Population Representation of the Confidence in a Decision in the Lateral Intraparietal Area of the Macaque

Ariel Zylberberg[1,2*] **and Michael N. Shadlen**[1,3,4,5]

[1]Mortimer B Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States;
[2]Virtual Confidence and Metacognition Laboratory;
[3]Department of Neuroscience, Columbia University, New York, United States;
[4]The Kavli Institute for Brain Science, Columbia University, New York, United States;
[5]Howard Hughes Medical Institute, Chevy Chase, United States

*For correspondence:
ariel.zylberberg@gmail.com

Abstract   Confidence in a decision is the belief, prior to feedback, that one's choice is correct. In the brain, many decisions are implemented as a race between competing evidence-accumulation processes. We ask whether the neurons that represent evidence accumulation also carry information about whether the choice is correct (i.e., confidence). Monkeys performed a reaction time version of the random dot motion task. Neuropixels probes were used to record from neurons in the lateral intraparietal (LIP) area. LIP neurons with response fields that overlap the choice-target contralateral to the recording site (T_{in} neurons) represent the accumulation of evidence in favor of contralateral target selection. We demonstrate that shortly before a contralateral choice is reported, the population of T_{in} neurons contains information about the accuracy of the choice (i.e., whether the choice is correct or incorrect). This finding is unexpected because, on average, T_{in} neurons exhibit a level of activity before the report that is independent of reaction time and evidence strength—both strong predictors of accuracy. This apparent contradiction is resolved by examining the variability in neuronal responses across the population of T_{in} neurons. While on average, T_{in} neurons exhibit a stereotyped level of activity before a contralateral choice, many neurons depart from this average in a consistent manner. From these neurons, the accuracy of the choice can be predicted using a simple logistic decoder. The accuracy of the choice predicted from neural activity reproduces the hallmarks of confidence identified in human behavioral experiments. Therefore, neurons that represent evidence accumulation can also inform the monkey's confidence.

Introduction

Choice, reaction time, and confidence are often considered the three pillars of choice behavior; a comprehensive model of decision-making should account for all three. However, the most prevalent models of binary decision making, Signal Detection Theory (SDT) and the Drift-Diffusion Model (DDM), can only account for two. In SDT, the choice depends on the sign of the difference between a sample of evidence and a decision criterion; confidence is a monotonically increasing function of the magnitude of that difference (Galvin et al., 2003; Kepecs and Mainen, 2012). Because SDT frames the decision as a categorization of one sample of evidence, SDT cannot account for reaction times, other than to posit that evidence closer to criterion might lead to slower choices, linked to uncertainty; (e.g., Carpenter and Williams, 1995). In the DDM, the decision is made by accumulating many samples of evidence over time (Ratcliff, 1978; Palmer et al., 2005; Shadlen et al., 2006; Ratcliff et al., 2016). The decision ends when the accumulated evidence exceeds an upper or lower bound. The model naturally accounts for choice and reaction time, but lacks a straightforward explanation of confidence. This is because at the moment of choice, the state of accumulated evidence is uninformative about the accuracy of the choice.

In the brain, simple binary decisions are implemented as a race between two competing evidence-accumulation processes (Gold and Shadlen, 2007; Hanks et al., 2015). The first process to reach an upper bound terminates the decision and determines the choice and reaction time (RT). A case in point is the random dot motion task, in which monkeys make binary decisions about the net direction of random dot motion and communicate their decision with a saccadic eye movement. Neurons in the lateral intraparietal area (LIP) and associated brain regions represent the accumulation of noisy momentary motion evidence for and against each response alternative. The decision process is well captured by race models of decision-making, which generalize drift-diffusion models by allowing the

competing evidence-accumulation processes to be imperfectly anticorrelated (Wang, 2002; Usher and McClelland, 2001).

In addition to being supported by neurobiology, race models offer the leading explanation of choice, reaction time, and confidence in simple binary decisions, guided by the accumulation of evidence. Confidence is often modeled under the *balance of evidence* hypothesis, which postulates that confidence is a function of the difference, at the moment of choice, between the evidence accumulated by the winning race and the losing race (Vickers, 1979). Since the evidence accumulated by the winning race at the moment of choice is at its upper bound, confidence is determined by the state of the accumulation that has not reached its upper bound (i.e., the losing race). The greater the distance is from its upper bound, the weaker the accumulated evidence is for the losing alternative, and the stronger the confidence is for the chosen alternative. Models that embrace the balance of evidence hypothesis are able to account for several behavioral regularities of confidence, including the relationship between confidence and evidence strength, reaction time, and accuracy (Vickers 1979; Kiani et al. 2014; van Den Berg et al. 2016; Brus et al. 2021; Smith and Vickers 1988; Hellmann et al. 2023; Moreno-Bote 2010; Rolls et al. 2010; Wei and Wang 2015; Vivar Lazo 2024; Vickers et al. 1985, but see Zylberberg et al. 2012; Comay et al. 2023).

We tested a key prediction of confidence models based on the balance of evidence hypothesis, namely, that the state of the losing race contains more information about the accuracy of the choice than the state of the winning race. We reanalyzed data from a recent study in which Steinemann, Stine et al. (2022) used high-density Neuropixels probes to record the spiking activity of a broad population of neurons in the lateral intraparietal area (LIP) of the macaque brain. The monkeys made simple decisions about the net direction—left or right—in a stochastic random dot motion display. LIP neurons with response fields that overlap the choice-target contralateral to the recording site (T_{in} neurons) represent the accumulation of evidence in favor of contralateral target selection. On trials where the monkey chooses the contralateral target (i.e., the one within the neurons' response field), the T_{in} neurons represent the *winning* race, whereas when the monkey chooses the ipsilateral target, the same neurons represent the *losing* race. We use the single-trial population response of T_{in} neurons to predict the accuracy of the choice.

Contrary to the prediction of balance of evidence models of confidence, we found that the winning race contains more information about the accuracy of the choice than the losing race. This finding is unexpected because the average firing rate of T_{in} neurons is known to reach a stereotyped level just before the monkey issues a saccadic eye movement to the contralateral choice target. The key insight is that the average firing rate of of the T_{in} neurons belie considerable heterogeneity of T_{in} responses across the population. We find that not all T_{in} neurons reach a common level of activity at the time of choice. Instead, some T_{in} neurons maintain a trace of evidence strength, while others maintain a representation of elapsed decision time. From the T_{in} neurons, it is possible to decode the probability that the choice is correct, and this probability exhibits hallmarks of confidence identified in human behavioral experiments.

Results

Task, behavior and neurophysiological recordings

We analyzed multineuron recordings previously published by Steinemann, Stine et al. (2022). In this study, two rhesus monkeys (*Macaca mulatta*) reported their decisions about the net direction of motion in a dynamic random dot display. The monkeys indicated their choices by redirecting their gaze from a central fixation point to a left or right choice target. Monkeys were allowed to indicate their decision when ready, thus giving rise to two behavioral measures, choice and reaction time[1] (RT) (Fig. 1A). The degree of difficulty was controlled by the *motion coherence*, defined as the probability that a dot displayed at time t will be redrawn in the direction of motion when replotted 40 ms later, as opposed to being randomly repositioned. On each trial, motion coherence was selected pseudorandomly from the list $\{\pm 0\%, \pm 3.2\%, \pm 6.4\%, \pm 12.8\%, \pm 25.6\%, \pm 51.2\%\}$. The sign of the motion coherence indicates the direction (positive for leftward). For the 0% coherence motion, the sign indicates the random direction to be rewarded on that trial. The proportion of leftward choices increases with motion coherence, and reaction

---

[1]we use reaction time and choice-response time synonymously, as both refer to the same latency: from onset of the random motion to the onset of the saccadic eye movement used to report the choice. We prefer reaction time to disambiguate behavioral and neural 'response' by reserving response for the latter category.

time shortens as a function of motion strength (the absolute value of motion coherence) (Fig. 1B). On about half of the trials, a brief (100 ms) pulse of motion, equivalent to a small change in motion coherence, was presented at a random time.

The relationship between choice, reaction time, and motion coherence is well captured by a race model in which two drift-diffusion processes compete until one of them reaches a threshold or bound (Fig. 1C). The first process that reaches its upper bound determines the choice and the decision time. The reaction time is the sum of the decision time and a non-decision time, which is assumed to be Normally distributed and independent of decision time. In our instantiation of the race model, the drift-diffusion processes cannot fall below a lower *reflective* bound, which realizes the constraint that firing rates are non-negative (Zylberberg and Shadlen, 2016).

In addition to the main task, the monkeys also made visually–guided and memory-guided saccade to peripheral targets after variable delays (see Methods) (Gnadt and Andersen, 1988; Mazzoni et al., 1996; Colby et al., 1996). In the memory-guided saccade task, a target is briefly presented in the periphery while the monkey maintains its gaze on a fixation point. When the fixation point is extinguished, the monkey makes an eye movement to the remembered location of the target. The task served to identify, *post hoc*, neurons that display persistent activity as the monkey plans a saccadic eye movement towards the choice target contralateral to the recording site. These neurons are referred to as T_{in} neurons. Monkeys also performed a passive motion viewing task in which they were rewarded for maintaining fixation while viewing random dot motion (higher motion strengths only; see Methods).

The neural data from Steinemann, Stine et al. (2022) and Stine et al. (2023) were recorded from area LIP using high-density NHP-neuropixels probes. Between 54 and 203 single neurons were recorded simultaneously over eight recording sessions (mean = 135.5 neurons/session) (Fig. 1E). Steinemann, Stine et al. (2022) showed that T_{in} neurons represent the drift-diffusion signal associated with stochastic choice and reaction time on single trials. On average, T_{in} neurons tend to ramp with positive slope on trials where the monkey chooses the target in the neurons' response field, and this slope is steeper as a function of motion strength (Fig. 1F)(Roitman and Shadlen, 2002). Importantly, the population of T_{in} neurons reach a common level of activity ~ 100 ms before a saccadic eye movement towards the contralateral choice target. These observations conform to the predictions of the race model illustrated in Fig. 1C, under the assumption that the T_{in} neurons represent the accumulation of evidence for a contralateral choice, and that another (unobserved) population of neurons represents the accumulation of evidence for the ipsilateral (rightward) choice (e.g., Usher and McClelland, 2001; Wong and Wang, 2006).

Choice accuracy decoded from LIP population activity

We investigate whether T_{in} neurons, which represent the accumulation of noisy evidence, or drift-diffusion, are also predictive of whether the choice would be the correct one and thus capable of informing the monkey's confidence (or reward prediction). Our approach is to ascertain whether the population activity shortly before the choice is reported was predictive of whether the choice would be correct or incorrect. We train a logistic decoder to predict the accuracy of each contraversive (i.e., left) choice from the activity of the T_{in} neurons:

$$\text{logit}[p_{\text{correct}}] = \boldsymbol{\beta_{\text{conf}}}^{\top} \tilde{\boldsymbol{S}}_{\textbf{Tin}} + \beta_0, \tag{1}$$

where $\boldsymbol{\beta_{\text{conf}}}$ is a column vector of regression coefficients with as many elements as there are T_{in} neurons in the session, and β_0 is a bias term. $\tilde{\boldsymbol{S}}_{\textbf{Tin}}$ contains the standardized (i.e., z-scored) number of spikes emitted by each T_{in} neuron in the time interval between 150 ms and 50 ms before the choice report; we refer to this time interval as the *presaccadic window* throughout the manuscript. The model is fit separately for each session and choice category, left or right. The logistic decoder outputs a probability that quantifies the confidence the decision-maker should have in the choice. The model is fit separately for each session using 10-fold cross-validation. Specifically, we divide the data into 10 groups of approximately equal number of trials, using one group as a prediction set, and the remaining 9 for training; we repeat this process 10 times so that confidence estimates for every trial are based on a prediction.

We use a receiver operating characteristic (ROC) analysis to assess how effectively the probability correct, predicted with Eq. 1, distinguishes between correct and incorrect choices. The approach is illustrated in Fig. 2A. The figure shows the predicted probability correct for trials with a contraversive (i.e., left) choice.
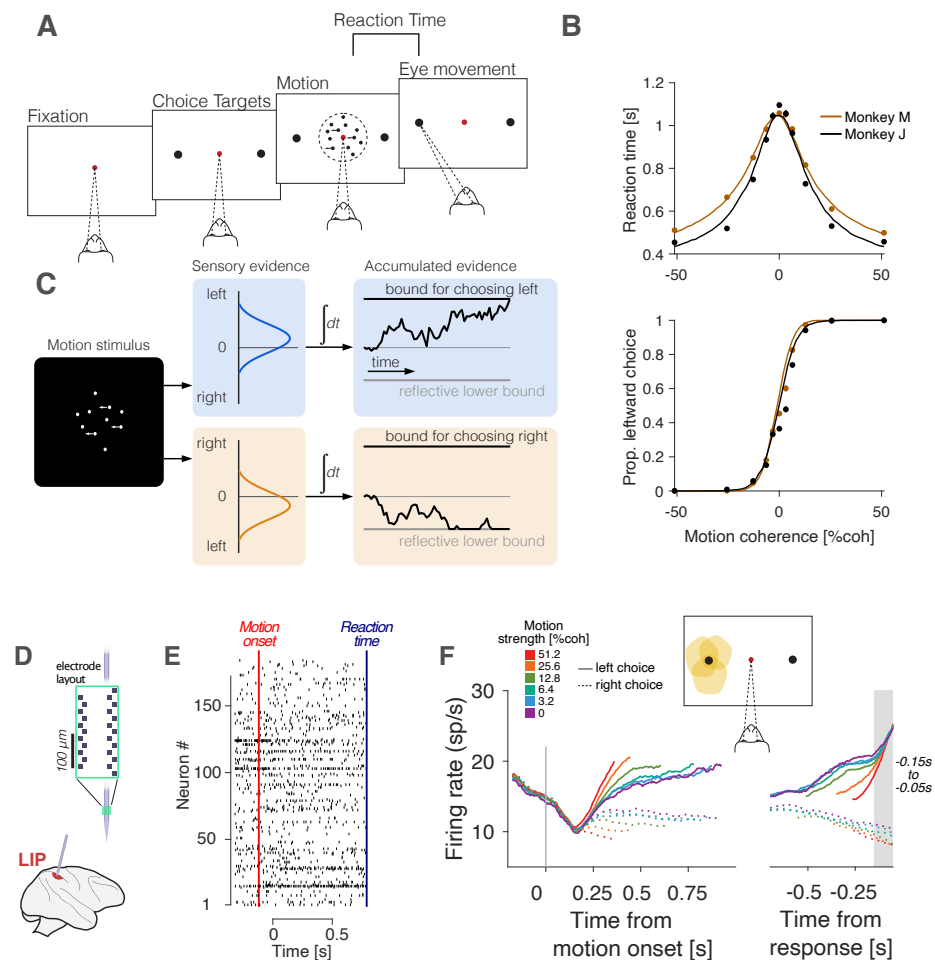
Figure 1. Large-scale recordings from LIP in a decision-making task
(A) Sequence of events in the random dot motion task. After the monkey fixates on a central spot, two choice targets are displayed, followed, after a random delay, by the random dot motion stimulus. The monkey is free to report its decision when ready by making a saccadic eye movement to one of the choice targets. The monkey is rewarded for choosing the left or right target for leftward or rightward motion, respectively. On noise-only (0% coherence motion) trials a reward is given with probability $\frac{1}{2}$. **(B)** Psychometric functions for the two monkeys studied by Steinemann, Stine et al. (2022). The average reaction time (top), and proportion of leftward choices (bottom) are plotted as a function of motion coherence. Positive and negative coherence indicate leftward and rightward motion, respectively. The solid lines are fits of the race model depicted in the next panel. **(C)** Sketch of the race model. The random dot motion stimulus provides sequential samples of momentary evidence for right minus left and left minus right, which are accumulated as a function of time to render two drift-diffusion processes. The samples are idealized as draws from Normal distributions with means proportional to motion coherence. The samples are anticorrelated ($\rho = -0.7$). The first process that reaches its positive bound terminates the decision and resolves the choice and decision time. A lower non-absorbing bound constrains the values of the negative accumulation. The reaction time is the sum of the decision time and a normally distributed non-decision time. **(D)** Schematic representation of the Neuropixels probe used to record neural activity from LIP in the right hemisphere (both monkeys). **(E)** Raster plot showing the spiking activity of 191 neurons simultaneously recorded during a representative trial. The red and blue vertical lines indicate the onset of motion and the choice report, respectively. **(F)** Average response of T_{in} neurons aligned with motion onset (left) and saccade initiation (right). Motion strength is indicated by color (legend). Solid lines indicate trials with leftward motion and dashed lines indicate trials with rightward motion. Only correct trials are included. The gray shading indicates the time between 150 ms and 50 ms before saccade initiation; most of our analyses focus on this time period.
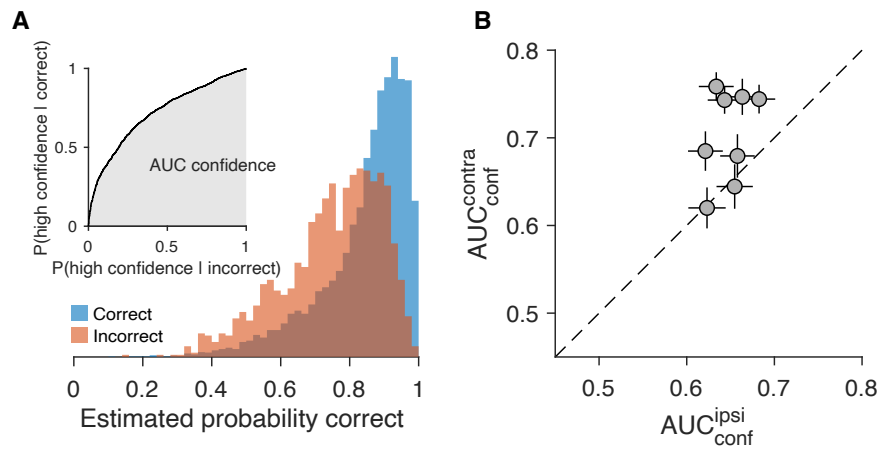
Figure 2. Confidence inferred from the population of T_{in} neurons
(A) Confidence estimates obtained with Eq. 1, shown separately for factually correct (blue) and incorrect (red) choices. These distributions are used to construct an ROC (inset) and to calculate the area under the ROC (AUC_{conf}; gray). Confidence estimates from the eight sessions are pooled. Only trials with contralateral (left) choices are included. **(B)** Area under the confidence ROC calculated from trials in which the ipsilateral (abscissa) or contralateral (ordinate) target was chosen. Each data point corresponds to a different session. Error bars indicate s.e. (bootstrap).

