Dynamical processing of orientation precision in the primary visual cortex

Hugo J. Ladret[1, 2, ✉], **Nelson Cortes**[2], **Lamyae Ikan**[2], **Frédéric Chavane**[1], **Christian Casanova**[2], **and Laurent U. Perrinet**[1]

[1] Institut de Neurosciences de la Timone, UMR 7289, CNRS and Aix-Marseille Université, Marseille, France
[2]School of Optometry, Université de Montréal, Montréal, QC H3C 3J7, Canada

In our daily visual environment, the primary visual cortex (V1) detects oriented features from a large variety of orientation distributions. The precision of these distributions is fluctuating from one part of the visual scenery to another, and it remains unclear how V1 can account for this range of input configurations in order to achieve robust orientation recognition performances. To answer this question, we used natural-like stimuli and characterised the responses of V1 neurons to quantified variations of orientation precision. Generally, decrements of orientation precision correlated with decreased single neuron tuning, but thirty percent of the recorded neurons remained robustly tuned to imprecise inputs. The existence of these resilient neurons, which could be explained by a model of recurrent cortical inhibition, allows to encode both the orientation and its precision, improving the robustness of the overall population response in V1. Altogether, our results support the notion that the processing of orientation involves precision-specific neurons and slow recurrent dynamics, which possibly constitutes an experimental signature of predictive processes occurring in V1.

V1 | orientation | precision | decoding | predictive coding
Correspondence: ***hugo.ladret@univ-amu.fr***

Introduction

Selectivity to the orientation of stimuli is an archetypal feature of neurons in the mammalian primary visual cortex (V1) (1). In more than 60 years of investigations, simple stimuli such as oriented gratings have uncovered many filter-like properties in V1 (2). Fundamentally, these low-complexity stimuli are limited in their capability to explain how the rich cortical dynamics involved in the vision of natural images (3, 4) are driven by the complex content of our everyday visual scenery (5). Rather than gratings or Gabor filters, natural inputs to the visual system are distributions of orientations (Figure 1). Yet, most of our understanding of orientation processing still comes from studies using single-orientation stimuli, i.e. made of an exaggeratedly precise orientation distribution. The effect of orientation precision on orientation selectivity is therefore a major gap that needs to be addressed to better understand the visual system.

Furthermore, the precision of a sensory variable is of central importance to empirical Bayesian inference, in which it weights the relative influence of the likelihood with respect to a prior when computing the estimation of the posterior distribution. In predictive coding (PC) terms, this corresponds to a modulation of the integration of sensory prediction and prediction error (6, 7), that is, a modulation of feedback context versus feedforward content computation. In that sense, precision allows the brain to balance between its own internal representations (predictions) and its new sensory inputs (prediction errors) to accurately predict the states in the environment. In practical terms, given a low-precision image of a lion moving through the savannah, a predictive coding system that seeks to avoid the lion should (relatively to a high-precision image) rely more on its "lion" internal prediction than on its "savannah grass" sensory input to produce an optimal behaviour. Hence, given the mounting evidence that PC captures the behaviour of the visual system (8–10), it is physiologically likely - and mathematically required by PC - that the first level of the visual hierarchy, V1, possesses a mechanism which can account for the precision of sensory inputs.

So far, our understanding of the influence of orientation precision on vision mainly come from the field of texture perception (11). Psychophysical experiments have shown that decreased orientation precision correlates non-linearly with lower orientation discrimination performances (12–14), which can be readily explained by models of recurrent cortical activity (15–17). Despite these converging observations, there has been only a handful of investigations on the neural correlates of the processing of orientation precision. Recent experimental results have shown that orientation precision causes tuning modulations of single V1/V2 neurons (18), coherently with previous psychophysical investigations. While this has led to the proposal that population variability can be used to encode sensory feature (im)precision (19), the detailed processes that would be involved in such a population code remain by and large unestablished.

Here, we used natural-like stimuli to study the encoding of the precision of oriented inputs in V1. We first quantified the precision-response relationships of single neurons, which shed light on two types of precision-modulated activity profiles. Some neurons exhibited "precision-resilient" activity, which, as opposed to "precision-vulnerable" neurons, remained well-tuned even at minimal input precision. These resilient neurons reached a peak spiking activity sometimes hundreds of milliseconds after stimulation onset, and had an optimal tuning delay which depended on the input's precision. Using a phenomenological recurrent neural network, we showed that varying the strength of lateral cortical inhibition can explain the existence of these two categories of neurons.

This compelling evidence for the involvement of recurrent connectivity in precision responses led us to investigate

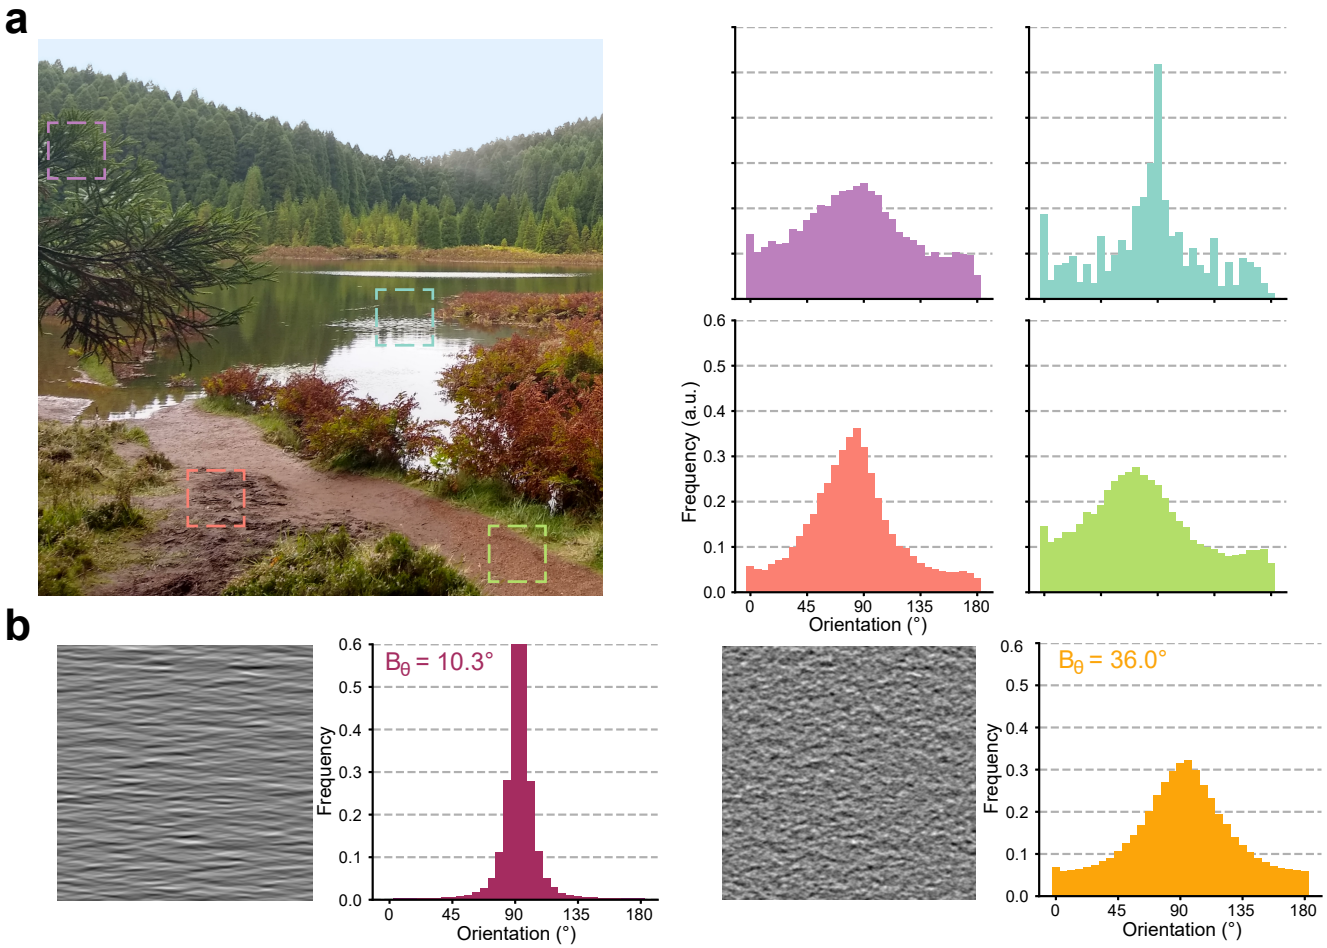

Figure 1. Orientation composition of natural images and natural-like stimuli. (**a**) Distributions of orientations retrieved from a natural image (São Miguel Island, Azores, Portugal, picture taken by H.J.L.) using an Histogram of Oriented Gradients (16x16 pixels cell size, 200x200 pixels image subsets, color-coded to the associated distributions). (**b**) Single frames from Motion Clouds stimulus, oriented at an angle of 90° relative to the vertical axis, of high ($B_\theta = 10.3°$) and lowest precision ($B_\theta = 36.0°$).

how precision modulates the overall population activity in V1. Using a biologically plausible neuronal decoder (20), we found that orientation could be robustly inferred from the population activity at each level of precision. While there was no evidence for an explicit encoding of precision in V1 neurons, we showed that precision and orientation are actually co-encoded by resilient neurons, which significantly improves the orientation encoding of the neural population. This encoding involves specific neural dynamics, in which the cortical processing time and the population variability are directly related to the input precision, resulting in temporally segregated neural representations of precisions within the same neural network. Taken together, these results suggest that the precision of sensory inputs is actively processed in V1. Additionally, slow non-canonical dynamics and neuronal sub-populations dedicated to the processing of input precision in V1 are compelling experimental traces of a predictive process occurring in our recordings.

Results

Single neuron modulations by orientation precision. Single-unit activity of 249 V1 neurons was recorded in three anesthetized cats. Orientation selectivity was measured using band-pass filtered white noise images called Motion Clouds (21), whose underlying generative framework allowed to finely control the orientation distribution of each visual stimulus. Here, we focused on varying two parameters of the Motion Clouds: the median orientation of the distribution, θ, and the bandwidth of the distribution B_θ (Figure 1b). As a measure of the standard deviation σ of the distribution, it is directly related to the precision $1/\sigma^2$. Hence, increasing B_θ corresponds to decreasing the precision of the oriented input to V1. One additional benefit from using Motion Clouds is their stationarity in the spatial domain, which removes any potential second-order correlation found in natural images, further allowing an isolation of the effect of precision on orientation selectivity (Figure 1b). Moreover, they conform to the statistics of natural images, namely the $1/f^2$ distribution of the power spectrum (22).

Here, we generated Motion Clouds at 12 evenly spaced values of θ (steps of 15°) and 8 linearly distributed values of B_θ (starting at $\approx 0°$, ending at 36°) and measured the orientation selectivity of single V1 neurons. As precision decreased (i.e. B_θ increased), the vast majority of tuning curves remained centered on the same preferred orientation (98.8% units, $p < 0.05$, Wilcoxon signed-rank test) but diminished in peak amplitude (95.1% units, $p < 0.05$, Wilcoxon signed-

