noticed the color change ("stimulus," S2). RT was measured as the elapsed time between stimulus and response. The different foreperiod delays provided categorically different levels of temporal expectation of stimulus arrival at the time of stimulus presentation (Luce, 1986; A. C. Nobre & Van Ede, 2018; Ollman & Billington, 1972; Salet et al., 2022). On short-delay trials, the stimulus was presented when there was uncertainty about whether the trial was a short- or long-delay trial, resulting in relatively low temporal expectation of stimulus arrival. On long-delay trials, the stimulus was presented when the trial could be identified unequivocally as a long-delay trial, resulting in relatively high temporal expectation of stimulus arrival (Fig. 1).

Dual behavioral signatures of anticipatory biases

The participants' RTs included endogenous variability for both delay conditions (median per-participant RT inter-quartile range=66.67 ms for short-delay and 66.75 for long-delay trials; examples are shown in Fig. 1A, Table S2), with two primary effects of anticipation that were consistent with previous findings (Luce, 1986; Nickerson, 1965; Noorani & Carpenter, 2016; Ollman & Billington, 1972). First, participants had faster RTs (paired *t*-test, $t(22)=5.57$, $p<0.001$) on long- versus short-delay trials ("RT bias"), albeit with substantial individual variability (mean RT range across participants=354–595 ms and 323–529 ms on short- and long-

delay trials, respectively, Fig. 1C). Second, participants had higher false-alarm rates on long- versus short-delay trials ("FA bias"; paired *t*-test, $t(22)=4.43$, $p<0.001$; range 1.7–58%, versus 0–7.8%, respectively, Fig. 1D). We did not observe a significant correlation with delay-related differences in RT and false-alarm rate ($p>0.3$, Fig. 1E) but did observe a correlation between RT and false-alarm rate when considering only long-delay trials ($r=0.47$, $p=0.02$).

We modeled these anticipatory effects as prestimulus modulations of an abstracted "rise-to-bound" motor-preparatory process (Fig. 1B; Noorani & Carpenter, 2016a). Specifically, we modeled each RT on short-delay trials as the time taken for a latent variable to rise from a fixed starting point to a fixed bound value to trigger a motor response ("rising process"). Trial-to-trial variability in the rate of rise accounts for endogenous RT variability and the characteristic (delay-independent) right-tailed RT distribution. For correct trials with RT>250 ms, we assumed that this rising process was triggered by the onset of the stimulus S2. In contrast, for trials with false alarms, we assumed that this rising process was triggered prior to the onset of S2, according to a stochastic process that occurred with uniform probability during the 500 ms preceding S2. We modeled each RT on long-delay trials as emerging from a rising process that was modulated by increased temporal anticipation. This anticipation took the form of an elevation of the baseline starting point of the rising process, bringing the process closer to the threshold value required to trigger a response. This baseline elevation increased the probability of triggering responses prior to S2 (FA bias) and reduced the time taken to generate a correct response after S2 (RT bias). We used separate parameters for anticipatory baseline elevation to trigger premature responses versus speed up correct RTs, to account for largely independent variability in RT and FA biases across participants.

This anticipatory starting-point model provided a good fit to participants RTs and anticipatory response biases (R^2 mean=0.85, range=0.62–0.97; Fig. S1 for individual model fits). For comparison, we tested two alternative models with the same number of free parameters but other mechanisms to explain anticipatory RT biases. One alternative model replaced modulations of the starting point with modulations of the variance of the rate of rise. This model produced poorer fits (R^2 mean=0.77, range=0.37–0.97). The other alternative model assumed that FA biases were based on trial-by-trial modulations of the starting point but that RT biases were based on trial-by-trial modulations of the mean rate of rise. This model produced fits that were similar to those produced by the starting-point-only model (R^2 mean=0.86, range=0.62–0.97), which highlights the difficulty in modeling specific algorithmic substrates of anticipatory behavior (Luce, 1986; Salet et al., 2022). Below we focus on the more parsimonious starting-point-only model to identify relevant neural mechanisms.

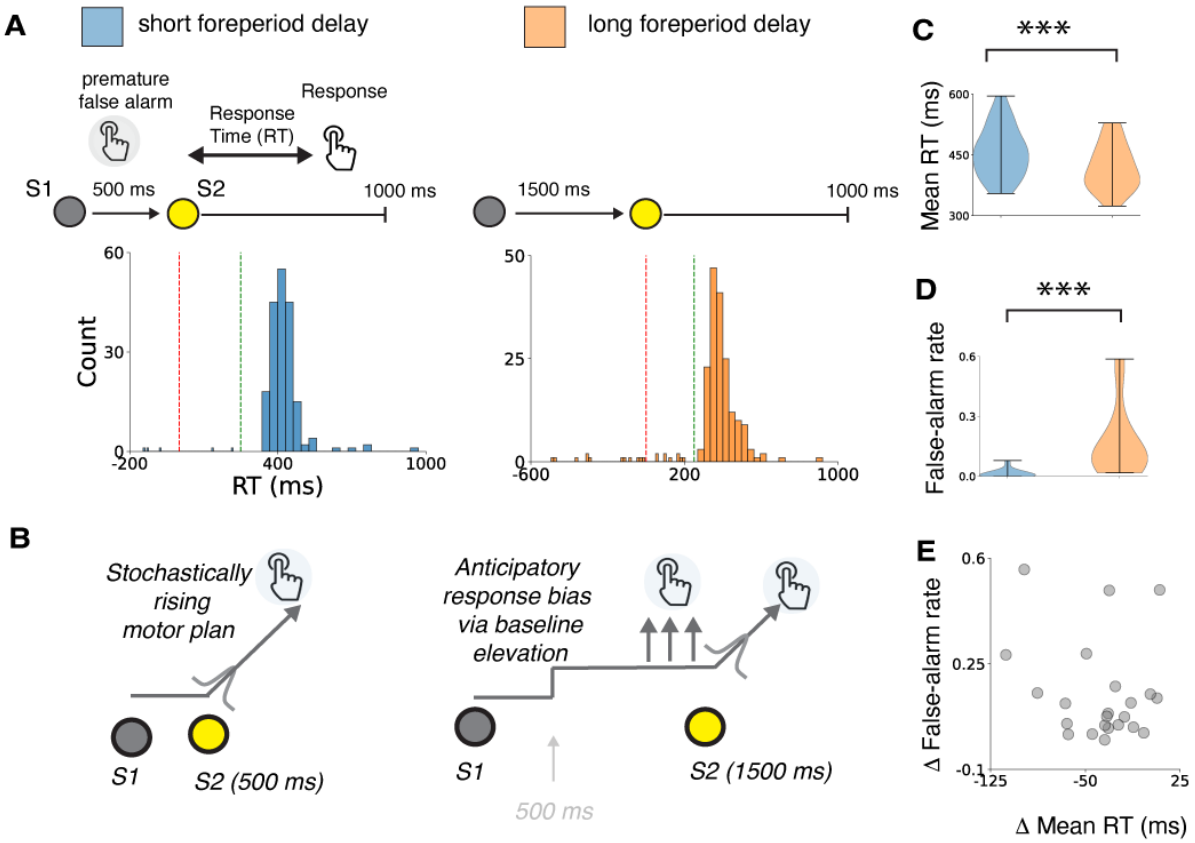

Figure 1. Two forms of anticipatory biases. (**A**) Task summary and RT distributions from an example participant. Red line indicates stimulus onset, green vertical line indicates 250 ms after stimulus onset (the fast-response threshold). Blue and orange histograms indicate timing of responses on short- and long-delay trials, respectively. Premature false alarms are responses that fall to the left of the red line. (**B**) Model schematic illustrating abstracted preparatory motor processes contributing to sensory-motor behavior. Anticipatory elevation of baseline activity can account for both a decrease in RT and an increase in false-alarm rate. (**C,D**) Violin plots showing distributions of mean RTs (**C**) and premature false-alarm rates (**D**) on short- (blue) and long- (orange) delay trials for all 23 participants. (**E**) Scatterplot showing covariance of delay-related changes in mean RT and premature false-alarm rate across participants. Each circle corresponds to data from a single participant.

Task- and RT-modulated neural responses were distributed widely in the brain

We obtained neural measurements from intraparenchymal depth electrodes implanted in participants with medically refractory epilepsy for clinical purposes (Fig. 2A). We focused on

high-frequency activity (HFA, 70–200 Hz power) recorded from bipolar pairs of electrodes, which provides a reliable surrogate of local (within ~3 mm) neural population spiking activity (Dubey & Ray, 2019; Leonard et al., 2023; Manning et al., 2009; Ray & Maunsell, 2011). In total, we studied recordings from 2,609 bipolar pairs of intraparenchymal depth electrodes distributed widely throughout the brain in 23 patients (mean=113.4 electrodes/participant). We localized these recordings to various intrinsic brain networks (Figs. 2B). For each electrode, we measured task-related activity of the nearby neural population time-locked to target onset and motor response in ~50 ms sliding time intervals, z-scored to an aggregate baseline from the entire recording session.

