Prefrontal Cortex Neural Correlates of the Contraction Bias in Frequency Discrimination

Luis Serrano-Fernández[a,b], Manuel Beirán[c], Ranulfo Romo [d], and Néstor Parga [a,b]

[a]Departamento de Física Teórica, Universidad Autónoma de Madrid, 28049 Madrid, Spain

[b]Centro de Investigación Avanzada en Física Fundamental, Universidad Autónoma de Madrid, 28049 Madrid, Spain

[c]Center for Theoretical Neuroscience, Zuckerman Institute, Columbia University, New York, NY, USA

[d]El Colegio Nacional, 06020 México DF, México.

Abstract

Perception is influenced by sensory stimulation, prior knowledge, and contextual cues, which collectively contribute to the emergence of perceptual biases. However, the precise neural mechanisms underlying these biases remain poorly understood. This study aims to address this gap by analyzing neural recordings from the prefrontal cortex (PFC) of monkeys performing a vibrotactile frequency discrimination task. Our findings provide empirical evidence supporting the hypothesis that perceptual biases can be reflected in the neural activity of the PFC. We found that the state-space trajectories of PFC neuronal activity encoded a warped representation of the first frequency presented during the task. Remarkably, this distorted representation of the frequency aligned with the predictions of its Bayesian estimator. The identification of these neural correlates expands our understanding of the neural basis of perceptual biases and highlights the involvement of the PFC in shaping perceptual experiences. Similar analyses could be employed in other delayed comparison tasks and in various brain regions to explore where and how neural activity reflects perceptual biases during different stages of the trial.

Keywords— Perceptual biases, Prefrontal Cortex, Bayesian Computations, Decision Making, State-space geometry

Introduction

Perception is shaped by various factors including the current and previous stimulation, prior knowledge, and contextual cues [1], [2]. Collectively, these factors can contribute to the emergence of perceptual biases [2]–[10]. These biases manifest in various forms, such as the central tendency observed in tasks involving magnitude or duration reproduction [11]–[15], the contraction bias observed in delayed comparison tasks [7], [10], [16]–[21], and the phenomenon of serial dependence observed in delayed response experiments [22]–[24]. A normative Bayesian account posits that perception emerges through the integration of a noisy representation of the stimulus with prior knowledge about the sensory environment [1], [19], [25]. As the level of sensory noise increases, Bayesian observers tend to rely more heavily on their prior knowledge when forming their estimations. In certain situations, this can result in a perceptual bias towards perceiving stimuli parameters that are closer to the center of their distribution. For instance, this Bayesian framework has been suggested to explain the central tendency effect [26]–[29], the contraction bias [10], [19], [20], [30] and serial dependence [31]–[33].

Our current understanding of the relationship between neural activity and the contraction bias is still quite limited. However, there are some notable studies that have shed light on this topic [7], [20], [21]. One such study by Akrami et al. (2018) highlighted the importance of working memory in the manifestation of the contraction bias. Additionally, another study by Sarno et al. (2022) explored how the contraction bias modulates dopamine activity in the midbrain. Benozzo and Genovesio (2023) analyzed the activity of prefrontal neurons to investigate the effects of the contraction bias on the decision-making process. While these studies are intriguing, they do not address the issue of how Bayesian computations leading to the contraction bias manifests in the activity of cortical neurons.

Recently, some progress has been done in that direction. In tasks involving time interval reproduction, where subjects report their perception of duration, it has been observed that the curvature of state-space trajectories during the presentation of the sample interval is a distorted representation of the interval's duration [14], [29]. This effect was also evident when the interval, instead of being reproduced, needs to be compared with another interval presented after a delay period [20]. Recurrent networks of spiking neurons trained on this task also displayed a curved manifold carrying a warped representation of the intervals during the presentation of the first stimulus. Throughout the delay period, the interval continued to be represented in a distorted manner by other geometric characteristics of the trajectories. Importantly, in both studies, these distorted representations aligned with the Bayesian estimation of the interval duration. This line of research provides valuable insights into the cognitive and neural processes involved in the formation of perceptual biases and their implications for decision-making. However, further studies are necessary to deepen our understanding of the specific computational mechanisms and population activity dynamics through which the biases emerge in cortical areas.

In this work, to investigate the neural representation of the contraction bias in cortical circuits and the underlying computational mechanisms, we examined the activity of prefrontal cortex (PFC) neurons during a vibrotactile frequency discrimination task performed by two monkeys [34], [35]. Our findings revealed that PFC neurons encoded information about the current stimulus but not about those in preceding trials. In contrast, we found evidence supporting the idea that PFC neurons possess knowledge about the structure of the stimulus set, a basic requirement for Bayesian computations. To

further investigate the compatibility of our findings with Bayesian principles, we developed a normative Bayesian model and successfully fitted it to behavioral data. Remarkably, the geometric characteristics of the state-space activity trajectories spanning from the presentation of the first frequency to the end of the delay period exhibited features that corresponded to the Bayesian estimator of that frequency. These results show that the dynamics of population activity in the PFC align with Bayesian concepts. This study opens the way for future research to investigate the underlying neural mechanisms and circuitry involved in the emergence of the contraction and other perceptual bias.

Results

Tasks. The research was conducted on recordings collected from two monkeys in previous experiments [34], [35] (SI Appendix, Methods). Both monkeys had been trained to compare two frequencies, f_1 and f_2, separated by a delay period. However, while the first monkey was required to respond shortly after the offset of f_2 (Fig. 1*A, Upper*), the second monkey had to wait for 3 s before communicating its decision at the moment of the Probe Up (PU) event (Fig. 1*A, Bottom*). Additionally, different stimulus sets were used; the first monkey was presented with the set 1, shown in Fig. 1*B, Left*, while the second monkey was presented with the set 2, depicted in Fig. 1*B, Right*. In both sets, frequency pairs (f_1,f_2) (stimulus classes) are organized along two diagonals, to be referred to as the upper and lower diagonals of the set. For later convenience classes were assigned a label (indicated inside the squares in Fig. 1*B*). Note that set 2 deliberately excluded any value of f_1 that would allow subjects to predict the appropriate motor response.

Behavior. Monkeys exhibited the contraction bias. The percentage of correct responses (Fig. 1*C*) revealed that performance is contingent on the class, even though in either all (Fig. 1*C*, *Left* and *Middle*) or a majority (Fig. 1*C, Right*) of the classes the difference between the two frequencies presented in the trial possesses the same absolute value ($\Delta f =\mid f_1 - f_2 \mid= 8$ Hz). This discrepancy can be attributed to the presence of a contraction bias, which impacted decision accuracy based on the specific pairing of stimuli. This internal process caused a shift in the perceived frequency of the base stimulus f_1 towards the center of the frequency range ($\langle f_1 \rangle = 22$ Hz and $\langle f_1 \rangle = 21.67$ Hz, in Figs. 1*B, Left* and *Right*, respectively). Consequently, when f_1 had a lower value, it was perceived as larger, whereas a higher value was perceived as smaller, as illustrated in Fig. 1*D*. Thus, the influence of the bias progressively hindered accurate evaluations for classes where $f_1 < f_2$ as they were encountered from right to left (upper diagonals in Fig. 1*B*). In contrast, the bias increasingly facilitated classes where $f_1 > f_2$ (lower diagonals). To uphold the bias structure, we designated the classes benefiting from the bias effects as positioned towards the extreme ends, while the central classes encountered restrictions imposed by the bias (Fig. 1*B*). In this classification arrangement, the classes positioned near the two ends, where the bias had a positive influence, demonstrated the highest level of decision accuracy. As already noticed in [10], when plotted against the class number, the resulting curves displayed V-shaped patterns, as depicted in Fig. 1*C*.

In this work we will focus on the effect that prior knowledge combined with observation on the

current trial may have on the occurrence of the contraction bias. The Bayesian explanation predicts that increasing the length of the delay period will lead to an increment in the magnitude of the bias [23]. This is because the maintenance in short-term memory of f_1 may lead to greater uncertainty about its value. To verify this prediction, we started by quantifying the contraction bias B as the root-mean-square deviation of the animal performance (SI Appendix, Methods). Calculating B for two different durations of this interval, $\Delta = 3$ s or $\Delta = 6$ s (monkey 1), yielded $B = 0.036$ and $B = 0.078$, respectively, confirming the Bayesian prediction.

PFC neurons carried information about the current but not about the previous trial f_1. A possible source for the contraction bias is sensory history [7], [20], [36], acting, e.g., through the persistence in the population firing activity of the representations of stimuli applied in previous trials. However, this does not imply that the activity of neurons in a given area should show such dependence, even if the neural activity in that area exhibits the contraction bias in some manner. To clarify the behavior of PFC neurons, we first calculated the mutual information (MI) that each neuron carried about the current trial f_1 value and its value from the previous trial (SI Appendix, Methods). Subsequently, we determined the fraction of neurons for which those pieces of information were found to be significant. In both monkeys, for the two choices, neurons exhibited significant information regarding the value of f_1 in the current trial (Supplementary Fig. 1). The population size conveying this information initially increased upon presentation of f_1, then decreased during the early phase of the delay period, and eventually recovered prior to the onset of f_2 (Fig. 2, blue lines). However, there was no population of neurons carrying significant information about the value of f_1 from the previous trial (Fig. 2, red lines). This result is particularly evident for monkey 2, as the proportion of neurons exhibiting MI about the current f_1 becomes nonsignificant already during the second delay period of the task (Fig. 2*B*). This finding contrasts with what has been observed in PPC neurons in rats [7]. In that work, the neuronal activity exhibited MI about the previous trial f_1 but not about its value in the current trial.