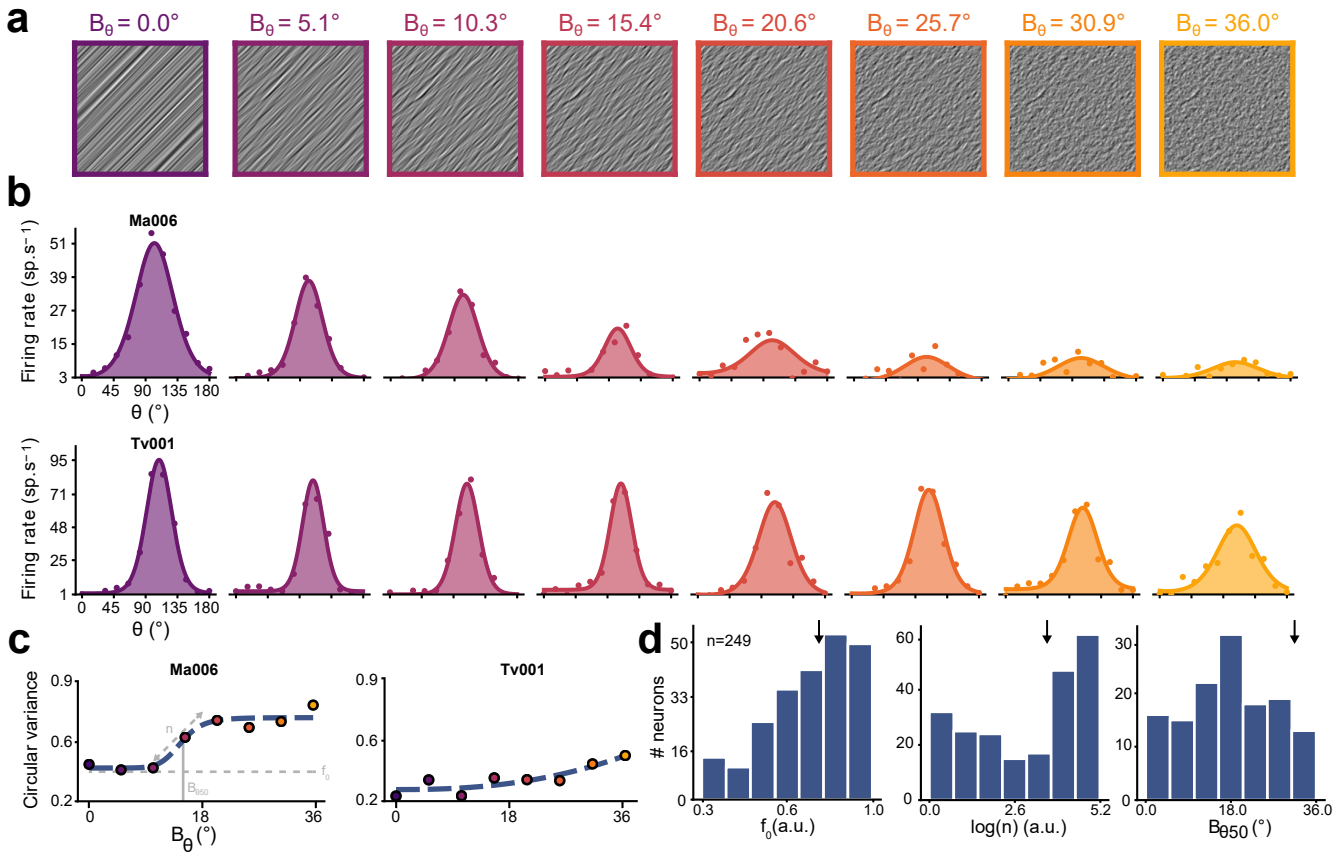

Figure 2. Single neuron responses to variations of orientation precision. Additional examples are provided in Supplementary Figure 1. (**a**) Examples of Motion Clouds, oriented at 45° relative to the vertical axis. Neural activity was elicited with Motion Clouds of 12 different orientations θ and 8 orientation precisions B_θ, drifting orthogonally to their main orientation. (**b**) Tuning curves of two examples neurons (labelled 'Ma006' and 'Tv001') stimulated with Motion Clouds of increasing B_θ from left to right, i.e. of decreasing orientation precision. Colored dots represent the mean firing rate across trials (baseline subtracted, 300 ms average), with lines indicating a fitted von Mises function. (**c**) Precision-response functions (PRFs), assessing the changes of orientation tuning as measured by the Circular Variance (CV, colored dots) as a function of B_θ, fitted with Naka-Rushton (NKR) functions (dashed curves). The effect of NKR parameters are illustrated (in grey) for neuron Ma006's PRF. The parameters for Ma006's PRF are $f_0 = 0.4$, $n = 8.4$, $B_{\theta 50} = 14.7°$. For Tv001, they are $f_0 = 0.3$, $n = 2.4$, $B_{\theta 50} = 36.0°$. (**d**) Distribution of the NKR parameters for all recorded neurons. Median values of the distributions are indicated by a black arrow, respectively 0.75, 3.6 and 31.5 for f_0, $\log(n)$ and $B_{\theta 50}$

rank test, 73.1% mean amplitude decrease). Only 28.5% of the recorded neurons were still tuned when orientation precision reached $B_\theta = 36.0°$ ($p < 0.05$, Wilcoxon signed-rank test). Hence, a reduction in the precision of oriented inputs to V1 caused a decrease in single neuron tuning. This precision modulation unfolded heterogeneously across neurons, as illustrated by two example tuning curves (Figure 2b). For instance, neuron Ma006 was no longer tuned to stimuli of $B_\theta = 36.0°$ ($p = 0.24$, Wilcoxon signed-rank test, firing rate of preferred vs orthogonal orientation), contrary to neuron Tv001 which remained orientation selective even at the lowest of input precision ($p = 10^{-6}$).

Those two distinctive neuronal behaviours were further depicted by measuring the variation of the goodness of tuning (circular variance) as a function of orientation precision B_θ. These "precision-response functions" (PRFs) showed various degrees of non-linearity between neurons (Figure 2c), which were readily captured by a Naka-Rushton curve (23) (Supplementary Figure 2). This allowed to characterize the single neuron response to variations of precision with three parameters: f_0, the goodness of orientation tuning for most precise inputs ; n, the degree of non-linearity of the PRF and $B_{\theta 50}$, the level of input precision at which a neuron's response transitioned from tuned to untuned state. Overall, the PRF parameters were highly heterogeneous through all the recorded neurons (Figure 2d).

Yet, in this heterogeneous population, what separates neurons that remain orientation tuned at lowest input precision from those that do not? We identified the properties of these types of neurons by separating the spectrum of PRF parameters into two diametrically opposed categories: "resilient neurons" (71 neurons, 28.51 % of recorded units), which remained significantly orientation tuned at lowest precision ($B_\theta = 36.0°$) and "vulnerable neurons", which were conversely no longer tuned at the same precision. Those two groups exhibited different dynamics: the maximum firing rate of resilient neurons was delayed compared to vulnerable ones, regardless of the input's precision (Figure 3e,a, $p = 8.10^{-5}$ for $B_\theta = 0.0°$ and $p = 0.01$ for $B_\theta = 36.0°$, Mann-Whitney U test). However, contrarily to vulnerable neurons, resilient neurons displayed a delayed time to reach maximum orientation tuning which changed with input precision (Figure 3f,b, $p = 0.002$, Mann-Whitney U test). Furthermore, both categories of neurons also varied by their PRFs. Resilient neurons showed significantly lower f_0 (Figure 3c, $p = 0.004$, Mann-Whitney U test) and n (Figure 3d, $p = 0.04$,

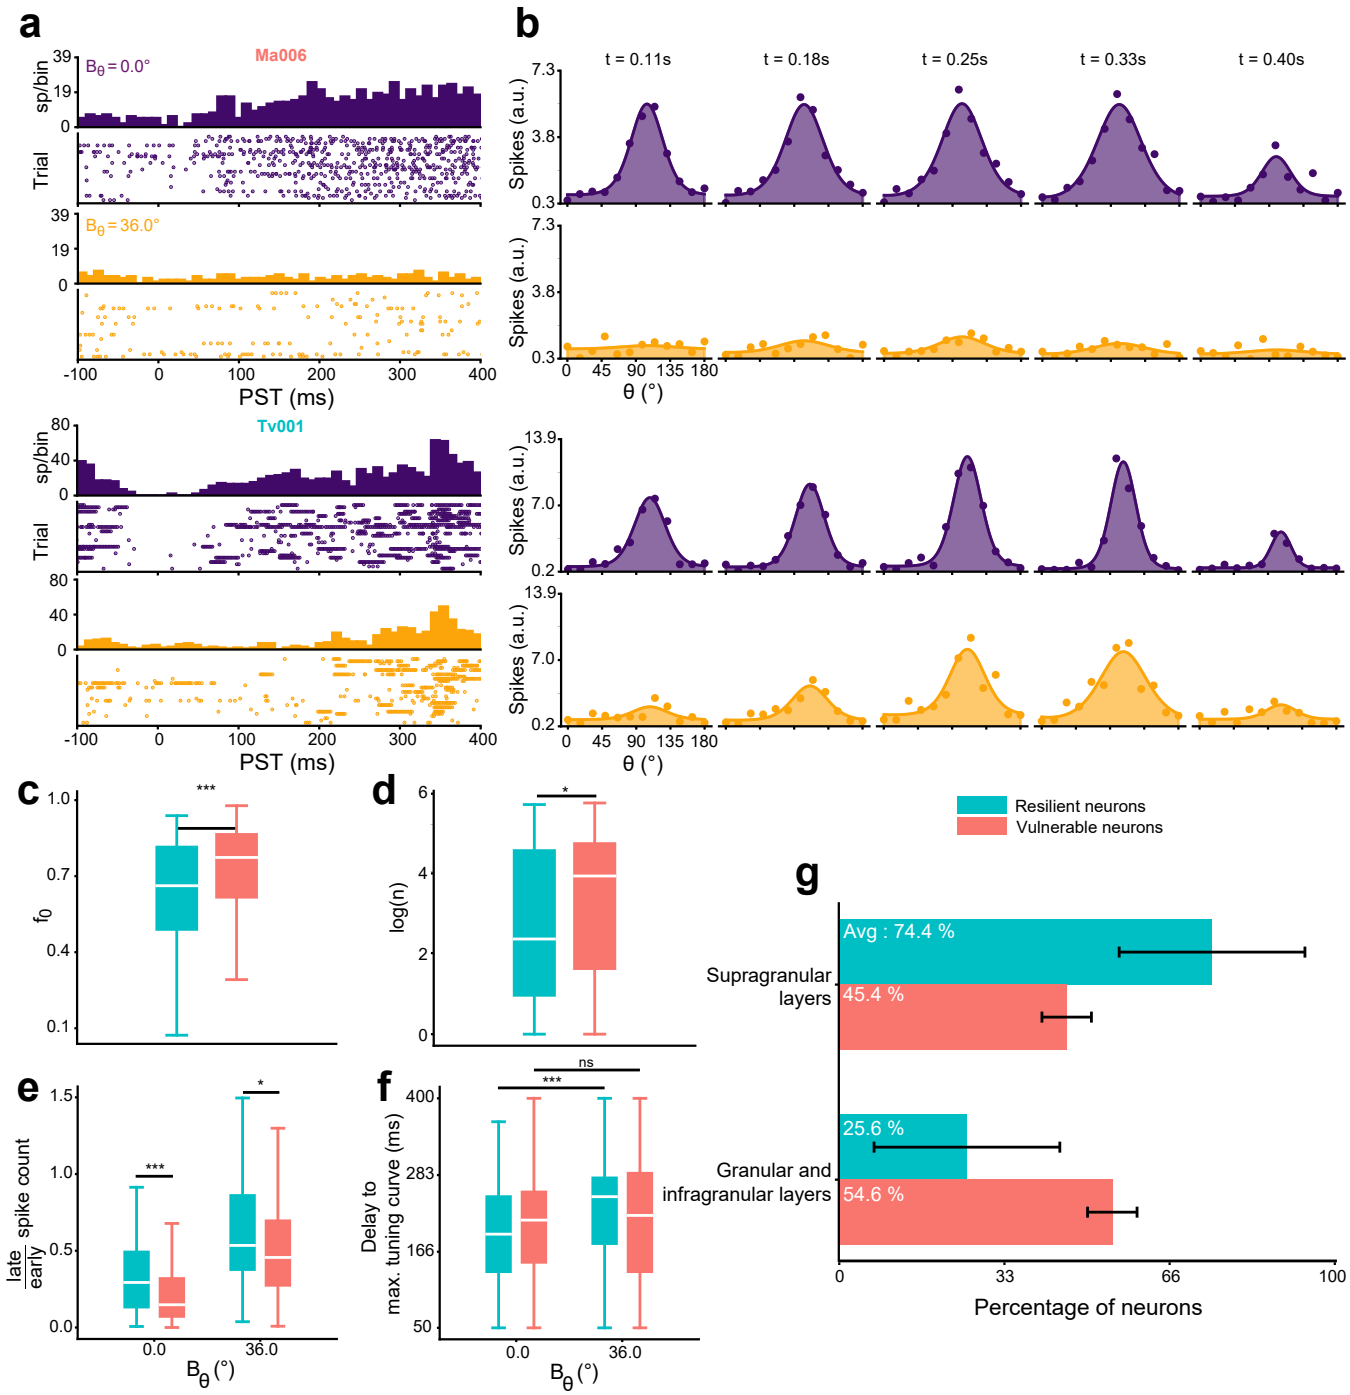

Figure 3. Dynamical properties of the neuronal responses. Neuron labels are colored in orange for neurons untuned to stimuli of $B_\theta = 36.0°$ ("vulnerable neurons"), and conversely colored in cyan when still significantly tuned to the same lowest-precision stimuli ("resilient neurons"). Additional examples are provided in Supplementary Figure 3. (**a**) Peri-stimulus time (PST) histogram and rasterplots of the two example neurons, for trials with $B_\theta = 0.0°$ (purple) and $B_\theta = 36.0°$ (yellow). (**b**) Temporal evolution of the tuning curves of the same neurons, taken in a 100 ms window starting at the labelled time. (**c**) Distributions of the baseline of the Naka-Rushton functions, f_0, for the two groups of neurons. (ns, not significant; $\star$, $p < 0.05$; $\star\star\star$, $p < 0.001$, Mann-Whitney U-test). (**d**) Distributions of the exponent of the Naka-Rushton functions, n. (**e**) Distributions of the spike count distributions of the two groups of neurons, measured as a late (last 100 ms of the trials) over early (first 100 ms of the trials) raw spike count ratio. (**f**) Distributions of the time to reach maximum tuning curve for the two groups of neurons. (**g**) Average proportion of resilient and vulnerable neurons locations in cortical layers (see Materials and Methods). Error bar represent the standard deviations across animals.