Correct and incorrect decisions are indicated in blue and red, respectively. The area under the ROC curve (AUC_{conf}) is a measure of how well the predicted probability correct for each trial discriminates correct from incorrect choices (Fig. 2A, inset). More specifically, it is the probability that given two choices, one correct and one incorrect, the predicted probability correct is greater for the correct choice. Using this simple metric we can test the prediction of the balance of evidence hypothesis mentioned above.

We compared the AUC_{conf} derived from the logistic model fit separately for contralateral and ipsilateral choices. As mentioned, T_{in} neurons represent the winning race for contralateral choices and the losing race for ipsilateral choices. Therefore, according to the balance of evidence hypothesis, the T_{in} neurons should contain more information about choice accuracy when monkeys select the ipsilateral (right) target, than when they select the contralateral (left) target. Fig. 2B shows the AUC_{conf} for each of the 8 recording sessions. Contrary to the balance of evidence hypothesis, T_{in} neurons contain more information about choice accuracy when the monkey chooses the contralateral target (i.e., when the T_{in} neurons represent the winning race) than when it chooses the ipsilateral target ($p = 0.008$, one-tailed *t-test*).

As mentioned earlier, this finding is surprising because the T_{in} neurons appear to reach a stereotyped level of activity before a contralateral choice, independent of motion strength and reaction time (Fig. 1F) (Roitman and Shadlen, 2002). Since motion strength and reaction time are strong predictors of accuracy, one would not expect T_{in} neurons to contain information about choice accuracy in the presaccadic window. We will address this tension after bolstering the claim that the choice accuracy that we infer from neural activity replicates the behavioral hallmarks of confidence identified in human psychophysical experiments.

Choice accuracy inferred from neural activity reproduces behavioral features of confidence

The monkeys did not report their confidence, so we cannot establish a correlation between the putative confidence signal and behavior. Instead, we ask whether a monkey exploiting this signal would mimic the regularities of confidence reports observed in humans performing a task similar to the one performed by the monkeys. van Den Berg et al. (2016) asked human participants to perform a variant of the random dot motion task in which they reported choice and confidence (high/low) simultaneously by moving a handle to one of four targets (Fig. 3A). Data from a representative participant is shown in Fig. 3A. The data show that: (*i*) confidence is greater for correct than for incorrect decisions, even when controlling for motion strength (Fig. 3A, left), (*ii*) for correct decisions, confidence increases as a function of motion strength (Fig. 3A, left), (*iii*) for incorrect decisions, confidence also increases as a function of motion strength (Fig. 3A, left), (*iv*) confidence decreases as a function of reaction time (for both correct and incorrect

decisions) (Fig. 3A, center), (v) for a given reaction time, confidence is lower for incorrect decisions than for correct decisions (Fig. 3A, center), (vi) for a given reaction time, confidence increases as a function of motion strength, even when controlling for accuracy (Fig. 3A, right).

The putative confidence signal reproduces these observations. To parallel the design of van Den Berg et al. (2016), we thresholded the putative confidence signal (obtained with Eq. 1) using a criterion set such that the proportion of high-confidence choices is equal to the proportion of high-confidence reports in the human experiment (61% high-confidence choices). Without any free parameters, the confidence signal qualitatively reproduces all the behavioral hallmarks of confidence observed by van Den Berg et al. (2016) (Fig. 3B). The similarity is not a consequence of the thresholding step in the analysis. The same qualitatve reproduction of the human regularities are also present without thresholding (Fig. S2).

In short, the population activity of T_{in} neurons measured just before a contralateral choice report contains information bearing on whether the choice is correct or incorrect. The confidence signal varies with motion strength, reaction time, and accuracy, in a similar manner to explicit confidence reports. Importantly, it is not necessary to consider the state of the losing race or the decision time to reproduce the behavioral features of confidence; the necessary information is contained in the population response of the T_{in} neurons.

Because of the tight link between confidence and reaction time (e.g., Henmon, 1911), we reasoned that the population activity of T_{in} neurons just before the response should contain information about RT. We fit a regression model to the reaction times for trials with a contralateral choice, again using the spike counts of the T_{in} neurons in the presaccadic window (Eq. 8). If all T_{in} neurons reach a stereotyped level of activity before the response, then there would be no information about reaction time just before the choice. Contrary to this prediction, we were able to distinguish longer from shorter reaction time (relative to the median) reliably from the population activity of T_{in} neurons (AUC= 0.85 ± 0.02; mean ± s.e. across sessions). Thus, shortly before the choice report T_{in} neurons contain information about the time required to make the decision. This information likely contributes to the ability of the accuracy decoder to reproduce the behavioral features of confidence that are thought to require an explicit representation of decision time.
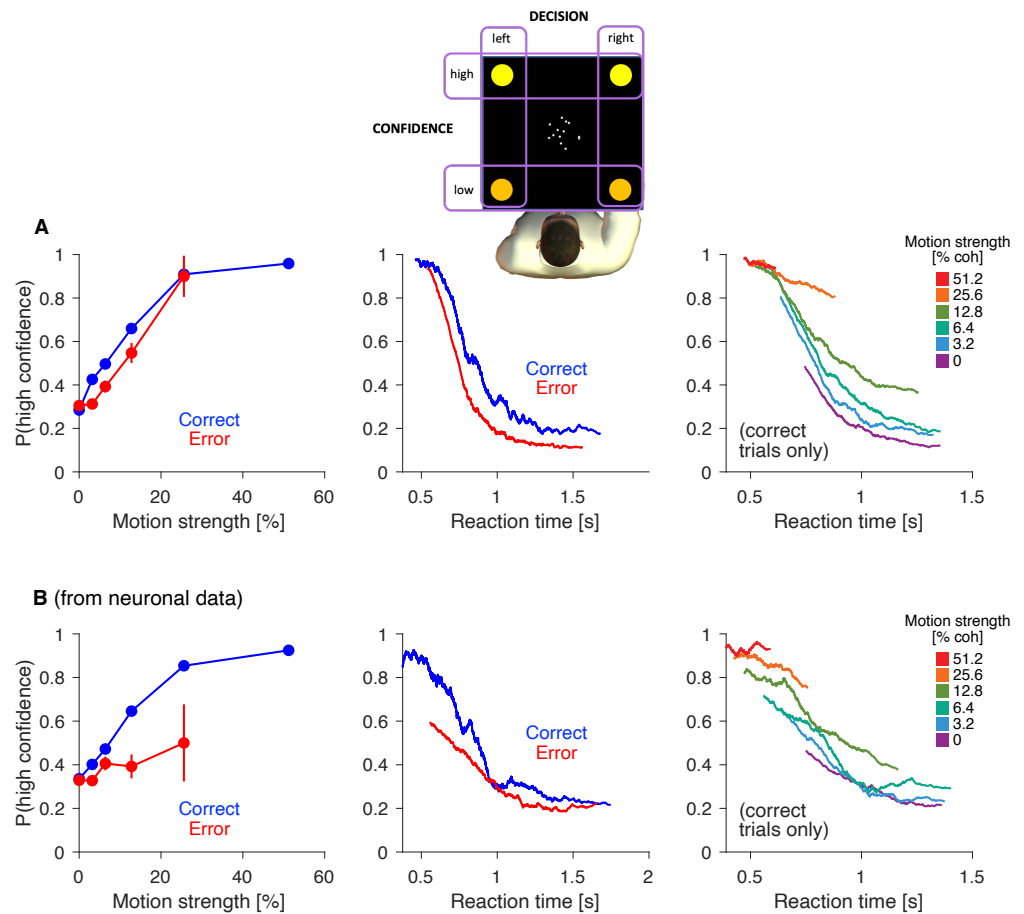
Heterogeneity of T_{in} responses underpins the representation of choice accuracy
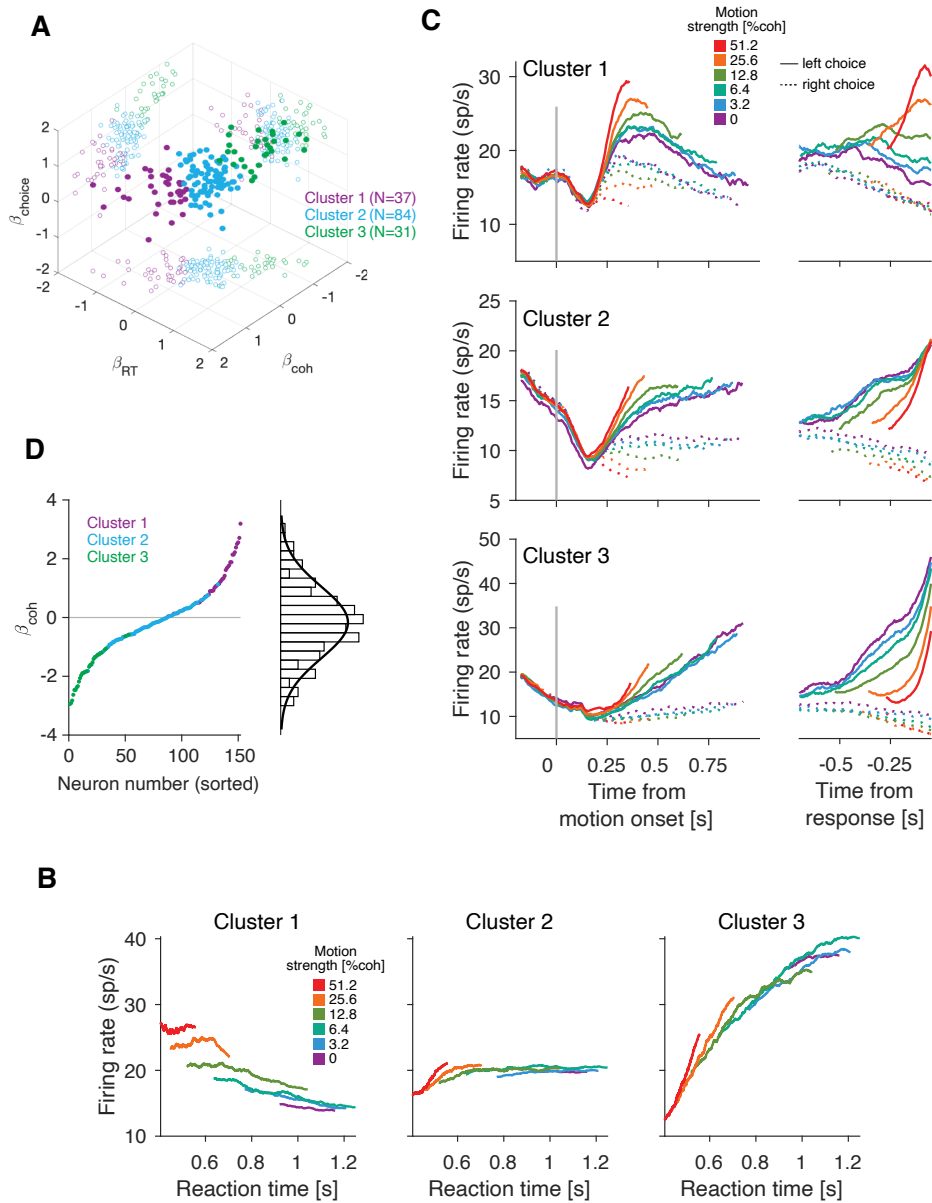
That T_{in} neurons contain information about the accuracy of a contralateral choice seems incompatible with the observation that these neurons achieve a stereotyped state at the end of the decision. However, this characterization is adduced from firing rates averaged over many trials and many neurons (e.g., Roitman and Shadlen, 2002; Steinemann, Stine et al., 2022). We thus considered the possibility that individual T_{in} neurons might contain information about choice accuracy that is not apparent in the firing rate averages.

To test this idea, we used a combination of linear regression and *k-means* clustering. The linear regression sets out to explain the spike counts in the presaccadic window for each T_{in} neuron on each trial, using three variables: the motion coherence, choice, and reaction time, plus an offset (Eq. 9). We fit the model independently for each neuron and applied *k-means* to assign the neurons to three clusters. This classification was based on the regression coefficients associated with the three variables (Fig. 4A).