PFC neurons have information on the set structure. Earlier we saw that, in agreement with the Bayesian hypothesis, the contraction bias B increases with the duration of the delay period. We now turn our attention to examining the extent to which the activities of neurons in the PFC reflect knowledge about the structure of the stimulus set, before f_2 is presented. We start by acknowledging that within set 1 (Fig. 1*B, Left*), classes 4-7 can be uniquely determined by the value of the first frequency. In other words, the information needed to make the correct choice is already available when f_1 is presented. Consequently, if the activity of PFC neurons has acquired this specific property of the set, it would be reasonable to anticipate the presence of neurons with information about the choice made at the end of the trial. To validate this hypothesis, we evaluated, with the data recorded from monkey one, the MI between each neuron's activity and the final choice (SI Appendix, Methods). As anticipated, the result confirmed that during stimulation with f_1 and the delay period a fraction of the neuronal population exhibits a significant MI on the choice (Fig. 3*A*). To provide additional support for this finding, we assessed the MI while considering classes 4-7 and classes 1-3, 8-10 separately. As

expected, we observed that for the latter set of classes, there was not a neural population exhibiting significant MI about the choice until the presentation of f_2 (compare Figs. 3*B* and 3*C*).

In the stimulus set employed with monkey two (Fig. 1*B, Right*), there is no specific value of f_1 that completely determines a particular class. Consequently, we anticipated that during stimulation with f_1 there would be no subpopulation of neurons exhibiting significant MI with respect to the choice, as confirmed by our calculations (Fig. 3*D*). Collectively, these findings suggest that the activity of PFC neurons reflects good prior knowledge of the structure of the stimulus set.

State-space Trajectories of Neural PFC activity. Our approach is based on the understanding that neural circuits encode and process information by utilizing the dynamics of neural populations within a state-space framework [5], [20], [37]–[42]. Then, to explore the relationship between behavior and neural firing activity, we conducted our analysis at the population level, using all the recorded neurons, and projected the population activity (z-score) onto the state-space directions that accounted for a minimum of 90% of the data variance (SI Appendix, Methods).

Our findings revealed that the population z-score evolved separately from each other during both the presentation of the f_1 (Figs. 4*A,D*) and the delay period (Figs. 4*B,E*). During the final stages of the delay interval, the trajectories displayed a decrease in their velocities (Figs. 4*C, Left*, and 4*F, Left*), ultimately converging towards terminal states that were dependent on the value of f_1. Notably, these terminal states exhibited an organized alignment (see the magenta line in Figs. 4*B,E*). As the delay period neared its end, there was an increase in the percentage of variance accounted for by the first principal component (PC) (Figs. 4*C, Right*, and 4*F, Right*), while the subsequent PC decreased. This finding is consistent with the notion that the trajectories approached a one-dimensional quasi-attractor region (magenta lines in Figs. 4*B,E*). The striking similarity in the trajectories observed in both monkeys (as depicted in Figs. 4*A,D*, and Figs. 4*B,E*) strongly supports the notion that the corresponding PFC neural populations are engaged in similar computational processes.

Nonetheless, the current finding does not specify which estimator of f_1 is upheld by the network during the period starting from the presentation of f_1 until it reaches the quasi-attractor. To investigate this matter and determine the geometric properties of the trajectories that encode the estimator, it is essential to establish a link between behavior and population activity. This aspect will be the focus of our upcoming analysis.

Bayesian Model and Fits of Behavioral Data. By leveraging the benefits of normative models in identifying underlying network computations associated with observed behavior, we developed a Bayesian model that incorporates detailed *a priori* knowledge (SI Appendix, Methods). We adopted the perspective of an ideal observer who makes imperfect measurements of the frequencies f_1 and f_2, possesses knowledge of the stimulus set (including the mean value and variance of the f_1 distribution), and can formulate beliefs regarding the presented stimuli. We assumed a Gaussian nature of the frequency measurements, with means f_1 and f_2, and standard deviations $\sigma_1 = W_1 f_1$ and $\sigma_2 = W_2 f_2$, respectively. Here, W_1 and W_2 represent the corresponding Weber fractions.

Since the activities of PFC neurons have no information about the frequency of the previous trial

(Fig. 2, red lines), to compare neuronal activity with behavior it will be sufficient to formulate a behavioral model in terms of the probability of observation of the current stimulus, without conditioning on classes in previous trials. In our model, in accordance with the results in Fig. 3, we assume that the ideal observer possesses good knowledge of the stimulus set. Additionally, we assume that it uses the long-term sensory history to estimate the moments of the distribution of f_1. When making an observation of the first frequency at the current trial n, this observer combines the noisy measurement of f_1 ($f'_{1,n}$) with the observation in the preceeding trial ($o_{1,n-1}$)

$$o_{1,n} = (1-\eta)\, f'_{1,n} + \eta o_{1,n-1}, \tag{1}$$

where if $\eta = 0$ there is no contribution from the long-term sensory history while if $\eta = 1$ the current stimulus does not contribute. From here it results that the variance of the observation is influenced by both the current measurement of f_1 and the stationary sensory history. In the Gaussian approximation, the likelihood $P\left(\mathcal{O}_1 | f_{1,n}\right)$ of an observation $\mathcal{O}_1$ has the mean and variance (see SI Appendix, Methods for details)

$$\mathcal{M} = f_1 + \eta(\langle f_1 \rangle - f_1), \tag{2}$$

$$\Sigma_1^2 = (1-\eta)^2\, W_1^2 f_1^2 + \eta^2 \Sigma_{1,st}^2, \tag{3}$$

$$\Sigma_{1,st}^2 = \frac{1-\eta}{1+\eta}\left(W_1^2 \left\langle f_1^2 \right\rangle + \sigma_{f_1}^2\right), \tag{4}$$

where $\Sigma_{1,st}^2$ is the variance of the stationary distribution of f_1 observations and $\left\langle f_1^2 \right\rangle$ and $\sigma_{f_1}^2$ are the second moment and variance of f_1 in the stimulus set. It should be noted that both $\mathcal{M}$ and Σ_1 play a role in the contraction bias phenomenon. The mean explicitly expresses how f_1 is pulled towards its own mean, and this contraction is proportional to the difference between the actual value of the interval and the mean of f_1. Qualitatively, a Bayesian contribution to the contraction bias occurs when the prior has a more substantial impact on the posterior of the first frequency compared to the second frequency [19], particularly when Σ_1 exceeds σ_2.

To fit the model parameters (η, W_1 and W_2), for each monkey we considered the corresponding accuracy curve (Figs. 1*C, Left* and *Right*), and used it to define the similarity between model and data (SI Appendix, Methods). The fits and the optimal parameter values for the two animals are shown in Fig. 5.

Relationship Between PFC State-Space Geometry and Behavior. We are now investigating the link between state-space geometry and behavior. In this task, the neural activity is expected to represent an estimate of the frequency f_1 during its stimulation and maintain a representation throughout the delay period until its conclusion. Consequently, we anticipate that during these epochs, an estimate of f_1 may be reflected in the geometric properties of the trajectories [14], [20], [29]. Our observations indicate that during the presentation of the stimulus, the trajectories evolve separately while maintaining finite relative distances between them (Fig. 4*A,D*), during the delay period still remain distant from each other (Fig. 4*B,E*), ultimately converging to a quasi-attractor region at the end of the delay epoch (Fig. 4*C,F*). Based on these findings, we hypothesize that estimations

of f_1 are represented: firstly, in the relative distances between trajectories as they evolve under the influence of the first stimulus, then, during the delay period, in the relative distances maintained by the trajectories, and finally, in the distances between consecutive terminal states on the quasi-attractor.

However, the specific nature of the estimator used by the cortical network remains unknown, as only the comparison choices are reported by the animals in the discrimination task. This differs from magnitude reproduction tasks, where the reproduced magnitude can be considered as a noisy version of the reported estimate [28]. Due to the limited information regarding the employed estimator for f_1, we explored three alternatives: the true f_1, the Bayesian estimator ($f_{1,Bayes}$), and the Maximum a Posteriori (MAP) estimator ($f_{1,MAP}$) (SI Appendix, Methods). We examined how these estimators were associated with the relative distances between neural trajectories during the stimulation and short-term memory periods, as well as with the distances between terminal states on the quasi-attractor at the end of the delay epoch.

To evaluate the estimator that most accurately captured the relative distances between neural trajectories during the stimulus presentation phase, we employed the following procedure. For monkey one, the relative distance $D_j^{f_1}(\mathrm{t})$ between each trajectory (indexed as $j = 1, \ldots, 7$) and the reference trajectory (f_1=22 Hz) was determined using the KiNeT methodology [13] (SI Appendix, Methods). The temporal evolution of these distances (Fig. 6*A*), as well as their temporal average values $D_j^{f_1} \equiv \langle D_j^{f_1}(t) \rangle$ (inset), appeared warped as a function of f_1. We subsequently investigated whether $D_j^{f_1}$ more effectively encoded the true f_1 values or their Bayesian or MAP estimators. To compare these alternatives, we individually fitted a linear expression ($aD_j^{f_1} + b$) to each of these three estimates (Fig. 6*B*) and assessed their root mean square errors (RMSEs) [14] (Fig. 6*C*). The results strongly supported the hypothesis favoring the Bayesian estimator. Subsequently, we conducted a comparable analysis focusing on the relative distances during the delay period, denoted as $D_j^d(\mathrm{t})$ ($j = 1, \ldots, 7$). Once again, the obtained results provided support for the hypothesis that favors the Bayesian estimator (Fig. 6*D-F*). In order to assess the performance of the estimators on the quasi-attractor, we calculated the distances δ_j between the terminal points of the trajectories within the attractor region. Similar to the previous analyses, we conducted a comparative assessment using the true f_1 value and the Bayesian and MAP estimators. Once again, the results strongly supported the Bayesian estimator (Figs. 6*G-I*). Collectively, the outcomes of these three tests consistently suggest that the representation of the first frequency is effectively captured by its Bayesian estimator. This conclusion was further supported by conducting a similar analysis on the data from monkey two (Supplementary Fig. 3).