Mann-Whitney U test). Respectively, this corresponds to a better tuning to precise oriented inputs and reduced variations of tuning when precision is decreasing, compared to vulnerable neurons. Altogether, resilient neurons exhibited sharper orientation tuning and slower dynamics, which are distinctive features of supragranular neurons (24, 25). This correlates with the location of the recorded units, which predominantly places resilient neurons in supragranular layers (Figure 3g). In terms of precision-responses, their more linear PRF gives them increased dynamic range, which corresponds to better orientation coding in across a wide range of precisions (26, 27).

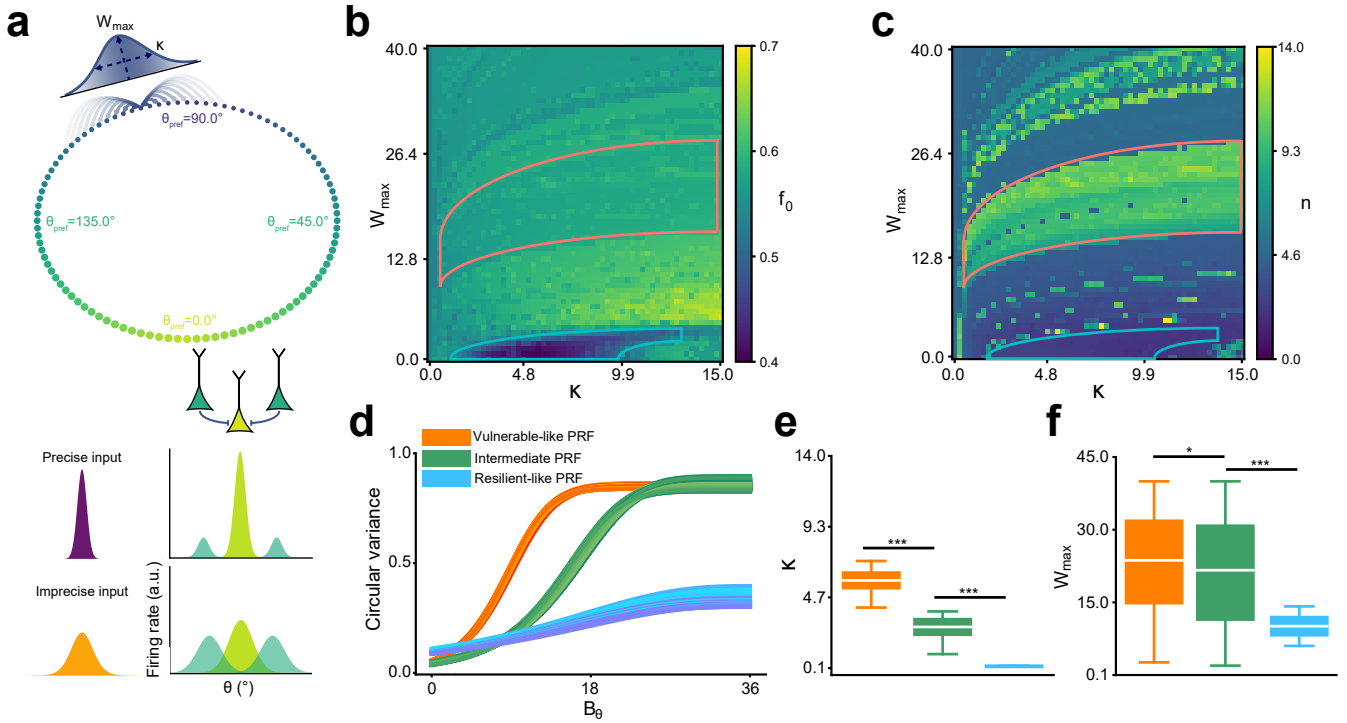

Figure 4. A model of recurrent inhibition can account for the presence of resilient and vulnerable neurons. (**a**) Top: Illustration of the model. A network of orientation-selective neurons tiling the orientation space is connected by recurrent inhibitory synapses. Two parameters control this connectivity: $W_{\max}$, the maximum strength of inhibition and κ, a measure of the concentration of the inhibition, with smaller values of κ causing a broader inhibition (see Materials and Methods). Bottom : Illustration of the effect of precision on V1 activity. Precise inputs (low B_θ) elicit an activity centered on few orientation-selective neurons. Conversely, imprecise inputs (high B_θ) elicit an activity that spreads on a greater number of neurons, causing an overlap in the activation of tuning curves, which drives - in our model - an inhibition on part of the tuning curves. (**b**) 2D map of the parameters $W_{\max}$ and κ and their effect on the neuron's PRF parameter f_0. Combinations of parameters resulting in the creation of a resilient-like or vulnerable-like PRFs are highlighted in cyan and orange, respectively. (**c**) 2D map of the parameters $W_{\max}$ and κ and their effect on the neuron's PRF parameter n. (**d**) Selected examples of PRFs with various parameters of simulation. Three different sets of parameters are shown, causing a low $B_{\theta 50}$ and high n (orange curves), an intermediate $B_{\theta 50}$ and high n (green curves) and a very high $B_{\theta 50}$ and low n (blue curves). (**e**) Distribution of the parameter κ for these three sets of stimulation. A higher value of κ causes a broader inhibition in the orientation space. ($\star\star\star$, p < 0.001, Mann-Whitney U-test). (**f**) Distribution of the parameter $W_{\max}$for these three sets of stimulation. ($\star$, p < 0.05, Mann-Whitney U-test).

Recurrent inhibition controls precision modulations. To explain these differences in single neuron activity, we turned to a computational approach in order to understand how resilient and vulnerable properties can emerge in V1. As resilient neurons exhibited some prototypical features of layer 2/3 neurons, which are driven by recurrent cortical activity (28, 29), we modeled the activity of the recorded neurons with a recurrent neural network organized in a "ring" topology (Figure 4a). Such models have been used to explain responses to simple mixtures of oriented inputs as the result of inhibitory activity, which causes non-linear sharpening and decrements of tuning curves, as observed in our data (Figure 2). Indeed, theoretical approaches have shown that intracortical inhibition plays a central role in shaping the response to stimuli containing multiple orientations, by reducing the activity of the off-median orientation content (16, 30). This is coherent with neurophysiological observations, which have proven that inhibitory processes between oriented components create responses to stimuli of multiple orientation which are smaller than the sum of their parts (31).

Because of this pivotal role of recurrent inhibition, and given the fact that both inhibition and excitation are balanced variables in cortical networks (32), we chose to focus only on inhibitory activity to model our results. Specifically, based on the notion that thalamocortical inputs of lower precision increase the spread (in orientation space) of activity (Figure 4a, bottom), we created a phenomenological population rate model of orientation-tuned neurons. In this model, neurons are tiling the orientation space and connect recurrently according to a von Mises distribution (Figure 4a, top), of which we varied two parameters: $W_{\max}$, the maximum strength of the inhibition and κ, a measure of the concentration of the inhibition in orientation space (decreasing κ broadens the inhibition, see also Materials and Methods). These parameters had a non-linear effect on the shape of a neuron's PRF. A specific range of parameters, namely κ ranging from ≈ 3.0 to ≈ 10 and $W_{\max}$ between ≈ 0.0 and ≈ 3.5 caused a neuron to display resilient-like characteristics (low f_0 and n, Figure 4b,c, orange highlight, similar to Figure 3c). To create a vulnerable neuron, a wider choice of parameters was available, namely $W_{\max}$ within the ≈ 12.5 to ≈ 26.4 range, without restriction on the spatial extent of inhibition κ (high f_0 and n, Figure 4b,c,blue highlight, similar to Figure 3d). This greater freedom on the conditions which created a vulnerable neuron could also explain their predominance in our data (71.5% of the recorded units are vulnerable). By selecting a restricted number of simulations (Figure 4d) which illustrated the prototypical PRFs observed in neuron Ma006 (Figure 2b, top) and Tv001 (Figure 2b, bottom), we found that the transition from the vulnerable to resilient PRF was