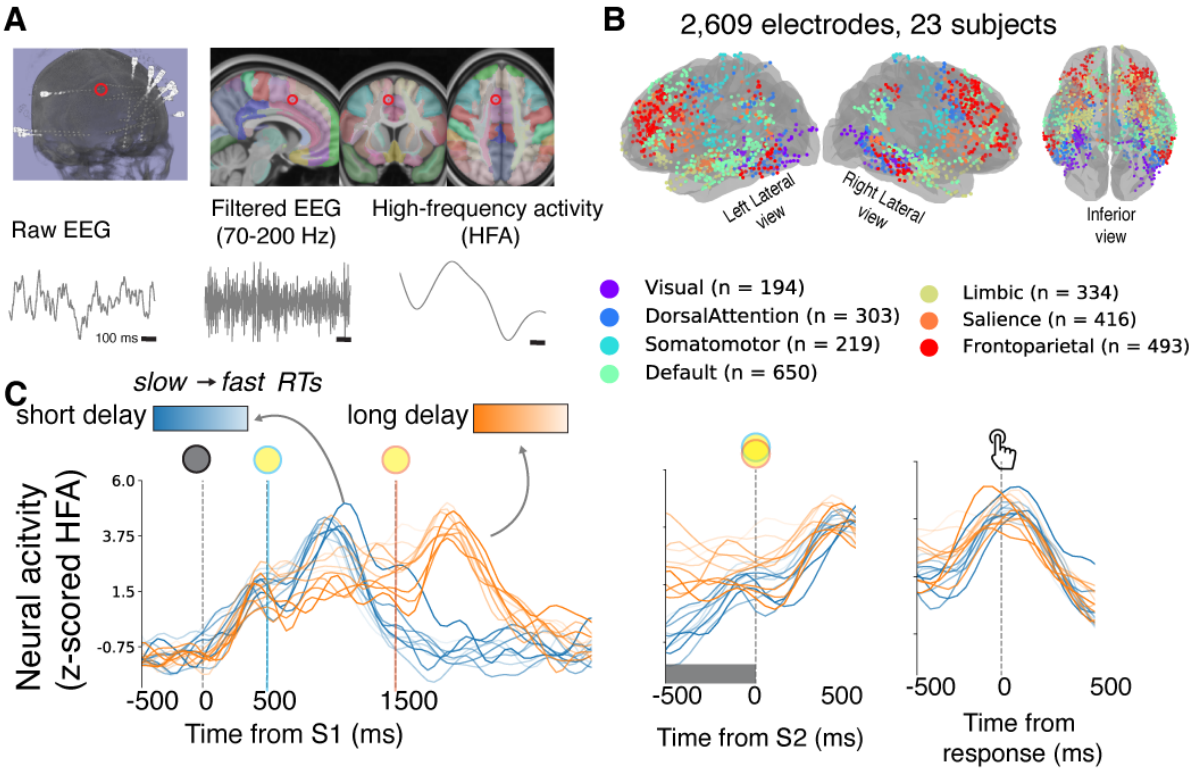

Figure 2. Intracranial recordings provide neural measurements with high resolution and broad anatomical coverage. (**A**) Extracting high-frequency activity (HFA, 70–200 Hz power) from intraparenchymal depth electrodes as an estimate of local spiking (red circle indicates example prefrontal electrode described in **C**). (**B**) Brain plot showing electrode locations from all participants in standard MNI coordinates. Colors indicate intrinsic brain networks based on a normative atlas (Yeo et al., 2011). (**C**) Task-driven responses of local neural activity (average z-scored HFA) measured at the electrode indicated in Fig. 2**A**, plotted separately for short- (blue) and long- (orange) delay trials and binned by stochastic RT percentile (10 bins; lighter shading indicates faster RTs). Left panel shows target-locked activity. Vertical lines indicate time of target onset (grey), short-delay color change (blue), and long-delay color change (orange). Middle and right panels show stimulus- and response-locked activity, respectively. Shaded gray box indicates a time interval during which we observed a correlation between neural activity and RT variability.

We identified task-related modulations and/or correlations with endogenous (delay-independent) RT variability at various time intervals throughout the trial, after accounting for transient sensory- and motor-driven responses, from 2,142 out of 2,609 electrodes (Fig. S2). To identify task-related changes in activity, we compared neural activity following the warning signal (500 ms following S1), stimulus onset (500 ms following S2), and response onset (1000 ms following the button press), relative to a baseline interval (500 ms prior to S1; paired *t*-tests, $p<0.05$). To relate neural activity at each electrode with endogenous, delay-independent RT variability, we used a multivariate model that included neural activity in various task-related time intervals (significance via non-parametric shuffle procedure $p<0.05$; Fig. S2).

An example electrode showing both task- and RT-related activity modulations is shown in Fig. 2C (the electrode location is indicated in Fig. 2A). This local neural population showed rising activity following the warning signal (S1) that peaked near the time of response and descended back to baseline. The prestimulus baseline activity was relatively higher during faster RT and long-delay trials, followed by a largely RT-independent rate of rise before the motor

response. In other words, this electrode showed modulations that were roughly consistent with "rise-to-bound" dynamics found in our and similar models.

However, this electrode was hardly unique in showing task and RT modulations, which we found throughout the brain (Fig. 3). Task modulations across all task epochs involved activity increases, rather than decreases, in response to task-relevant sensory and/or motor events (any interval; two-sided binomial test $p<0.001$, 95% CI: 83–86%; expected 50%), and rarely included both task-driven increases and decreases in different task epochs measured at the same electrode (n=65/2,609, 3%). Task-related increases were generally uniformly distributed across the brain, but neural populations in the salience network showed more frequent task-driven increases than expected (two-tailed binomial test, FDR-corrected $p<0.001$, 95% CI: 74–85%; Fig 3).

RT modulations were also distributed widely, but with different spatial patterns. We observed a regional intermixing of effects that included both positive correlations, such that increased activity corresponded to slower RTs, and negative correlations, such that decreased activity correlated with faster RTs. Neural activity showed more frequent correlations with RT than chance across the brain (n=638/2609, one-tailed binomial test $p<0.001$, 95% CI>23%, where chance=5%) and within each intrinsic brain network (corrected ps<0.001, 95% CIs >15–30%), even when separately considering only positive or negative correlations (corrected $p<0.03$, except limbic populations rarely showed positive correlations, corrected $p>0.5$) We rarely observed neural populations at a single electrode that showed both positive and negative RT correlations during different task epochs (n=60/2,609, 2%). These RT modulations were not distributed uniformly across the brain: visual neural populations showed positive RT correlations more frequently than expected (two-tailed binomial test, corrected $p=0.04$; 17–30%), whereas

somatomotor neural populations showed negative RT correlations more frequently than expected (corrected p=0.02; 24–37%).

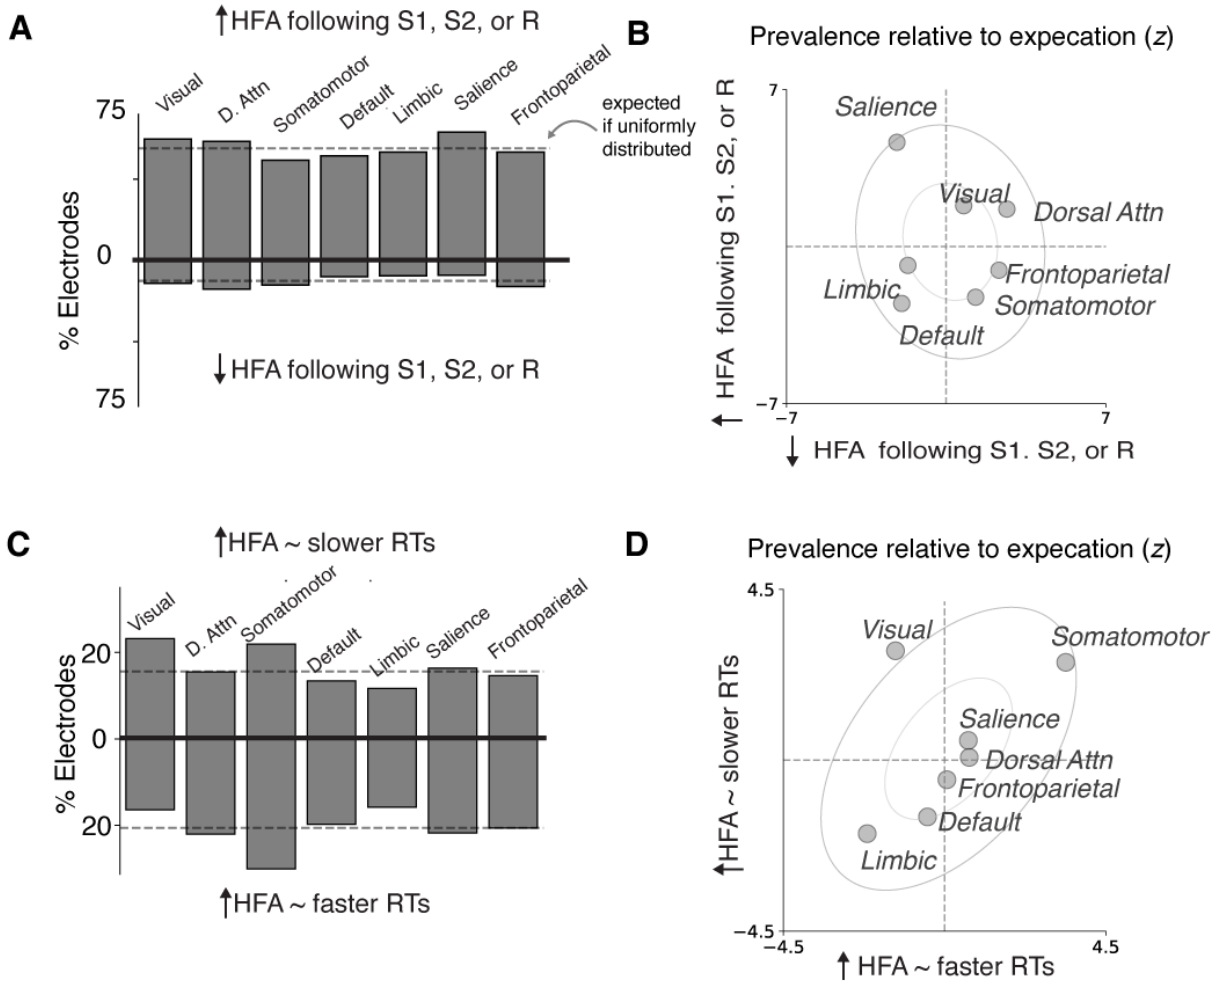

Figure 3. Task-related neural activity changes were regionally widespread and intermixed (**A**) Percentages of electrodes in each region with activity that showed task-related activity changes following S1, S2, or R. Positive values indicate task-related increases in activity; negative values indicate task-related decreases in activity. (**B**) Scatterplot showing the relative frequency of electrodes with positive (ordinate) and negative (abscissa) changes in activity to task-relevant events in each intrinsic brain network relative to their overall (expected) frequency across the brain (z-scores). Positive values indicate increased relative frequency; negative values indicate decreased relative frequency. Inner and outer ellipses indicate 1σ and 2σ confidence intervals derived from the joint distribution, respectively. (**C,D**) Same as **A,B**, but for positive and negative correlations with endogenous, trial-to-trial RT variability during any time interval. Negative values in **C** indicate increased activity with faster RTs; positive values indicate increased activity with slower RTs.

Data-driven clustering of neural activity patterns

To better understand these broadly distributed, diverse activity patterns, we used a data-driven hierarchical clustering algorithm to group electrodes that showed similarities in task-related activity modulations and endogenous, delay-independent RT correlations. We selected a clustering level (4) that maximized the number of clusters that that included data from all of the participants (we also excluded clusters with <200 electrodes) and exhibited distinct patterns of modulations by task events and RT. Cluster 0 showed task-related increases without RT modulation. Cluster 1 showed task-related increases with negative RT correlations. Cluster 2 showed task-related decreases without RT modulation. Cluster 3 showed task-related increases with positive RT correlations. These clusters were distributed widely across the brain (see Fig. S3).