The T_{in} neurons cluster into groups that exhibit distinct response characteristics. Figure 4B shows the firing rates of neurons within each cluster, calculated within the presaccadic window, as a function of reaction time and split by motion strength. The analysis includes only correct contralateral choices. Just before the choice is reported, neurons in cluster 1 ($T_{in}^{k=1}$) exhibit higher firing rates for strong motion and faster choices. The traces corresponding to different motion strengths do not converge when conditioned on reaction time, indicating that the activity of these neurons at the moment of choice is informative about both reaction time and motion strength. In contrast, neurons in cluster 2 ($T_{in}^{k=2}$) appear to be largely unaffected by reaction time or motion strength, consistent with expectations for neurons that reach a stereotyped level of activity prior to response. For the neurons in cluster 3 ($T_{in}^{k=3}$) the activity increases strongly with reaction time and is only slightly influenced by motion strength.

We repeated the clustering analysis using only the odd-numbered trials or only the even-numbered trials

Figure 3. Choice accuracy inferred from neural activity reproduces behavioral signatures of confidence
(A) Random dot motion task with simultaneous choice and confidence reports (from van Den Berg et al., 2016). The two left and right targets are used to indicate leftward and rightward motion. In alternating blocks, either the top two targets or the bottom two targets were used to indicate high-confidence choices, and the remaining two targets were used to indicate low-confidence choices. Data correspond to a representative participant from van Den Berg et al. (2016) (N = 9,024 trials; proportion of high-confidence choices: 61%). Data from the other participants in van Den Berg et al. (2016) are shown in Fig. S1. *left*. Proportion of high confidence choices as a function of motion strength, shown separately for correct and incorrect choices. Conditions with fewer than 4 trials were excluded. *center*. Proportion of high confidence choices as a function of reaction time, shown separately for correct and incorrect choices. Trials were sorted by reaction time and smoothed with a boxcar filter (N = 300 trials). *right*. Proportion of high confidence choices as a function of reaction time, for correct trials only, plotted separately for each motion strength. Trials were sorted by reaction time and smoothed with a boxcar filter (N = 300 trials). **(B)** Same analyses as in panel A, but for the Steinemann, Stine et al. (2022) data, using as confidence the probability of correct inferred from the population of T_{in} neurons using logistic regression (Eq. 1). The continuous confidence estimate was thresholded so that the proportion of trials with high confidence was the same as in the behavioral data (panel A). A version with non-thresholded confidence estimates is included as Fig. S2.

Figure 4. Distinct response characteristics of T_{in} neurons

The T_{in} neurons vary in their representation of motion coherence, reaction time and choice. We distinguished three clusters, based on a linear regression model of these explanatory variables on each T_{in} neuron's response. **(A)** Regression coefficients associated with choice, motion coherence, and reaction time for each neuron. Each filled symbol shows the 3-tuple of one neuron. Color shows the cluster membership assigned by *K-means* procedure instructed to identify three clusters. Open circles are the 2D projections of the regression coefficients. The number of T_{in} neurons in each cluster is shown in parenthesis. **(B)** The average firing rate of the neurons within each cluster, calculated within the presaccadic window, is plotted as a function of reaction time. Traces are calculated separately for each motion strength, including only correct contralateral choices. Traces are smoothed with a boxcar filter (N = 500 trials). **(C)** Average response of T_{in} neurons for each cluster. Same conventions as in Figure 1F. **(D)** Regression coefficients associated with motion coherence, obtained from a linear regression using motion coherence (plus an offset) to explain the spike counts (z-scores) emitted by each neuron in the presaccadic window. The colors indicate the cluster to which each neuron belongs. The histogram of regression weights (right) is well described by a Normal distribution (thick black trace) with a mean and standard deviation of -0.17 and 1.2, respectively.

from each session. The regression weights assigned to motion coherence, RT and choice were highly consistent across both analyses (Fig. S3). Applying *k-means* to the coefficients derived from odd and even trials, we observed that most neurons (89%) were assigned to the same cluster in both realizations, indicating that the neuron clustering is robust.

Fig. 4C compares and contrasts the response properties of the three clusters. The figure shows the firing rate of neurons within each cluster, aligned to motion onset and reaction time. Neurons from all three clusters are choice selective, as can be seen by comparing the dashed and solid lines. Neither $T_{in}^{k=1}$ nor $T_{in}^{k=3}$ neurons reach a common level of activity prior to the response. $T_{in}^{k=1}$ neurons do not show the ramping activity usually associated with T_{in} neurons. Instead, the traces for the different motion coherences are largely parallel to each other, similar to the representation of momentary evidence in upstream visual areas (e.g., Britten et al., 1992). For contralateral choices, $T_{in}^{k=3}$ neurons appear to increase their firing rate more rapidly than $T_{in}^{k=2}$ or $T_{in}^{k=1}$ neurons. The firing rate of $T_{in}^{k=3}$ is higher just before choice in difficult (i.e., low motion strength) decisions compared to easier ones.

The latency to choice selectivity was similar across clusters. We calculated this latency independently for each neuron using the CUSUM method (Ellaway, 1978; Lorteije, Zylberberg et al., 2015). The average latency to direction selectivity was 0.22±0.03, 0.22±0.02 and 0.23±0.03 s from motion onset for neurons in clusters 1, 2, & 3, respectively. No significant differences in latency were observed between clusters (Fig. S4) ($p>0.3$ for all three pairwise comparisons, *t-test*).

The heterogeneity of neuronal responses across the population of T_{in} neurons can also be observed by analyzing individual neurons, without clustering. We used a linear regression model to characterize the relationship between motion coherence and the neuronal activity during the presaccadic window. In the regression model, motion coherence for each trial (plus an offset) was used to explain the standardized (z-scored) spike counts in the presaccadic window. Only trials with a contralateral choice were included. The regression coefficient associated with motion coherence varies substantially between neurons (Fig. 4D). For some neurons, activity increases with motion coherence ($\beta_{coh}>0$), while for others, it decreases ($\beta_{coh}<0$). The values of β_{coh} are approximately normally distributed (Fig. 4D). Because the mean is close to zero, the activity of T_{in} neurons just before the response appears to be unaffected by motion coherence when averaged over many neurons (Fig. 1F; Roitman and Shadlen 2002).

The heterogeneity of neuronal responses is not evident in control tasks

The memory-guided saccade task was used to identify the T_{in} neurons and historically to elucidate the their hallmark visual, memory and perisaccadic responses (Gnadt and Andersen, 1988; Mazzoni et al., 1996; Colby et al., 1996). We wondered whether the signs of the heterogeneity we identified in the random dot motion task would be evident in the memory-guided saccade task. Fig. S5A shows the response of neurons from the three clusters for memory saccades to the contralateral and ipsilateral targets. The three clusters show similar persistent activity during the delay period. We calculated the spike counts in the last 200 ms before the fixation point is extinguished, and computed the average difference in standardized counts between trials with memory-saccades to the left and right target. This measure fails to differentiate the three clusters ($p_{max}>0.2$, Wilcoxon rank–sum test).

The monkeys also performed a passive motion viewing task in which they were presented with strong ($c = \pm 51.2\%$) left and right motion and were rewarded for maintaining fixation. We expected neurons in cluster 1 that retain information about motion strength to have response fields that overlap the random dot motion stimulus, but there is no sign of this. Neurons in all three clusters maintained low activity during the passive viewing task (Fig. S5B). While firing rates were significantly greater for leftward than for rightward motion (p = 0.047, 0.0026, and 0.0145 for clusters 1, 2, and 3, respectively; one-sided Wilcoxon signed-rank test), decoding accuracy was poor and similar so for all clusters. We calculated the spike rate during motion viewing (excluding the initial 200 ms), and computed the average difference in standardized rates between trials with leftward and rightward motion. This measure does not distinguish between the three clusters ($p_{max}>0.25$, Wilcoxon rank–sum test). The capacity of this measure to distinguish between leftward and rightward motion was low (AUC =0.55 ± 0.01, mean ± s.e. across neurons). As an additional control, we repeated the decoding and clustering analyses after excluding the T_{in} neurons that significantly discriminated between leftward and rightward motion in the passive motion–viewing task. Excluding these neurons (N = 45) yielded qualitatively similar results Fig. S6. We conclude that

the response features that distinguish the three clusters of T_{in} neurons are not elucidated by memory saccades or passive motion viewing. Instead, the discriminating feature is more likely associated with the decision mechanism.

Clusters 1 and 3 are the most informative about the accuracy of the choice

Because confidence is informed by evidence strength and decision time, and because these variables are represented by the population of cluster-1 and cluster-3 neurons, we reasoned that neurons from these clusters should be the most informative about choice accuracy. We repeated the logistic regression analysis used to decode accuracy from population activity (Eq. 1), using signals from one cluster at a time. Cluster-1 neurons and cluster-3 neurons are more predictive of choice accuracy than cluster-2 neurons (Fig. 5A) ($p_{max} < 10^{-8}$, bootstrap).

The average spike counts across the population of cluster-1 neurons, combined with the average across the population of cluster-3 neurons, predict choice accuracy with fidelity. We trained the accuracy decoder (Eq. 1) using two independent variables: the average of the spike counts from cluster–1 neurons and cluster–3 neurons within the presaccadic window (plus an offset). The AUC_{conf} from this regression model is statistically indistinguishable from the model using all T_{in} neurons (p = 0.16, bootstrap; Fig. 5B). We also trained the accuracy decoder using the within-trial average of the spike counts across all T_{in} neurons. This model yields a much lower AUC_{conf} than that obtained from the mean activity of neurons from clusters 1 and 3 (p < 10^{-8}, bootstrap. We conclude that the activity of cluster–1 and cluster–3 neurons (i.e., two numbers per trial) accounts for most of the information about choice accuracy contained in all recorded T_{in} neurons.

We chose to use three clusters for analytical convenience. The number is not guided by a biological or computational principle. Indeed, the regression coefficients used for clustering exhibit continuous variation (Fig. 4A). Nonetheless, three seems to be the right number. We repeated the analysis using between two and six clusters. For each number, we calculated the AUC_{conf}. We averaged the number of spikes in the presaccadic window across neurons belonging to the same cluster and used these averaged spike counts to predict choice accuracy (Eq. 6). We found a significant difference in AUC_{conf} only between two clusters and more than two clusters, but not between three and four clusters or three and six clusters (Fig. S7). Three clusters may be adequate for our purposes as they are the minimum number required to capture the central tendency and both signs of diversity.

Separable contributions of T_{in} neurons to the decoding of choice and accuracy

We assessed whether the decoding of choice accuracy and the decoding of the choice itself relies on a common weighting of the activity of the T_{in} neurons. To this end, we fit a regression model similar to the one we used to predict decision accuracy (Eq. 1), but here the variable to predict is choice (left/right) (Eq. 7). We used the spike counts of the T_{in} neurons in the presaccadic window to derive the best-fitting regression weights, β_{choice}. The regression weights define a coding direction (CD) in the state space, where each T_{in} neuron represents a different dimension. Unsurprisingly, choice can be decoded with high precision from the population of T_{in} neurons ($AUC_{choice} = 0.96 \pm 0.014$; mean ± s.e. across sessions). More interestingly, the cosine similarity between the directions defined by the regression weights on choice, β_{choice}, and the regression weights on accuracy, β_{conf}, is low: the average (across sessions) absolute value of the cosine similarity is 0.25 ± 0.05, indicating that the two directions in state space are closer to orthogonal than similar (Fig. 5C). Indeed, the projection of the population activity on the choice-CD is barely informative about accuracy (AUC_{conf} = 0.57 ± 0.024; mean ± s.e. across sessions) and significantly less informative than the projection on the confidence CD (p = 0.0009, *t-test*; Fig. 5D).

We reasoned that the separable contribution of T_{in} neurons to decoding choice and accuracy may be evident at the level of clusters. We repeated the logistic regression analysis, again using signals from one cluster at a time (as in Fig. 5A) but now training the decoder to predict choice. As shown in Fig. 5E, the neurons from cluster 2 are more predictive of the choice than neurons from clusters 1 and 3 ($p_{max} < 10^{-8}$, bootstrap).
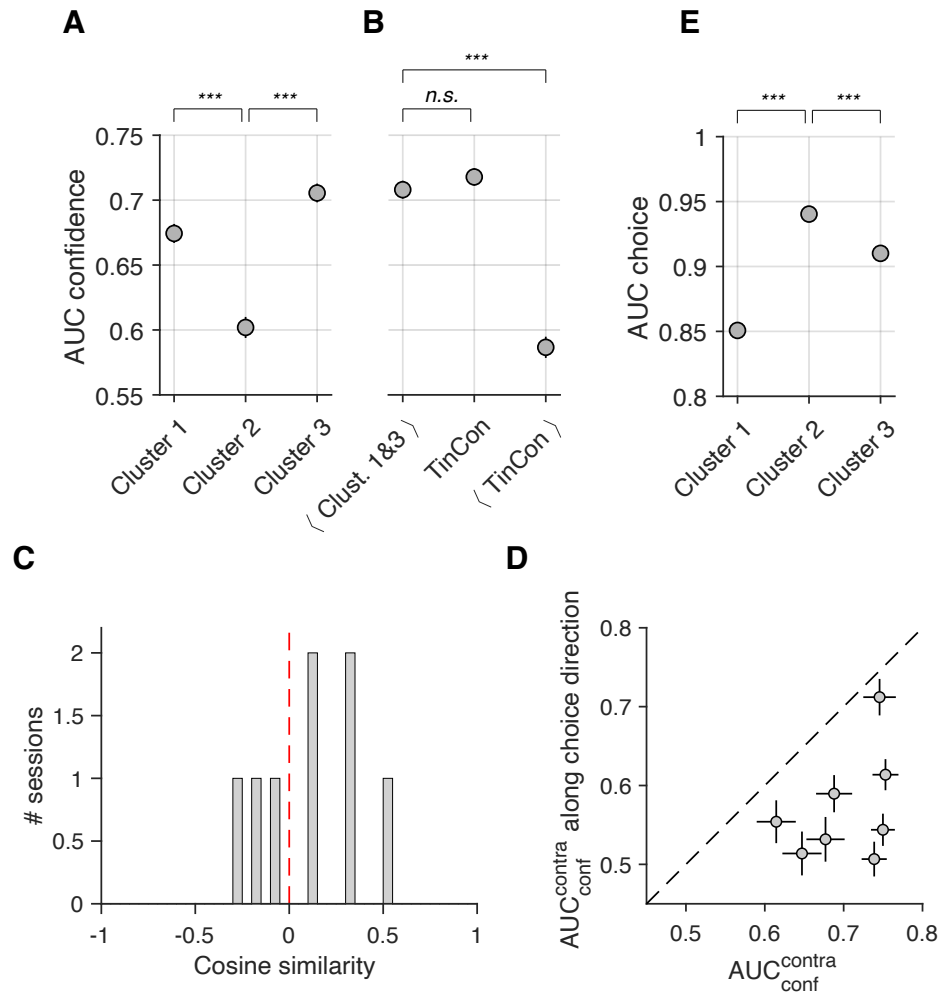
Figure 5. Information about choice accuracy differs across clusters

(A) Neurons from cluster 3 are the most informative about choice accuracy. Symbols show the area under the receiver-operator curve (AUC_{conf}) which quantifies the separation between the distribution of spike counts in the presaccadic window on correct vs. incorrect left choices. The accuracy predictions used to calculate the AUC_{conf} are based on the activity of the T_{in} neurons assigned to each of the three clusters (abscissa). Error bars indicate s.e. (bootstrap). Bootstrapping was used to assess statistical significance (n.s.: $p > 0.05$; ∗∗∗: $p < 10^{-8}$). **(B)** AUC_{conf} values obtained from the accuracy decoders trained on either (*i*) the average activity of neurons from cluster 1 and cluster 3, (*ii*) the individual T_{in} neurons, or (*iii*) the average (across neurons) of the activity of the T_{in} neurons. Error bars indicate s.e. (bootstrap). **(C)** Cosine similarities between the weights assigned by the regression to choice, β_{choice}, and the regression to accuracy, β_{conf}, across sessions. The set of weights define coding directions in the neuronal state space. **(D)** We project the neural activity onto the coding directions defined by β_{choice} and by β_{conf}, and compute the information about choice accuracy, measured by the AUC_{conf}, contained in these two projections. Information about choice accuracy is larger for the projection onto the direction defined by β_{conf} (abscissa) than onto the one defined by β_{choice} (ordinate). Each data point corresponds to a different session. Error bars indicate s.e. (bootstrap). **(E)** AUC_{choice} derived from three separate regression analyses, each including neurons from a single cluster. Error bars indicate s.e. (bootstrap). Neurons from cluster 2 are the most informative about choice.