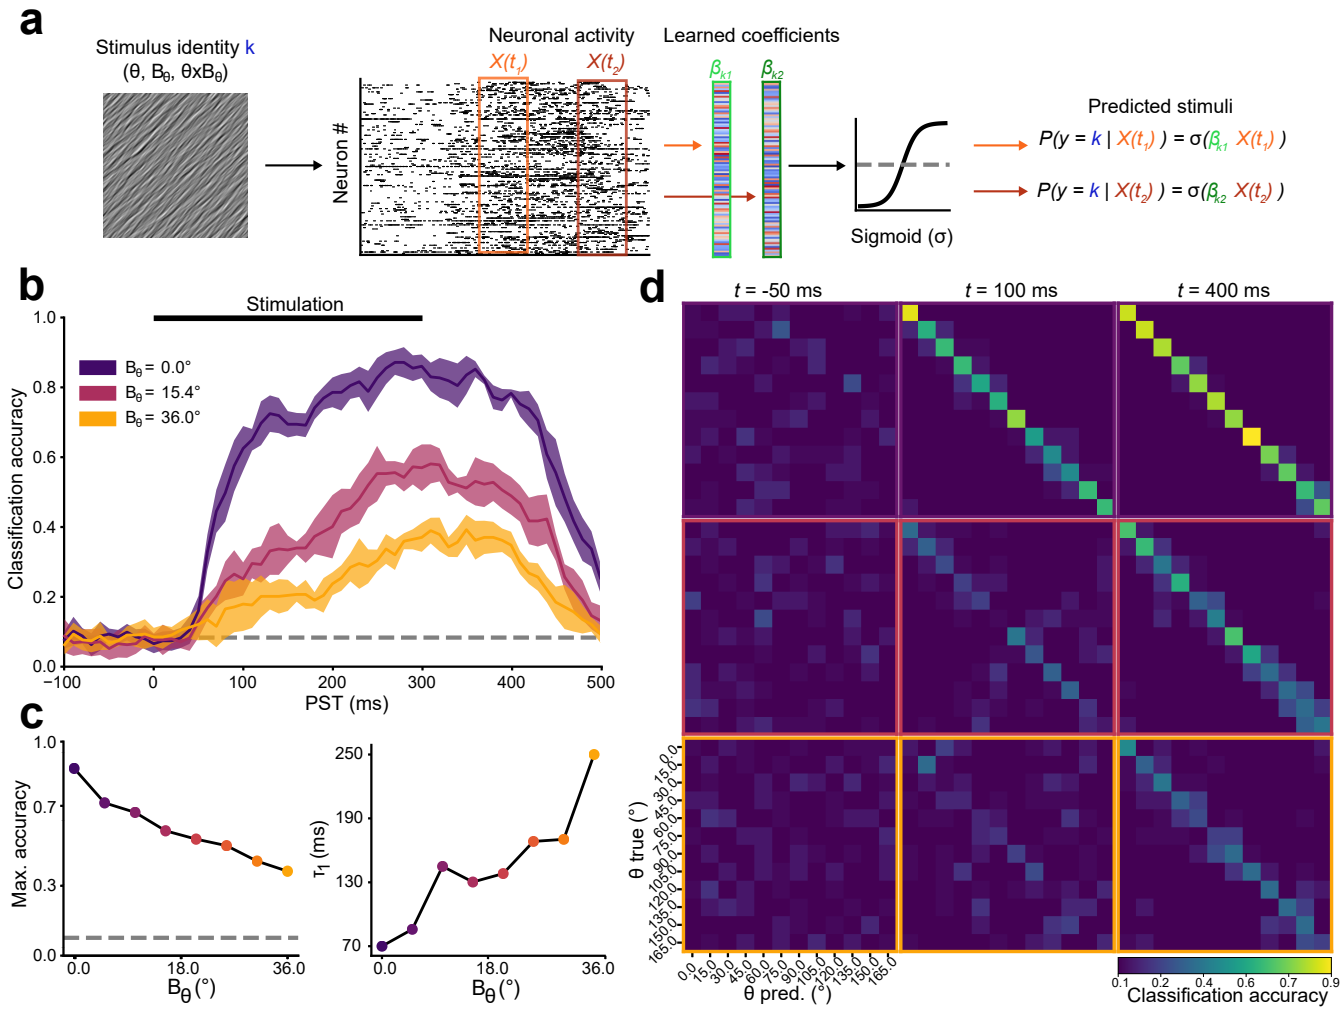

Figure 5. Decoding orientation from population activity. (**a**) Illustration of the decoding model. For each trial, spikes from all neurons in a given time window were pooled in a single neural activity vector $X(t)$. A multinomial logistic classifier predicts the identity k of the stimulus in a given trial by estimating the probability of each hypothesis y. This is achieved by learning a weight vector β_k, which maximizes the accuracy of the prediction $P(y = k|X(t)) \approx \sigma(\beta_k^T X(t))$, where σ is the sigmoid function. The resulting accuracy of this classifier serves as a proxy for the discriminability of the population's neural activity within feature space. (**b**) Time course around the peri-stimulus time (PST) of the accuracy of three independent decoders trained to predict the orientation θ of Motion Clouds for three given orientation precisions B_θ. Solid dark line represent the mean accuracy of a 5-fold cross validation and filled contour the SD. Decoding at chance level (here, $\frac{1}{12}$) is represented by a gray dashed line. (**c**) Parameters of the decoders' time courses, estimated by fitting a difference of two sigmoids, with τ_1 the time constant of the first sigmoid (i.e the time constant of the rising exponential phase of the accuracy curve, see Materials and Methods) (**d**) Confusion matrices of the decoders at three time points around stimulation onset. Values on the diagonals of each matrix represent the instances of correct decoding, that is, where $\theta(t)_{\text{pred}} = \theta_{\text{true}}$. Each row of matrices is framed with the color code of the corresponding decoder in (**b**).

caused by an increase of the broadness of inhibition in the orientation space (Figure 4e) and a decremented maximum strength of inhibition (Figure 4f).

The strength of recurrent inhibition within a population of orientation-selective neurons can thus account for the presence of vulnerable and resilient neurons. Resilient neurons integrate inhibitory information over a larger extent of the orientation space, so they should receive a more varied range of orientation-tuned input, which would position them near the pinwheel centers of orientation maps (33). The properties of vulnerable neurons aren't tied to a specific distance of inhibition, but rather to a specific inhibitory strength which is higher than the one of resilient neurons, fitting the fact that their activity decreases with precision (Figure 2, 3). Based on this modelling approach, it is clear that precision modulations are influenced by multi-neuronal activity and it thus becomes necessary to understand how V1 populations respond to variations of orientation precision.

Decoding orientation from population activity. Given the compelling evidences for the involvement of recurrent activity to produce precision-modulated responses, we next moved to study the effect of precision at the level of larger neural networks in V1. To understand what parameters of the stimuli (orientation or precision) were encoded in this population activity, we used a neuronal decoder that probed for population codes in V1. Briefly, we trained a multinomial logistic regression classifier, which is a probabilistic model that classifies data with multiple outcomes (34) (see materials and methods). Here, this classifier was fed the firing rate of all recorded neurons in a given time window and learned, for each neuron, a coefficient that best predict the identity (θ, B_θ or $\theta \times B_\theta$) of the stimulus (Figure 5a). We first used this method to study the changes in the orientation (θ) code in the population activity, and repeated this decoding for each

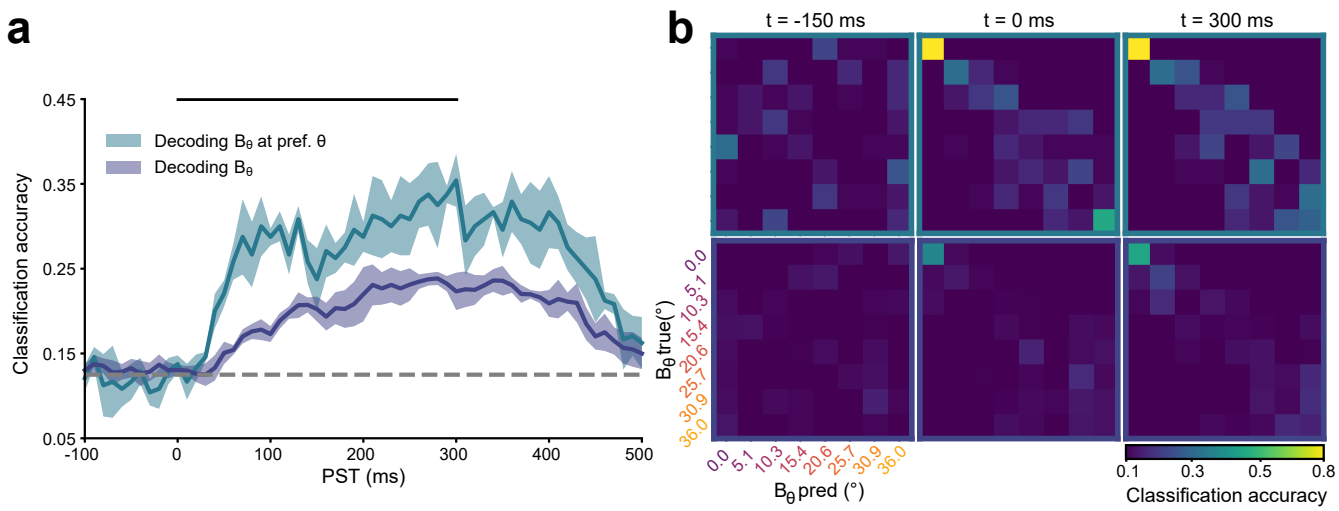

Figure 6. Decoding orientation precision from population activity. (**a**) Time course of two independent decoders, trained to retrieve the orientation precision B_θ of Motion Clouds either with (teal) or without (dark blue) orientation θ prior. Solid dark line represent the mean accuracy of a 5-fold cross validation and filled contour the SD. Decoding at chance level (here, $1/8$) is represented by a gray dashed line. (**b**) Confusion matrices of the two decoders, for the decoding task with orientation prior (top) and without (bottom).

different level of input precision (B_θ). For this, the dataset of trials were separated for each orientation precision, such that 8 independent orientation precision-specific decoders were learned, with optimal parametrization (Supplementary Figure 4a). These decoders were then able to predict the correct stimulus orientation well above the chance level of 1 out of 12 possible orientations. The temporal evolution of decoding performance for these decoders (Figure 5b) showed that optimal orientation encoding is reached faster for stimuli with higher precisions (270 and 370 ms post-stimulation for $B_\theta = 0.0°$ and $B_\theta = 36.0°$, respectively). It should be noted that as the decoders are trained independently in each time window, the accumulative process observed here is occurring in V1, rather than in the decoder itself. Interestingly, the sustained neural activity present hundreds of milliseconds after the stimulation onset can be used to decode orientation for each precision (Figure 5b,c). Further description of these decoders' performances is given by confusion matrices (Figure 5d), which represent the accuracy of the decoder for all possible combinations of true and predicted stimuli (see Materials and Methods). The maximal diagonalization (number of significant $y_{pred} = y_{true}$ instances) of the $B_\theta = 0°$ decoder (upper row), observed as early as 100 ms after the stimulation onset, is close to the diagonalization of the $B_\theta = 36.0°$ decoder (lower row) observed hundreds of ms later (12/12 versus 11/12 classes significantly decoded above chance level, respectively, $p < 0.05$, 1000 permutations test).

Thus, orientation is encoded in the V1 population at all levels of precision, albeit with increasingly slow dynamics in response to low precision stimulus. The short delay required to process precise inputs is congruent with the feedforward processing latency of V1 (20), while the increased time required to reach maximum accuracy for low precision orientations would hint at the presence another slower mechanism. As was put forward by the model, this could mean a transition from a dominantly feedforward regime towards a more recurrent mechanism when processing inputs of low precision (Figure 4).

Given the fact that not all neurons are still active at low precision (Figure 2), we sought to perform orientation decoding on resilient and vulnerable neurons to see if, and how, their contribution to the orientation population code differed. Using the same approach as previously, we trained a precision-specific orientation decoder on the spikes of either resilient or vulnerable neurons. For precise inputs (Supplementary Figure 6a, left), both groups of neurons performed with about the same maximum accuracy, which was below the combined performance of decoding orientation using both types of neurons. Unsurprisingly, when precision is low (Supplementary Figure 6a, right), the activity of resilient neurons can serve to explain the entirety of the decoding. However these resilient neurons, which have delayed responses (Figure 3e,f), displayed shorter population time constants compared to vulnerable neurons. Both results are not incompatible: if resilient neurons integrate information from more cortical neurons than vulnerable ones, then each spikes can potentially carry more information about the population activity, which results in a faster increment of decoder performance.

Decoding orientation precision. As the encoding of orientation in V1 was impacted by orientation precision, we next sought to understand if orientation precision itself was encoded in the population activity. A decoder was trained to retrieve the precision (B_θ) of stimuli (Figure 6a,b, dark blue) but could not decode with satisfying accuracy the precision of orientation from the population activity. The confusion matrix indicates that while precise orientation seem to have a distinctive signature (Figure 6b, lower row), it is impossible to infer the precision of the stimuli once B_θ starts to increase, regardless of the neuron category (Supplementary Figure 6c). One likely explanation for this is the fact that the decoder takes into account spikes from any possible θ, and thus does not have any information on the orientation of the