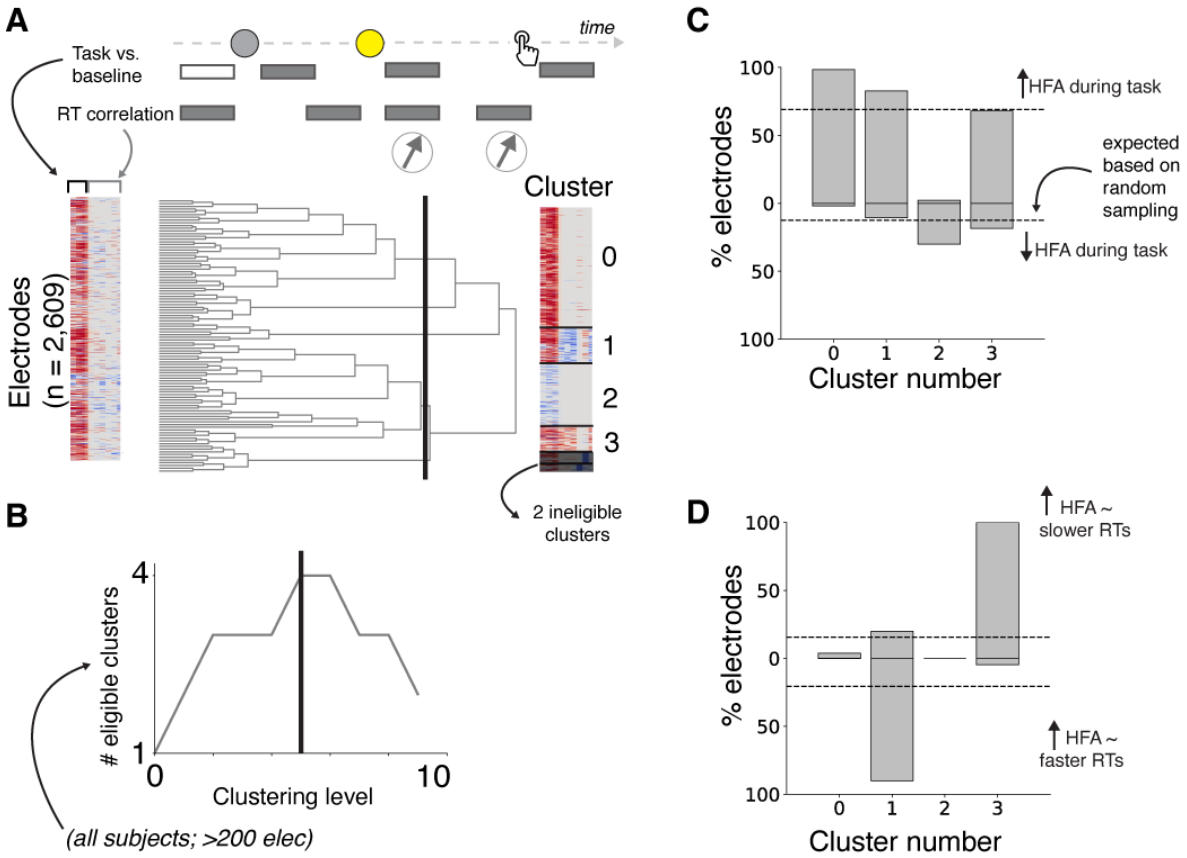

Figure 4. Data-driven clustering of task-related neural activity. Illustration of hierarchical clustering of electrodes based on similar task-related activity modulations. Top row: time intervals used to measure task modulation and RT correlation for each electrode as detailed in main text. Bottom row: left, colormap representing a feature matrix across all electrodes, where each row represents an electrode, and each column represents a feature (task or RT modulation); middle: dendrogram representing similarities between electrodes; right: feature matrix re-organized based on similar task and RT modulations at the clustering level indicated by the black line in middle panel, which we used in these analyses. (**B**) We identified a level of clustering (black vertical line in dendrogram) based on an objective function that maximized the number of clusters that were well sampled in our dataset (i.e., each eligible cluster was observed in all participants and consisted of at least 200 electrodes). (**C**) Percentage of electrodes that showed task-related increases and decreases (positive and values, respectively as in Fig. 3A). (**D**) Same **C** but for RT correlations. Negative values indicate increased activity with faster RTs; positive values indicate increased activity with slower RTs (as in Fig. 3C).

Dual neural signatures of anticipatory biases in prestimulus activity

Prestimulus activity in two distinct neural clusters encoded participant-to-participant variability in the two main anticipatory biases we identified from behavior: RT bias in Cluster 3, and FA bias in Cluster 1 (Fig. 5). Specifically, we measured prestimulus activity in each cluster on long-delay trials (ranging from 250 ms prior to S1 to 50 ms prior to S2, excluding trials with FA and RTs<250 ms, averaged within participants) and related these participant-wise measures to anticipatory RT and FA biases (as shown in Fig 1E). We found that increased prestimulus activity in Cluster 1 correlated with increased FA bias (ρ=0.53, corrected p=0.04; partial correlation controlling for RT bias). In contrast, increased prestimulus activity in Cluster 3 correlated with RT bias (ρ=0.62, corrected p=0.002). We further detail the nature of prestimulus modulations in these two different clusters below.

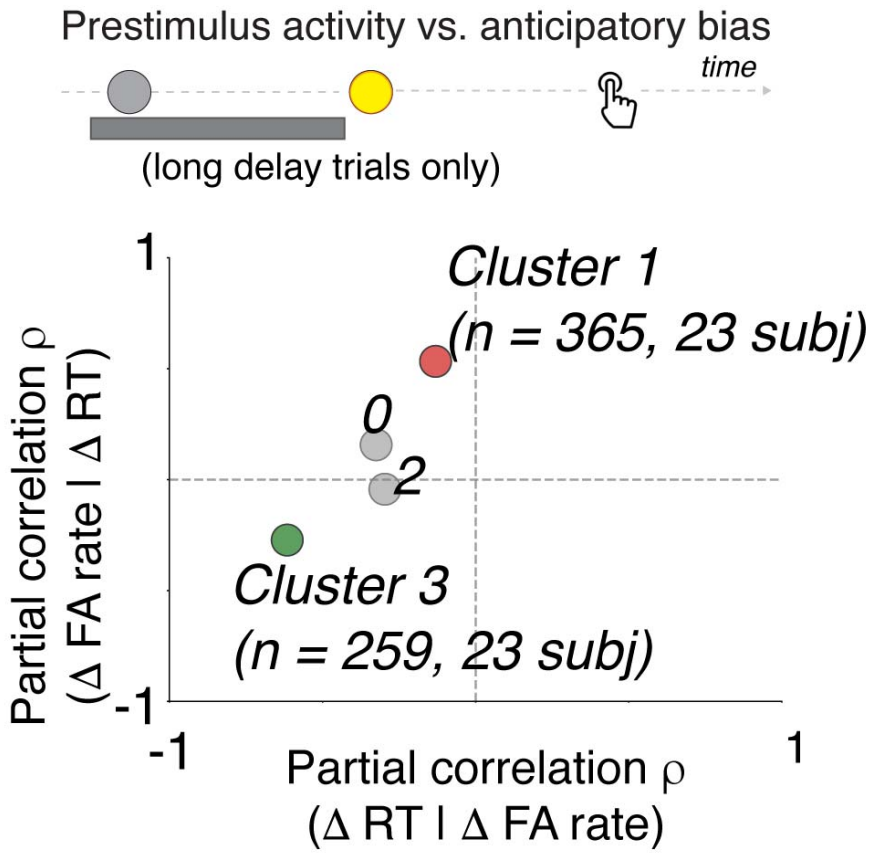

Figure 5. Dissociable prestimulus neural correlates of anticipatory biases. The top schematic shows the time interval used to compute prestimulus activity in each cluster. The scatterplot shows partial correlation coefficients for RT bias (controlling for FA bias) on the abscissa and FA bias (controlling for RT bias) on the ordinate. The green circle indicates a significant relationship with RT bias, and the red circle indicates a significant relationship with FA bias (corrected $p<0.05$).

Anticipatory FA biases were encoded by prestimulus modulations in Cluster 1 neural activity patterns, which resembled a preparatory motor process (Fig. 6). These activity patterns, on average, tended to rise after the warning signal, stay elevated during the foreperiod delay, and peak at the time of the motor response on both correct and false-alarm trials (Fig. 6A,B). There also were reliable modulations by trial-to-trial RT variability, such that increased activity corresponded to faster RTs (Fig. 6A). A time-resolved partial-correlation analysis relating prestimulus activity in Cluster 1 with FA bias (Fig. 6C; 250 ms sliding advanced 10 ms steps, controlling for RT bias) showed a peak correlation strength immediately preceding S2, when expectation of stimulus arrival should be relatively high (1180—1430 ms following S1, highlighted in Fig. 6B; a scatterplot of this correlation is shown in Fig. 6D).

Anticipatory RT biases were encoded by prestimulus modulations in Cluster 3 neural activity patterns, which were activated more strongly by sensory than motor events (Fig. 7). These activity patterns, on average, tended to show a transient rise in activity after the onset of the warning signal and did not show rising activity preceding response, including on FA trials (Fig 7B). They showed positive correlations with endogenous RT variability such that increased activity (including prior to S1) corresponded to slower RTs (Fig. 7A). A time-resolved partial correlation analysis relating prestimulus activity in Cluster 3 with RT bias (7C; 250 ms sliding advanced 10 ms steps, controlling for FA bias) showed a peak correlation strength immediately

near the time of S1 onset, when participant's presumably saccade to target (120 ms prior to and 30 ms following S1, highlighted in Fig 7B; a scatterplot of this correlation is shown in Fig. 7D).