It is noteworthy that the dynamics of population activity directly transformed the true value of f_1 into its Bayesian estimator. The computations necessary to integrate prior knowledge and current sensory information were not explicitly carried out. This phenomenon resembles what is observed in time interval estimation experiments [14], [29] and in simulated recurrent neural networks trained to discriminate between the durations of two time intervals [20].

Discussion

Previous studies have suggested that the contraction bias, a phenomenon in perception, can be attributed to the integration of prior knowledge with noisy sensory observations [10], [19], [20], [43]. In

the current work, we aimed to investigate the validity of this proposition and explore how Bayesian integration manifests in the population activity of PFC neurons. To accomplish this, we examined neural recordings from monkeys engaged in a discrimination task involving two tactile frequencies within the flutter range [34], [35]. Throughout the task, we observed distinct geometric characteristics in the state-space trajectories, which consistently represented the first frequency in a compressed manner. This compression was evident during both the stimulus presentation and the subsequent delay period, as reflected in the relative distances between trajectories. Towards the end of the delay period, these compressions converged towards a one-dimensional quasi-attractor. Importantly, these observed compressions conform with the predictions of a Bayesian model, indicating that they underlie the manifestation of the contraction bias in the activity of PFC neurons.

Our findings can be compared to recent experimental and computational studies investigating neural activity in other tasks that exhibit similar perceptual biases. In a temporal reproduction task, it was found that a curved one-dimensional manifold exhibited a compressed representation of the time interval, which was consistent with the compression predicted by its Bayesian estimator [13], [14], [29]. Computational studies have demonstrated that the contraction bias can also occur in time interval discrimination tasks [20]. To investigate the underlying mechanisms of this bias, the authors trained recurrent neural networks composed of spiking neurons on such tasks. In the trained networks, the bias arose from the integration of prior knowledge and the current observation. Interestingly, similar to the findings in the temporal reproduction task [29], a curved one-dimensional manifold displayed a compressed representation of the first time interval, aligning with the predictions of the Bayesian estimator. Moreover, during the delay interval, the state-space trajectories exhibited similar compressions in their relative distances.

An intriguing finding emerged from our study, revealing that neurons in the PFC encode information specifically about the current first stimulus, while remaining insensitive to stimuli presented in previous trials. This stands in stark contrast to observations in the posterior parietal cortex of rodents, where the opposite pattern was observed [7]. The contraction bias may manifest in various cortical areas in different ways, which could reflect how it is constructed and expressed in behavior. To gain a comprehensive understanding of the emergence of this bias, it is crucial to extend the current study to other implicated brain areas and examine the relationship between neuronal population activity and behavior in each case.

In contrast to a previous study focusing on the decisional stage of PFC neurons [21], our research investigates the earlier stages of the task, namely the stimulation and maintenance epochs. Our findings reveal that the Bayesian estimator is consistently reflected in the state-space geometry from the presentation of the first stimulus until the end of the delay period. This challenges the hypothesis that the estimator is obtained by sequentially combining prior information with the current stimulus but aligns with the findings observed in other tasks [29]. Furthermore, our investigation of PFC neural activity demonstrates that the contraction bias is already evident upon the presentation of the first stimulus, suggesting that it may originate from a sensory area where it begins to take shape. This observation is in line with the findings of Preuschhof et al. (2010) [30], supporting the integration of prior information in early brain areas such as S2. Overall, our study provides new insights into the temporal dynamics and neural mechanisms underlying the contraction bias, contributing to a deeper understanding of this phenomenon.

References

[1] D. C. Knill and W. Richards, *Perception as Bayesian inference*. Cambridge University Press, 1996.

[2] C. Summerfield and F. P. De Lange, "Expectation in perceptual decision making: Neural and computational mechanisms," *Nature Reviews Neuroscience*, vol. 15, no. 11, pp. 745–756, 2014.

[3] F. Carnevale, V. de Lafuente, R. Romo, and N. Parga, "Internal signal correlates neural populations and biases perceptual decision reports," *Proceedings of the National Academy of Sciences*, vol. 109, no. 46, pp. 18 938–18 943, 2012.

[4] K. M. Bausenhart, O. Dyjas, and R. Ulrich, "Temporal reproductions are influenced by an internal reference: Explaining the vierordt effect," *Acta psychologica*, vol. 147, pp. 60–67, 2014.

[5] F. Carnevale, V. de Lafuente, R. Romo, O. Barak, and N. Parga, "Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty," *Neuron*, vol. 86, no. 4, pp. 1067–1077, 2015.

[6] Z. Shi and D. Burr, "Predictive coding of multisensory timing," *Current Opinion in Behavioral Sciences*, vol. 8, pp. 200–206, 2016.

[7] A. Akrami, C. D. Kopec, M. E. Diamond, and C. D. Brody, "Posterior parietal cortex represents sensory history and mediates its effects on behaviour," *Nature*, vol. 554, no. 7692, pp. 368–372, 2018.

[8] F. P. De Lange, M. Heilbron, and P. Kok, "How do expectations shape perception?" *Trends in cognitive sciences*, vol. 22, no. 9, pp. 764–779, 2018.

[9] A. Hermoso-Mendizabal, A. Hyafil, P. E. Rueda-Orozco, S. Jaramillo, D. Robbe, and J. De la Rocha, "Response outcomes gate the impact of expectations on perceptual decisions," *Nature communications*, vol. 11, no. 1, pp. 1–13, 2020.

[10] S. Sarno, M. Beirán, J. Falcó-Roget, *et al.*, "Dopamine firing plays a dual role in coding reward prediction errors and signaling motivation in a working memory task," *Proceedings of the National Academy of Sciences*, vol. 119, no. 2, e2113311119, 2022. DOI: https://doi.org/10.1073/pnas.2113311119.

[11] K. Thurley, "Magnitude estimation with noisy integrators linked by an adaptive reference," *Frontiers in Integrative Neuroscience*, vol. 10, p. 6, 2016.

[12] M. Jazayeri and M. N. Shadlen, "A neural mechanism for sensing and reproducing a time interval," *Current Biology*, vol. 25, no. 20, pp. 2599–2609, 2015.

[13] E. D. Remington, D. Narain, E. A. Hosseini, and M. Jazayeri, "Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics," *Neuron*, vol. 98, no. 5, pp. 1005–1019, 2018.

[14] S. W. Egger, E. D. Remington, C.-J. Chang, and M. Jazayeri, "Internal models of sensorimotor integration regulate cortical dynamics," *Nature neuroscience*, vol. 22, no. 11, pp. 1871–1882, 2019.

[15] M. Beiran, N. Meirhaeghe, H. Sohn, M. Jazayeri, and S. Ostojic, "Parametric control of flexible timing through low-dimensional neural manifolds," *Neuron*, 2023.

[16] H. L. Hollingworth, "The central tendency of judgment," *The Journal of Philosophy, Psychology and Scientific Methods*, vol. 7, no. 17, pp. 461–469, 1910.

[17] O. Dyjas, K. M. Bausenhart, and R. Ulrich, "Trial-by-trial updating of an internal reference in discrimination tasks: Evidence from effects of stimulus order and trial sequence," *Attention, Perception, & Psychophysics*, vol. 74, no. 8, pp. 1819–1841, 2012.

[18] O. Dyjas and R. Ulrich, "Effects of stimulus order on discrimination processes in comparative and equality judgements: Data and models," *Quarterly Journal of Experimental Psychology*, vol. 67, no. 6, pp. 1121–1150, 2014.

[19] P. Ashourian and Y. Loewenstein, "Bayesian inference underlies the contraction bias in delayed comparison tasks," *PloS one*, vol. 6, no. 5, e19551, 2011.

[20] L. Serrano-Fernandez, M. Beiran, and N. Parga, "Emergent perceptual biases from state-space geometry in spiking recurrent neural networks trained to discriminate time intervals," *bioRxiv*, pp. 2022–11, 2022.

[21] D. Benozzo, L. Ferrucci, and A. Genovesio, "Effects of contraction bias on the decision process in the macaque prefrontal cortex," *Cerebral Cortex*, vol. 33, no. 6, pp. 2958–2968, 2023.

[22] J. Fischer and D. Whitney, "Serial dependence in visual perception," *Nature neuroscience*, vol. 17, no. 5, pp. 738–743, 2014.

[23] C. Papadimitriou, A. Ferdoash, and L. H. Snyder, "Ghosts in the machine: Memory interference from the previous trial," *Journal of neurophysiology*, vol. 113, no. 2, pp. 567–577, 2015.

[24] D. P. Bliss, J. J. Sun, and M. D'Esposito, "Serial dependence is absent at the time of perception but increases in visual working memory," *Scientific reports*, vol. 7, no. 1, p. 14 739, 2017.

[25] W. J. Ma and M. Jazayeri, "Neural coding of uncertainty and probability," *Annual review of neuroscience*, vol. 37, pp. 205–220, 2014.

[26] Z. Shi, R. M. Church, and W. H. Meck, "Bayesian optimization of time perception," *Trends in cognitive sciences*, vol. 17, no. 11, pp. 556–564, 2013.

[27] F. H. Petzschner, S. Glasauer, and K. E. Stephan, "A bayesian perspective on magnitude estimation," *Trends in cognitive sciences*, vol. 19, no. 5, pp. 285–293, 2015.

[28] M. Jazayeri and M. N. Shadlen, "Temporal context calibrates interval timing," *Nature neuroscience*, vol. 13, no. 8, pp. 1020–1026, 2010.

[29] H. Sohn, D. Narain, N. Meirhaeghe, and M. Jazayeri, "Bayesian computation through cortical latent dynamics," *Neuron*, vol. 103, no. 5, pp. 934–947, 2019.

[30] C. Preuschhof, T. Schubert, A. Villringer, and H. R. Heekeren, "Prior information biases stimulus representations during vibrotactile decision making," *Journal of Cognitive Neuroscience*, vol. 22, no. 5, pp. 875–887, 2010.