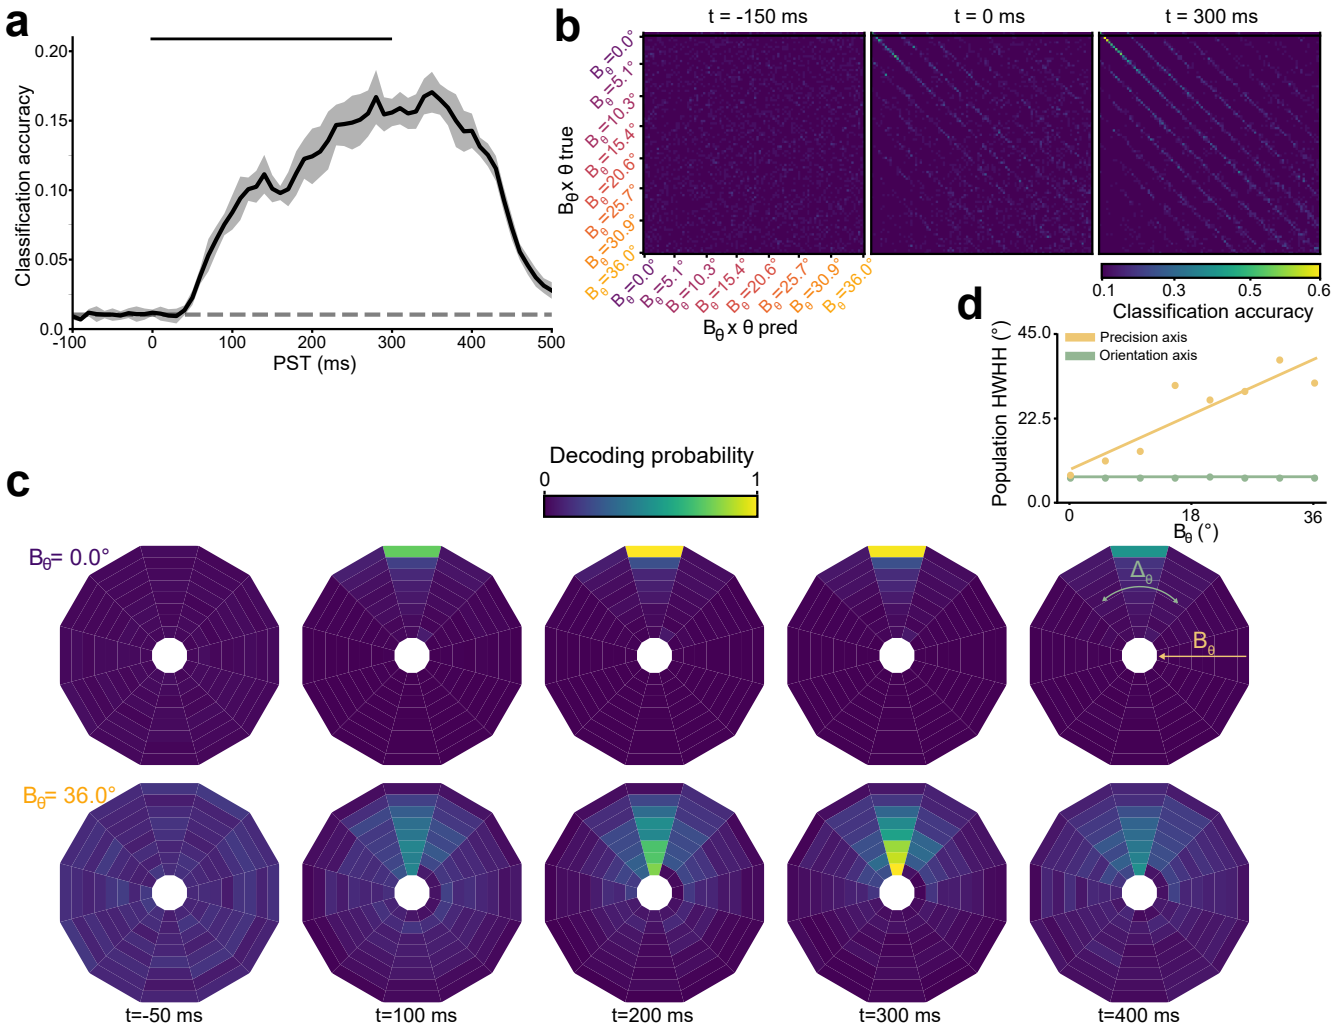

Figure 7. Joint decoding of orientation and precision. (**a**) Time course of a decoder trained to retrieve both the orientation and precision $\theta \times B_\theta$ of Motion Clouds. The solid dark line represents the mean accuracy of a 5-fold cross validation and the filled contour the SD. Decoding at chance level (here, $1/96$) is represented by the gray dashed line. (**b**) Confusion matrices of the decoder. In each section labeled by a given B_θ, the local matrix organization is identical as in Figure 5, with $\theta = 0.0°$ located in the upper left corner of the sub-matrix. (**c**) Polar plots of the decoded neural representations. The probabilistic output of the decoder is color-coded in a 2D polar matrix, where the error Δ_θ on the θ identity of the stimulus is represented as the angle of each bin from the upper vertical, while the eccentricity corresponds to the inferred B_θ (lower precisions towards the center). The temporal evolution of the probability from the decoder is normalized by row. See Supplementary Figure 7 for a detailed explanation. (**d**) Spreading of population representations. By fitting a von Mises curves on either the orientation (azimuth) or precision (eccentricity) axis of the probabilistic output of the decoder, the Half-Width at Half-Height (HWHH) measures the concentration of the population code around the correct input class. Values of HWHH are fitted with a linear regression, with $y = .81x + 9.85$, $r = .9$, p-value = 2.10^-3 for the precision axis and $y = 7.58$, $r = 0.08$, p-value = 0.84 for the orientation axis.

stimuli. As observed in the progressive flattening of tuning curves (Figure 2b, Supplementary Figure 1a), it is impossible to distinguish the firing rate at the orthogonal orientation when $B_\theta = 0.0°$ from the one at the preferred orientation when $B_\theta = 36.0°$. Thus, it is challenging, at least here, to encode orientation precision without any information on the orientation itself. In that sense, adding an orientation prior to the decoder by learning only from spikes coming from each neuron's preferred orientation significantly improved precision decoding (Figure 6a,b, teal color). While this decoder remained rather inaccurate, this was partly due to the reduction in the amount of data available to train the decoder (using only θ reduced the size of the dataset by 12). This is an intrinsic limit of our experimental approach, and it is possible that V1 signals orientation prior with a much greater efficiency, possibly via orientation-specific synaptic connectivity in orientation maps (35).

Decoding orientation and its precision. Although precision affects orientation encoding without itself being encoded in the population, an alternative hypothesis is that both orientation and its precision are jointly encoded in V1. We thus trained a $\theta \times B_\theta$ decoder, which could retrieve the identity of Motion Clouds with a maximum accuracy of about 16 times the chance level $(1/96)$ in 350 ms (Figure 7a). This delay to maximum accuracy is similar to the delay observed for low precision decoding (Figure 5b,c). Unsurprisingly, the confusion matrix of this $\theta \times B_\theta$ decoder showed greater accuracy for patterns of higher precision (Figure 7b, 44% average θ decoding for $B_\theta = 0.0°$ and 12% for $B_\theta = 36.0°$). A confounding factor of the decoder was observed in these confusion matrices, in the form of secondary, parallel diagonals (Figure 7b), which are instances where the decoder correctly predicted θ with no exact certainty on the associated B_θ (only elements within the main diagonals are correct decoding oc-

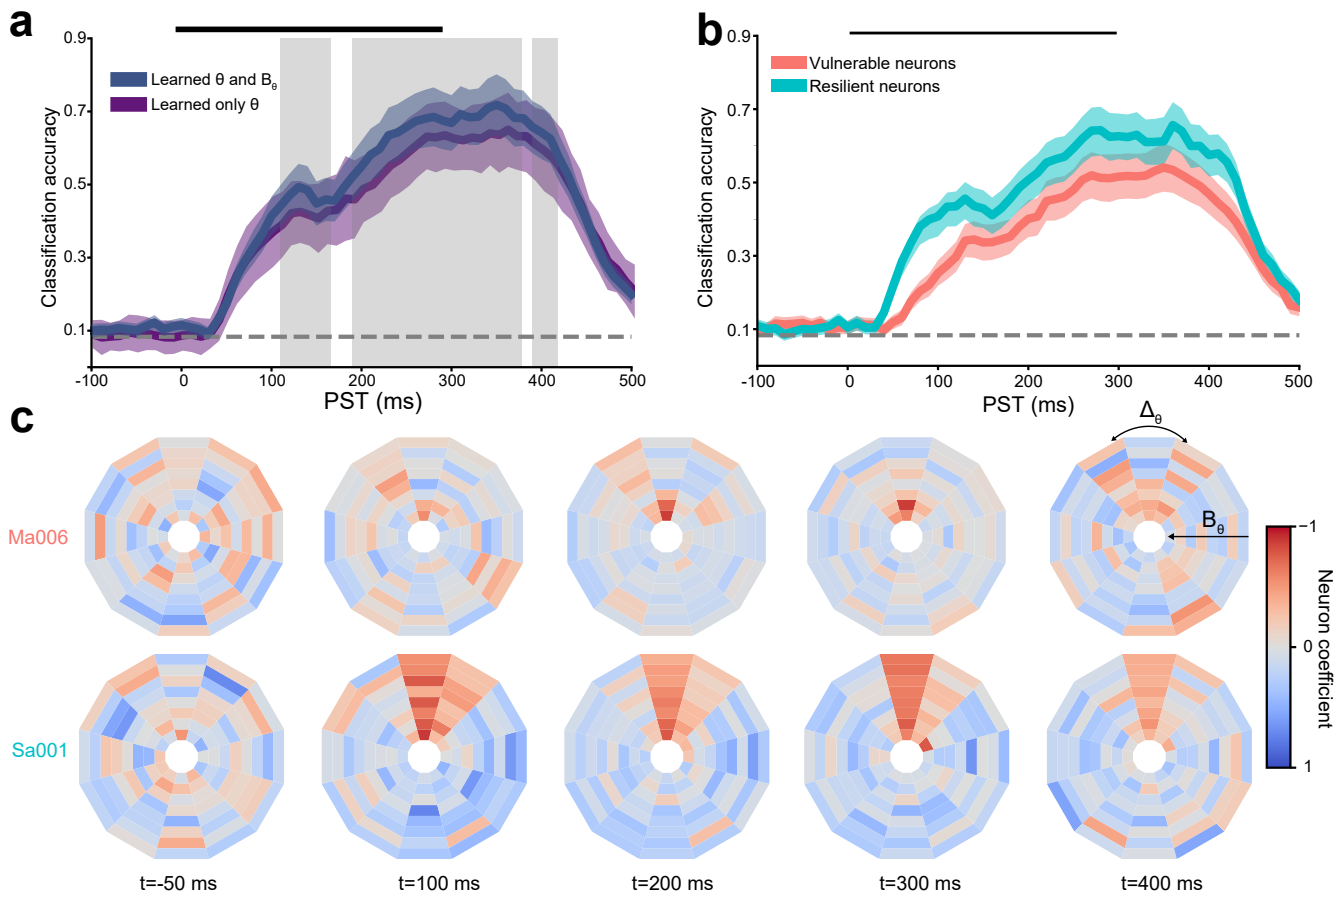

Figure 8. Integrating orientation and precision in a single neural code. (**a**) Time courses of a two independent decoders trained to retrieve orientation θ, either by being precision agnostic (red) or by having prior information on precision (blue). The decoder retrieving orientation with prior precision information is a marginalization of the $\theta \times B_\theta$ decoder trained in Figure 7a,b. Solid dark line represent the mean accuracy of a 5-fold cross validation and filled contour the SD. Regions where the marginalized decoder performed significantly better than the precision-agnostic decoder are indicated in grey (1000 permutations test, n=6, threshold p < 0.05). (**b**) Time courses of two marginalized decoder, trained either from vulnerable or resilient neurons data. (**c**) Polar plots of the weight matrices from a vulnerable (top row) and a resilient (bottom row) neuron. As in Figure 7, the angle of each bin represents the error on the θ identity of the stimulus, Δ_θ, and the eccentricity corresponds to the weight for each B_θ (lower precisions towards the center). The temporal evolution of the weights from the decoder are normalized for each neurons by row. Additional examples are provided in Supplementary Figure 8.

currences, where $y_{pred} = y_{true}$).

This ambiguity on the precision identity of the stimulation is linked to the previous observation that the B_θ decoders fail to retrieve precision from the population activity (Figure 6b). When reporting the performance of a multinomial logistic regression classifier, using the accuracy as a metric can hide part of the neural code's actual structure. Indeed, accuracy is the argmax of the decoder's learned probabilities and the full structure of this probability matrix can sometime shed additional light on the composition of the neural code. Thus, we looked in more details at the structure of the $\theta \times B_\theta$ code by analyzing the probabilistic structure of the decoder (see materials and methods). In such radially organised representation shown in Figure 7c, the error of the orientation classification, Δ_θ is represented as the angle of each bin from the upper vertical, while the eccentricity of each bin from the center is representing the precision probability. For instance, in the first row, precise stimulations ($B_\theta = 0.0°$) elicit right away a probabilistic neural representation of a sharp orientation (low Δ_θ) and precision (no extent on the eccentricity axis) code. By measuring the Half-Width at Half-Height (HWHH) of this probabilistic matrix on either the precision (eccentricity) axis or the orientation (angle) axis, one can observe that as precision decreases (Figure 7c, third row), the spread of the precision code increases, but not the orientation code (Figure 7d). Rather than being totally unavailable as the accuracy results would suggest (Figure 6), precision information is encoded in the form of a gradient in the probabilistic representation of the system. Hence, the precision of stimuli is encoded in V1, but in the form of a $\theta \times B_\theta$ population code.