Multiple time-scales of evidence accumulation represented by T_{in} neurons

We considered the possibility that the different response characteristics of T_{in} neurons might be explained by differences in how long (or persistently) the momentary motion evidence affects neuronal activity. We evaluate this idea by analyzing the influence of the short (100 ms) motion pulses introduced into the random dot motion stimulus. For each T_{in} neuron i and trial j, we count the spikes emitted between t and $(t+100)$ ms, where t is the time from the onset of the motion pulse, and standardize the counts separately for each motion coherence and 100 ms time window. We refer to the standardized counts as $\tilde{R}_i(t, j)$. We then calculate the average difference, $\Delta \tilde{R}_i(t)$, between trials with a leftward pulse and trials with a rightward pulse.

The difference $\Delta \tilde{R}_i(t)$ is a measure of the influence of the motion pulse on neuronal activity. We averaged this difference across neurons belonging to the same cluster and fit a function (Eq. 13) to these averages (Fig. 6A). The function implements two assumptions: (*i*) the pulse affects the neural response with a variable latency, and (*ii*) its effect dissipates exponentially with time constant α^{-1} (Lorteije, Zylberberg et al., 2015). The best-fitting α values are 100, 3.1, & -0.17 for neurons of clusters 1, 2, and 3 respectively, consistent with a more persistent effect on cluster–2 neurons than on cluster–1 neurons, and a more persistent effect on cluster–3 neurons than on cluster–2 neurons (Fig. 6A). The negative α (cluster 3) reflects the monotonic increase in the impact of the pulse as a function of time (Fig. 6A, *bottom*). We used a bootstrap analysis to create a distribution of α values for neurons of the three clusters. The probability of observing in the bootstrap distributions an α value greater in cluster 1 than in cluster 2 was 86.8 %, and a value greater in cluster 3 than in cluster 2 was 96.6 %. While the probability of observing this rank ordering by chance is not negligible, the analysis is consistent with the hypothesis that the time constant of evidence accumulation is different for neurons from the three clusters.

We further substantiate this interpretation with an analysis of the pairwise correlations between neurons belonging to different clusters. We counted the spikes of neurons from cluster k, $S_{Tin}^k(t)$, in 25ms windows. We formed pairs $\{x, y\}$, where $x = \tilde{S}_{Tin}^k(t_x)$ and $y = \tilde{S}_{Tin}^j(t_y)$, with k and j representing different clusters. The tilde in these expressions indicates the use of standardized residual values, for each motion coherence. The heat maps in Fig. 6B-D illustrate the correlation of these residuals across trials, for different pairs of clusters. Fluctuations in the activity of cluster–1 neurons predict—at later times— fluctuations in the activity of cluster–2 neurons, and the activity of cluster–3 neurons (both $p < 10^{-8}$, permutation test; Fig. 6B & C). Similarly, a noise correlation analysis between the activity residuals of cluster–2 and cluster–3 neurons reveals that the fluctuations in both clusters are largely positively correlated, and that the fluctuations in activity of cluster–2 neurons at time t predict the fluctuations in activity of cluster–3 neurons for times $t' > t$ ($p < 10^{-8}$, permutation test; Fig. 6D). The observations suggest that the time constant of integration increases with ascending cluster number (or that the degree of integration *leak* decreases with ascending cluster number).

Representation of momentary motion evidence by T_{in} neurons

The transient effect of the motion pulses on the activity of cluster-1 neurons made us wonder whether these neurons resemble other neurons in LIP that mimic responses of direction selective neurons in area MT (Freedman and Assad, 2006). The Steinemann, Stine et al. (2022) dataset also contains such neurons, termed M_{in}—for motion in response field—neurons. Fig. 7A shows firing rate averages from leftward and rightward preferring M_{in} neurons in the random dot motion task. For the M_{in} neurons, the firing rate traces associated with different motion coherences are predominantly parallel. They do not resemble the ramp-like dynamics associated with evidence accumulation. Some M_{in} neurons show selectivity for contraversive motion (M_{in}^{left}) and others for ipsiversive motion (M_{in}^{right}) (Fig. 7A–C, top and bottom panels respectively).

In contrast to T_{in} neurons, M_{in} neurons do not show persistent activity in the memory-guided saccade task (Fig. 7B) but show direction selectivity in the passive motion viewing task (Fig. 7C). However, during the random dot motion task, M_{in} and $T_{in}^{k=1}$ neurons display similar responses in that they are direction selective but represent neither evidence accumulation nor decision termination. Note the similarity of the traces between the M_{in} neurons (Fig. 7A) and the $T_{in}^{k=1}$ neurons (Fig. 4D, *top*).

We wondered if the $T_{in}^{k=1}$ neurons would share signals with the M_{in} neurons, despite the different locations of their response fields. We tested this idea using a noise correlation analysis similar to that shown in
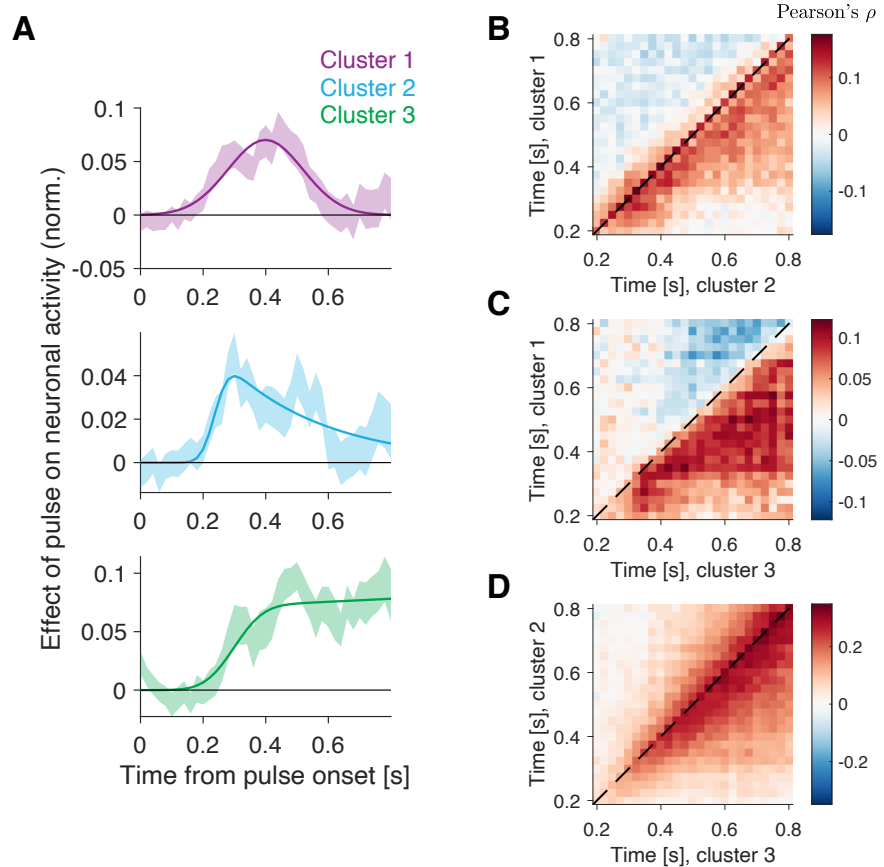
Figure 6. Different time constants of evidence integration by T_{in} neurons
(A) Influence of a brief motion pulse on the activity of cluster–1 (top), cluster–2 (middle) and cluster–3 (bottom) neurons (Eqs. 10 and 11). Error bars indicate s.e. (bootstrap). Solid lines are fits of a function (see Methods). The pulses affect the neural response with latency ≈ 200 ms. **(B–D)** Noise correlations between neurons from clusters 1 and 2 (panel B), clusters 1 and 3 (panel C), and clusters 2 and 3 (panel D). Time is relative to pulse onset. The correlation coefficients were calculated in non-overlapping 25 ms windows, using only trials with contralateral choices.
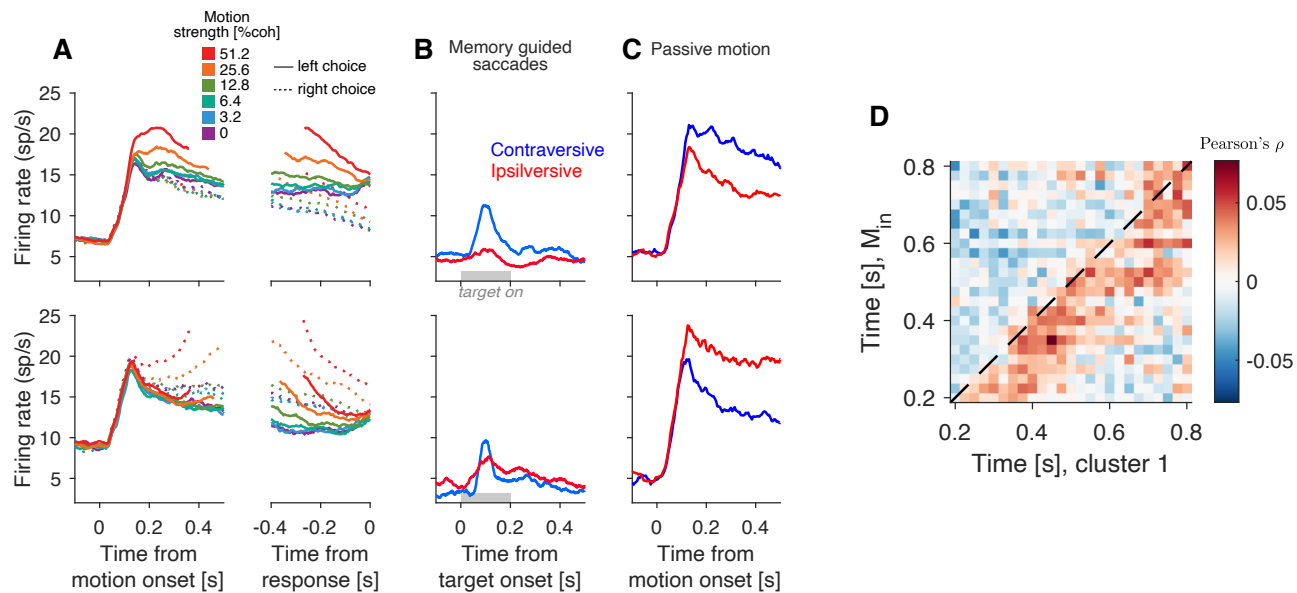
Figure 7. Momentary motion evidence in LIP

(A–C) Firing rate of the M_{in} neurons in the random dot motion task (**A**), the memory-guided saccade task (**B**), the passive motion viewing task (**C**). The upper (lower) row represents M_{in} neurons that prefer leftward (rightward, respectively) motion. **(D)** Noise correlations between the motion-selective neurons with the motion stimulus on their response field (ordinate), and the $T_{in}^{k=1}$ neurons (abscissa). Same conventions as in Fig. 6B–D.

Fig. 6B-D, but where the correlations are calculated between $T_{in}^{k=1}$ and M_{in} neurons. That is, $x = \tilde{S}_{Tin}^{k=1}(t_x)$ and $y = \tilde{S}_{Min}^{left}(t_y) - \tilde{S}_{Min}^{right}(t_y)$. As before, the tilde in these expressions indicates the use of standardized residual values, for each motion coherence. The heat map (Fig. 7D) illustrates the correlation of these residuals across trials. Correlations are stronger for off-diagonal elements where the activity of the cluster-1 neurons lags the activity of the M_{in} neurons by approximately 100 ms. Correlations are significantly higher for $t_x > t_y$ than for $t_y > t_x$ ($p < 10^{-8}$; permutation test). That is, fluctuations in the activity of M_{in} neurons predict changes in $T_{in}^{k=1}$ neurons at later times. This result is consistent with the idea that the signals that drive the M_{in} neurons are reflected later in the activity of the $T_{in}^{k=1}$ neurons.

Information about choice and accuracy evolves over time

So far we have mainly focused our analyses on a time window just before the choice report. Here we look for choice and accuracy signals outside the presaccadic window to characterize the time course of information about choice and accuracy. We construct two population signals by projecting the neural activity in the coding directions defined by β_{choice} and β_{conf}. Neural activity is obtained by binning the spike counts of T_{in} neurons in sliding windows of 100 ms. At each time t, we compute the area under the ROC curve (AUC) obtained from the projections onto the choice and accuracy coding directions. The AUC values indicate how well the projections discriminate between left and right choices and between correct and incorrect choices, respectively.

Both AUC values peak near the time of reporting (Fig. 8). Unsurprisingly, the decoding of choice is more veridical than the decoding of choice accuracy (i.e., the choice predictions better distinguishes left from right choices than the accuracy predictor distinguishes correct from errors). We assessed whether accuracy information lags behind choice, using a latency analysis based on fitting a bilinear "dogleg" function (Lorteije, Zylberberg et al., 2015) to the time course of AUC values. Information about choice diverges from baseline at $0.165 \pm 0.02 s$ from motion onset; for choice the divergence occurs $0.187 \pm 0.04 s$ from motion onset (Fig. 8). This difference is not significantly different from zero (in a bootstrap analysis, choice lags confidence in 24.8% of samples). This suggests that information about choice and accuracy are practically contemporaneous.
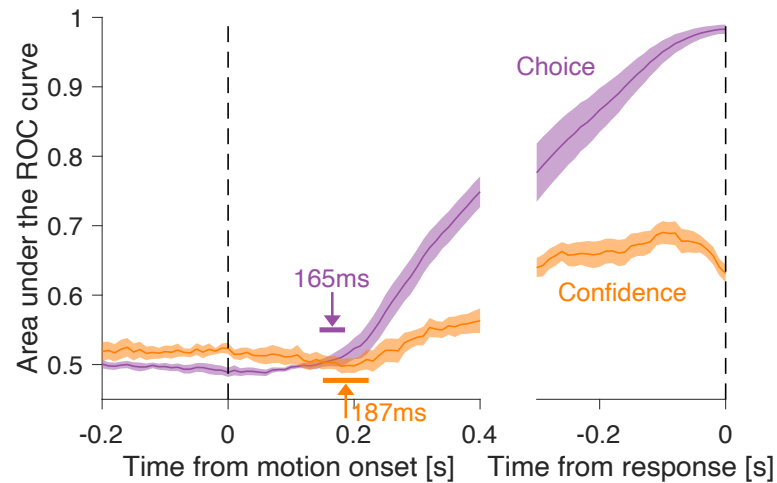
Figure 8. Contemporaneous decoding of choice and accuracy
Time-course of the AUC values obtained from the projection of the neuronal activity along the directions defined by β_{choice} (purple) and β_{conf} (orange). Shading indicates s.e. across sessions. Projections were calculated in 100 ms windows in steps of 10 ms. The arrows indicate the time when the traces first deviate from baseline (see Methods), and the associated horizontal bars are the s.e. of these estimates.