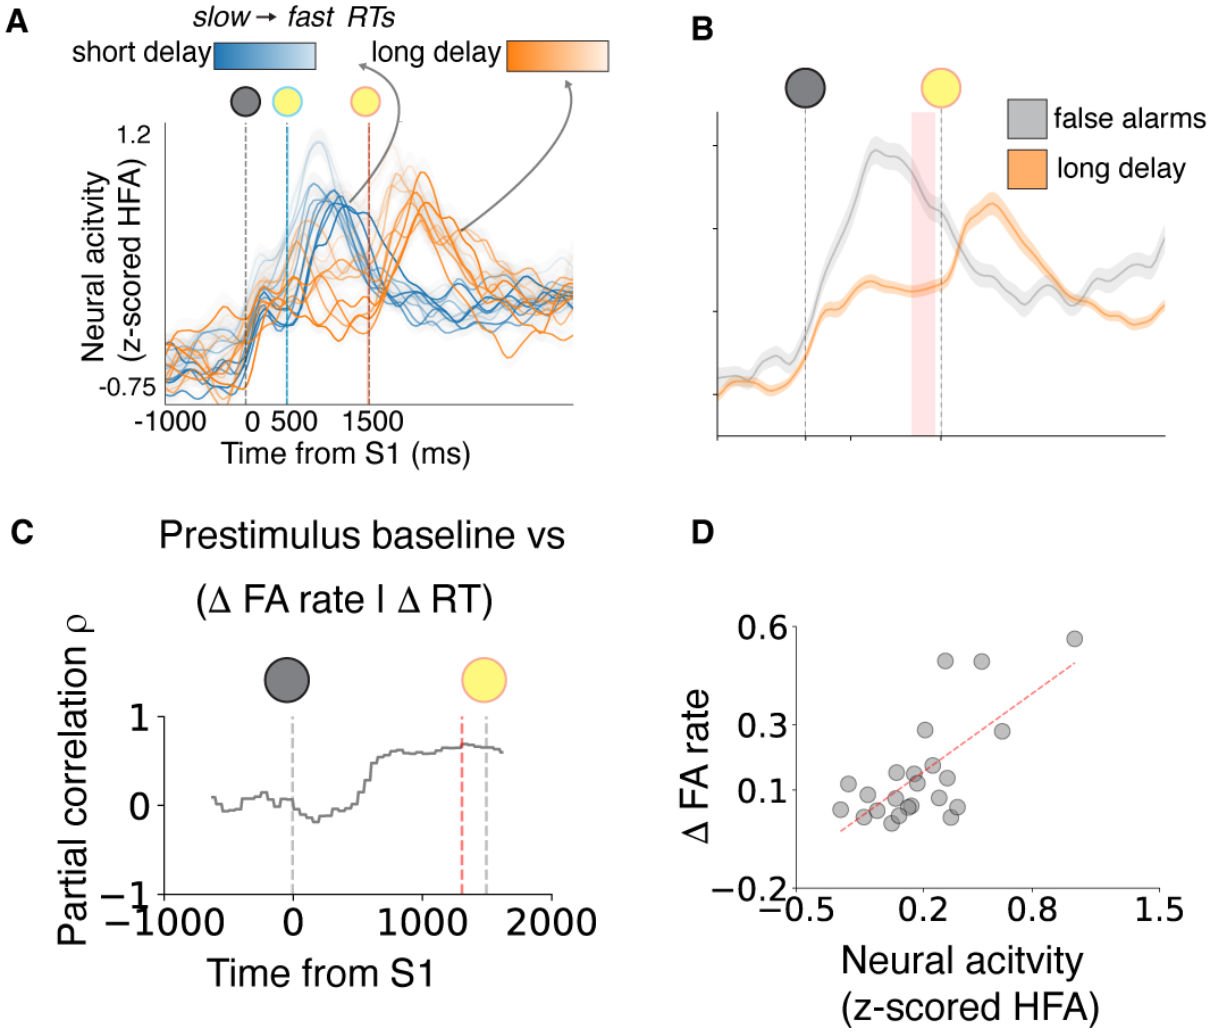

Figure 6. Neural correlates of anticipatory FA bias. (**A**) Average task-related neural responses for Cluster 1 electrodes (same format as Fig. 2C, left panel). (**B**) Average neural activity locked to target onset (S1) for long-delay trials associated with correct responses (RT>250; orange) and false alarms (gray). (**C**) Time resolved partial correlation coefficient relating neural activity with FA bias across participants, controlling for RT bias (vertical red line indicates peak correlation, corresponding to highlighted time interval in B). (**D**) Scatterplot showing across-participant correlation between FA bias and prestimulus neural activity (averaged within the 1180-1430ms window relative to S1, as highlighted in B.).

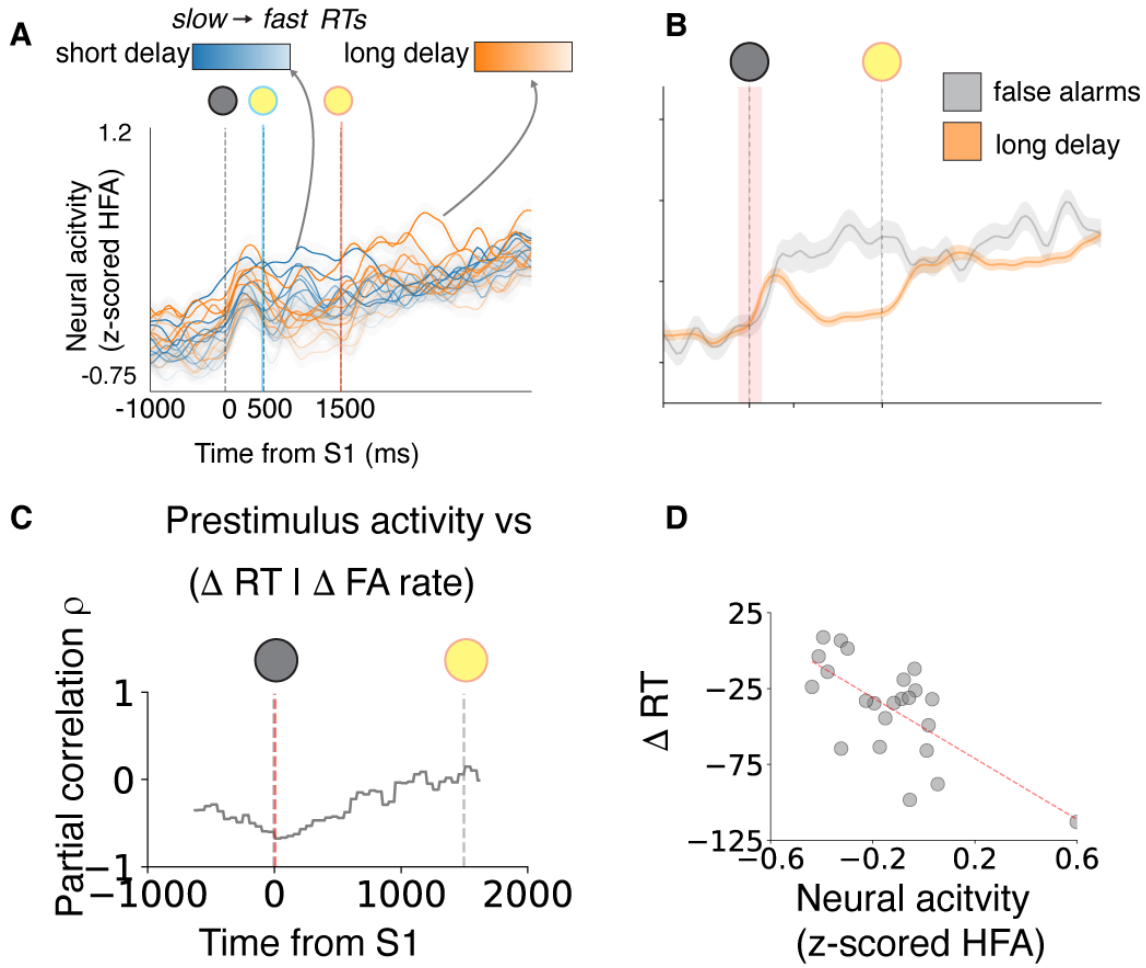

Figure 7. Neural correlates of anticipatory RT bias. (**A-B**) Same as Fig. 6, but for Cluster 3 electrodes (**C**) Time resolved partial correlation coefficient relating neural activity with RT bias across participants, controlling for FA bias (vertical red line indicates peak correlation). (**D**) Scatterplot showing across-participant correlation between RT bias and prestimulus neural activity (averaged within the -120-230ms window relative to S1, as highlighted in B).

These results suggest a functional heterogeneity in prestimulus activity related to anticipatory processing, such that Cluster 1 activity prior to S2 encodes FA bias, whereas Cluster 3 activity around the time of S1 encodes RT bias. We tested for such a functional dissociation using a Linear Mixed Model (LMM) analysis, in which we related prestimulus activity (z-scored HFA) to RT bias (interacting with “cluster,” and “time window,” defined based on peak

correlations from the previous analysis) and FA bias (also interacting with "cluster" and "time window"), with "participant" and "intrinsic brain network" as random effects (varying intercepts; see Materials and Methods). We observed a significant three-way interaction between RT bias, cluster, and time window ($F = 3.93$, $p = 0.005$), and between FA bias, cluster, and time window ($F=19.69$, $p<0.001$), suggesting a functional heterogeneity in prestimulus activity modulations during anticipatory processing.

Cluster 1 and 3 electrodes were widely distributed and regionally intermixed across the brain, but each showed non-uniform distributions (Fig. 8A–C). Cluster 1 electrodes were found more frequently in somatomotor networks (two-tailed binomial test, corrected $p=0.02$; found 15–27%, expected 14%), and less frequently in visual and limbic networks (corrected ps<0.03; 03–13%), than expected. Cluster 3 electrodes were found more frequently in visual networks (two-tailed binomial test, corrected $p=0.02$; found 12–23%, expected 10%) than expected. The visual network was an outlier in showing an increased probability of containing Cluster 3 electrodes, and a decreased probability of containing Cluster 1 electrodes (Fig. 8C).