To understand what role does this co-encoding serve, we compared the performances of two orientation decoders: the first learned to decode orientation θ without any prior knowledge on precision B_θ, while the second was the marginalization of the $\theta \times B_\theta$ decoder over B_θ, i.e. an orientation decoder with prior knowledge of the precision (Figure 8a). When marginalized, the $\theta \times B_\theta$ decoder performed significantly better than the precision-agnostic decoder ($12/12$ orientation significantly better decoded, $p < 0.05$, 1000 permutations test). Prior knowledge on the orientation precision associated with an oriented input thus yielded better performances of orientation decoding. Training such a marginalized decoder on vulnerable and resilient neurons showed that the latter are significantly better at performing a co-encoding

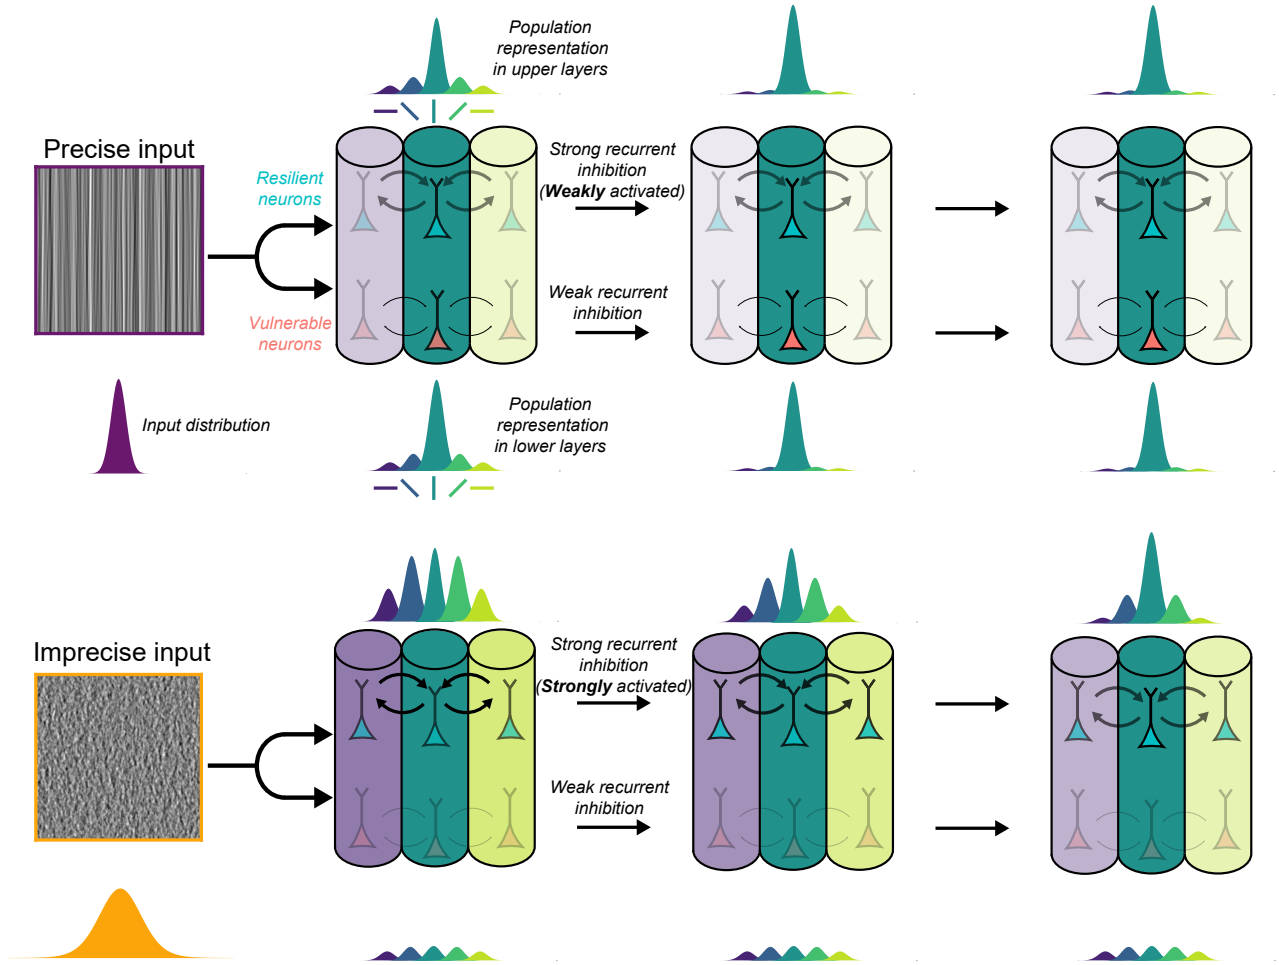

Figure 9. Schematic interpretation of the processing of orientation precision in V1. The color of cortical columns represents their preferred orientation while the intensity of the color represents the degree of their contribution to the population response. Inter-layer connectivity and neuronal types are are not shown nor considered here. The underlying distribution of the input as well as the cortical population representations are shown as Gaussian-like curves. Top row : For high precision inputs, the orientation content is highly concentrated around a median orientation, and can be easily read out by V1 with classical feedforward processing, both by resilient (cyan) and vulnerable neurons (orange). Bottom row : For low precision inputs, the distribution is enlarged: retrieving the median orientation requires a competition between multiple components that each needs orientation processing, implicating recurrent connectivity and additional processing time by the supragranular resilient neurons. As processing time goes on, a consensus emerges in the population, which allows to maintain a stable orientation code through multiple level of input precision.

of orientation and precision (Figure 8b). The visualization of the decoder's weight (Figure 8c, Supplementary Figure 8) showed how the $\theta \times B_\theta$ is implemented in single neurons. These weights, which represents the contributions of each neurons towards the population code, showed that the activity of resilient neurons was used to encode orientation and precision, as observed by the extent of the bins in the eccentricity (B_θ) axis (Figure 8c, bottom). Conversely, the decoding process yielded an orientation-only information from the activity of vulnerable neurons (Figure 8c, top).

Taken together, these results can be interpreted in terms of an interplay between cortical populations with dissimilar tuning properties (Figure 9). Indeed, the accumulative dynamics of the decoders (Figure 5) are directly tied to the notion that neurons are performing incrementally better separation of features within orientation space (34). In other terms, this means that as time elapses, a dominant representation of orientation emerges within the population, which can serve to decode the median orientation contained in the stimuli. Based on the data shown here, we propose that recurrent inhibitory activity, which mediates precision resilience behaviours (Figure 4), is the biological support of this competition (Figure 9). Given the single neurons properties (Figure 3) and the feature they encode (Figure 8), it is likely that resilient neurons are the one mainly involved in this recurrent process. This transition from a feedforward to a recurrent state within the same cortical circuit supports an important property of the population code. As precision of the inputs decreases, the robustness of this encoded precision does decrease, but the population code remains largely invariant in the orientation domain (Figure 7).

Discussion

The precision of oriented inputs impacts orientation selectivity and we have sought to understand how V1 processes this essential visual feature. Using naturalistic stimuli, we showed that specific pool of neurons responses support a population activity that co-encodes orientation and precision through differential dynamical processes. These results suggest that recurrent population-level mechanisms in V1 can

serve to encode orientation across the varying degrees of precision found in natural images.

While we do not presently pinpoint the neural substrate for precision processing, three likely candidates can be put forward, based on theoretical considerations and biological requirements. Precision, by mathematical definition $1/\sigma^2$, requires summing the activity of multiple neurons to compute σ at the population level. As shown by the decoding results, this population activity extends over a wide range of time scales, from tens of milliseconds for precise stimuli to few hundreds milliseconds for imprecise ones.

By taking the classical view of the visual system as a hierarchical network, one first possible substrate of precision processing is the feedforward/feedback connectivity. In Bayesian formulation, changing the precision of likelihoods (i.e. sensory inputs) drives the posterior closer or further away to the priors (i.e. predictions), with imprecise likelihoods causing the posterior to shift towards the prior (7). One could thus expect that inputs with decreasing levels of precision are encoded by increasingly hierarchically high areas, which would explain the long delays observed for low precision stimuli. Furthermore, even under anesthesia, feedback from extrastriate areas does not change the content of the orientation response in cat V1 (36–41), but may provide contextual modulation (42), which is arguably similar to precision weighting of cortical activity (43). However, this explanation raises a theoretical issue, since the computation of precision in a cortical area of hierarchical level N would become non-local, since precision would then be processed by another area of level $N+1$. This poses a quickly-scaling problem of regulating cortical flow of information between multiple areas to finely communicate precision.

If precision processing does take place across multiple cortical areas, a more straightforward candidate exists, in the form of the pulvinar nucleus. This higher-order thalamic nucleus establishes cortico-thalamo-cortical connectivity with the visual hierarchy, providing bi-directional shortcuts between visual areas (44). Interestingly, the pulvinar's action upon the cortex is a function of the hierarchical level of each area (45), regulating and increasing the efficiency of the information transmission across non-linear response functions (45–47). The temporal scale of this pathway matches the observed delays as well as the feedback would. Considering the pulvinar as the seat of precision modulations offers the advantage of simplifying the theoretical considerations, as pulvinar can then be a nucleus partly dedicated to processing precision of cortical areas (43). However, computing such precision matrix of any cortical area of size n neurons requires an order of n^2 computations, which would be prohibitively costly for the pulvinar (48). It is more likely that the pulvinar could compute an approximation of this precision matrix, which would be then refined within each cortical area. This matches nicely the specific anatomy of transthalamic connectivity, in which pulvino-cortical synapses targets supragranular recurrent layers (45).

This would support intra-area cortical recurrence as implementing precision modulation in the early visual cortex. This interpretation has the advantage of being supported by predictive coding theories, namely by the notion of precision of predictions errors being canonically assigned to supragranular neurons (6, 49). In line with our results, recurrent synaptic connectivity is heterogeneous in V1, which sustains a resilient orientation tuning (33) that could reflect the diversity of single-neurons precision modulations. Temporal scales of recurrent connectivity, namely slowly-conducted horizontal waves in an orientation map (29), fit the view of precision processing being the result of multiple iterations of a canonical precision computation implemented by recurrent interactions. Furthermore, it would explain the slower time constant of vulnerable neurons (Supplementary Figure 6) as the result of a propagation of activity between granular and infragranular layers. This general interpretation of the result also fits our proposal of low-precision input recruiting more columns of different preferred orientations (Figure 9). Inhibitory suppression of off-median orientation progressively increases the accuracy of a population decoder (Figure 5) that reads out the activity of these columns, which also explains the decreases in single neurons response (Figure 2).

While this last consideration is an easier fit to our results, we have preferred a conservative approach in the present modelling approach by considering orientation space rather than actual mapping to neurobiology, namely due to the fact that very few other works have studied the encoding of orientation precision in V1. A notable exception is from Goris *et al.* (18), who, using stimuli similar to Motion Clouds, have reported that heterogeneously tuned populations are better suited to encode the mixtures of orientations found in natural images. Consistent with their findings, we also reported a variety of single-neuron modulations by the precision of oriented inputs, which then serves to decode orientation mixtures of varying precisions at the population level — a requirement for encoding natural images in V1 (50). Their results have further shown that variance of the population activity is directly tied to the (im)precision of the stimulation (19). Similarly, we observed that tuning variability (measured as circular variance in Figure 2 and Figure 3) increases for lower precisions. While we do not directly infer precision of stimulus from population variance, we do show that the population variability is a signature of precision, but not of orientation (Figure 7d).

Overall, we have shown that precision, an intrisic feature of natural images, is actively processed in the primary visual cortex. While we have highlighted previously unknown population mechanisms, the actual functional organization of precision processing remains an open question. All three hypothesis we have put forward make easily testable experimental predictions, which would simply mandate additional electrophysiological recordings in extrastriate areas. This would result in a better understanding of the neural basis of slow cortical dynamics observed here, which would be a crucial step to properly depart from the fast, feedforward-centric interpretation of visual processes. By using a generative stimulus framework, as presented here (21), precision processing can be investigated for any sensory modality, at

any hierarchical level. Manipulating precision within higher-level paradigms, like oddball effects (51), effectively allows fine control over top-down predictions and bottom-up predictions errors, which could yield pivotal new insights into our understanding of predictive processes in the brain.

ACKNOWLEDGEMENTS
This article is dedicated to the memory of Umit Keysan (1992-2019), a dear colleague and wonderful friend. The authors would like to thank Genevieve Cyr for her technical assistance, Bruno Oliveira Ferreira de Souza and Visou Ady for advices regarding experimental procedures as well as Louis Eparvier and Jean-Nicolas Jérémie for their comments on the manuscript and Jonathan Vacher for fruitful exchanges on the formalization of the generation of synthetic images and for his contributions to the analysis of neurophysiological recordings. This work was supported by ANR project "Horizontal-V1" ANR-17-CE37-0006 and by a CIHR grant to C.C (PJT-148959). H.J.L. was funded by an Ecole Doctorale 62 PhD grant.

DATA AVAILABILITY
The data generated in the present study are available from the corresponding author, H.J.L., upon reasonable request. No publicly available data was used in this study.

CODE AVAILABILITY
Data was analyzed using custom Python code, available at https://github.com/hugoladret, using libraries SciPy (52), scikit-learn (53), numpy (54), PyTorch (55), lmfit (56) and Matplotlib (57).