[31] R. S. Van Bergen and J. F. Jehee, "Probabilistic representation in human visual cortex reflects uncertainty in serial decisions," *Journal of Neuroscience*, vol. 39, no. 41, pp. 8164–8176, 2019.

[32] M. Fritsche, E. Spaak, and F. P. De Lange, "A bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception," *Elife*, vol. 9, e55389, 2020.

[33] G. M. Cicchini, K. Mikellidou, and D. C. Burr, "The functional role of serial dependence," *Proceedings of the Royal Society B*, vol. 285, no. 1890, p. 20 181 722, 2018.

[34] R. Romo, C. D. Brody, A. Hernández, and L. Lemus, "Neuronal correlates of parametric working memory in the prefrontal cortex," *Nature*, vol. 399, no. 6735, pp. 470–473, 1999.

[35] C. D. Brody, A. Hernández, A. Zainos, and R. Romo, "Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex," *Cerebral cortex*, vol. 13, no. 11, pp. 1196–1207, 2003.

[36] O. Raviv, M. Ahissar, and Y. Loewenstein, "How recent history affects perception: The normative approach and its heuristic approximation," *PLoS computational biology*, vol. 8, no. 10, e1002731, 2012.

[37] S. Vyas, M. D. Golub, D. Sussillo, and K. V. Shenoy, "Computation through neural population dynamics," *Annual Review of Neuroscience*, vol. 43, p. 249, 2020.

[38] S. Saxena and J. P. Cunningham, "Towards the neural population doctrine," *Current opinion in neurobiology*, vol. 55, pp. 103–111, 2019.

[39] R. B. Ebitz and B. Y. Hayden, "The population doctrine in cognitive neuroscience," *Neuron*, vol. 109, no. 19, pp. 3055–3068, 2021.

[40] J. Wang, D. Narain, E. A. Hosseini, and M. Jazayeri, "Flexible timing by temporal scaling of cortical responses," *Nature neuroscience*, vol. 21, no. 1, pp. 102–110, 2018.

[41] F. Mastrogiuseppe and S. Ostojic, "Linking connectivity, dynamics, and computations in low-rank recurrent neural networks," *Neuron*, vol. 99, no. 3, pp. 609–623, 2018.

[42] M. Jazayeri and S. Ostojic, "Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity," *Current opinion in neurobiology*, vol. 70, pp. 113–120, 2021.

[43] I. Lieder, V. Adam, O. Frenkel, S. Jaffe-Dax, M. Sahani, and M. Ahissar, "Perceptual bias reveals slow-updating in autism and fast-forgetting in dyslexia," *Nature neuroscience*, vol. 22, no. 2, pp. 256–264, 2019.

[44] R. Romo, C. D. Brody, A. Hernández, and L. Lemus, "Single-neuron spike train recordings from macaque prefrontal cortex during a somatosensory working memory task," *CRCNS.org*, 2016.

Methods

Experimental Data. Data from two monkeys, previously trained to perform a somatosensory vibrotactile frequency discrimination task [34], [35], [44], were utilized in this study. The monkeys, named RR13 and RR15 (monkey 1 and monkey 2, respectively), underwent training on distinct sets of vibrotactile frequency classes. The task was executed with a high level of accuracy, and our analysis specifically focused on correct trials. Both monkeys were recorded with a 3-second delay interval, while additional recordings were obtained from monkey RR13 with a 6-second delay interval. However, only monkey two had to wait for 3 s before reporting its decision. 474 neurons from monkey one and 170 cells from monkey two were considered for this study. Monkey one neurons were recorded over 116 experimental sessions while monkey two units proceeded from 56 sessions. In each session, there were around 10 blocks of trials. Each block had as many trials as there are classes in the set and classes were randomly chosen without replacement. A total of 10,540 trials of neural recording were considered for monkey one and 5,972 trials for monkey two.

Data analyses: Behavior. *Accuracy curve.* In order to examine the discrimination properties, we plotted the percentage of correct trials as a function of the index of enumerated classes (Fig. 1*C*).
Phenomenological measure of the contraction bias. Denoting the probability of deciding f_1 smaller (S) given a presented class C as $P(S|C)$, a phenomenological measure B of the contraction bias was defined as the root-mean-square deviation of $P(S|C)$ on all classes, relative to classes 3 and 8 in case of monkey one (Fig. 1*B, Left*)

$$B_{M1}^2 = \frac{1}{10}\left[\sum_{uC=1}^{5}\Big(P(S|uC) - P(S|3)\Big)^2 + \sum_{lC=6}^{10}\Big(P(S|lC) - P(S|8)\Big)^2\right], \tag{5}$$

and relative to classes 3 and 9 in case of monkey two (Fig. 1*B, Right*)

$$B_{M2}^2 = \frac{1}{12}\left[\sum_{uC=1}^{6}\Big(P(S|uC) - P(S|3)\Big)^2 + \sum_{lC=7}^{12}\Big(P(S|lC) - P(S|9)\Big)^2\right], \tag{6}$$

where uC and lC denote classes on the upper and the lower diagonals of the stimulus set, respectively. Note that the value of B^2 corresponding to the minimum possible bias equals 0, while for the maximum possible bias it is 3/5 or 1/2 in case of monkeys one or two, respectively.

Data analyses: Neural Activity. *Firing Rates.* The trial-averaged firing rates were estimated by conditioning on f_1. The analysis window, which encompassed the presentation of the first stimulus and the delay period, slid in uniform steps. The analysis intervals were initially positioned at the onset and the offset of f_1, respectively. The temporal window size was set to 250 ms, with a sliding interval of 25 ms, except during the analysis of the attractor (explained below), where a window of 100 ms with a sliding interval of 10 ms was employed.
Mutual Information (MI). The information that neurons carried about the first stimulus and monkey's choice ($f_1 < f_2$ or $f_1 > f_2$) was measured using MI [45]. This analysis let us to extract information from the single neuron response, quantified as the number of spikes, r_t, fired in the t-th time bin. The quantities $r = \{r_1, ..., r_L\}$, where L is the total number of bins, were computed using a temporal

window of size 250 ms, sliding in steps of 25 ms. R collects all values of r. If the information sought is about the first stimulus, all possible values of f_1 are collected in F_1. With this, the MI about f_1 carried by a neuron is

$$I(F_1|R) = \sum_{f_1} P(f_1) \sum_{r} P(r|f_1) \log_2 \frac{P(r|f_1)}{P(r)}, \tag{7}$$

here $P(f_1)$ is the f_1 probability and $P(r|f_1)$ is the conditional probability of observing a response r given that the presented stimulus was f_1. The evaluation of the MI about particular sets of values of f_1 or choice follows an analogous procedure. Equation (7) was computed using a MATLAB-based toolbox given by [46]. This toolbox introduces a correction to a bias, due to the limited-sampling of the neural response [47], by means of an extrapolation quadratic in the number of trials, N_{trial},

$$bias = \frac{a}{N_{trial}} + \frac{b}{(N_{trial})^2}, \tag{8}$$

here, a and b are free parameters.

Also, this toolbox analysed the MI significance from a bootstrap procedure that pairs randomly stimuli and responses in order to destroy their correlation [46]. A single neuron's MI at a specific time was considered significant (Supplementary Figs. 1 and 2) if exceeded the sum of its corresponding bootstrapped MI and the standard deviation obtained from the permutations used in the bootstrap procedure. Subsequently, the percentage of neurons exhibiting significant MI over time was calculated. (Figs. 2 and 3). To eliminate causality between responses of consecutive trials, we employed a second shuffling procedure. This involved pooling together the firing rates of all neurons from all trials and carrying out 25 random permutations without replacement. Using this shuffled dataset, we calculated, bin-by-bin, the mean and confidence intervals of the percentage of neurons with significant MI (horizontal shadings in Figs. 2 and 3).

Data analyses: State-space Analysis and Geometry. *Principal Component Analysis.* We employed principal component (PC) analysis to analyze the firing rates averaged across trials and identify a low-dimensional space that captured a minimum of 90% of the variance in the data. The neural activity, conditioned on the first stimulus, exhibited curved trajectories within this reduced space. The speed associated with these trajectories was computed as the distance covered by a sliding time window of 250 ms, with a 25 ms increment (Fig. 4*C,F*).

KiNeT Method for relative distances. We utilized the Kinematic analysis of Neural Trajectory (KiNeT) methodology [13] to determine the relative distance between a pair of trajectories, denoted as i and j. In this analysis, one trajectory was chosen as the reference (ref=i), and the corresponding states between the reference trajectory and the other trajectory (j) were defined. Specifically, at each time point t (aligned to SF1), if the reference trajectory was in state $s_{ref}(t)$, the corresponding state $s_j(t)$ on trajectory j was identified as the one with the minimum Euclidean distance to $s_{ref}(t)$. This allowed us to compute the relative distance $D_j(t)$ between the two trajectories. For monkey one, the reference trajectory was set to ref=22 Hz, while for monkey two, it was ref=24 Hz. The relative distances were computed in a PC space capturing at least 90% of the total variance during the presentation of f_1 (Fig. 6*A* and Supplementary Fig. 3*A*) and the delay period (Fig. 6*D* and Supplementary Fig. 3*D*). To estimate the standard errors, we employed bootstrap resampling with replacement, generating 100

resamples. The mean ± 95% confidence intervals (CI) were then calculated from the PC analysis and the corresponding KiNeT analysis performed on each resample.