The presaccadic confidence signal accommodates an informative prior

In the random dot motion task, confidence is influenced not only by reaction time and motion strength, but also by the prior probability (base rate) of the different response alternatives (Zylberberg et al., 2018). We ask whether the neural representation of choice accuracy that we identified is also sensitive to manipulations of prior probability. We reanalyzed data from Hanks et al. (2011) in which the prior probability that the motion is rightward or leftward was varied in blocks of ∼400 trials.

We decoded choice accuracy using the same approach that we used for the Neuropixels data. We select trials with a contralateral choice and predict whether the choice is correct or incorrect using the neuronal data recorded on that session, again focusing on the presaccadic window. In the experiment of Hanks et al. (2011), only one T_{in} neuron was recorded per session, hence the predictions are less veridical. Nevertheless, the predicted choice accuracy is greater for trials in which the monkey chose the target with the greater base rate (Fig. 9). This holds for each level of motion strength and for both correct and incorrect decisions (Fig. 9 A & B, respectively), consistent with behavioral observations (Zylberberg et al., 2018). Therefore, the neural representation of choice accuracy that we identified is not only informed by motion strength and reaction time (Fig. 3), but also by the prior probability of the chosen option (Fig. 9), thus furthering the idea that the T_{in} neurons support the computation of confidence.

Discussion

We show that neurons in parietal Area LIP that represent an evolving decision variable (Roitman and Shadlen, 2002; Steinemann, Stine et al., 2022) also carry information about the probability that the decision is correct—or the probability that it will be rewarded. Information about choice accuracy is present in LIP even though the monkeys were not required to report confidence or use it to inform a subsequent decision. Consistent with human imaging studies (Lebreton et al., 2015), the present finding supports the view that the calculation of confidence is automatic: an obligate component of decisions, whether or not it is put to use.

The putative confidence signal reproduces features of confidence reports from humans in tasks similar to the one performed by the monkeys. These features include the relationship between confidence and decision difficulty (as determined by motion strength), prior probability, accuracy, and reaction times (Kiani et al., 2014; Zylberberg et al., 2018). Accounting for these features is thought to require knowledge of the decision time and/or the accumulated evidence for the unchosen alternative (Vickers, 1979; Kiani et al., 2014; van Den Berg et al., 2016; Brus et al., 2021; Smith and Vickers, 1988; Zylberberg et al., 2018, 2012; Hellmann et al., 2023; Moreno-Bote, 2010; Rolls et al., 2010; Wei and Wang, 2015). However, all the information needed to account for these behavioral features is contained in the activity of the
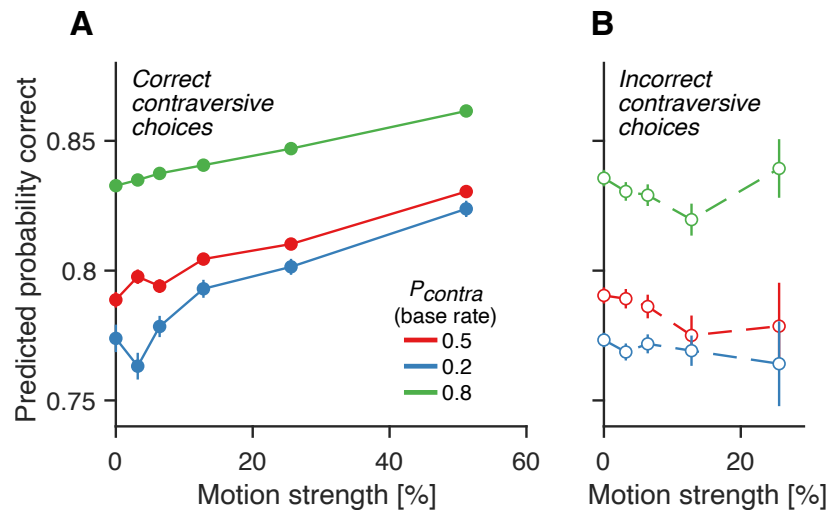
Figure 9. The confidence signal accommodates an informative prior
(A) Probability correct inferred from individual LIP neurons from the experiment of Hanks et al. (2011). Monkeys performed blocks of trials in which the prior probability that the target in the neuron's response field is the one to be rewarded (p_{contra}) was either 0.5, 0.2 or 0.8. The predicted probability is shown separately for trials of the different blocks (colors). The predicted probability correct increases with motion strength and with the strength of the prior supporting the choice. **(B)** Same as A, but for incorrect choices. In both A and B, the analysis includes only trials in which the monkey chose the target contralateral to the recording site (i.e., the target in the neurons' response field). Error bars indicate s.e.m. across trials.

population of LIP neurons with response fields that overlap the chosen target.

The finding is surprising because such LIP neurons have been shown to reach a stereotyped level of activity at the end of the decision that is independent of the strength of evidence and of reaction time—two strong predictors of accuracy. We replicate this property in our data when we examine the firing rates averaged over neurons, as in previous studies. However these averages belie heterogeneity across the population. Not all of these neurons reach a stereotyped level of activity. Importantly, the way each neuron departs from the average is consistent across multiple decisions. The heterogeneity thus reflects systematic differences between functional subtypes of T_{in} neurons, and we show that one of the functional distinctions manifests as different degrees of leaky integration (Fig. 6). Such variation might be supported by different levels of recurrence in the neural circuit (Usher and McClelland, 2001; Wong and Wang, 2006; Lange et al., 2021; Zylberberg et al., 2009). This heterogeneity endows the T_{in} neurons with signals that could be used to support a confidence judgment or reward prediction. An alternative interpretation of heterogeneity, in general, is that it confers robustness to a unidimensional signal. It would be of great interest to determine whether the projections of LIP to the SC and the orbitofrontal cortex comprise groupings that support decoding of choice and reward prediction, respectively. Alternatively, the downstream areas might receive the same signals but weight them differently to extract the choice or confidence signal. Simultaneous recordings from functionally related populations could provide an answer to this question in the near future. It bears on a broader question what heterogeneity achieves along the spectrum from a specific feature to robustness of the broader population of T_{in} neurons to variation in a property.[2]

Neurophysiological studies in non-human animals have identified neural correlates of confidence in several brain areas, including the superior colliculus (Odegaard et al., 2018), the pulvinar (Komura et al., 2013), the orbitofrontal cortex (Kepecs et al., 2008; Lak et al., 2014; Masset et al., 2020), the lateral intraparietal area (Kiani and Shadlen, 2009; Vivar Lazo, 2024), the supplementary eye fields (So and Stuphorn, 2016; Middlebrooks and Sommer, 2012), the visual cortex (Fetsch et al., 2014; Zylberberg et al., 2016; Boundy-Singer et al., 2024), and the midbrain (Lak et al., 2017). Kiani and Shadlen (2009) examined a version of the random dot motion task where the experimenter controlled the stimulus duration. Monkeys had the option to opt out of making decisions to receive a small but guaranteed reward. Kiani

---

[2]For example, heterogeneity of speed tuning in visual cortex could confer coding of speed or robustness of the direction-selective signal to variation in velocity.

& Shadlen suggested that the monkey's decision to choose or waive the sure bet depends on (*i*) the average activity of the T_{in} neurons at the end of the evidence stream, approximately 200 ms after random dot motion offset, and (*ii*) the random dot motion duration from onset to offset. They did not suggest a mechanism to combine these factors but summarized the process as a time-dependent criterion applied to the magnitude of the average. It is conceivable that the diversity of T_{in} neurons could be involved in implementing this computation.

We used logistic decoders to identify coding-directions (CDs) in state space that are most informative about choice and accuracy. The cosine similarities between the two coding directions are low (Fig. 8), indicating that the two signals are potentially distinguishable by downstream structures. The existence of distinct signals might explain why confidence and choice can be dissociated by certain lesions, inactivations, or behavioral manipulations (Komura et al., 2013; Del Cul et al., 2009; Rounis et al., 2010; Miyamoto et al., 2018; Peters et al., 2017; Maniscalco et al., 2016; Koizumi et al., 2015; Samaha and Denison, 2020; Aitchison et al., 2015; Zylberberg et al., 2012, 2014; Rahnev and Denison, 2018; Dou et al., 2024).

The presaccadic confidence signal we identified may offer insights into key features of confidence, such as positive evidence bias (PEB), that were not explored in this study. The PEB refers to the observation that confidence is more strongly influenced by the evidence supporting the chosen option than by the evidence supporting the non-chosen option (Zylberberg et al., 2012; Aitchison et al., 2015; Peters et al., 2017; Maniscalco et al., 2016; Mazor et al., 2023; Samaha and Denison, 2020; Sepulveda et al., 2020; Vivar Lazo, 2024; Mazor et al., 2023). Our results provide a tentative explanation for the PEB. Our data suggest that confidence is represented by neurons that constitute the "winning" race. Thus, the evidence that drove the winning race is expected to contribute more to confidence than the evidence that drove the losing race, leading to a PEB. This prediction has to be tested in a task in which the evidence for one alternative is not necessarily interpreted by the decision maker as evidence against the other one (e.g. Zylberberg et al., 2012). Other features of confidence that need to be explored in suitable tasks include the increase in confidence with response time when stimulus duration is controlled by the experimenter (Irwin et al., 1956; Vickers et al., 1985; Kiani and Shadlen, 2009; Zylberberg et al., 2016) and when decisions are made under time pressure (Vickers and Packer, 1982).

A key contribution of our study is the realization that confidence in a decision can be estimated at decision termination simply as a weighted average of the activity of LIP neurons whose response fields overlap with the chosen target. That said, the surprising fact that this information is present does not imply that the monkey exploits it. Additional experiments would be needed to evaluate such an assertion (Ritchie et al., 2019). In our case, this concern is somewhat mitigated by the fact that we: (*i*) used simple, linear decoders, which implies that it is easy for downstream areas to read confidence from the activity of T_{in} neurons, (*ii*) validated the prediction against many behavioral features of confidence, and (*iii*) focused on a very specific moment in the trial when these neurons are presumably communicating important information (e.g., where to move the eyes) to downstream areas. While our results need to be validated in tasks with confidence reporting, it is enticing to think that a group of neurons that share a functional characterization, contain enough information in the final 100 ms of decision formation to resolve choice, reaction time and confidence.

Methods

Behavioral tasks

Random dot motion task

In the main task, the monkeys had to decide the net direction (leftward or rightward) of a circular patch of limited-lifetime, dynamic random dots and report their choice when ready by making a saccadic eye movement from the central fixation to the left or right choice target. The monkey initiates at trial by directing the gaze to a central fixation point. After 0.25–0.7 s (truncated exponential with time constant λ=0.15 s), two red choice targets (diameter 1 dva; degrees of visual angle) are presented in the left and right visual fields. After a random delay (0.25-0.7 s, λ=0.4 s), the random dot motion stimulus is displayed until the monkey initiates a saccadic eye movement to report its choice.

The random dot motion comprises limited lifetime dots displayed within a circular area (diameter 5 dva) centered on the fixation point. The dot density is 16.7 dots $\cdot$ dva$^{-2} \cdot$ s^{-1}. The direction and strength of the motion are chosen pseudorandomly on each trial, such that the coherence, $C \in \{\pm 0\%, \pm 3.2\%, \pm 6.4\%, \pm 12.6\%, \pm 25.6\%, \pm 51.2\%\}$. The sign of C determines the direction of motion; positive values indicate leftward motion. For $C = 0\%$, the sign indicates the random direction that is rewarded on that trial. The absolute value $|C|$ establishes the motion strength: the probability that a dot displayed in video frame n is displaced by Δx in frame $n + 3$ (i.e., 40 ms later). Otherwise the dot is replaced by a new dot at a random position. The displacement, $\Delta x = \pm 0.2$ dva, is consistent with apparent motion speed of 5 dva per s (see Roitman and Shadlen, 2002, for further details). Monkeys are rewarded for making a saccadic eye movement to the correct choice target. On trials with 0% motion coherence, either saccadic choice is rewarded with a probability of 0.5. Incorrect responses are penalized by increasing the inter-trial interval by up to 3 seconds (see Stine et al., 2023, for further details). On approximately half of the trials, the motion coherence is incremented or decremented for 100 ms by 4% coherence for monkey M and 3.2% for monkey J (Stine et al., 2023). The onset time of the *pulse* is chosen randomly from a truncated exponential distribution: t_{min} =0.1 s to t_{max} =0.8 s from motion onset ($\lambda = 0.4$ s). Monkey M completed 9,684 trials across five sessions, while Monkey J completed 8,142 trials in three sessions.

Control tasks

Monkeys also completed two additional tasks in each session: a passive motion viewing task and a memory-guided saccade task. In the passive motion viewing task, the monkey views $\pm 51.2\%$ coherent motion for 0.5 s (and for 1 s on a small number of trials in one session). The task matches the main task but without choice targets. The monkey is rewarded for maintaining fixation during the motion presentation.

In the memory-guided saccade task (Hikosaka and Wurtz, 1983; Gnadt and Andersen, 1988), a target was briefly flashed (200 ms) at a pseudo-random location in the visual field. After a variable delay (0.2-0.9 seconds for monkey M, $\lambda = 0.3$ seconds; 0.3-1.3 seconds for monkey J, $\lambda = 0.2$ seconds), the fixation point was extinguished and the monkey had to make a saccadic eye movement to the remembered location of the target. The monkey was rewarded if the gaze was within ±2.5 degrees of visual angle of the target location.

Biased prior probability task

The analysis of prior probabilities makes use of previously published single neuron recordings from two other monkeys that performed the same task as Steinemann, Stine et al. (2022). However, Hanks et al. (2011) varied the prior probability that the rewarded choice was left or right. In blocks of trials, the sign of the motion coherence was biased in favor of positive or negative. They also included blocks with a neutral (i.e., uninformative) prior. In blocks with neutral priors, both targets had an equal 50% chance of being correct. In biased conditions, one direction had an 80% probability of being correct and the other had a 20% probability of being correct, except for a small number of sessions in one monkey where a 67-33% prior was used (these data were not included in our study nor in Hanks et al. 2011).