AUTHOR CONTRIBUTIONS
L.U.P., C.C. F.C., N.C., and H.J.L. designed the study. H.J.L., N.C. and L.K. collected the data. H.J.L., N.C. and L.U.P. analyzed the data. H.J.L. and L.U.P. wrote the original draft of the manuscript. All authors reviewed and edited the manuscript.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

References

1. Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology* **148**, 574 (1959).
2. Priebe, N. J. Mechanisms of orientation selectivity in the primary visual cortex. *Annual review of vision science* **2**, 85–107 (2016).
3. Fiser, J., Chiu, C. & Weliky, M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* **431**, 573–578 (2004).
4. Yoshida, T. & Ohki, K. Natural images are reliably represented by sparse and variable populations of neurons in visual cortex. *Nature communications* **11**, 1–19 (2020).
5. Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annual review of neuroscience* **24**, 1193–1216 (2001).
6. Friston, K. A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences* **360**, 815–836 (2005).
7. Friston, K. & Kiebel, S. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences* **364**, 1211–1221 (2009).
8. Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* **2**, 79–87 (1999).
9. Schwiedrzik, C. M. & Freiwald, W. A. High-level prediction signals in a low-level area of the macaque face-processing hierarchy. *Neuron* **96**, 89–97 (2017).
10. Boutin, V., Franciosini, A., Chavane, F., Ruffier, F. & Perrinet, L. Sparse deep predictive coding captures contour integration capabilities of the early visual system. *PLoS computational biology* **17**, e1008629 (2021).
11. Julesz, B. Textons, the elements of texture perception, and their interactions. *Nature* **290**, 91–97 (1981).
12. Phillips, G. C. & Wilson, H. R. Orientation bandwidths of spatial mechanisms measured by masking. *JOSA A* **1**, 226–232 (1984).
13. Heeley, D., Timney, B., Paterson, I. & Thompson, R. Width discrimination for band-pass stimuli. *Vision research* **29**, 901–905 (1989).
14. Heeley, D. W. & Buchanan-Smith, H. M. The influence of stimulus shape on orientation acuity. *Experimental Brain Research* **120**, 217–222 (1998).
15. Keeble, D., Kingdom, F., Moulden, B. & Morgan, M. Detection of orientationally multimodal textures. *Vision Research* **35**, 1991–2005 (1995).
16. Beaudot, W. H. & Mullen, K. T. Orientation discrimination in human vision: Psychophysics and modeling. *Vision research* **46**, 26–46 (2006).
17. Baudot, P. *et al.* Animation of natural scene by virtual eye-movements evokes high precision and low noise in V1 neurons. *Frontiers in neural circuits* **7**, 206 (2013).
18. Goris, R. L., Simoncelli, E. P. & Movshon, J. A. Origin and function of tuning diversity in macaque visual cortex. *Neuron* **88**, 819–831 (2015).
19. Hénaff, O. J., Boundy-Singer, Z. M., Meding, K., Ziemba, C. M. & Goris, R. L. Representation of visual uncertainty through neural gain variability. *Nature communications* **11**, 1–12 (2020).
20. Berens, P. *et al.* A fast and simple population code for orientation in primate V1. *Journal of Neuroscience* **32**, 10618–10626 (2012).
21. Leon, P. S., Vanzetta, I., Masson, G. S. & Perrinet, L. U. Motion clouds: model-based stimulus synthesis of natural-like random textures for the study of motion perception. *Journal of neurophysiology* **107**, 3217–3226 (2012).
22. Field, D. J. Relations between the statistics of natural images and the response properties of cortical cells. *Josa a* **4**, 2379–2394 (1987).
23. Naka, K. & Rushton, W. A. S-potentials from colour units in the retina of fish (cyprinidae). *The Journal of physiology* **185**, 536–555 (1966).
24. Ringach, D. L., Shapley, R. M. & Hawken, M. J. Orientation selectivity in macaque V1: diversity and laminar dependence. *Journal of Neuroscience* **22**, 5639–5651 (2002).
25. Ringach, D. L., Hawken, M. J. & Shapley, R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature* **387**, 281–284 (1997).
26. Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung c* **36**, 910–912 (1981).
27. Kinouchi, O. & Copelli, M. Optimal dynamical range of excitable networks at criticality. *Nature physics* **2**, 348–351 (2006).
28. Douglas, R. J., Martin, K. A. & Whitteridge, D. A canonical microcircuit for neocortex. *Neural computation* **1**, 480–488 (1989).
29. Chavane, F. *et al.* Lateral spread of orientation selectivity in v1 is controlled by intracortical cooperativity. *Frontiers in systems neuroscience* **5**, 4 (2011).
30. Carandini, M. & Ringach, D. L. Predictions of a recurrent model of orientation selectivity. *Vision research* **37**, 3061–3071 (1997).
31. Gizzi, M. S., Katz, E., Schumer, R. A. & Movshon, J. A. Selectivity for orientation and direction of motion of single neurons in cat striate and extrastriate visual cortex. *Journal of neurophysiology* **63**, 1529–1543 (1990).
32. Denève, S. & Machens, C. K. Efficient codes and balanced networks. *Nature neuroscience* **19**, 375–382 (2016).
33. Monier, C., Chavane, F., Baudot, P., Graham, L. J. & Frégnac, Y. Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* **37**, 663–680 (2003).
34. Bishop, C. M. *Pattern recognition and machine learning* (springer, 2006).
35. Ohki, K. *et al.* Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* **442**, 925–928 (2006).
36. Martinez-Conde, S. *et al.* Effects of feedback projections from area 18 layers 2/3 to area 17 layers 2/3 in the cat visual cortex. *Journal of Neurophysiology* **82**, 2667–2675 (1999).
37. Wang, C., Waleszczyk, W., Burke, W. & Dreher, B. Modulatory influence of feedback projections from area 21a on neuronal activities in striate cortex of the cat. *Cerebral Cortex* **10**, 1217–1232 (2000).
38. Galuske, R. A., Schmidt, K. E., Goebel, R., Lomber, S. G. & Payne, B. R. The role of feedback in shaping neural representations in cat visual cortex. *Proceedings of the National Academy of Sciences* **99**, 17083–17088 (2002).
39. Huang, L., Chen, X. & Shou, T. Spatial frequency-dependent feedback of visual cortical area 21a modulating functional orientation column maps in areas 17 and 18 of the cat. *Brain research* **998**, 194–201 (2004).
40. Shen, W., Liang, Z. & Shou, T. Weakened feedback abolishes neural oblique effect evoked by pseudo-natural visual stimuli in area 17 of the cat. *Neuroscience letters* **437**, 65–70 (2008).
41. Wang, C., Huang, J. Y., Bardy, C., FitzGibbon, T. & Dreher, B. Influence of 'feedback'signals on spatial integration in receptive fields of cat area 17 neurons. *Brain research* **1328**, 34–48 (2010).
42. Ito, M. & Gilbert, C. D. Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* **22**, 593–604 (1999).
43. Kanai, R., Komura, Y., Shipp, S. & Friston, K. Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences* **370**, 20140169 (2015).
44. Merabet, L., Desautels, A., Minville, K. & Casanova, C. Motion integration in a thalamic visual nucleus. *Nature* **396**, 265–268 (1998).
45. de Souza, B. O. F., Cortes, N. & Casanova, C. Pulvinar modulates contrast responses in the visual cortex as a function of cortical hierarchy. *Cerebral Cortex* **30**, 1068–1086 (2020).
46. Cortes, N. H. & Van Vreeswijk, C. The role of pulvinar in the transmission of information in the visual hierarchy. *Frontiers in computational neuroscience* **6**, 29 (2012).
47. Cortes, N. & Vreeswijk, C. v. Pulvinar thalamic nucleus allows for asynchronous spike propagation through the cortex. *Frontiers in computational neuroscience* **9**, 60 (2015).
48. Millidge, B., Seth, A. & Buckley, C. L. Predictive coding: a theoretical and experimental review. *arXiv preprint arXiv:2107.12979* (2021).
49. Bastos, A. M. *et al.* Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
50. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).
51. Schindel, R., Rowlands, J. & Arnold, D. H. The oddball effect: Perceived duration and predictive coding. *Journal of vision* **11**, 17–17 (2011).
52. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
53. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).
54. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
55. Paszke, A. *et al.* Pytorch: An imperative style, high-performance deep learning library. In Wallach, H. *et al.* (eds.) *Advances in Neural Information Processing Systems 32*, 8024–8035 (Curran Associates, Inc., 2019).
56. Newville, M. *et al.* Lmfit: Non-linear least-square minimization and curve-fitting for python. *ascl* ascl–1606 (2016).
57. Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* **9**, 90–95 (2007).
58. Villeneuve, M. Y. & Casanova, C. On the use of isoflurane versus halothane in the study of visual response properties of single cells in the primary visual cortex. *Journal of neuroscience methods* **129**, 19–31 (2003).
59. Siegle, J. H. *et al.* Open ephys: an open-source, plugin-based platform for multichannel electrophysiology. *Journal of neural engineering* **14**, 045003 (2017).
60. Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M. & Harris, K. D. Fast and accurate

spike sorting of high-channel count probes with kilosort. In *Advances in neural information processing systems*, 4448–4456 (2016).
61. Rossant, C. *et al.* Spike sorting for large, dense electrode arrays. *Nature neuroscience* **19**, 634–641 (2016).
62. Binzegger, T., Douglas, R. J. & Martin, K. A. A quantitative map of the circuit of cat primary visual cortex. *Journal of Neuroscience* **24**, 8441–8453 (2004).
63. Peirce, J. *et al.* Psychopy2: Experiments in behavior made easy. *Behavior research methods* **51**, 195–203 (2019).
64. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nature Reviews Neuroscience* **13**, 51–62 (2012).
65. Movshon, J. A., Thompson, I. & Tolhurst, D. Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *The Journal of physiology* **283**, 101–120 (1978).
66. Smith, M. A., Majaj, N. J. & Movshon, J. A. Dynamics of motion signaling by neurons in macaque area mt. *Nature neuroscience* **8**, 220–228 (2005).
67. Swindale, N. V. Orientation tuning curves: empirical description and estimation of parameters. *Biological cybernetics* **78**, 45–56 (1998).
68. Brodersen, K. H., Ong, C. S., Stephan, K. E. & Buhmann, J. M. The balanced accuracy and its posterior distribution. In *2010 20th International Conference on Pattern Recognition*, 3121–3124 (IEEE, 2010).
69. Quiroga, R. Q., Reddy, L., Koch, C. & Fried, I. Decoding visual inputs from multiple neurons in the human temporal lobe. *Journal of neurophysiology* **98**, 1997–2007 (2007).
70. Guitchounts, G., Masis, J., Wolff, S. B. & Cox, D. Encoding of 3d head orienting movements in the primary visual cortex. *Neuron* **108**, 512–525 (2020).

Materials and Methods

Animal procedures. Experiments were performed on 3 adult cats (3.6 - 6.0 kg, 2 males). All surgical and experimental procedures were carried out in compliance with the guidelines of the Canadian Council on Animal Care and were approved by the Ethics Committee of the University of Montreal (CDEA #20-006). Animals were first administered atropine (0.1 mg/kg) and acepromazine (Atravet, 1mg/kg) subcutaneously to reduce the parasympathetic effects of anesthesia and provoke sedation, respectively. Anesthesia was induced with 3.5% Isoflurane in a 50:50 (v/v) mixture of O_2 and N_2O. Isoflurane concentration was maintained at 1.5% during surgical procedures. A tracheotomy was performed and animals were immobilized using an intravenous injection of 2% gallamine triethiodide. Animals were then artificially ventilated and a 1:1 (v/v) solution of 2% gallamine triethiodide (10 mg/kg/h) in 5% of dextrose in lactated ringer solution was continuously administered to maintain muscle relaxation. Throughout the experiment, expired level of CO_2 was maintained between 35 and 40 mmHg by adjustement of the tidal volume and respiratory rate. Heart rate was monitored and body temperature was maintained at 37°C using a feedback controlled heated blanket. Dexamethasone (4mg) was administered intramuscularly every 12h to reduce cortical swelling. Pupils were dilated using atropine (Mydriacyl) and nictitaing membranes were retracted using phenylephrine (Midfrin). Rigid contact lenses of appropriate power were used to correct the eyes' refraction and eye lubricant was used to avoid corneal dehydratation. Lidocaine hydrochloride (2%) was used in all incisions and pressure points. A craniotomy was performed between Horsley-Clarke coordinates 4-8 P ; 0.5 - 2 L to access the area 17 (V1) contralateral to the stimulation side. Small durectomies were performed for each electrode penetration. A 2% agar solution in saline was applied over the exposed regions to stabilize recordings and avoid the drying of cortical surface.

Electrophysiological recordings. During recording sessions, anesthesia was changed to 0.5-1% Halothane, as isoflurane has been shown to yield a depression of visual responses (58). Extracellular activity was recorded using 32 channel linear probes (≈1 MΩ, 1×32-6 mm-100-177, Neuronexus) and acquired at 30 KHz using an Open Ephys acquisition board (59). Single units were isolated using Kilosort (60) and manually curated using the phy software (61). Clusters of low amplitude templates or ill-defined margins were excluded from analysis. Further exclusion was carried if the firing rate of a cluster dropped below 5 spikes.s^{-1} for more than 30 seconds or its tuning curve was poorly fitted with a von Mises distribution ($r^2 < .75$), leaving 249 putative neurons for further analysis. Clusters were characterised as belonging to supragranular layers if they were located below a fixed distance of 500µm (62) from the top of the electrode (i.e. 5 contacts).