Bayesian Model. *Formulation and Likelihood function.* Behavioral data were fitted using a Bayesian model. We assumed that at the current trial n the Bayesian observer makes an observation of the first frequency ($o_{1,n}$), given by a combination of the current-trial noisy measurement of f_1 ($f'_{1,n}$) and the observation in the preceding trial ($o_{1,n-1}$)

$$o_{1,n} = (1-\eta)\, f'_{1,n} + \eta o_{1,n-1}. \tag{9}$$

Here η is a parameter that measures the contribution of $o_{1,n-1}$ (if $\eta = 0$ there is no contribution from sensory history while if $\eta = 1$ the current stimulus does not contribute).
We assumed that the noisy measurement of f_1 can be described by a normal distribution

$$f'_{1,n} \sim \mathcal{N}\left(f_{1,n}, W_1^2 f_{1,n}^2\right), \tag{10}$$

where W_1 is the Weber fraction, a parameter of the Bayesian model to be determined. In addition, we assumed that all previous trials constitute a Gaussian stationary regime (the long-term sensory history). The mean and variance of the stationary variable are obtained by self-consistency arguments

$$M_{1,st} = \langle f_1 \rangle, \tag{11}$$

$$\Sigma_{1,st}^2 = \frac{1-\eta}{1+\eta}\left(W_1^2 \left\langle f_1^2 \right\rangle + \sigma_{f_1}^2\right), \tag{12}$$

where $\langle f_1 \rangle$ and $\sigma_{f_1}^2$ are the mean and variance of the prior probability of f_1. Within this Gaussian approximation the observation in the current trial is also a Gaussian variable, $\mathcal{O}_1$, with mean $\mathcal{M}$ and variance Σ_1^2,

$$\mathcal{O}_1 \sim \mathcal{N}\left(\mathcal{M}, \Sigma_1^2\right). \tag{13}$$

The distribution of this variable defines the likelihood of an observation $\mathcal{O}_1$, that we denote by $P\left(\mathcal{O}_1 | f_{1,n}\right)$. The mean and variance of the observation probability are given by

$$\mathcal{M} = f_1 + \eta(\langle f_1 \rangle - f_1), \tag{14}$$

$$\Sigma_1^2 = (1-\eta)^2 W_1^2 f_1^2 + \eta^2 \Sigma_{1,st}^2. \tag{15}$$

Beliefs. The computation of posterior probabilities closely aligns with a Bayesian model used for tactile frequency discrimination (Sarno et al., 2022). However, the key distinction lies in the likelihoods, which are determined by the derived expressions mentioned earlier. In our analysis, we represent the discrete prior probabilities of the frequency f_1 as $P(f_1^i)$, where the index i corresponds to the possible values of f_1 within the stimulus set (see Fig. 1*B*). For monkey one, i ranges from 1 to 10, while for monkey two, it ranges from 1 to 12. According to Bayes' theorem, the posterior distribution $P(f_1^i|\mathcal{O}_1)$ is obtained by combining the noisy observation with the prior distribution

$$b_1(f_1^i) \equiv P(f_1^i|\mathcal{O}_1) \propto P(\mathcal{O}_1|f_1^i)P(f_1^i). \tag{16}$$

The left hand side indicates that this probability is also known as the belief $b_1(f_1^j)$ about f_1. Likewise, the posterior probability $P(d_2^l|\mathcal{O}_2)$, or belief $b_2(d_2^l)$ about f_2 is

$$b_2(f_2^l) \equiv P(f_2^l|\mathcal{O}_2) \propto P(\mathcal{O}_2|f_2^l)P(f_2^l). \tag{17}$$

where $P(f_2^l)$ is the discrete prior probability of f_2 (l labels the possible values of f_2 in the stimulus set, Fig. 1*B*).

Let us denote class k as $C^k = (f_1^k, f_2^k)$, where $k = 1, \ldots, 10$ or 12 (for monkey one or two, respectively) is a class label (Fig. 1*B*). The joint posterior probabilities, or joint beliefs, indicating the likelihood that class C^k was presented in the current trial given the observations $\mathcal{O}_1$ and $\mathcal{O}_2$, can be expressed as follows [10]

$$B(C^k|\mathcal{O}_1, \mathcal{O}_2) = \frac{P(\mathcal{O}_2)}{P(\mathcal{O}_2|\mathcal{O}_1)} \frac{P(f_2 = f_2^k|f_1 = f_1^k)}{P(f_2^k)} b_1(f_1^k) b_2(f_2^k). \tag{18}$$

The initial quotient is a normalization factor. $P(f_2 = f_2^k|f_1 = f_1^k)$ represents the transition probability matrix indicating the probability of f_2 being presented after f_1. If the animals possess a good understanding of the set's structure, the transition matrix elements that are non-zero exclusively correspond to the classes within the set.

The Bayesian observer utilizes the relative beliefs to make a decision between the choices "$f_1 < f_2$" and "$f_1 > f_2$". These beliefs are calculated by aggregating the class beliefs $B(C^k|\mathcal{O}_1^*, \mathcal{O}_2)$ for classes k associated with a particular choice.

$$\begin{cases} b(L) = \sum_{f_1 > f_2} B(C^k|\mathcal{O}_1, \mathcal{O}_2) & \text{belief about } f_1 \text{ Larger (L="}f_1 > f_2\text{")}, \\ b(S) = 1 - b(L) & \text{belief about } f_1 \text{ Smaller (S="}f_1 < f_2\text{")}. \end{cases} \tag{19}$$

Probability of deciding "$f_1 < f_2$". Given the information $\mathcal{O}_1$ and $\mathcal{O}_2$, when the current class is $C^k = (f_1^k, f_2^k)$, the Bayesian observer selects the option "$f_1 < f_2$" with the probability

$$P(S|C^k) = \sum_{\mathcal{O}_1, \mathcal{O}_2} P(\mathcal{O}_1|f_1^k) P(\mathcal{O}_2|f_2^k) G(b(S)), \tag{20}$$

where $G(b(S))$ equals 1 if $b(S) > 0.5$ and and it equals zero otherwise [10], [19].

Fitting the Bayesian Model to Data. The model was fitted by minimizing mean squared differences between the accuracy curve derived from Eq. (20) and the corresponding behavioral curve obtained from experimental data (Fig. 1*C*, left and right). Minimizations were done using MATLAB's *simulannealbnd* function which implements a simulated annealing algorithm.

Estimation of f_1 and its relationship with state-space geometry. To assess whether geometric features of the trajectories encoded an estimate of the first stimulus, we conducted separate comparisons between three estimators: the true f_1, the Bayesian estimator $f_{1,Bayes}$, and the maximum a posteriori (MAP) estimator $f_{1,MAP}$. We evaluated these estimators using measures of relative distances during

both the presentation of f_1 and the delay period, as well as the spatial structure of the attractor. To accomplish this, we computed the estimators in single trials by simulating the model (using eqs. (14, 15)). Then, we obtained the means $f_{1,Bayes}$ and $f_{1,MAP}$ as corresponding averages over trials.

Bayesian and MAP estimators. For a given trial n, the Bayesian estimator of $f_{1,n}$ ($f_{1,n}^{Bayes}$) results from a loss function; taking it as the squared error

$$(f_{1,n}^{Bayes} - f_{1,n})^2,$$

the Bayesian estimator is the mean of f_1 weighted with the posterior probability $b_1(f_{1,n}^i)$,

$$f_{1,n}^{Bayes} = \sum_{i=1}^{N_1} f_1^i \cdot b_1(f_{1,n}^i), \tag{21}$$

where N_1 represents the number of possible values for f_1, which is equal to 7 for monkey one and 6 for monkey two.

The MAP estimator ($f_{1,n}^{MAP}$) is the value of f_1 that maximizes the belief,

$$f_{1,n}^{MAP} = \arg\max_{f_1} b_1(f_{1,n}^i). \tag{22}$$

After determining the model parameters, we conducted simulations consisting of 10,000 trials and computed the corresponding Bayesian estimator (Eq. (21)) and MAP estimator (Eq. (22)). Taking the trial averages of these estimators, we obtained $f_{1,Bayes}$ and $f_{1,MAP}$, respectively.

Mean relative distances. To determine which of the three estimators provided the best explanation for the relative distances between trajectories (see KiNeT section above), we compared the following three hypotheses during the presentation of the first stimulus (Fig. 6*B* and Supplementary Fig. 3*B*) and the delay period (Fig. 6*E* and Supplementary Fig. 3*E*):

$$f_1;\ f_{1,Bayes};\ f_{1,MAP} = aD_j + b. \tag{23}$$

Here D_j ($j = 1, \ldots, N_1$. Again, $N_1 = 7$ in case of monkey one and $N_1 = 6$ in case of monkey two) is the temporal mean of the relative distances [14] during either the presentation of the first stimulus, $D_j^{f_1}$, (inset from Fig. 6*A* and Supplementary Fig. 3*A*) or the delay period, D_j^d, (inset from Fig. 6*D* and Supplementary Fig. 3*D*). For both intervals, we utilized an analysis window of 250 ms with a sliding interval of 25 ms. We retained the number of PCs necessary to explain a minimum of 90% of the variance.

Distances between terminal states on the attractor. To examine which estimator is encoded by the attractor, we evaluated the Euclidean distances δ_j ($j = 1, \ldots, N_{f_1}$) between the final states of each f_1 trajectory and the final state of the trajectory with the smallest f_1 (Fig. 6*G* and Supplementary Fig. 3*G*) and made hypotheses analogous to those in Eq. (23). This test (Fig. 6*H* and Supplementary Fig. 3*H*) took place in the last 100 ms before the end of the delay period. The attractor was considered to be embedded in a 2-dimensional subspace, which accounted for 84% of variance in monkey one and 93% of variance in monkey two.

In the three tests the goodness of the fits was compared using their Root Mean Square Errors (RMSEs; Figs. 6 *C,F,I* for monkey one and Supplementary Figs. 3 *C,F,I* for monkey two). Significance was assessed using a bootstrap procedure. We generated 100 resamples by randomly sampling with replacement from the original neural population. Subsequently, we conducted the tests on each resample and calculated the respective dispersions.

Code and Data Availability

Matlab codes, as well as the raw data necessary for full replication of the figures will be available upon publication.

References

[34] R. Romo, C. D. Brody, A. Hernández, and L. Lemus, "Neuronal correlates of parametric working memory in the prefrontal cortex," *Nature*, vol. 399, no. 6735, pp. 470–473, 1999.