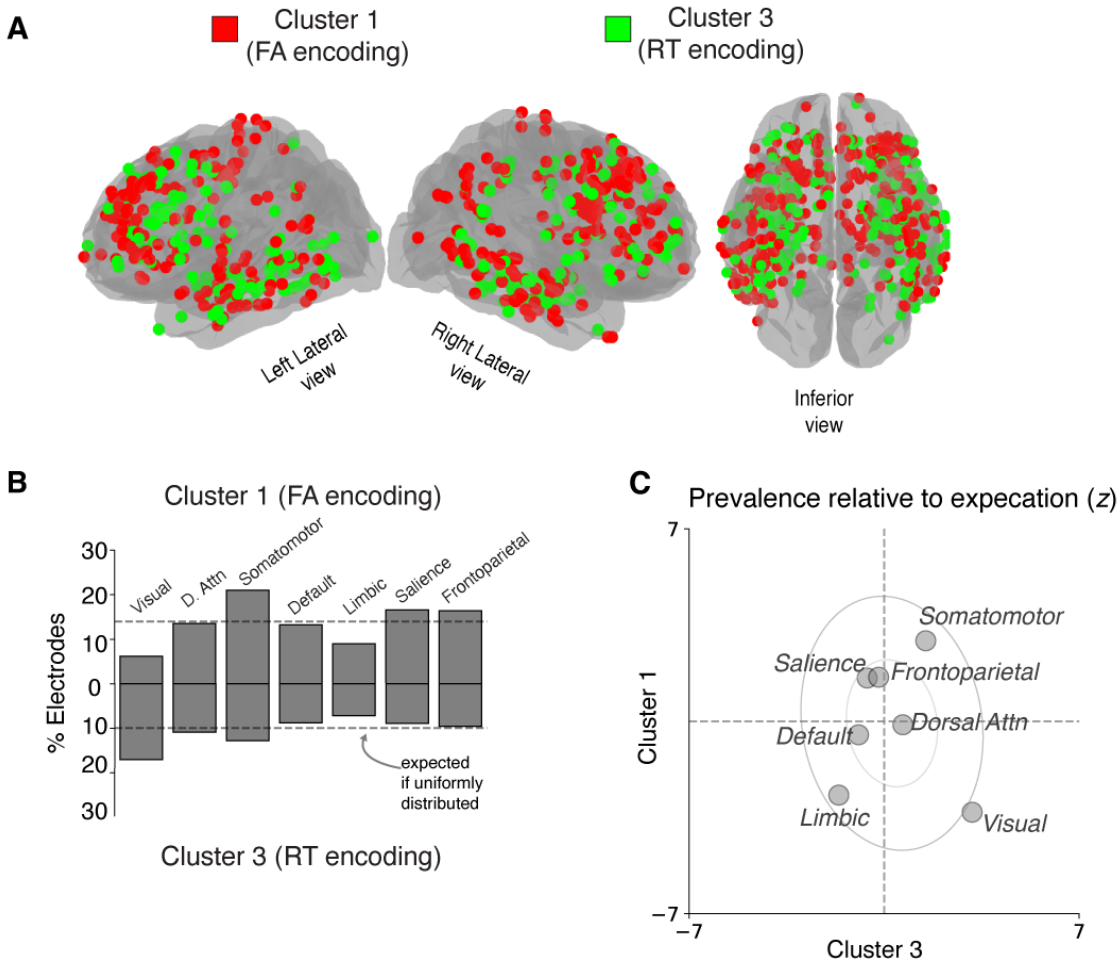

Figure 8. Anatomical distribution of dual distributed networks encoding anticipatory biases. (**A**) Brain plot showing anatomical distribution of Cluster 1 (red) and Cluster 3 (green), which encoded FA and RT biases, respectively. (**B**) Percentage of electrodes in each intrinsic brain network assigned to Cluster 1 (top bars), and Cluster 3 (bottom bars). Horizontal dashed line is expected percentage assuming a uniform anatomical distribution. (**C**) Scatterplot showing the relative frequency of Cluster 1 electrodes (ordinate) and Cluster 3 electrodes (abscissa) in each intrinsic brain network relative to their overall (expected) frequency across the brain (z-scores; same format as Fig. 3B)

Discussion

We identified neural correlates in the human brain of two behaviorally distinguishable effects of anticipation on a simple sensory-motor behavior. The first effect, which we refer to as a false-alarm (FA) bias, was characterized behaviorally by an increase in premature responses under conditions of higher (more certain) anticipation – akin to pressing the car brakes sooner than necessary when you think the car in front of you is about to stop. As has been reported previously, these increased false alarms tended to be accompanied by faster RTs, thus reflecting a form of speed-accuracy trade-off (Green & Swets, 1966; Luce, 1986). We found that these FA biases were encoded by prestimulus neural activity near the end of the foreperiod delay, when the probability of the stimulus was relatively certain, in widespread neural populations that tended to ramp up just preceding the motor response and showed increased activity when RTs were faster, and were prevalent in somatomotor networks, consistent with a role in response generation (Hanes & Schall, 1996a).

The second effect, which we refer to as an RT bias, was characterized behaviorally by faster RTs under conditions of higher (more certain) anticipation – akin to pressing the car breaks faster than usual when you expect, and then see, the brake lights on the car in front of you go on. This bias is consistent with well-established relationships between stimulus uncertainty and mean RT (Klemmer, 1957; Niemi & Näätänen, 1981), including for very similar task designs using randomly interleaved foreperiod delays. Under these conditions, responses are thought to be suppressed while waiting the estimated duration of the short delay (when it is unknown whether it is a short- or long-delay trial) and then facilitated around the estimated end of the long delay (Luce, 1986; A. Nobre et al., 2007; A. C. Nobre & Van Ede, 2018; Ollman & Billington, 1972; Salet et al., 2022). We found that these RT biases were encoded by prestimulus

neural activity at the beginning of the foreperiod delay interval, which is a reference moment for implicit time estimation, in widespread neural that tended to have transient visual responses and build gradually throughout each trial, and showed increased activity when RTs were slower, consistent with a role in visuospatial attention and response inhibition (Houghton & Tipper, 1984; Neill et al., 1995; van Moorselaar et al., 2020).

Previous work also identified many instances of anticipatory modulations of neural activity (e.g., Nobre & Van Ede, 2018). Our work provides new insights into those findings, leveraging the unique combination of high spatio-temporal resolution and broad anatomical sampling of iEEG measurements (Parvizi & Kastner, 2018) to show that such anticipatory-driven modulations of neural activity: 1) are not limited to particular, spatially restricted sensory and/or motor neural populations, as might be inferred from animal electrophysiology studies that typically target spatially restricted recording sites (e.g., Ghose & Maunsell, 2002; Janssen & Shadlen, 2005)□; 2) are more heterogeneous than might be inferred from scalp EEG studies that report only signals that reflect neural activity patterns that have been aggregated across large cortical areas (Miniussi et al., 1999; Rohenkohl & Nobre, 2011; Walter et al., 1964)□; and 3) have more complex temporal dynamics than might be inferred from functional neuroimaging studies with relatively low temporal resolution (Coull & Nobre, 1998; A Vallesi, 2010; Antonino Vallesi et al., 2009). Our results also build on recent work (Paraskevopoulou et al., 2021) showing that that spontaneous fluctuations in neural activity that underlie intertrial variability in human behavior are far more widespread than shown by prior functional neuroimaging studies (Fox et al., 2007).

One of our key findings was that the "rise-to-bound" dynamics of our model, which have previously been used to link algorithmic descriptions of behavior to their implementations in the brain (Gold & Shadlen, 2007; Hanes & Schall, 1996a; Heitz & Schall, 2012; O'Connell et al., 2012b), had a relatively direct mapping to just a subset of the task-relevant neural signals that we identified. Instead, we found two distributed networks with task-relevant modulations of neural activity, corresponding roughly two distinct classes of algorithmic models of anticipatory behavior: those that feature motor preparation (Los et al., 2014; Noorani & Carpenter, 2016; Salet et al., 2022) and those that feature visuospatial attention (Janssen & Shadlen, 2005; Summerfield & Egner, 2009). More generally, our results suggest that the mapping between algorithmic models fit to behavior and brain dynamics is not as straightforward as suggested by certain local and aggregate neural signals, as has been noted previously (Hauser et al., 2018; Heitz & Schall, 2012; O'Connell et al., 2012b, 2018). In particular, the mapping appears to involve neural signals that have diverse forms and locations, which we characterized though our clustering analysis. Results from these analyses are in-line with the emerging view that although the cortex can be segmented into distinct regions based on structural features (Desikan et al., 2006; Fischl et al., 2004), information processing is largely distributed across parallel, interleaved processing streams ("intrinsic brain networks"), with localized information processing limited to certain primary sensory and motor regions (Glasser et al., 2016; Van Essen et al., 1992; Yeo et al., 2011).

Our study has several limitations. First, it is possible that some of these results are particular to our patient population. In general, epilepsy patients can show additional forms of inter-individual variability in their brain networks related to their pathology (Bettus et al., 2008)

that might also manifest in certain categorical behavioral differences compared to healthy controls (Bruhn & Parsons, 1977). However, we sought to mitigate such concerns by focusing on neural signals associated with specific, highly controlled behaviors. Under these kinds of conditions, it has been shown that neural findings from intracranial EEG studies in patients with epilepsy can generalize to healthy controls populations(Long et al., 2014). Nevertheless, more work is needed to fully understand the relationship between behavioral and neural variability across individuals (Genon et al., 2022), which can have broad evolutionary (Bechara et al., 2000), developmental (Tenenbaum et al., 2011), and functional (Yang & Wang, 2020) causes. Second, our clustering analysis was intended to identify distinct functional neural population that were evident across all subjects we studied. A more granular clustering of distinct functional profiles may be possible with a larger dataset. Third, we focused on within-trial anticipatory biases for this study but recognize that across-trial biases may also occur (Salet et al., 2022). Further work is needed to understand if and how within- and across-trial anticipatory biases relate to each other on a neural level.

Despite these limitations, our study provided new insights into the distributed neural processes in the human brain that support anticipatory influences on sensory-motor behaviors. Despite relatively simple algorithmic explanations underlying anticipatory behavior, we found that anticipatory computations had heterogeneous, distributed, and regionally intermixed neural correlates in the human brain. These results reflect the fundamental role of anticipation in higher brain function and can help explain why simple sensory-motor processing engages widespread brain networks (Gonzalez-Castillo et al., 2012; Mesulam, 1998). These results illustrate how high-resolution neural measurements in the human brain can complement algorithmic models in

illuminating cognitive processes underlying human behavior. Moreover, these results motivate the need for more high-resolution and large-scale neural recordings combined with sophisticated computational models of brain activity and behavior to help bridge gaps in our understanding of the computational underpinnings of behavior and their complex neural substrates (Bogacz & Gurney, 2007; Wong & Wang, 2006; Yang & Wang, 2020).

Materials and Methods

Participants

We studied 23 patients with medically refractory epilepsy who underwent surgical implantation of intracranial electrodes for seizure localization (Table S1). Patients provided informed consent to perform cognitive testing as part of our research study while they were admitted to the hospital. Our study was approved by the University of Pennsylvania Institutional Review Board. Clinical circumstances alone determined the number and placement of implanted electrodes.

Stimulus-detection task

We used a stimulus-detection task with a variable foreperiod delay (Luce, 1986; Nickerson, 1965; Niemi & Näätänen, 1981). Participants viewed visual stimuli on a laptop computer screen and responded by pressing a button on a game controller with their right thumb. Each trial began with the presentation of a small white box at a randomized location on the screen as a fixation target (one of nine locations on a 3 x 3 grid). The stimulus changed color to yellow after one of two randomly selected foreperiod delays: 1) short delay=500 ms, or 2) long delay=1500 ms. We measured response time (RT) as the time between color change (stimulus) and button press (response). If a response was provided within a predefined response interval

(1000 ms after color change), visual feedback was provided as follows: 1) for RTs ≤ 150 ms, salient positive feedback was shown, consisting of a smiley face with text below ("wow!"); 2) for RTs > 150 ms, the RT was shown, color-coded by binned values (green for 150–300 ms, yellow for 300–600 ms, red for 600–1000 ms). If a response occurred before the color change of the stimulus ("premature false alarm") or if no response was provided within 1000 ms after color change, the participant was shown a blank screen for 2500–2750 ms. During each session, the participant performed blocks of 9 trials each. After six blocks, the participant was provided with an option of performing additional blocks of trials. On average, participants performed a mean of 149 trials (i.e., 1.34 sessions). Two participants performed additional trials with a 1000 ms foreperiod delay, which we excluded from our analyses.