Sessions typically began with a block of 200-400 trials under a neutral prior. The monkeys then completed 300-600 trials in which the prior favored one of the targets, with the most-likely target being the one chosen least often during the neutral prior block. To signal to the monkeys which target was more likely, each biased block was preceded by 20 trials of 100% coherent motion toward the more likely target. These trials were not included in our analysis. In some sessions, monkeys completed an additional block

with a prior favoring the opposite target. See Hanks et al. (2011) for details.

Combined choice-confidence task in humans

van Den Berg et al. (2016) asked participants to discriminate the direction of motion of a random dot motion display similar to that of Steinemann, Stine et al. (2022) and Hanks et al. (2011). Subjects held the handle of a vBOT manipulandum used to record the position of the handle at 1,000 Hz (Howard et al., 2009). A horizontal mirror projecting a downward facing CRT monitor prevented subjects from seeing their arm. A chin rest ensured that the viewing distance was approximately 40 cm.

Participants reported choice and confidence simultaneously by moving the handle to one of four circular targets displayed at the corners of a 17 cm x 17 cm square. The two targets on the left corresponded to a leftward motion choice, and the two on the right corresponded to a rightward motion choice. In half of the blocks, the two top targets corresponded to a high-confidence choice and the bottom targets to a low-confidence choice; in the other half, the mapping was reversed such that the bottom targets corresponded to high-confidence and the top targets to high-confidence. To motivate participants to make calibrated confidence reports, the high- and low-confidence targets had different payoffs for correct and incorrect choices. The low confidence targets gave a 1 point reward for a correct choice and a 1 point loss for an incorrect choice. The high-confidence target gave 2 points for a correct choice and a loss of 3 points for an incorrect choice.

Fig. 3A reproduces data from a representative participant (Subject 2 in Figure 2 of van Den Berg et al. 2016), who completed 9,023 trials over 12 experimental sessions. Data from the other participants is shown in Fig. S1. Details of the experimental procedure should be sought in the original publication (van Den Berg et al., 2016).

Race model of decision making

In the race model, two drift-diffusion processes, $x_L(t)$ and $x_R(t)$, compete until one of them reaches an upper bound. The first to reach the upper bound determines choice and RT. x_R accumulates evidence for right minus left, and x_L accumulates evidence for left minus right. The dynamics of the decision variables is described by the following difference equations:

$$x_L^{(t+1)} = x_L^{(t)} + \kappa \Delta t (C + C_0) + u^{(t+1)} + \eta_L^{(t+1)} \sqrt{\Delta t}, \quad (2)$$

$$x_R^{(t+1)} = x_R^{(t)} - \kappa \Delta t (C + C_0) + u^{(t+1)} + \eta_R^{(t+1)} \sqrt{\Delta t} \quad (3)$$

where κ is a measure of the signal-to-noise, $\Delta t = 0.005$s is the time step, C is the motion coherence (positive for leftward motion), C_0 is a bias term and η is zero-mean normally distributed noise with unit variance. $\eta_L(t)$ and $\eta_R(t)$ are sampled from a bivariate Normal distribution such that the correlation between them is $\rho = -0.7$. At time $t = 0$, $x_L = x_R = 0$.

The urgency signal $u^{(t+1)}$ decreases the amount of evidence needed to trigger a response as time progresses (Hanks et al., 2011). For times $t < d$, the value of $u^{(t+1)}$ is zero, indicating no urgency. For times greater than d, $u^{(t+1)}$ assumes a constant value equal to $a\Delta t$, where a is a parameter that represents the linear rate of rise of the urgency signal.

The decision variables in the model cannot drop below a lower reflective bound. If an update to a decision variable would result in a value lower than B_{reflect}, the variable's value is set to B_{reflect} for that time step. That is:

$$x_L^{(t+1)} \leftarrow \max(x_L^{(t+1)}, B_{\text{reflect}}), \quad (4)$$

$$x_R^{(t+1)} \leftarrow \max(x_R^{(t+1)}, B_{\text{reflect}}). \quad (5)$$

The lower, non-absorbing bound simply instantiates the fact that firing rates must be ≥ 0.

The decision terminates when one of the races reach an upper bound at B. The choice is leftward (rightward) if x_L (x_R) reaches the bound first. The decision time is the time taken to reach the bound. The reaction time is the sum of the decision time and a non-decision time, which is normally distributed with a mean of μ_{nd} and standard deviation σ_{nd}.

The free parameters of the model are $\Theta = \{\kappa, B, a, d, C_0, \mu_{nd}, \sigma_{nd}\}$. We use simulations to fit them to data. Data from each monkey were fit separately. For a given set of parameters, we simulated 10 times as

many trials as were completed by the monkey. For each combination of choice and motion coherence, we fit the distribution of decision times obtained from the model with an Epanechnikov (parabolic) kernel to obtain a smooth probability density function of decision times. The distribution of decision times is convolved with the distribution of non-decision times to obtain a probability density function of reaction times. We compute a separate p.d.f. of RTs for each combination of choice and motion coherence, and use them to calculate the likelihood of the parameters given the single-trial choice and reaction time data. We use BADS (Acerbi and Ma, 2017) to find the maximum-likelihood parameters. The best-fitting parameters are shown in Table 1.

	κ	B	$a[s^{-1}]$	$d[s]$	ρ	$\mu_{nd}[s]$	$\sigma_{nd}[s]$	C_0	B_{rectif}
Monkey M	14.86	1.73	1.63	0.13	-0.7	0.28	0.07	0.01	-1
Monkey J	12.97	0.88	0.57	0.62	-0.7	0.3	0.03	0	-1

Table 1. Best-fitting parameters of the race model. ρ and B_{rectif} were not fit but fixed to predefined values.

Neurophysiological recordings (LIP)

Main task

Steinemann, Stine et al. (2022) used a prototype "alpha" version of Neuropixels1.0-NHP45 probes (developed by IMEC and HHMI-Janelia) to record multiple single-unit activities in the ventral part of area LIP (LIPv). Steinemann, Stine et al. (2022) used anatomical MRI to localize LIPv and used single-neuron recordings (Thomas Recording GmbH) to verify that the activity conformed to known physiological properties of LIPv before proceeding with multi-neuron recordings. The Neuropixels probes are equipped to record from 384 of the 4,416 available electrical contacts distributed along their 45 mm long shaft. Data was only recorded from the 384 contacts closest to the probe tip (Bank 0), covering 3.84 mm. The reference and ground signals were directly connected to each other and to the monkey's headpost. A total of 1,084 neurons were recorded over eight sessions, with each session yielding between 54 and 203 neurons (see Table 2 of Steinemann, Stine et al. 2022 for details).

Biased prior probability task

Hanks et al. (2011) recorded fifty-two neurons from the LIP area of two rhesus monkeys. Recordings were made using standard methods for extracellular recording of action potentials from single neurons (Roitman and Shadlen, 2002). See Hanks et al. (2011) for details.

Data analysis

Preprocessing of neuronal data

Our study focuses on T_{in} neurons, i.e., those that show persistent activity during saccade planning to the target contralateral to the recording site. For the Neuropixels data, T_{in} neurons were identified post hoc using a memory-guided saccade task (Steinemann, Stine et al., 2022; Stine et al., 2023). Hanks et al. (2011) used the same task to identify neurons with spatially selective persistent activity and to place targets within the response field of these neurons.

Unless otherwise stated, neurophysiological analyses are based on the number of spikes emitted by each T_{in} neuron in the 100 ms epoch that ends 50 ms before the initiation of the saccadic eye movement used to report the choice. We refer to this time interval as the presaccadic window.

Accuracy decoder

In each of the 8 sessions, we trained two decoders to predict the accuracy of the monkey's choice, using only contralateral (left) or ipsilateral (right) choices, respectively. We trained the decoder using simple logistic regression:

$$\text{logit}[p_{\text{correct}}] = \boldsymbol{\beta_{\text{conf}}}^{\top} \boldsymbol{S_x} + \beta_0 \quad , \tag{6}$$

where $\boldsymbol{S_x} = \tilde{\boldsymbol{S}}_{\mathbf{Tin}}$ is the standardized spike count of each T_{in} neuron in the interval from 150 ms to 50 ms before saccade onset. Standardization (i.e., z-score) is applied to each neuron independently. The fitted coefficients, $\boldsymbol{\beta_{\text{conf}}}$, establish a vector the same size as the number of simultaneously recorded T_{in} neurons. We refer to this vector as a coding direction in neuronal state space. β_0 is a constant that captures the accuracy rate over all stimulus conditions, which is typically much better than chance (typically, $p_{\text{correct}} > 0.7$).

We apply the same strategy to decode from subpopulations of neurons by redefining S_x. For example, Fig. 5A shows the result of applying the regression model (Eq. 6) to the subset of T_{in} neurons belonging to each cluster. For the analysis of the data from Hanks et al. (2011), S_x contains the standardized spike counts of the single T_{in} neurons recorded in separate experiments. We use the same presaccadic window as in the analysis of the Neuropixels data.

To derive the confidence estimates, the decoders are trained using 10-fold cross-validation. The data are divided into 10 groups, each containing an approximately equal number of trials selected randomly. One group is used as the prediction set, while the remaining nine groups are used for training. This process is repeated 10 times, ensuring that the confidence estimates for each trial are based on a prediction. For the analyses shown in Fig. 8, we derive $\beta_{\mathbf{choice}}$ and $\beta_{\mathbf{conf}}$ using all trials instead of using cross-validation. This approach allows us to obtain a single set of regression weights per session, rather than 10 sets as produced by the cross-validation method.

The confidence estimates are used to calculate the area under the ROC curve (AUC). Fig. 2A exemplifies the distributions that are used for the ROC analysis. For the analyses shown in Fig. 2B and Fig. 5D, AUC values were calculated separately for each session. The statistical comparisons use a one-tailed paired *t-test* applied to the logit-transformed AUC values from each session. Elsewhere, AUCs are not computed per session, but rather the confidence estimates from all sessions are pooled before calculating a single AUC. Standard errors were computed using bootstrapping ($N = 5,000$ samples).

We also use bootstrap samples to determine if there is a significant difference between two AUC values. For instance, to determine if the AUC_{conf} obtained using only the cluster–1 neurons is larger than that obtained using only the cluster–2 neurons (Fig. 5A), we bootstrap to obtain $N = 5,000$ AUC_{conf} values for each cluster. We then compare the AUC values for all $N = 25 \times 10^6$ pairwise comparisons and determine significance as the proportion of comparisons for which the AUC value from cluster–1 neurons is larger than that from cluster–2 neurons.

Choice decoder

We use the same spike-count standardization (z-scoring), analysis time window, and cross-validation method to predict the monkey's choice on each trial:

$$\mathrm{logit}[p_{\mathrm{left}}] = \beta_{\mathbf{choice}}^{\top} \tilde{S}_{\mathbf{Tin}} + \beta_0. \tag{7}$$

Here, β_0 reflects the monkeys bias for or against a left choice, not the monkey's accuracy. Unlike the accuracy decoder, the choice decoder is trained on trials with both contralateral and ipsilateral choices.

Reaction time decoder

We also use a logistic decoder to assess whether T_{in} neurons contain, just before the response, information about reaction time. The model is:

$$\mathrm{logit}[p_{\mathrm{fast}}] = \beta_{\mathbf{RT}}^{\top} \tilde{S}_{\mathbf{Tin}} + \beta_0. \tag{8}$$

It was fit independently for each session using only trials with contraversive (left) choices. Fast and slow responses were defined relative to the median RT. We use the same cross-validation method that we used for the accuracy and choice decoders.

Latency analysis

We used the cumulative sum (CUSUM) method to determine the latency of direction-selective responses in T_{in} neurons (Ellaway, 1978).

Receiver operating characteristic (ROC) analysis was used to assess the directional selectivity of each T_{in} neuron. The area under the curve (AUC) in this analysis represents the degree separation—from 0.5 (complete overlap) to 1 (complete separation)—between the spike count distributions for leftward and rightward motion in single trials. The AUC was calculated from spike counts in the interval 100–400 ms after motion onset. We restricted the analysis to correct trials with reaction times greater than 450 ms and motion coherence greater than 10%.

For neurons with an AUC greater than 0.55, we calculated the difference in spike counts (in 25 ms bins) between leftward and rightward choice trials. These differences were then added cumulatively over time,

as required by the CUSUM method. Typically, the cumulative difference remains around zero prior to the onset of direction selectivity, and then gradually increases or decreases depending on the neuron's preferred choice.

To determine the onset of direction selectivity, we fit a "dogleg" function to the cumulative spike sum. This function starts with a flat line from $t_0 = 0$ and transitions to a linear increase or decrease starting at $t_1 > t_0$. The end of the flat portion of the fit, which occurs between 0 and 500 ms after the onset of motion, was considered the latency.

Using cumulative sums of spikes to estimate latencies helps reduce the effect of neural noise. The fitting step further reduces the influence of the number of trials on latency estimates, providing an advantage over traditional methods such as t-tests in moving windows (e.g., Lorteije, Zylberberg et al., 2015).

The significance of the difference in latency between neurons from different clusters was assessed with a two-tailed *t-test*.

We conducted a similar analysis to estimate the latencies depicted in Fig. 8. We construct bootstrap samples combining the time course of AUC values from individual sessions (N = 5,000 bootstrap samples). A dogleg function was fit to each bootstrap sample, resulting in 5,000 latency estimates. The arrows in Fig. 8 identify the mean latency across the bootstrap samples. The p-value we report is the proportion of bootstrap samples for which the choice signal deviates from baseline later than the confidence signal.

Clustering

We use linear regression to explain the spike counts of each T_{in} neuron in the presaccadic window. As independent variables, we used the motion coherence (C), the choice and reaction time (RT), and an offset:

$$\tilde{S}_{Tin}^{(i)} = \beta_0^{(i)} + \beta_1^{(i)} C + \beta_2^{(i)} \mathrm{RT} + \beta_3^{(i)} \mathrm{choice}. \tag{9}$$

$\tilde{S}_{Tin}^{(i)}$ is the standardized spike count of neuron i in the aforementioned time interval. The regression analysis was performed separately for each T_{in} neuron, including correct trials only.

We then apply *k-means* clustering, using 3 clusters, to the regression coefficients associated with motion coherence, RT and choice. The cluster labels (1–3) were chosen so that the average of β_1 is smallest for cluster–1 neurons, intermediate for cluster 2 and largest for cluster 3. We confirmed the robustness of the neuron cluster assignments by independently deriving them using either the odd or even trials (Fig. S3). We also fit a regression model similar to Eq. 9, but without the choice and RT terms and considering only correct contralateral choices. The regression coefficients associated with motion coherence are shown in Fig. 4D.

Motion pulses

We estimate the time course of the effect of the brief motion pulses on neuronal activity by aligning the spike times of each T_{in} neuron to the onset of the motion pulse. For each time $-60 < t < 800$ ms relative to the onset of the motion pulse, we calculate the number of spikes emitted in the time epoch between t and $(t+100)$ ms. Time t is advanced in steps of 20 ms. $R_i(t, j)$ contains the spike counts for neuron i emitted at time t from pulse onset, in trial j. Only trials where the reaction time is at least 150 ms greater than t, and those with a motion pulse ($\approx \frac{1}{2}$) are included in the analysis. We then standardize the spike counts independently for each motion coherence, to obtain $\tilde{R}_i(t, j)$. We average $\tilde{R}_i(t, j)$ across trials with a leftward (contraversive) pulse, and trials with a rightward (ipsiversive) pulse, and calculate the difference, left minus right, between these averages, to obtain $\Delta \tilde{R}_i(t)$.