[35] C. D. Brody, A. Hernández, A. Zainos, and R. Romo, "Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex," *Cerebral cortex*, vol. 13, no. 11, pp. 1196–1207, 2003.

[44] R. Romo, C. D. Brody, A. Hernández, and L. Lemus, "Single-neuron spike train recordings from macaque prefrontal cortex during a somatosensory working memory task," *CRCNS.org*, 2016.

[45] C. E. Shannon, "A mathematical theory of communication," *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.

[46] C. Magri, K. Whittingstall, V. Singh, N. K. Logothetis, and S. Panzeri, "A toolbox for the fast information analysis of multiple-site lfp, eeg and spike train recordings," *BMC neuroscience*, vol. 10, pp. 1–24, 2009.

[47] A. Treves and S. Panzeri, "The upward bias in measures of information derived from limited data samples," *Neural Computation*, vol. 7, no. 2, pp. 399–407, 1995.

Acknowledgements

This work was supported by the grant PGC2018-101992-B-I00 from the Spanish Ministry of Science, Innovation and Universities.

Author contributions

N.P. and R.R. designed research. L.S.-F. implemented the numerical computations and analyzed the data. M.B. and N.P. developed the Bayesian model. L.S.-F. and N.P. wrote the paper.

Figures

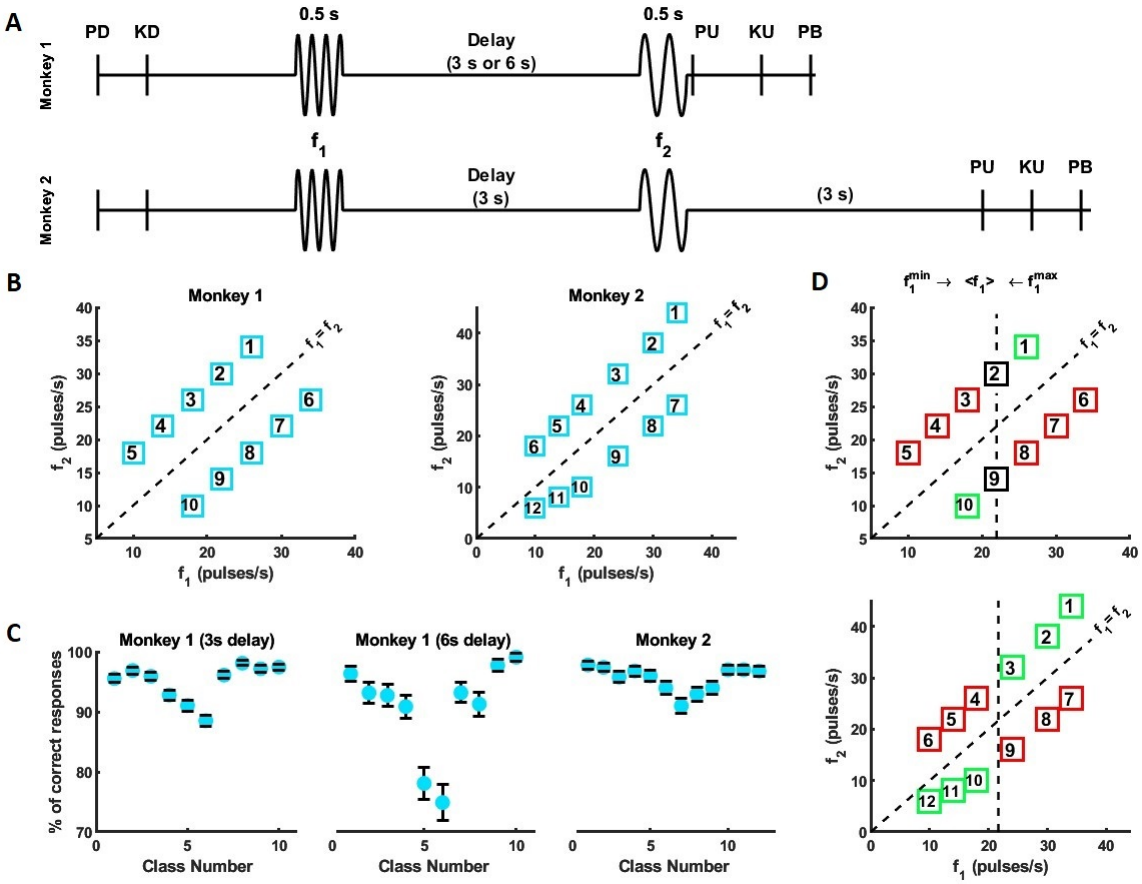

Fig. 1. Task, stimulus sets and behavior. *(A)* Schemes with the timeline of the frequencies discrimination task for monkey one *(Upper)* and monkey two *(Bottom)*. A mechanical probe (Probe Down, PD) is applied to the smooth skin of a hairless finger on a restrained hand. To initiate each trial, the monkey presses a key with its free hand (Key Down, KD), indicating its readiness. A 500 ms long mechanical vibration (first stimulus, f_1), followed by a delay period of 3 or 6 seconds, is then delivered. Subsequently, another 500 ms long mechanical vibration (second stimulus, f_2) is applied. Once the mechanical probe is lifted (Probe Up, PU), the monkey is allowed to report its choice regarding the comparison between the two vibrations. At this point, the monkey's free hand is no longer on the readiness key (Key Up, KU), and it proceeds to press one of two push-buttons (PB) to indicate its choice. The only difference between both schemes is a 3 s delay period between the offset of f_2 and PU in the task of monkey two. *(B)* Stimulus set of possible classes or pairs (f_1, f_2) for monkey one (set 1; *(Left)*) and monkey two (set 2; *(Right)*). *(C)* Percentage of correct trials from monkey one with 3s delay *(Left)* and with 6s delay *(Middle)* and from monkey two *(Right)*. Error bars were computed with 1,000 bootstrap resamples. *(D)* Contraction bias refers to the tendency of the first frequency (f_1) to be coded in a contracted manner towards its distribution mean. This bias leads to errors in certain classes (red squares), while favoring others (green squares). *Upper* panel corresponds to set 1, while *bottom* panel corresponds to set 2.

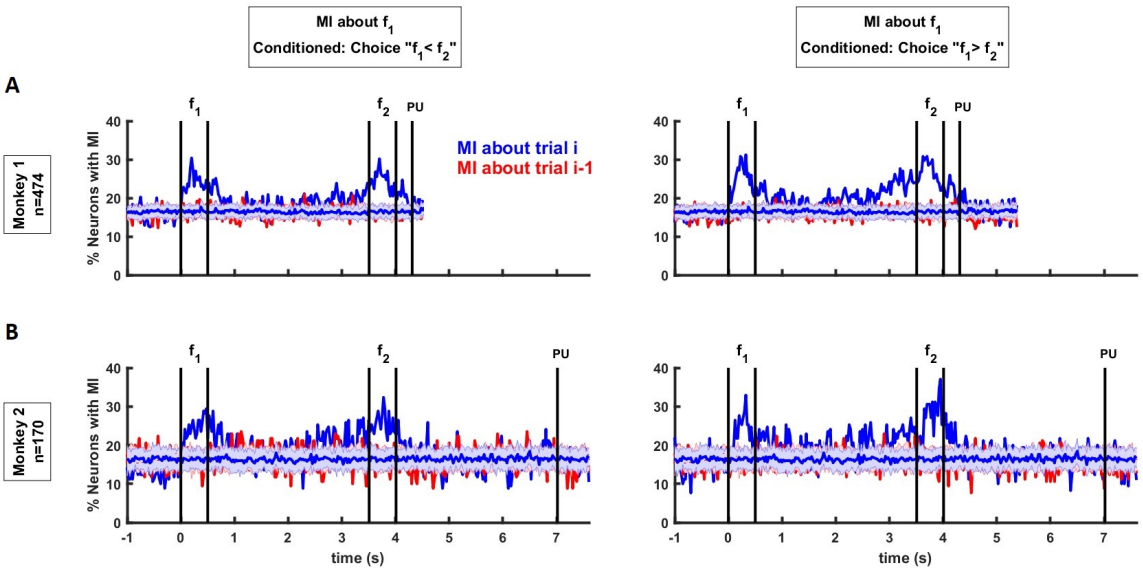

Fig. 2. PFC neurons exhibit significant mutual information (MI) about current trial f_1 but not on the previous trial f_1. *(A)* Percentage of neurons from monkey one with significant MI about the current (blue line) and previous (red line) first stimulus, conditioning on correct trials and on choice (*Left*: choice "$f_1 < f_2$"; *Right*: choice "$f_1 > f_2$"). Horizontal blue and red shadings refer to the mean and standard deviation computed from shuffled data (see SI Appendix, Methods) for current and previous trial, respectively. *(B)* Same MI analysis for monkey two.