Extended LATER model of RT

All model-fitting procedures described below were implemented by finding the minimum of the appropriate objective function using the SciPy Optimization package in Python. We modeled RT distributions using an extended version the LATER model(Carpenter & Williams, 1995; Noorani & Carpenter, 2016). This rise-to-bound model describes RTs as measuring the time it takes for a latent variable to rise linearly from a starting value (S_0) at stimulus presentation to a threshold value (S_t) at response. The model accounts for trial-to-trial RT variability by assuming that the rate-of-rise of this process varies stochastically from trial-to-trial per a Gaussian process. The basic form of the LATER model describes an RT distribution using two free parameters, as follows:

$$RT \sim \frac{1}{N(\mu, \sigma^2)}$$

Where N represents a Gaussian distribution with mean μ and standard deviation σ. Note that this version assumes that the threshold value S_t=1; an equivalent formulation assumes that σ=1, and S_t is a free parameter. These parameters are estimated by fitting a Gaussian probability density function to reciprocal RTs, which typically resembles a Gaussian distribution (Noorani & Carpenter, 2016).

We extended this basic model to account for premature false alarms and RTs<250 ms, which we modeled as a stochastic process that occurred with uniform probability during the 500 ms preceding S2. For these and other RT analyses, we assumed that this false alarm-generating process also triggered responses within 250 ms after stimulus onset ("fast responses"), because these fast responses are thought to arise from processes distinct from those driving slower RTs (we chose 250 ms as a conservative cutoff to ensure that our neural analyses do not conflate these two processes; Luce, 1986; Noorani & Carpenter, 2016).

We first fit this extended LATER model to behavior on the short-delay trials. We then fit a modified model to behavior on long-delay trials that included two additional free parameters representing delay-related changes in: 1) the starting point of the LATER unit, and 2) the uniform probability of generating a premature or fast response. We evaluated the overall model fits by computing an R^2 value comparing the (*z*-scored) model-predicted probability distributions and an empirically estimated probability distribution based on a Gaussian-smoothed reciprocal RT histograms (including all responses from 500 ms prior to stimulus onset to 1000 ms following stimulus onset, with an offset such that all RTs on anticipatory false alarms were positive values and RTs on short and long-delay trials were aligned relative to stimulus onset),

and also the cumulative probability of observing a premature false alarm based on the uniform distribution (see Fig. S1 for individual model fits for each participant).

Intracranial neural recordings via intraparenchymal electrodes

Patients were implanted exclusively with intraparencymal depth electrodes ("stereo EEG," Ad-tech, 1.1 diameter, 4 contacts spaced 5 mm apart), except in one patient who also had subdural grid electrodes (Participant #142, Ad-tech, 4 mm contacts, spaced 10 mm apart). iEEG was recorded using a Natus recording system. Based on the amplifier and the discretion of the clinical team, signals were sampled at either 512 or 1024 Hz. Signals were converted to a bipolar montage by taking the difference of signals between each pair of immediately adjacent contacts on the same electrode. The resulting bipolar signals were treated as new virtual electrodes (henceforth, "electrodes"), originating from the midpoint between each electrode pair (Burke et al., 2014; Ramayya et al., 2015, 2021). Digital pulses synchronized the electrophysiological recordings with task events. We excluded electrodes that recorded prominent 60 Hz electrical line noise, defined as electrodes that showed greater spectral power in the 58–62 Hz range compared to the 18–22 Hz range, or electrodes that were disconnected (standard deviation=0). We excluded trials with prominent noise artifacts (e.g., if voltage data were not recorded due to saturation, or if the mean or standard deviation of voltage was >10 standard deviations of all trials). We did not specifically exclude electrodes based on epileptic activity because our analyses focused on behaviorally linked neural activity, which should not be influenced systematically by epileptic networks (Liu & Parvizi, 2019). We analyzed data from 2,609 electrodes.

Anatomical localization of electrodes

Intracranial electrodes were identified manually on each post-operative CT scan and contact coordinates recorded using custom software based on the center of density of the radiodense contacts on thresholded images. To obtain contact locations in each patient's native anatomic space as well as a common reference space (MNI coordinates), we used Advanced Normalization Tools (Avants et al., 2011)to register the post-operative CT to the pre-operative MRI, and the MRI to the Montreal Neurological Institute (MNI) average brain. We assigned each electrode to various canonical intrinsic brain networks ("7 network model") using a volumetric atlas in MNI coordinates (Yeo et al., 2011). We refer to the "ventral attention" network as the "salience" network based on its resemblance to behaviorally defined networks important for emotion (Seeley et al., 2007), but otherwise use terminology as reported in the original study (Yeo et al., 2011).

Extracting high-frequency activity (HFA)

We extracted 5000 ms segments of iEEG data around the following task events: 1) 2000 ms prior to target onset, 2) 5000 ms after target onset, 3) the stimulus color change, and 4) onset of the motor response. For each segment, we extracted spectral power using five complex-valued Morlet wavelets (wave number 3, to increase temporal resolution) with logarithmically spaced center frequencies from 70 to 200 Hz. We squared and log-transformed the wavelet convolutions to obtain power estimates at each time sample. We removed 1000 ms buffers at the beginning and end of each segment to avoid contamination from edge artifacts. We averaged these power estimates across the 5 wavelets, resulting in a single power value for each time sample. We convolved each power-time series with a Gaussian kernel (half-width of 75 ms), resulting in a continuous representation of high-frequency activity (HFA) surrounding each task event. We z-

scored HFA by the mean and standard deviation of the distribution of HFA values obtained from randomly selected segments of iEEG data recorded from that session clips (matched to the number of total task events) to not bias values towards any particular task-related event (Burke et al., 2014; Ramayya et al., 2015, 2021). We refer to *z*-scored HFA as "HFA."

Measuring and preprocessing electrode-specific activation functions

We quantified the activity pattern recorded by each electrode by measuring the average HFA around target onset during, measured separately on short-delay and long-delay trials (subsequently referred to as "activation functions"). We measured target-locked activation functions from 1000 ms prior to, until 4000 ms following target onset.

To ensure that all subsequent analyses focused on modulations of the activation functions that were separate from the immediate sensory and motor responses, we preprocessed the data as follows (all curve fits were performed using the *SciPy* Optimization package in *Python* ("cuve_fit"):

First, we estimated the overall (linear) trends in HFA that occurred over the course of the trial independent of changes that occurred immediately following stimulus presentation using the target-locked activation function (averaged across both trial types). Specifically, we fit a line to estimate how HFA changed over time using during two short segments of this activation function: 1) 1000 ms to -500 before target onset, and 2) 1000 ms after color change occurred on long-delay trial (2500 ms to 3000 ms after target onset). We subtracted these lines from each target-locked activation function to obtain detrended functions.

Second, we estimated the changes in activity that occurred immediately after target onset across both trial types. Specifically, if the detrended target-locked activation function (averaged

570 across both trial types, from target onset to 500 ms afterwards) contained any activity peaks or
571 troughs that were outside of the electrode's baseline activity range, we fit a Gaussian function to
572 model these activity excursions. If both a peak and trough were detected in a given electrode's
573 activation function, we modeled whichever one was larger. We then removed the post-target
574 Gaussian from each target-locked activation function. For long-delay trials, we also fit a
575 Gaussian to the 500–1500 ms interval after target onset and subtracted it from the activation
576 function to remove any expectation-related changes in activity that occurred in response to the
577 target not changing color.

578 Third, we modeled activity locked to the color change separately for short- and long-
579 delay trials. Specifically, we first fit a Gaussian to a short time segment after the color change
580 (0–500 ms post-stimulus) to capture any changes in activity that occurred immediately after
581 color change. We also fit any residual peaks or troughs over a longer time segment (0–1000 ms
582 post-stimulus). We then subtracted these fit Gaussians from the activation function from each
583 trial.

584 Fourth, we modeled activity locked to the motor response separately for short- and long-
585 delay trials. Specifically, we fit separate Gaussians to the 500 ms preceding the response, to
586 capture any pre-response ramping activity or any smeared stimulus-locked activity from the
587 preceding color change, and to any residual peaks or troughs during the 500 ms post-response
588 interval, to capture further changes in activity related to either the response or feedback
589 presentation. We then subtracted these fit Gaussians from the activation function from each trial.

In summary, these procedures resulted in residual trial-by-trial HFA measures that were time-locked to four task events: 1) target-locked, short delay; 2) target-locked, long delay; 3) response-locked, short delay; and 4) response-locked, long delay.

Relating residual neural activity to endogenous RT variability

We measured endogenous (non delay-related) RT variability for each participant as follows. First, we transformed RTs during correct responses (excluding "fast response" RTs < 250 ms) by taking the negative reciprocal ($-1/RT$), which transforms right-tailed RTs into an approximately Gaussian distribution (the negative sign is applied for convenience such that long RTs are still associated with larger values; Noorani & Carpenter, 2016). Then, we removed delay-related RT variability by z-scoring reciprocal RTs within each delay condition.

We assessed whether endogenous RT variability was related to neural activity at a given electrode on a trial-by-trial basis, after removing stimulus- and response-locked components of the activation function, including ramping components (Fig. S2). We considered several trial-by-trial neural features locked to trial events: 1) baseline interval prior to target onset (500 ms), 2) baseline interval prior to stimulus (500 ms), 3) post-stimulus activity (250 ms), 4) post-stimulus buildup rate (slope of a line fit to HFA trend 250 ms following stimulus), 5) pre-response buildup rate (slope of a line fit to HFA 250 ms prior to response), and 6) pre-response activity (250 ms).

We performed an omnibus test to assess for any relationship between neural activity and stochastic RT variability using a multi-variate linear regression model. The dependent variable was stochastic RT variability across all trials, and the independent variables were each of the neural features described above. We measured the predictive power of this model using the sum

squared error of residuals (SSE). We also assessed specific relationships between each neural feature and stochastic RT variability using Spearman correlations.

We assigned non-parametric *z*-statistics and *p* values for each of these tests by comparing each of these "true" statistics (SSE of multi-variate model, and Spearman's *rho* for each neural features) to null distributions generated for each electrode by misaligning RTs and neural data (using a circular shift procedure, 1000 iterations, to account for any autocorrelation in RTs). For the omnibus test, we assigned a one-tailed *p* value (i.e., if the true SSE was <5% of null SSE values, we assigned a p value of 0.05). For Spearman correlations, we assigned a two-tailed *p* value (i.e., if *rho* was >2.5% of null *rho* values, we assigned a *p* value of 0.05).

Hierarchical clustering of neural populations based on functional properties.