$\Delta \tilde{R}_i(t)$ approximates for each T_{in} neuron i and time t, the influence of the motion pulse on neuronal activity. We average $\Delta \tilde{R}_i(t)$ across the N_K neurons belonging to the same cluster K,

$$\Delta \tilde{R}_K(t) = \frac{1}{N_K} \sum_{i \in K} \Delta \tilde{R}_i(t) \quad , \tag{10}$$

and normalize the average subtracting a baseline,

$$D_K(t) = \Delta \tilde{R}_K(t) - \mathrm{baseline} \quad , \tag{11}$$

where baseline is the average of $\Delta \bar{R}_K(t)$ for times t between 0 and $0.2s$. Fig. 6A shows $D_K(t)$ for the three clusters.

We use a curve-fitting approach to estimate the rate at which the effect of the motion pulse dissipates over time. We fit $D_K(t)$ using a function $f(x)$ constructed on the following two assumptions: (i) the onset latency of the effect of the motion pulse on neuronal activity follows a Normal distribution, (ii) the effect dissipates exponentially. Given these assumption, the differential equation describing the time course of $f(t)$ is

$$\frac{\partial f(t)}{\partial t} = -\alpha f(t) + \mathcal{N}(t|\mu, \sigma) \quad , \tag{12}$$

where μ and σ are the mean and standard deviation of the Normal distribution ($\mathcal{N}$), and α is the reciprocal of the time constant of the dissipation. The solution to this equation is:

$$f(t) = d \cdot \exp(\mu\alpha + \frac{1}{2}\sigma^2\alpha^2 - \alpha t) \cdot \Phi(t|m, \sigma) \quad , \tag{13}$$

where d is a scaling parameter, and $\Phi(\cdot|m, \sigma)$ is the cumulative Gaussian distribution with mean $m = \mu + \sigma^2\alpha$ and standard deviation σ. The fit function has four parameters: $\{d, \mu, \sigma, \alpha\}$. We fit the parameters to minimize the sum over time points of the squared differences between $f(t)$ and $D_K(t)$.

We compared the best-fitting dissipation parameter (α) across clusters. We generated bootstrap samples (N=5,000) by selecting with replacement from the pool of neurons that belong to a given cluster. For each of these samples, we compute the parameter α. We evaluate the significance of the difference in α values in the data by the proportion of bootstrap samples that result in a difference in α values as extreme as the one we observed in the data.

Cross-correlation analysis

The analysis depicted in Fig. 6B–D is based on the spike counts of neurons from clusters 1–3. Spike counts were calculated in 25 ms windows, aligned to motion onset, up to 50 ms before the reaction time. We compute standardized residuals separately for each time bin, motion coherence and session. Standardized residuals were combined across sessions. The processed signals are referred to as $\tilde{S}_{\text{Tin}}^{k=1}$, $\tilde{S}_{\text{Tin}}^{k=2}$ and $\tilde{S}_{\text{Tin}}^{k=3}$. We then calculated the Pearson correlation coefficient between every pair of signals (Fig. 6B–D), for every pair of time steps between 0.2 and 0.8s.

We used permutation tests to asses statistical significance. We define two regions of interest based on the time from stimulus onset in the x and y dimensions (Fig. 6B). ROI1 is defined by $t_x > t_y$, and ROI2 is defined by $t_y > t_x$, for time time points shown in Fig. 6B. If y causally affects x, or if y and x receive a common input but the integration time constant is greater for y than for x, then the pairwise correlations between x and y should be greater in ROI1 than in ROI2. We calculated the difference in correlations between two groups, $\langle \rho_{\text{ROI1}} \rangle - \langle \rho_{\text{ROI2}} \rangle$, where the expectation is calculated over the time bins within each region of interest (ROI). This difference was then contrasted with those obtained after randomly shuffling the order of trials for one of the dimensions ($N_{\text{shuffles}} = 200$). Significance was determined by the probability of achieving a difference as extreme as the one observed in the experimental data.

The same procedure was applied to the cross-correlation analysis shown in Fig. 7D, but with $x = \tilde{S}_{\text{Tin}}^{k=1}(t_x)$ and $y = \tilde{S}_{\text{Min}}^{\text{left}}(t_y) - \tilde{S}_{\text{Min}}^{\text{right}}(t_y)$, where $\tilde{S}_{\text{Min}}^{\text{left}}(t)$ and $\tilde{S}_{\text{Min}}^{\text{right}}(t)$ are the standardized residuals obtained from the activity of the M$_{\text{in}}$ neurons preferring contraversive and ipsiversive motion, respectively.

Acknowledgments

Author contributions

A.Z. conceived, designed and conducted the research presented in this study. A.Z. wrote the original draft. A.Z. and M.N.S. revised and edited the manuscript.

Funding

References

Acerbi L, Ma WJ. Practical Bayesian optimization for model fitting with Bayesian adaptive direct search. arXiv preprint arXiv:170504405. 2017; .

Aitchison L, Bang D, Bahrami B, Latham PE. Doubly Bayesian analysis of confidence in perceptual decision-making. PLoS computational biology. 2015; 11(10):e1004519.

Boundy-Singer ZM, Ziemba CM, Goris RLT. Sensory population activity reveals confidence computations in the primate visual system. bioRxiv. 2024; https://www.biorxiv.org/content/early/2024/08/04/2024.08.01.606172, doi: 10.1101/2024.08.01.606172.

Britten KH, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. Journal of Neuroscience. 1992; 12(12):4745–4765.

Brus J, Aebersold H, Grueschow M, Polania R. Sources of confidence in value-based choice. Nature communications. 2021; 12(1):7337.

Carpenter RH, Williams M. Neural computation of log likelihood in control of saccadic eye movements. Nature. 1995; 377(6544):59–62.

Colby CL, Duhamel JR, Goldberg ME. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. Journal of neurophysiology. 1996; 76(5):2841–2852.

Comay NA, Della Bella G, Lamberti P, Sigman M, Solovey G, Barttfeld P. The presence of irrelevant alternatives paradoxically increases confidence in perceptual decisions. Cognition. 2023; 234:105377.

Del Cul A, Dehaene S, Reyes P, Bravo E, Slachevsky A. Causal role of prefrontal cortex in the threshold for access to consciousness. Brain. 2009; 132(9):2531–2540.

van Den Berg R, Anandalingam K, Zylberberg A, Kiani R, Shadlen MN, Wolpert DM. A common mechanism underlies changes of mind about decisions and confidence. Elife. 2016; 5:e12192.

Dou W, Martinez Arango LJ, Castaneda OG, Arellano L, Mcintyre E, Yballa C, Samaha J. Neural Signatures of Evidence Accumulation Encode Subjective Perceptual Confidence Independent of Performance. Psychological Science. 2024; p. 09567976241246561.

Ellaway P. Cumulative sum technique and its application to the analysis of peristimulus time histograms. Electroencephalography and clinical neurophysiology. 1978; 45(2):302–304.

Fetsch CR, Kiani R, Newsome WT, Shadlen MN. Effects of cortical microstimulation on confidence in a perceptual decision. Neuron. 2014; 83(4):797–804.

Freedman DJ, Assad JA. Experience-dependent representation of visual categories in parietal cortex. Nature. 2006; 443(7107):85–8. http://www.ncbi.nlm.nih.gov/pubmed/16936716, doi: 10.1038/nature05078.

Galvin SJ, Podd JV, Drga V, Whitmore J. Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. Psychonomic bulletin & review. 2003; 10:843–876.

Gnadt JW, Andersen RA. Memory related motor planning activity in posterior parietal cortex of macaque. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale. 1988; 70(1):216–20. http://www.ncbi.nlm.nih.gov/pubmed/3402565.

Gold JI, Shadlen MN. The neural basis of decision making. Annual review of neuroscience. 2007; 30.

Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature. 2015; 520(7546):220–223.

Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. Journal of Neuroscience. 2011; 31(17):6339–6352.

Hellmann S, Zehetleitner M, Rausch M. Simultaneous modeling of choice, confidence, and response time in visual perception. Psychological review. 2023; .

**Henmon VAC**. The relation of the time of a judgment to its accuracy. Psychological review. 1911; 18(3):186.

Hikosaka O, Wurtz RH. Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. Journal of Neurophysiology. 1983; 49(5):1268–1284. doi: 10.1152/jn.1983.49.5.1268.

Howard IS, Ingram JN, Wolpert DM. A modular planar robotic manipulandum with end-point torque control. Journal of neuroscience methods. 2009; 181(2):199–211.

Irwin FW, Smith W, Mayfield JF. Tests of two theories of decision in an" expanded judgment" situation. Journal of experimental psychology. 1956; 51(4):261.

Kepecs A, Mainen ZF. A computational framework for the study of confidence in humans and animals. Philosophical Transactions of the Royal Society B: Biological Sciences. 2012; 367(1594):1322–1337.

Kepecs A, Uchida N, Zariwala HA, Mainen ZF. Neural correlates, computation and behavioural impact of decision confidence. Nature. 2008; 455(7210):227–231.

Kiani R, Shadlen MN. Representation of confidence associated with a decision by neurons in the parietal cortex. Science. 2009; 324(5928):759–764. doi: 10.1126/science.1169405.

Kiani R, Corthell L, Shadlen MN. Choice certainty is informed by both evidence and decision time. Neuron. 2014; 84(6):1329–1342.

Koizumi A, Maniscalco B, Lau H. Does perceptual confidence facilitate cognitive control? Attention, Perception, & Psychophysics. 2015; 77:1295–1306.

Komura Y, Nikkuni A, Hirashima N, Uetake T, Miyamoto A. Responses of pulvinar neurons reflect a subject's confidence in visual categorization. Nature neuroscience. 2013; 16(6):749–755.

Lak A, Costa GM, Romberg E, Koulakov AA, Mainen ZF, Kepecs A. Orbitofrontal cortex is required for optimal waiting based on decision confidence. Neuron. 2014; 84(1):190–201.

Lak A, Nomoto K, Keramati M, Sakagami M, Kepecs A. Midbrain dopamine neurons signal belief in choice accuracy during a perceptual decision. Current Biology. 2017; 27(6):821–832.

Lange RD, Chattoraj A, Beck JM, Yates JL, Haefner RM. A confirmation bias in perceptual decision-making due to hierarchical approximate inference. PLoS Computational Biology. 2021; 17(11):e1009517.

Lebreton M, Abitbol R, Daunizeau J, Pessiglione M. Automatic integration of confidence in the brain valuation signal. Nature neuroscience. 2015; 18(8):1159–1167.

Lorteije, Zylberberg, Ouellette BG, De Zeeuw CI, Sigman M, Roelfsema PR. The formation of hierarchical decisions in the visual cortex. Neuron. 2015; 87(6):1344–1356.

Maniscalco B, Peters MA, Lau H. Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. Attention, Perception, & Psychophysics. 2016; 78:923–937.

Masset P, Ott T, Lak A, Hirokawa J, Kepecs A. Behavior-and modality-general representation of confidence in orbitofrontal cortex. Cell. 2020; 182(1):112–126.

Mazor M, Maimon-Mor RO, Charles L, Fleming SM. Paradoxical evidence weighting in confidence judgments for detection and discrimination. Attention, Perception, & Psychophysics. 2023; p. 1–30.

Mazzoni P, Bracewell RM, Barash S, Andersen RA. Motor intention activity in the macaque's lateral intraparietal area. I. Dissociation of motor plan from sensory memory. Journal of neurophysiology. 1996; 76(3):1439–1456.

Middlebrooks PG, Sommer MA. Neuronal correlates of metacognition in primate frontal cortex. Neuron. 2012; 75(3):517–530.

Miyamoto K, Setsuie R, Osada T, Miyashita Y. Reversible silencing of the frontopolar cortex selectively impairs metacognitive judgment on non-experience in primates. Neuron. 2018; 97(4):980–989.

Moreno-Bote R. Decision confidence and uncertainty in diffusion models with partially correlated neuronal integrators. Neural computation. 2010; 22(7):1786–1811.

Odegaard B, Grimaldi P, Cho SH, Peters MA, Lau H, Basso MA. Superior colliculus neuronal ensemble activity signals optimal rather than subjective confidence. Proceedings of the National Academy of Sciences. 2018; 115(7):E1588–E1597.

Palmer J, Huk AC, Shadlen MN. The effect of stimulus strength on the speed and accuracy of a perceptual decision. Journal of vision. 2005; 5(5):1–1.

Peters MA, Thesen T, Ko YD, Maniscalco B, Carlson C, Davidson M, Doyle W, Kuzniecky R, Devinsky O, Halgren E, et al. Perceptual confidence neglects decision-incongruent evidence in the brain. Nature human behaviour. 2017; 1(7):0139.

Rahnev D, Denison RN. Suboptimality in perceptual decision making. Behavioral and brain sciences. 2018; 41:e223.

Ratcliff R. A theory of memory retrieval. Psychological review. 1978; 85(2):59.

Ratcliff R, Smith PL, Brown SD, McKoon G. Diffusion decision model: Current issues and history. Trends in cognitive sciences. 2016; 20(4):260–281.

Ritchie JB, Kaplan DM, Klein C. Decoding the brain: Neural representation and the limits of multivariate pattern analysis in cognitive neuroscience. The British journal for the philosophy of science. 2019; .

Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. Journal of neuroscience. 2002; 22(21):9475–9489.

Rolls ET, Grabenhorst F, Deco G. Choice, difficulty, and confidence in the brain. Neuroimage. 2010; 53(2):694–706.

Rounis E, Maniscalco B, Rothwell JC, Passingham RE, Lau H. Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. Cognitive neuroscience. 2010; 1(3):165–175.

Samaha J, Denison R. The positive evidence bias in perceptual confidence is not post-decisional. BioRxiv. 2020; p. 2020–03.

Sepulveda P, Usher M, Davies N, Benson AA, Ortoleva P, De Martino B. Visual attention modulates the integration of goal-relevant evidence and not value. Elife. 2020; 9:e60705.

Shadlen MN, Hanks TD, Churchland AK, Kiani R, Yang T. The speed and accuracy of a simple perceptual decision: a mathematical primer. Bayesian brain: Probabilistic approaches to neural coding. 2006; p. 209–37.

Smith PL, Vickers D. The accumulator model of two-choice discrimination. Journal of Mathematical Psychology. 1988; 32(2):135–168.

So N, Stuphorn V. Supplementary eye field encodes confidence in decisions under risk. Cerebral Cortex. 2016; 26(2):764–782.

Steinemann, Stine, Trautmann, Zylberberg, Wolpert, Shadlen. Direct observation of the neural computations underlying a single decision. bioRxiv. 2022; p. 2022–05.

Stine GM, Trautmann EM, Jeurissen D, Shadlen MN. A neural mechanism for terminating decisions. Neuron. 2023; .

Usher M, McClelland JL. The time course of perceptual choice: the leaky, competing accumulator model. Psychological review. 2001; 108(3):550.

Vickers D. Decision processes in visual perception. Academic Press; 1979.

Vickers D, Packer J. Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. Acta psychologica. 1982; 50(2):179–197.

Vickers D, Smith P, Burt J, Brown M. Experimental paradigms emphasising state or process limitations: II effects on confidence. Acta Psychologica. 1985; 59(2):163–193.