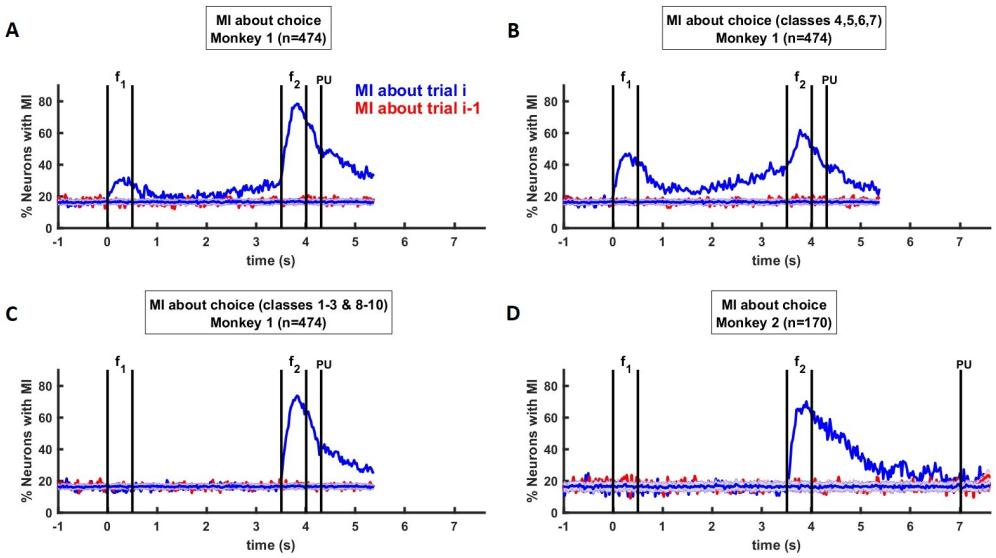

Fig. 3. PFC neural activity has information on the set structure. *(A-D)* Percentage of neurons with significant mutual information (MI) about the current (blue line) and previous (red line) choice. The horizontal blue and red shadings represent the mean and standard deviation computed from shuffled data (see SI Appendix, Methods) for current and previous trial, respectively. *(A)* MI computed exclusively using correct trials from monkey one. *(B)* MI computed from correct trials of monkey one for classes 4, 5, 6 and 7 (its first stimulus identifies the correct choice). *(C)* MI computed from correct trials of monkey one for classes 1, 2, 3, 8, 9 and 10 (correct choice is unknown if only the first time interval is known). *(D)* MI computed using correct trials from monkey two.

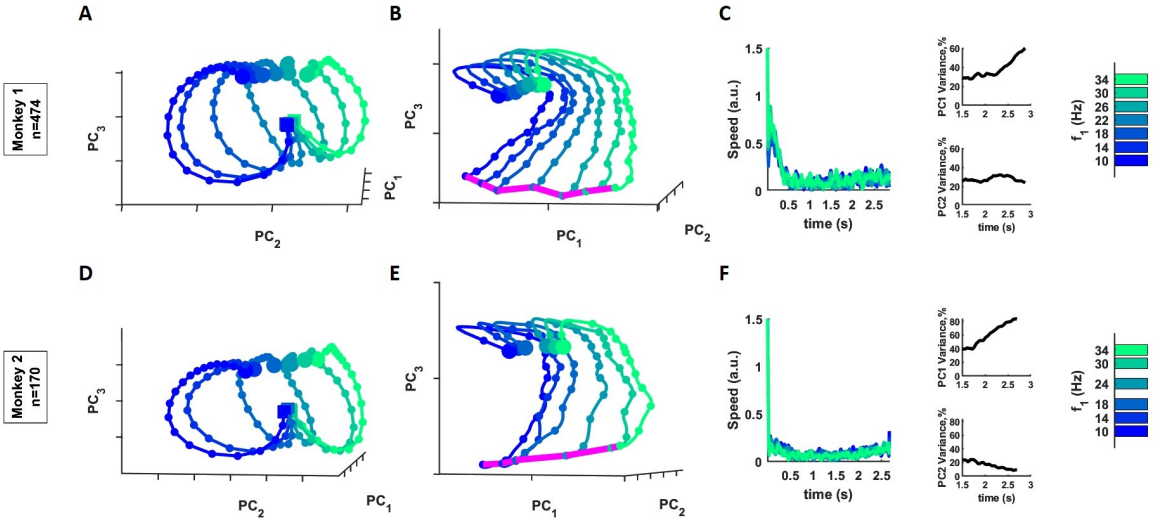

Fig. 4. State-space trajectories and attractor. *(A-C)* Monkey one. *(A)* Projection of neural trajectories onto the first three PCs during the presentation of f_1. The used time window had a size of 25 ms with a sliding window of 25 ms. Squares and big circles represent the onset and the offset of f_1, respectively. Small circles indicate neural states every 25 ms intervals. It is important to note that the distance between consecutive small circles reflects the speed of neural trajectories. *(B)* State-space neural trajectories during the 3 s delay period. The used analysis window had a size of 250 ms with a sliding window of 25 ms. Big circles indicate the offset of f_1. Small circles represent neural states at 250 ms increments. *(C, Left)* Speed of neural trajectories over time during the delay period. Magenta line connects the terminal states at the end of the delay period indicating the quasi-attractor. *(C, Right)* Percentage of variance explained by the first (*Upper*) and second (*Bottom*) PCs at the end of the delay period. Towards the end of the delay period, both variances sum around 90%. *(D-F)* Same state-space analysis but for monkey two.

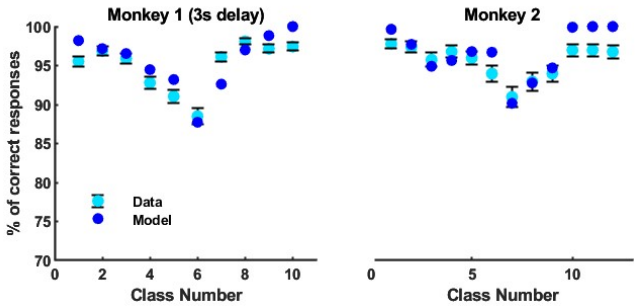

Fig. 5. Bayesian fits of the Accuracy curves. Cyan dots with error bars correspond to behavioral data while dark blue circles represent the fit using Eqs. (14-15). *(Left)* Monkey one. The adjusted model parameters are the Weber fractions of the first ($W_1 = 0.064$) and second ($W_2 = 0.088$) frequency and the sensory history parameter ($\eta = 0.509$). *(Right)* Monkey two. The fit parameters are $W_1 = 0.054$, $W_2 = 0.108$ and $\eta = 0.430$. Note that the relevant quantity to determine the Bayesian contribution to the bias is not W_1 but Σ_1. In both monkeys, the condition $\Sigma_1 > \sigma_2$ was satisfied for the majority of the frequency pairs (f_1, f_2).

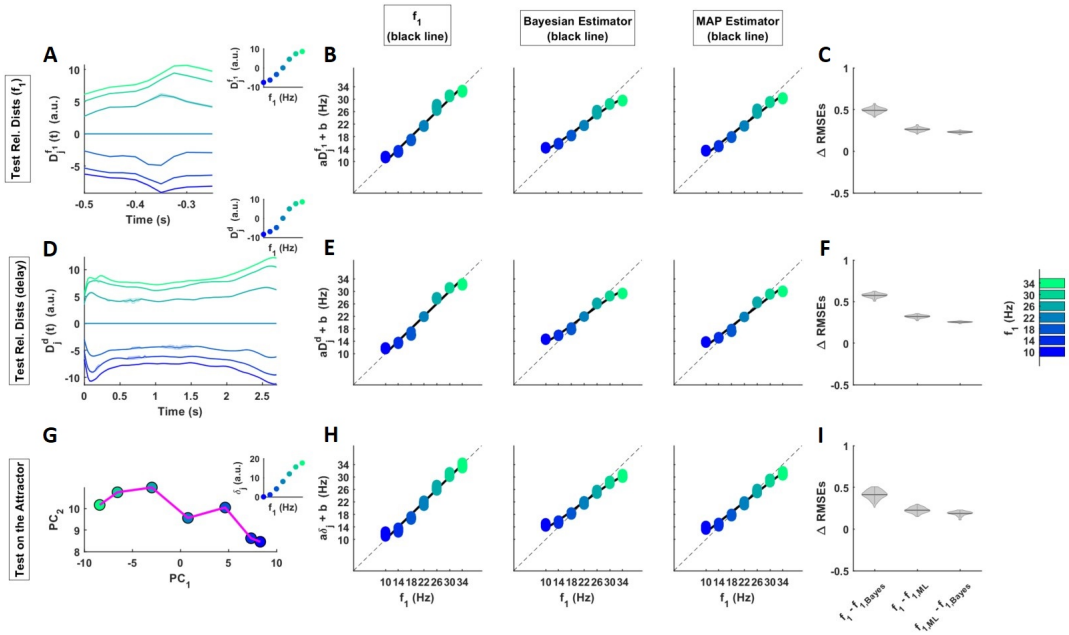

Fig. 6. Several geometric features from monkey one code the Bayesian estimator. *(A-C)* KiNeT trajectories and activity-behavior relationship during the presentation of the first frequency. *(A)* Relative distance, $D_j^{f_1}(t)$, between each trajectory $j = 1, ..., 7$ and the reference trajectory f_1=22Hz using KiNeT analysis. Alignment is located at the offset of the first stimulus. The state space that embeds these neural trajectories is constructed using 6 PCs, which account for 92% of the variance. Shadings were built from the mean ± 95% CIs of 100 bootstrap resamples. *Inset* displays D_j, the temporal means of relative distances as function of f_1. *(B)* Linear regression of the mean relative distances ($aD_j^{f_1}+b$) to the true f_1-value *(Left)*, the Bayesian estimator *(Middle)* and the MAP estimator *(Right)* as a function of f_1. Colored dots represent multiple estimates of the regression obtained through bootstrapping (100 resamples). Black solid lines indicate the true f_1 (*Left*), its Bayesian estimator *(Middle)* and its MAP estimator *(Right)*. The unity line is represented by a black dotted line. *(C)* Violin plots display the kernel probability density of shuffled RMSE differences among the three hypotheses: f_1, Bayesian estimator, and MAP estimator. The width of the gray area represents the density of the shufflings in that range. The horizontal line within each violin plot indicates the mean of the distribution. The chosen hypothesis is consistently the Bayesian estimator. *(D-F)* KiNeT trajectories and activity-behavior relationship during the delay period. Same format as *A-C*, as well as same alignment at the offset of f_1. Activity analysis with KiNeT was build from a state space of 30 PCs explaining 90% of variance. Again, Bayesian hypothesis is chosen by the 100 resamples. *(G-I)* Attractor and activity-behavior relationship at the end of the delay period. *(G)* Attractor region (magenta line) emerges from the terminal neural states (circles) located at the end of the delay period. This region was embedded within a 2-dimensional subspace (84% of variance). *Inset* shows the Euclidean distances, δ_j, between each f_1 terminal state and the smallest f_1 (taken as reference) on the attractor region, as function of f_1. *(H-I)* activity-behavior relationship in the same format as *B-C*. One more time, Bayesian estimator is preferred.