We used data-driven unsupervised clustering to group all electrodes based on similar task-driven neural activation patterns. This approach allowed us to study task-relevant neural representations that were possibly distributed across many regions. For each electrode, we defined a multi-variate feature vector representing the magnitude of (linear) relationships between: 1) each of the five features of the activation function listed above (i.e., baseline, post-stimulus activity, post-stimulus buildup rate, pre-response buildup rate, pre-response activity) and stochastic RT variability (i.e., z-scored RT computed separately for short- and long-delay trials, thus removing any mean difference in RT for the two delay conditions), and 2) task-related modulations relative to pre-target baseline activity (Fig. 3). We characterized this feature vector as a multinomial distribution indicating the presence and direction of a significant effect ($p<0.05$; 0 indicates no relation, or no significant task-related modulation; 1 indicates relatively increased activity with long RTs or task-related increases in activity; and -1 indicates relatively increased

activity with short RTs or task-related decreases in activity). We used an agglomerative clustering algorithm (*scikit learn, Python*) to group all electrodes using Euclidean distance as a measure of pairwise similarity and a linkage function that merges clusters to minimize the variance within all clusters. We identified 4 clusters that were grouped by similar functional profiles using an objective function maximized the number of clusters that were observed in all participants and contained at least 200 electrodes.

Statistical tests

We performed non-parametric statistical tests as described above, when appropriate. Otherwise, we performed *t*-tests to compare continuous distributions and binomial to compare categorical distributions (counts data). We used partial correlation analysis to assess across-participant correlations between neural activity and anticipatory biases (delay-related differences in mean RT and false alarm rate). Based on our partial correlation analysis, we specifically assessed for a functional dissociation between prestimulus activity modulations in Cluster 1 and Cluster 3 using a linear mixed-effects model, as follows:

$$Y \sim 1 + (\Delta\mathrm{RT} \times time \times cluster) + (\Delta\mathrm{FA} \times time \times cluster) + (1|subj) + (1|roi)$$

Where Y is a continuous variable that represents neural activity (z-scored HFA) averaged within each subject for a specific cluster (Cluster 1 or 3), and a specific time window (-120-130ms, "early" or 1180-1430ms, "late" relative to S1, as identified by the partial correlation analyses, Figs. 6, 7). The Fixed Effects are two separate three-way interactions between ΔRT (continuous variable indicating delay-related difference in mean RT for each participant), time (categorical variable indicating early or late time window), and cluster (categorical variable indicating Cluster 1 or 3), and between ΔFA (continuous variable indicating delay-related difference in FA

655 rate for each participant), time, and cluster. The Random Effects are subject ("subj") and
656 intrinsic brain network ("roi").
657 We used False Discovery Rate (FDR) correction for multiple comparisons (Benjamini &
658 Hochberg, 1995) and considered an FDR-corrected *p* value <0.05 to be statistically significant.
659 We report FDR-corrected *p* values when indicated, otherwise reported *p* values are uncorrected.
660 We performed most of our analyses using Python using both custom code and publicly available
661 packages (e.g., *NumPy* for numerical computing, *SciPy* for statistics and signal processing, *MNE*
662 for spectral analyses, *statsmodels* for regression modeling, *pingouin* for partial correlation
663 analysis). We fit linear mixed effects models using in R using the lme4 and lmeTest packages
664 (Bates et al., 2015; Kuznetsova et al., 2017).
665 Data sharing and Code Accessibility
666 Behavioral and neural data used for this study will be made publicly available and have
667 been submitted to Dryad. The Python code used to process these data are available online
668 (https://github.com/ashwinramayya/code_RamaEtal_AntReact)
669
670
671

Acknowledgments: We thank Drs. Michael Kahana, Michael Beauchamp, Tahra Eissa, Matthew Nassar, Long Ding, Kareem Zaghloul, Daniel Yoshor, Alex Vaz, and William Songjun Li for helpful feedback on the manuscript. We thank Cameron Brandon, Everett Prince, John Bernabei, Jacqueline Boccanfuso, and Joel Stein for technical assistance in data collection.

Funding: National Institutes of Health grant 6T32NS091006 (AGR)

691 References:

692 Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., & Gee, J. C. (2011). A
693 reproducible evaluation of ANTs similarity metric performance in brain image registration.
694 *NeuroImage*, *54*(3), 2033–2044. https://doi.org/10.1016/j.neuroimage.2010.09.025

695 Barack, D. L., & Krakauer, J. W. (2021). Two views on the cognitive brain. *Nature Reviews*
696 *Neuroscience*, *22*(June). https://doi.org/10.1038/s41583-021-00448-6

697 Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, Decision Making and the
698 Orbitofrontal Cortex. *Cerebral Cortex*, *10*(3), 295–307.
699 https://doi.org/10.1093/CERCOR/10.3.295

700 Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and
701 Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B*
702 *(Methodological)*, *57*(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

703 Bettus, G., Wendling, F., Guye, M., Valton, L., Régis, J., Chauvel, P., & Bartolomei, F. (2008).
704 Enhanced EEG functional connectivity in mesial temporal lobe epilepsy. *Epilepsy*
705 *Research*, *81*(1), 58–68. https://doi.org/10.1016/j.eplepsyres.2008.04.020

706 Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision
707 making between alternative actions. *Neural Computation*, *19*(2), 442–477.
708 https://doi.org/10.1162/neco.2007.19.2.442

709 Bruhn, P., & Parsons, O. A. (1977). Reaction Time Variability in Epileptic and Brain-Damaged
710 Patients. *Cortex*, *13*(4), 373–384. https://doi.org/10.1016/S0010-9452(77)80018-X

711 Burke, J. F., Zaghloul, K. A., Jacobs, J., Williams, R. B., Sperling, M. R., Sharan, A. D., &
712 Kahana, M. J. (2013). Synchronous and asynchronous theta and gamma activity during
713 episodic memory formation. *The Journal of Neuroscience : The Official Journal of the*
714 *Society for Neuroscience*, *33*(1), 292–304. https://doi.org/10.1523/JNEUROSCI.2057-
715 12.2013

716 Callaway, E., & Yeager, C. L. (1960). Relationship between reaction time and
717 electroencephalographic alpha phase. *Science*, *132*(3441), 1765–1766.
718 https://doi.org/10.1126/science.132.3441.1765

719 Carpenter, R. H. S., & Williams, M. L. L. (1995). Neural computation of log likelihood in
720 control of saccadic eye movements. *Nature*, *377*(6544), 59–62.
721 https://doi.org/10.1038/377059a0

722 Cattell, J. M. (1886). The time taken up by cerebral operations. *Mind*, *11*(42), 220–242.

723 Churchland, M., Byron, M., Cunningham, J., Sugrue, L., Cohen, M., Corrado, G., Newsome, W.,
724 Clark, A., Hosseini, P., Scott, B., & Bradley, D. (2010). Stimulus onset quenches neural
725 variability: a widespread cortical phenomenon. *Nature Neuroscience*, *13*(3).
726 https://doi.org/10.3389/conf.neuro.06.2009.03.295

727 Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: The neural systems for
728 directing attention to spatial locations and to time intervals as revealed by both PET and
729 fMRI. *Journal of Neuroscience*, *18*(18), 7426–7435. https://doi.org/10.1523/jneurosci.18-
730 18-07426.1998

731 Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R.
732 L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., & Killiany, R. J. (2006). An
733 automated labeling system for subdividing the human cerebral cortex on MRI scans into
734 gyral based regions of interest. *NeuroImage*, *31*(3), 968–980.
735 https://doi.org/10.1016/j.neuroimage.2006.01.021

736 Dubey, A., & Ray, S. (2019). Cortical electrocorticogram (Ecog) is a local signal. *Journal of*
737 *Neuroscience*, *39*(22), 4299–4311. https://doi.org/10.1523/JNEUROSCI.2917-18.2019

738 Dustman, R. E., & Beck, E. C. (1965). Phase of alpha brain waves, reaction time and visually
739 evoked potentials. *Electroencephalography and Clinical Neurophysiology*, *18*(5), 433–440.
740 https://doi.org/10.1016/0013-4694(65)90123-9

741 Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience. In
742 *Neuron* (Vol. 109, Issue 19, pp. 3055–3068). Cell Press.
743 https://doi.org/10.1016/j.neuron.2021.07.011

744 Eichenbaum, H. (2014). Time cells in the hippocampus: A new dimension for mapping
745 memories. In *Nature Reviews Neuroscience* (Vol. 15, Issue 11, pp. 732–744).
746 https://doi.org/10.1038/nrn3827

747 Fischl, B., Van Der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., Busa, E.,
748 Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., & Dale, A.
749 M. (2004). Automatically Parcellating the Human Cerebral Cortex. *Cerebral Cortex*, *14*(1),
750 11–22. https://doi.org/10.1093/cercor/bhg087

751 Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic Fluctuations within
752 Cortical Systems Account for Intertrial Variability in Human Behavior. *Neuron*, *56*(1), 171–
753 184. https://doi.org/10.1016/j.neuron.2007.08.023

754 Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews*
755 *Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Genon, S., Eickhoff, S. B., & Kharabian, S. (2022). Linking interindividual variability in brain structure to behaviour. *Nature Reviews Neuroscience*. https://doi.org/10.1038/s41583-022-00584-7

Ghose, G. M., & Maunsell, J. H. R. (2002). Attentional modulation in visual cortex depends on task timing. *Nature*, *419*(6907), 616–620. https://doi.org/10.1038/nature01057

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., & Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. https://doi.org/10.1038/NATURE18933

Goel, A., & Buonomano, D. V. (2014). Timing as an intrinsic property of neural networks: Evidence from in vivo and in vitro experiments. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 369, Issue 1637). https://doi.org/10.1098/rstb.2012.0460

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*(1), 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Gonzalez-Castillo, J., Saad, Z. S., Handwerker, D. A., Inati, S. J., Brenowitz, N., & Bandettini, P. A. (2012). Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(14), 5487–5492. https://doi.org/10.1073/pnas.1121049109

Hanes, D., & Schall, J. (1996). Neural control of voluntary movement initiation. *Science*, *274*(5286), 427--430.

Hauser, C. K., Zhu, D., Stanford, T. R., & Salinas, E. (2018). Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. *ELife*, *7*. https://doi.org/10.7554/eLife.33456

Heitz, R. P., & Schall, J. D. (2012). Neural Mechanisms of Speed-Accuracy Tradeoff. *Neuron*, *76*(3), 616–628. https://doi.org/10.1016/j.neuron.2012.08.030

Helmholtz, H. Von. (1866). Treatise on Physiological Optics. In *Book* (Issues 3d ed, 1910).