Vivar Lazo M. Cortical Population Dynamics Underlying Choice, Reaction Time, and Confidence. PhD thesis, Johns Hopkins University; 2024.

Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 2002; 36(5):955–968.

Wei Z, Wang XJ. Confidence estimation as a stochastic process in a neurodynamical system of decision making. Journal of neurophysiology. 2015; 114(1):99–113.

Wong KF, Wang XJ. A recurrent network mechanism of time integration in perceptual decisions. Journal of Neuroscience. 2006; 26(4):1314–1328.

Zylberberg A, Barttfeld P, Sigman M. The construction of confidence in a perceptual decision. Frontiers in integrative neuroscience. 2012; 6:79.

Zylberberg A, Dehaene S, Mindlin GB, Sigman M. Neurophysiological bases of exponential sensory decay and top-down memory retrieval: a model. Frontiers in computational neuroscience. 2009; 3:341.

Zylberberg A, Fetsch CR, Shadlen MN. The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. Elife. 2016; 5:e17688.

Zylberberg A, Roelfsema PR, Sigman M. Variance misperception explains illusions of confidence in simple perceptual decisions. Consciousness and cognition. 2014; 27:246–253.

Zylberberg A, Shadlen MN. Cause for pause before leaping to conclusions about stepping. BioRxiv. 2016; p. 085886.

Zylberberg A, Wolpert DM, Shadlen MN. Counterfactual reasoning underlies the learning of priors in decision making. Neuron. 2018; 99(5):1083–1097.

Supplemental information
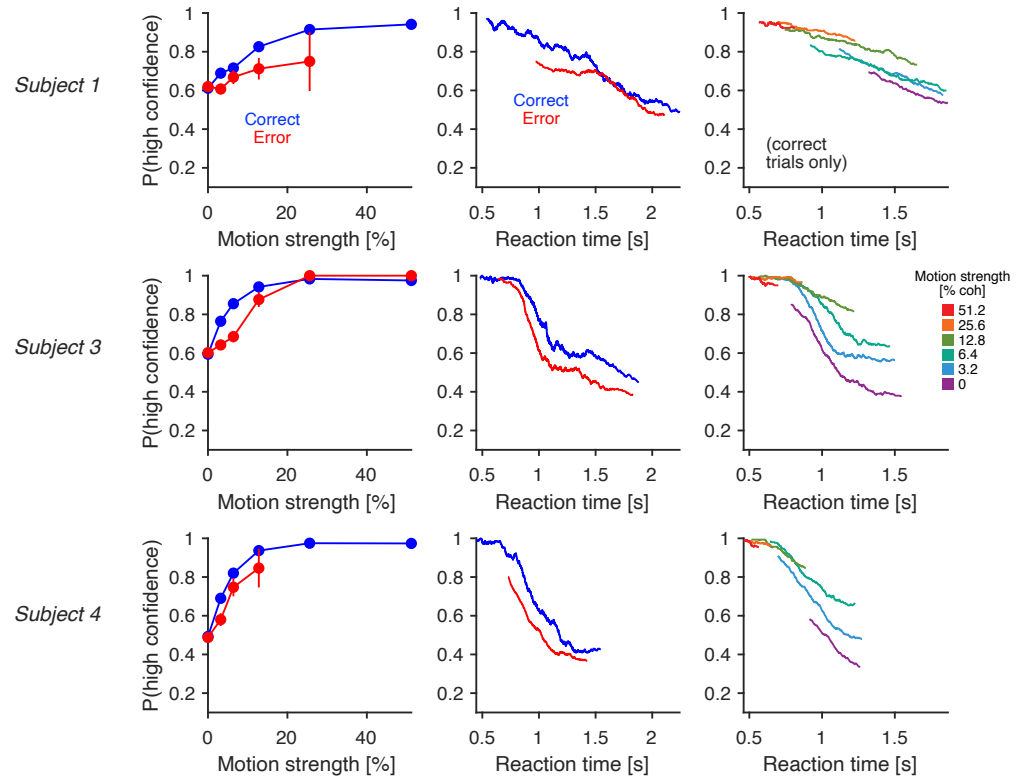
Figure S1. Signatures of confidence in the data of van Den Berg et al. (2016)
Same as Fig. 3A for the other three participants in van Den Berg et al. (2016).
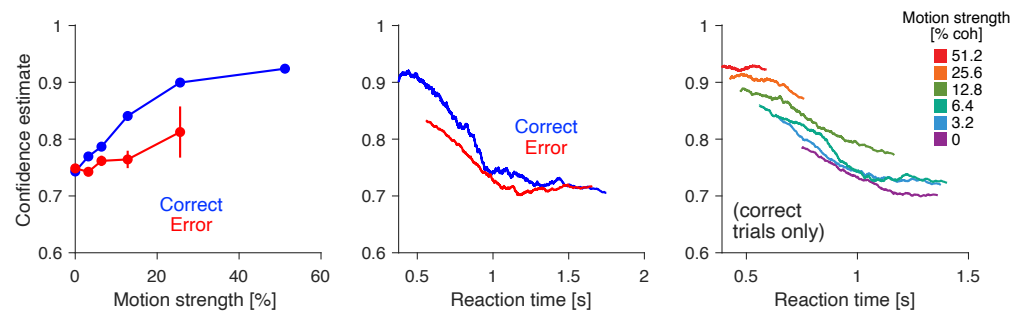
Figure S2. Choice accuracy inferred from neural activity

Same as Fig. 3B, except that the confidence estimates obtained from the accuracy decoder are not thresholded into high and low confidence categories. Results are qualitatively similar to those obtained from the behavioral data (Fig. 3A).
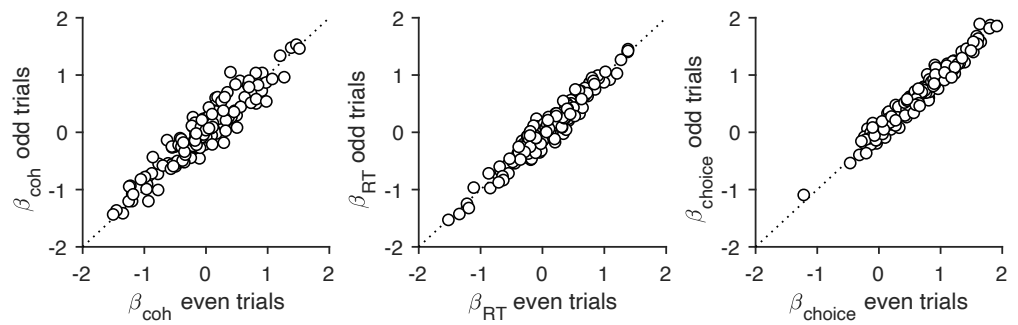
Figure S3. Consistency of the neural representation of motion coherence, reaction time and choice
We use only the odd or even trials to explain the standardized (z-scored) spike counts of each neuron in the presaccadic window as a function of motion coherence, RT, and choice. The figure shows the best-fitting regression coefficients plotted against each other. Each data point corresponds to a different T_{in} neuron. Panels 1-3 correspond to the best-fitting regression coefficients for motion coherence, RT, and choice, respectively. The regression coefficients are highly consistent across independent regression analyses.
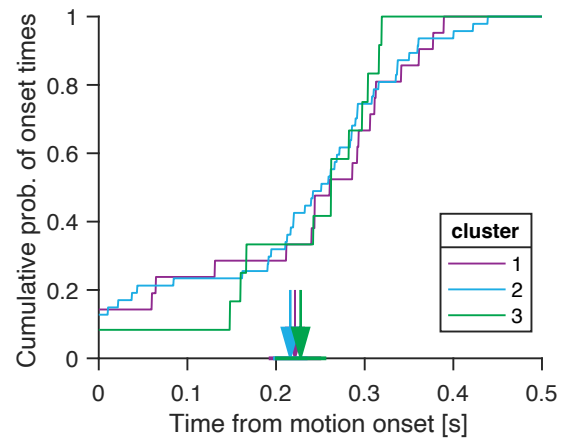
Figure S4. Latency to motion selectivity of individual T_{in} neurons
For each T_{in} neurons separately, we calculate the latency to motion direction selectivity using the CUSUM method (Ellaway, 1978). The cumulative distribution of onset times is shown separately for neurons belonging to the three clusters. The vertical arrows indicate the mean onset time for each cluster, and the horizontal line indicates the standard error of the mean.
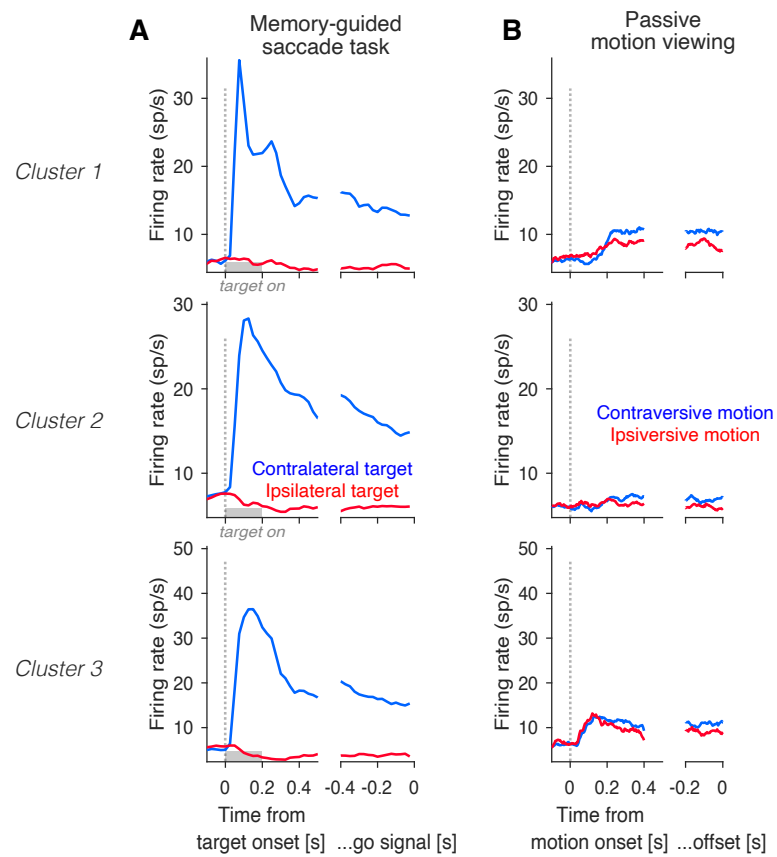
Figure S5. Mean response of neurons from each cluster in the memory-guided saccade and passive motion–viewing tasks

(A) memory-guided saccade task. Blue and red traces correspond to saccades to the target located contralaterally and ipsilaterally, respectively. The horizontal gray bar indicates the time of target presentation. In the panels on the right, neural activity is aligned to the go signal (i.e., the offset of the fixation point). The average time from target onset to the go signal is 0.82 seconds. **(B)** Passive motion–viewing task. Blue and red traces correspond to contraverise (leftward) and ipsiversive (rightward) motion, respectively. Traces are aligned to the onset (left) and offset (right) of the random dot motion stimulus.
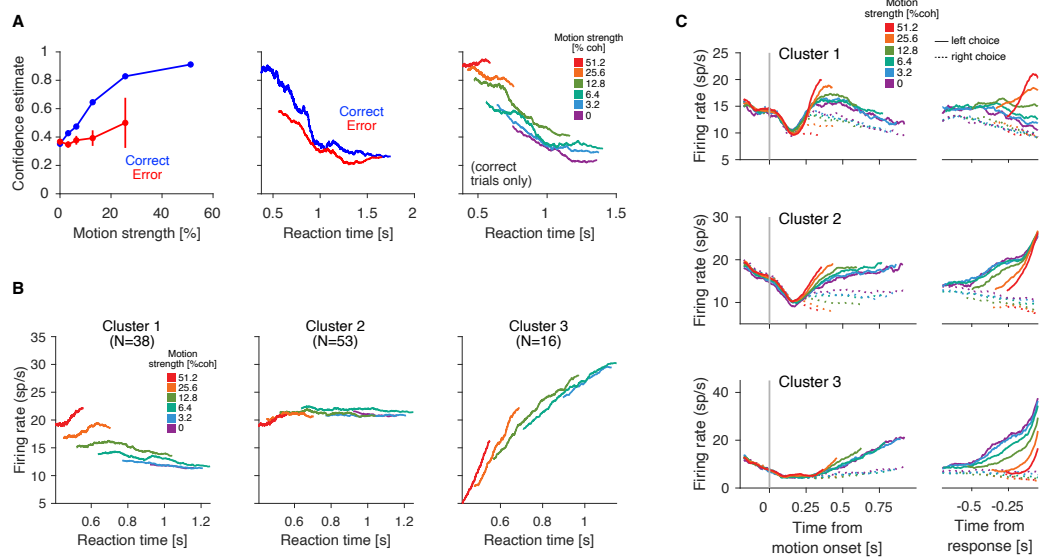
Figure S6. Decoding and clustering without the T$_{in}$ neurons with significant modulation in the passive motion-viewing task.

We repeated the decoding and clustering analyses without the T$_{in}$ neurons that significantly discriminated between leftward and rightward motion in the passive motion viewing task. Significance was assessed using a Wilcoxon rank–sum test comparing spike rates on leftward and rightward motion trials. Spike rates were calculated for each trial in the epoch between 0.2s after motion onset and motion offset. Neurons with p-values lower than 0.05 were deemed significant. The results of the decoding and clustering analyses are qualitatively similar to those obtained without excluding these neurons. **(A)** Analysis equivalent to that shown in Fig. 3B. **(B)** Analysis equivalent to that shown in Fig. 4B. **(C)** Analysis equivalent to that shown in Fig. 4C.
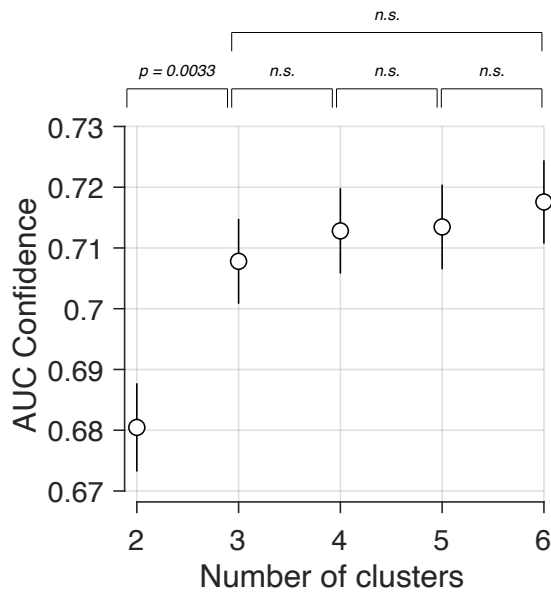
Figure S7. AUC_{conf} versus the number of clusters used in *k-means*
We repeat the clustering analysis applied to the regression coefficients obtained from Eq. 9, using different numbers of clusters (abccissa). We average the spike counts in the presaccadic window of neurons belonging to the same cluster. From these averaged spike counts, we calculate the AUC_{conf} (ordinate) using Eq. 6, just as we do elsewhere. We only find a significant difference in the AUC_{conf} values between $N=2$ clusters and $N>2$ clusters (bootstrap). P-values were higher than 0.15 for all non-significant comparisons (*n.s.*).