Supplementary Figures

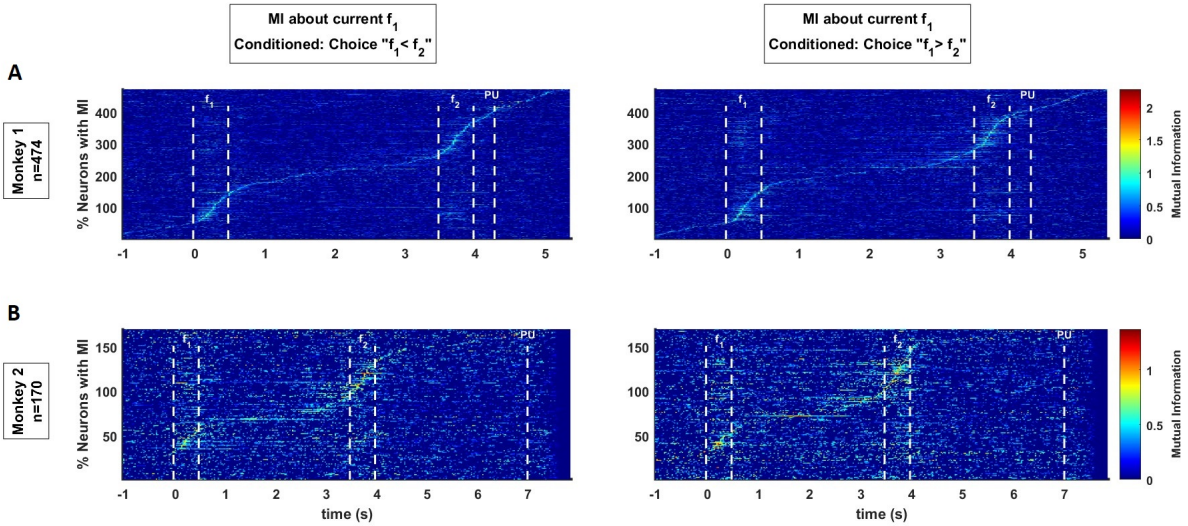

Supplementary Fig. 1. Mutual information (MI) about f_1 in the current trial. MI as a function of time from either *(A)* the 474 cells of monkey one or *(B)* the 170 cells of monkey two and conditioning on rewarded trials either with choice "$f_1 < f_2$" *(Left)* or with choice "$f_1 > f_2$" (*Right)*. Dark blue indicates the non-significant MI values (see SI Appendix, Methods).

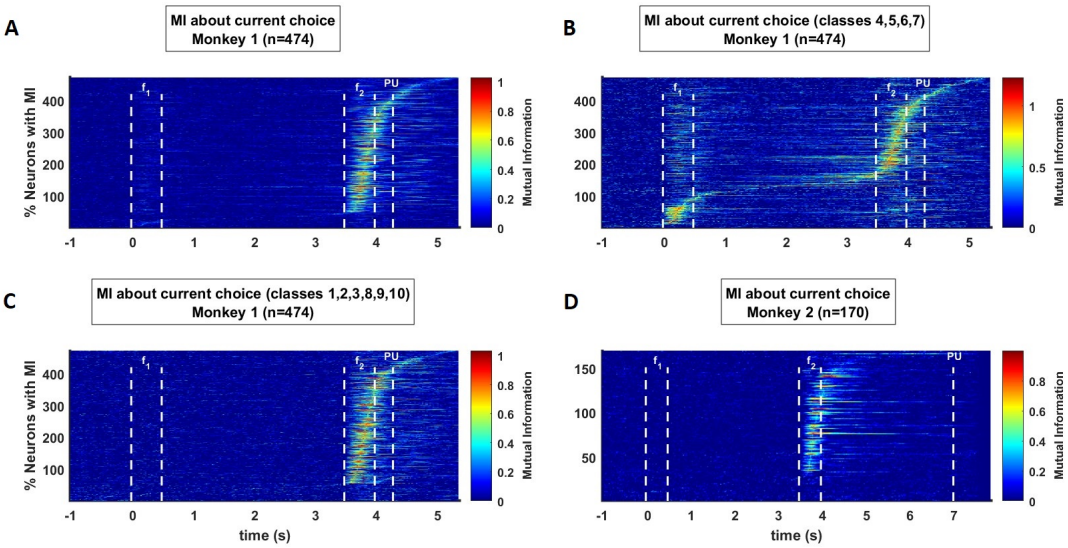

Supplementary Fig. 2. Mutual information (MI) about choice in the current trial. *(A)* MI as a function of time from the 474 cells of monkey one conditioning on rewarded trials. *(B)* MI as a function of time from monkey one conditioning on rewarded trials with classes whose first frequency identifies the correct response (classes 4, 5, 6 and 7). *(C)* MI as a function of time from monkey one conditioning on rewarded trials with classes whose first frequency does not identify the correct response (classes 1, 2, 3, 8, 9 and 10). *(D)* MI as a function of time from the 170 cells of monkey two, conditioning on rewarded trials. Dark blue indicates the non-significant MI values (see SI Appendix, Methods).

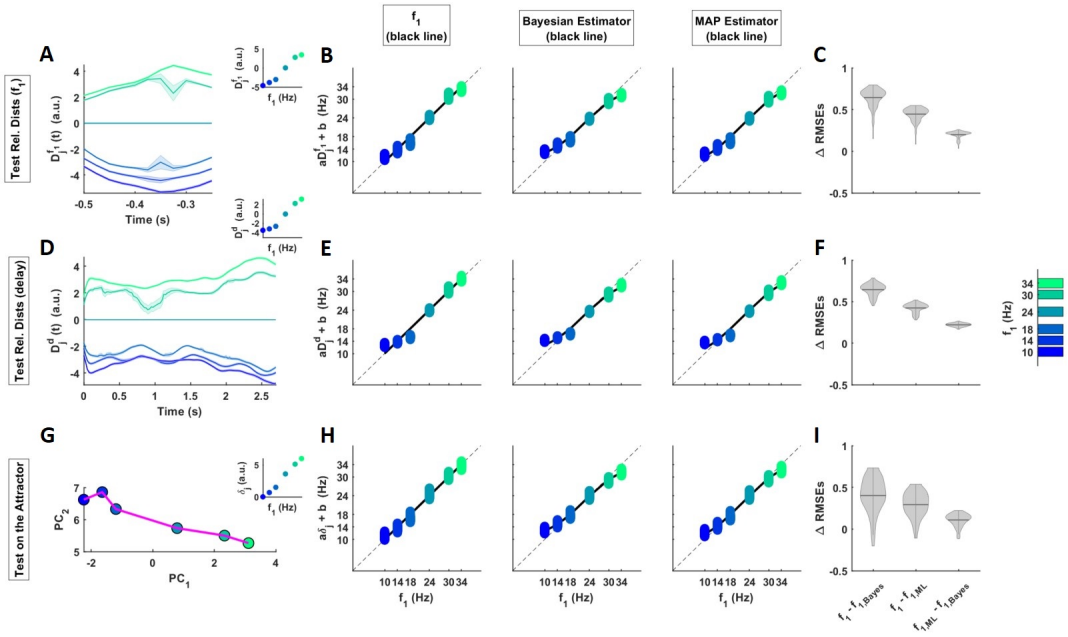

Supplementary Fig. 3. Several geometric features from monkey two code the Bayesian estimator. *(A-C)* KiNeT trajectories and activity-behavior relationship during the presentation of the first frequency. *(A)* Relative distance $D_j^{f_1}(t)$, between each trajectory $j = 1, ..., 6$ and the reference trajectory f_1=24Hz using KiNeT analysis. Alignment is located at the offset of the first stimulus. The state space that embeds these neural trajectories is constructed using 5 PCs, which account for 91% of the variance. Shadings were built from the mean ± 95% CIs of 100 bootstrap resamples. *Inset* displays $D_j^{f_1}$, the temporal means of relative distances as function of f_1. *(B)* Linear regression of the mean relative distances ($aD_j^{f_1} + b$) to the true f_1-value *(Left)*, the Bayesian estimator *(Middle)* and the MAP estimator *(Right)* as a function of f_1. Colored dots represent multiple estimates of the regression obtained through bootstrapping (100 resamples). Black solid lines indicate the true f_1 *(Left)*, its Bayesian estimator *(Middle)* and its MAP estimator *(Right)*. The unity line is represented by a black dotted line. *(C)* Violin plots display the kernel probability density of shuffled RMSE differences among the three hypotheses: f_1, Bayesian estimator, and MAP estimator. The width of the gray area represents the density of the shufflings in that range. The horizontal line within each violin plot indicates the mean of the distribution. The chosen hypothesis is consistently the Bayesian estimator. *(D-F)* KiNeT trajectories and activity-behavior relationship during the delay period. Same format as *A-C*, as well as same alignment at the offset of f_1. Activity analysis with KiNeT was , as well as same alignment at the offset of $f_1$build from a state space of 14 PCs explaining the 90% of variance. Again, Bayesian hypothesis is chosen by the 100 resamples. *(G-I)* Attractor and activity-behavior relationship at the end of the delay period. *(G)* Attractor region (magenta line) emerges from the terminal neural states (circles) located at the end of the delay period. This region is embedded within a 2-dimensional subspace (93% of variance). *Inset* shows he Euclidean distances δ_j, between each f_1 terminal state and the smallest f_1 (taken as reference) on the attractor region, as function of f_1. *(H-I)*. activity-behavior relationship in the same format as *B-C*. Here the 90% of resamples choose Bayesian hypothesis.