Howard, M. W., Shankar, K. H., Aue, W. R., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, *122*(1), 24–53. https://doi.org/10.1037/a0037840

Humphries, M. D. (2021). Strong and weak principles of neural dimension reduction. *Neurons, Behavior, Data Analysis, and Theory*. https://doi.org/10.51628/001c.24619

Janssen, P., & Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience*, *8*(2), 234–241. https://doi.org/10.1038/nn1386

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*(NOV). https://doi.org/10.3389/neuro.06.004.2008

Landman, B., & Warfield, S. (2012). MICCAI workshop on multi-atlas labeling. *Medical Image Computing and Computer Assisted Intervention Conference.*

Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, *349*(6244), 184–187. https://doi.org/10.1126/SCIENCE.AAA4056/SUPPL_FILE/LATIMER-SM.PDF

Liu, S., & Parvizi, J. (2019). Cognitive refractory state caused by spontaneous epileptic high-frequency oscillations in the human brain. *Science Translational Medicine*, *11*(514), 1–14. https://doi.org/10.1126/scitranslmed.aax7830

Long, N. M., Burke, J. F., & Kahana, M. J. (2014). Subsequent memory effect in intracranial and scalp EEG. *NeuroImage*, *84*, 488–494. https://doi.org/10.1016/j.neuroimage.2013.08.052

Los, S. A., Kruijne, W., & Meeter, M. (2014). Outlines of a multiple trace theory of temporal preparation. In *Frontiers in Psychology* (Vol. 5, Issue SEP). https://doi.org/10.3389/fpsyg.2014.01058

Los, S. A., & Van Den Heuvel, C. E. (2001). Intentional and unintentional contributions to nonspecific preparation during reaction time foreperiods. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(2), 370–386. https://doi.org/10.1037/0096-1523.27.2.370

Luce, R. D. (1986). Response Times: Their Role in Inferring Elementary Mental Organization. In *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195070019.001.0001

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontiguous events. *Neuron*, *71*(4), 737–749. https://doi.org/10.1016/j.neuron.2011.07.012

Manning, J. R., Jacobs, J., Fried, I., & Kahana, M. J. (2009). Broadband Shifts in Local Field Potential Power Spectra Are Correlated with Single-Neuron Spiking in Humans. *Journal of Neuroscience*, *29*(43), 13613–13620. https://doi.org/10.1523/JNEUROSCI.2041-09.2009

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., … Mazoyer, B. (2001). A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *356*(1412), 1293–1322. https://doi.org/10.1098/rstb.2001.0915

Mesgarani, N., Cheung, C., Johnson, K., & Chang, E. F. (2014). Phonetic feature encoding in human superior temporal gyrus. *Science*, *343*(6174), 1006–1010. https://doi.org/10.1126/science.1245994

Mesulam, M. M. (1998). From sensation to cognition. *Brain*, *121*(6), 1013–1052.

Miniussi, C., Wilding, E. L., Coull, J. T., & Nobre, A. C. (1999). Orienting attention in time. Modulation of brain potentials. *Brain*, *122*(8), 1507–1518. https://doi.org/10.1093/brain/122.8.1507

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, *16*(5), 1936–1947. https://doi.org/10.1523/jneurosci.16-05-01936.1996

Natraj, N., Silversmith, D. B., Chang, E. F., & Ganguly, K. (2022). Compartmentalized dynamics within a common multi-area mesoscale manifold represent a repertoire of human hand movements. *Neuron*, *110*(1), 154-174.e12. https://doi.org/10.1016/j.neuron.2021.10.002

Nickerson, R. S. (1965). Response time to the second of two successive signals as a function of absolute and relative duration of intersignal interval. *Perceptual and Motor Skills*, *21*(1), 3–10. https://doi.org/10.2466/pms.1965.21.1.3

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, *89*(1), 133–162. https://doi.org/10.1037/0033-2909.89.1.133

Nobre, A. C., & Van Ede, F. (2018). Anticipated moments: Temporal structure in attention. In *Nature Reviews Neuroscience* (Vol. 19, Issue 1, pp. 34–48). https://doi.org/10.1038/nrn.2017.141

Noorani, I., & Carpenter, R. H. S. (2016). The LATER model of reaction time and decision. In *Neuroscience and Biobehavioral Reviews* (Vol. 64, pp. 229–251). https://doi.org/10.1016/j.neubiorev.2016.02.018

O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735. https://doi.org/10.1038/nn.3248

O'Connell, R. G., Shadlen, M. N., Wong-Lin, K. F., & Kelly, S. P. (2018). Bridging Neural and Computational Viewpoints on Perceptual Decision-Making. *Trends in Neurosciences*, *41*(11), 838–852. https://doi.org/10.1016/j.tins.2018.06.005

Ollman, R. T., & Billington, M. J. (1972). The Deadline model for simple reaction times. *Cognitive Psychology*, *3*(2), 311–336. https://doi.org/10.1016/0010-0285(72)90010-2

Paraskevopoulou, S. E., Coon, W. G., Brunner, P., Miller, K. J., & Schalk, G. (2021). Within-subject reaction time variability: Role of cortical networks and underlying neurophysiological mechanisms. *NeuroImage*, *237*. https://doi.org/10.1016/j.neuroimage.2021.118127

Parvizi, J., & Kastner, S. (2018). Promises and limitations of human intracranial electroencephalography. *Nature Neuroscience*, *21*(4), 474–483. https://doi.org/10.1038/s41593-018-0108-2

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*(4), 1113–1143. https://doi.org/10.1037/a0020311

Ramayya, A. G., Pedisich, I., & Kahana, M. J. (2015). Expectation modulates neural representations of valence throughout the human brain. *NeuroImage*, *115*, 214–223. https://doi.org/10.1016/j.neuroimage.2015.04.037

Ramayya, A. G., Yang, A. I., Buch, V. P., Burke, J. F., Richardson, A. G., Brandon, C., Stein, J. M., Davis, K. A., Chen, H. I., Proekt, A., Kelz, M. B., Litt, B., Gold, J. I., & Lucas, T. H. (2021). Theta Synchrony Is Increased near Neural Populations That Are Active When Initiating Instructed Movement. *ENeuro*, *8*(1). https://doi.org/10.1523/ENEURO.0252-20.2020

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Ratcliff, R., Philiastides, M. G., Sajda, P., & Shiffrin, R. M. (2009). *Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG*. *106*(16), 6539–6544. https://doi.org/10.1073/pnas.0812589106

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. In *Trends in Cognitive Sciences* (Vol. 20, Issue 4, pp. 260–281). https://doi.org/10.1016/j.tics.2016.01.007

Ray, S., & Maunsell, J. H. R. (2011). Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLoS Biology*, *9*(4), e1000610. https://doi.org/10.1371/journal.pbio.1000610

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning and the effectiveness of reinforcement and non-reinforcement. In *Classical Conditioning. 2. Current Research and Theory*. (pp. 64–69). https://ci.nii.ac.jp/naid/10024275851/

Rohenkohl, G., & Nobre, A. C. (2011). Alpha oscillations related to anticipatory attention follow temporal expectations. *Journal of Neuroscience*, *31*(40), 14076–14084. https://doi.org/10.1523/JNEUROSCI.3387-11.2011

Salet, J., Kruijne, W., van Rijn, H., Los, S. A., & Meeter, M. (2022). fMTP: A unifying computational framework of temporal preparation across time scales. In *Psychological Review*.

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., & Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *Journal of Neuroscience*, *27*(9), 2349–2356. https://doi.org/10.1523/JNEUROSCI.5587-06.2007

Shadlen, M. N., Kiani, R., Newsome, W. T., Gold, J. I., Wolpert, D. M., Zylberberg, A., Ditterich, J., De Lafuente, V., Yang, T., & Roitman, J. (2016). Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making." *Science (New York, N.Y.)*, *351*(6280), 1406. https://doi.org/10.1126/SCIENCE.AAD3242

Shankar, K. H., & Howard, M. W. (2012). A scale-invariant internal representation of time. In *Neural Computation* (Vol. 24, Issue 1, pp. 134–193). https://doi.org/10.1162/NECO_a_00212

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*. https://www.sciencedirect.com/science/article/pii/S1364661309001442?casa_token=EMs14nsV5yUAAAAA:drBcxAV1eu4AZBYurLNwKW0_VDmrxgm9A88W5GYzHUr4zObJzwyJkeQnETpgeSESLzzT5CHcDg

915 Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind:
916 Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285.
917 https://doi.org/10.1126/SCIENCE.1192788/SUPPL_FILE/TENENBAUM.SOM.REV1.PD
918 F

919 Umakantha, A., Morina, R., Cowley, B. R., Snyder, A. C., Smith, M. A., & Yu, B. M. (2021).
920 Bridging neuronal correlations and dimensionality reduction. *Neuron*, *109*(17), 2740-
921 2754.e12. https://doi.org/10.1016/j.neuron.2021.06.028

922 Vallesi, A. (2010). Neuro-anatomical substrates of foreperiod effects. In *Attention and Time*.
923 https://books.google.com/books?hl=en&lr=&id=LBoDED__q3gC&oi=fnd&pg=PA303&dq
924 =vallesi+neuroanatomical+substrates&ots=8sKa_kv7Ih&sig=v1E1sodk_uTkVt9O1ELS-
925 mW_v2E

926 Vallesi, Antonino, McIntosh, A. R., Shallice, T., & Stuss, D. T. (2009). When time shapes
927 behavior: fMRI evidence of brain correlates of temporal monitoring. *Journal of Cognitive*
928 *Neuroscience*, *21*(6), 1116–1126. https://doi.org/10.1162/jocn.2009.21098

929 Van Essen, D. C., Anderson, C. H., & Felleman, D. J. (1992). Information processing in the
930 primate visual system: an integrated systems perspective. In *Science* (Vol. 255, Issue 5043,
931 pp. 419–423). https://doi.org/10.1113/jphysiol.2010.201798

932 Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through Neural
933 Population Dynamics. In *Annual Review of Neuroscience* (Vol. 43, Issue 1, pp. 249–275).
934 https://doi.org/10.1146/annurev-neuro-092619-094115

935 Walter, W. G., Cooper, R., Aldridge, V. J., McCallum, W. C., & Winter, A. L. (1964).
936 Contingent negative variation□: An electric sign of sensori-motor association and
937 expectancy in the human brain. *Nature*, *203*(4943), 380–384.
938 https://doi.org/10.1038/203380a0

939 Warrington, S., Bryant, K. L., Khrapitchev, A. A., Sallet, J., Charquero-Ballester, M., Douaud,
940 G., Jbabdi, S., Mars, R. B., & Sotiropoulos, S. N. (2020). XTRACT - Standardised
941 protocols for automated tractography in the human and macaque brain. *NeuroImage*, *217*,
942 116923. https://doi.org/10.1016/j.neuroimage.2020.116923

943 Wong, K. F., & Wang, X. J. (2006). A recurrent network mechanism of time integration in
944 perceptual decisions. *Journal of Neuroscience*, *26*(4), 1314–1328.
945 https://doi.org/10.1523/JNEUROSCI.3733-05.2006

946 Yang, G. R., & Wang, X. J. (2020). Artificial Neural Networks for Neuroscientists: A Primer.
947 *Neuron*, *107*(6), 1048–1070. https://doi.org/10.1016/j.neuron.2020.09.005

948 Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M.,
949 Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fisch, B., Liu, H., & Buckner, R. L.
950 (2011). The organization of the human cerebral cortex estimated by intrinsic functional
951 connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165.
952 https://doi.org/10.1152/JN.00338.2011

953

954

955