Visual Stimulation. Visual stimuli were generated using Psychopy (63) and were projected monocularly with a PROPixx projector (VPixx Technologies Inc., St-Bruno, QC, Canada) onto an isoluminant screen (Da-Lite© screen) located 57 cm from the animal's eye, covering $104° \times 79°$ of visual angle with a mean luminance of 25 cd/m^2. The stimuli used here, Motion Clouds (21), are a class of band-pass filtered white-noise textures (64) that provide parametric control over the content of the stimuli while retaining the statistics of natural images (22). The envelope of the filters in Fourier space are Gaussians in the coordinates of the relevant axis, in which they are described by their mean and precision. As such, a Motion Cloud (MC) is defined as:

$$\mathrm{MC} = \mathcal{F}^{-1}(G(f_0, B_f) \cdot O(\theta, B_\theta)) \quad \textbf{(1)}$$

where $\mathcal{F}$ is the Fourier transform, G the spatial frequency envelope and O the orientation envelope. The spatial frequency envelope follows a log-normal distribution:

$$G(f_0, B_f)(f_x, f_y) = \frac{1}{f_r} \exp\left\{ -\frac{1}{2} \left(\frac{\ln(\frac{f_r}{f_0})}{\ln(\frac{f_0+B_f}{f_0})} \right)^2 \right\} \quad \textbf{(2)}$$

where $f_r = \sqrt{f_x^2 + f_y^2}$ is the radial frequency, f_0 is the mean spatial frequency and B_f is the precision of the spatial frequency distribution, both in cycles per degree. The orientation envelope is a von Mises distribution:

$$O(\theta, B_\theta)(f_x, f_y) = \exp\left\{ \frac{\cos(2(\theta_f - \theta))}{4 \cdot B_\theta^2} \right\} \quad \textbf{(3)}$$

where θ_f is the angle of f_x, f_y in the Fourier plane, θ is the mean orientation and B_θ is the variance of the orientation distribution. For small B_θ, this distribution is close to a Gaussian and B_θ measures its standard deviation. For all stimuli, the spatial frequency parameters were set at $f_0 = B_f = 0.9$ cpd and orthogonal drift speed was set to 10°/s, within the response range of area 17 neurons (65). θ was varied in 12 even steps from 0 to π rad and B_θ in 8 even steps from $\pi/5$ to ≈ 0. All stimuli were displayed at 100% contrast, for 300 ms, interleaved with the presentation of a mean luminance screen for 150 ms. Trials were fully randomized and each stimulus was presented 15 times.

Single Neuron Analysis. Tuning curves were constructed by selecting a 300ms window maximizing spike-count variance (66), in which firing rate was averaged and baseline substracted. Tuning curves were averaged across drift directions and a von Mises distribution (67) was fitted to the data:

$$f(\theta_k) = R_0 + (R_{\text{max}} - R_0) \cdot \exp\{\kappa \cdot (\cos(2(\theta_k - \theta)) - 1)\} \quad \textbf{(4)}$$

where R_{max} is the response at the preferred orientation θ, R_0 the response at the orthogonal orientation, κ is a measure of concentration and θ_k the orientation of the stimuli. A global measure of the orientation tuning was assessed by computing the circular variance (CV) of the raw tuning curve data:

$$\text{CV} = 1 - \left| \frac{\sum_k R(\theta_k) \exp\{2i\theta_k\}}{\sum_k R(\theta_k)} \right| \quad \textbf{(5)}$$

where $R(\theta_k)$ is the response of a neuron to a stimuli of angle θ_k. CV varies from 0 for exceptionally orientation-selective neurons to 1 for orientation-untuned neurons (24). The significance of orientation tuning was assessed by comparing the firing rate of the neuron at the preferred and orthogonal orientations for all trials using a Wilcoxon signed-rank test, corrected for continuity. Changes of preferred orientation were assessed as the difference of $\arg\max f(\theta_k)$ between trials where $B_\theta = 0.0°$ and trials where B_θ was great enough so that a neuron was no longer significantly orientation tuned. Variations of peak amplitude was measured by comparing the firing rate at preferred orientation for both these conditions.

Precision-Response Functions. The variation of circular variance as a function of B_θ was assessed with a Naka-Rushton function:

$$f(B_\theta) = f_0 + f_{\text{max}} \frac{B_\theta^n}{B_\theta^n + B_{\theta 50}^n} \quad \textbf{(6)}$$

where f_0 is the base value of the function, f_{max} its maximal value, $B_{\theta 50}$ the orientation precision at half f_{max} and n a strictly positive exponent of the function (23).

Computational model. A phenomenological population rate model of a recurrent network was used to reproduce the activity of single neurons. A population of 100 neurons tiling the orientation space in even steps between 0 and πrad was modeled as an homogeneous set of orientation-selective units. The orientation distribution of Motion Clouds (Figure 1) with 20 precisions from $B_\theta = \pi/5$ to ≈ 0 was used to drive the activity of these neurons, according to the equation:

$$f(\theta_i) = \text{MC}_{\theta i} - w(\theta_i) \quad \textbf{(7)}$$

where $f(\theta_i)$ is the activity of a neuron with preferred orientation θ_i, $\text{MC}_{\theta i}$ the value of the orientation distribution of the input on this orientation and $w(\theta_i)$ the weight of the neighbouring inhibitory synapses, described by a von Mises function:

$$w(\theta_i) = w_{\text{max}} \frac{\exp(\kappa \cos(\theta_i - \theta))}{2\pi \text{I}_0(\kappa)} \quad \textbf{(8)}$$

where $\text{I}_0(\kappa)$ the Bessel function of order 0 and, θ_i is a neighbouring neuron with index i, w_{max} the maximum synaptic weight and κ the measure of concentration of the function. The effect of w_{max} and κ were explored by varying both parameters with 60 values, between $0.0; 40.0$ and $0.0; 15.0$, respectively. Furthermore, a non-linear threshold was added to the response of each neuron, such that neuron who fell beneath this threshold failed to produce a spiking activity. This parameter was not varied, and an empirical value of $\approx 10\%$ of the maximum input to the model (when $B_\theta \approx 0$) was used.

Population Decoding. The stimulus identity was decoded using a multinomial logistic regression classifier (34), which can be interpreted as an instantaneous, neurally plausible, readout of the population activity (20). For a given stimulus, this population activity was a vector $X(t) = \begin{bmatrix} X_1(t) & X_2(t) & \cdots & X_{249}(t) \end{bmatrix}$, where $X_i(t)$ is the spike count of neuron i in time window $[t; t + \Delta T]$, where t is the time bin and ΔT the size of the integration window, usually 100 ms. Time t was varied from -200 ms to 400 ms, relative to the stimulation onset, in steps of 10 ms. Each time bin was labeled as the end of the time window and was hence reported as $t + \Delta T$.

Practically, the multinomial logistic regression is an extension of the binary logistic regression (34), which was here trained to classify the neural activity vector between K possible classes. The probability of any such vector to belong to a given class is:

$$P(y = k | X(t)) = \frac{\exp\{\langle \beta_k, X(t) \rangle\}}{\sum_{k'=1}^{K} \exp\{\langle \beta_{k'}, X(t) \rangle\}} \quad \textbf{(9)}$$

where $\langle \cdot, \cdot \rangle$ is the scalar product over the different neurons, $k = 1, \ldots, K$ is the class out of K possible values and β_k are the vectors of learned coefficients of the classifier. We trained several such classifiers, to decode orientation θ ($K = 12$), orientation precision B_θ ($K = 8$) or both ($K = 12 \times 8 = 96$). All meta-parameters (the integration window size, penalty type, regularization strength and train/test split size) were controlled, showing that the decoder is optimally parameterized and that its performances result from experimental data and not decoder parametrization (Supplementary Figure 4).

Decoding accuracy was reported as the average accuracy across classes, also known as the balanced accuracy score (68), which accounts for possible imbalances in the learning or testing datasets. Decoding performance was also reported in the form of confusion matrices, in which the values on each row i and column j represents the normalized number of times a stimuli of class $k = i$ is predicted to belong to the class $k = j$. Hence, a perfect decoder would produce a perfectly diagonal confusion matrix (unit matrix). We reported the significativity of decoders by splitting the population activity in separate training and testing data sets, then performing 6 different such splits on different groups of neurons, comparing the resulting confusion matrices to class-shuffled decoder matrices. Previous publications have reported population decoding by merging neural activity across electrodes or experiments (69, 70), which we validated in our data (Supplementary Figure 5).

Decoding time course. To estimate the temporal parameters of the decoders, the evolution of the population decoding accuracy through time was fitted with a difference of two sigmoid functions:

$$\sigma_1 - \sigma_2 = \max_{\text{acc}} \left(\frac{1}{1+e^{-k_1\tau_1}} - \frac{1}{1+e^{-k_2\tau_2}} \right) + \min_{\text{acc}} \tag{10}$$

where $\max_{\text{acc}}$ and $\min_{\text{acc}}$ are respectively the maximum and minimum accuracies of the decoder, k the steepness and τ the time constant of the function.

Supplementary Figures

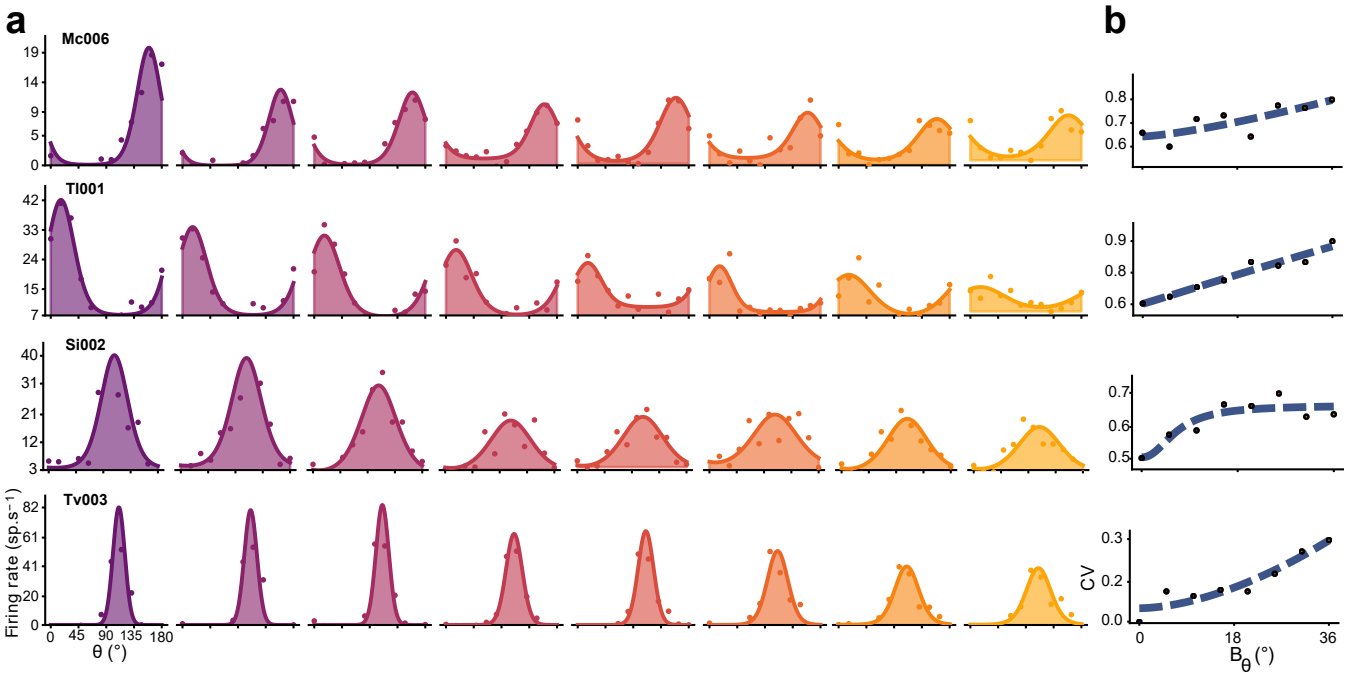

Supplementary Figure 1. Additional examples of single neuron tuning curves and PRFs. (**a**) Tuning curves of 4 additional neurons in response to Motion Clouds of increasing B_θ (left to right). Colored points represent the mean firing rate across trials (baseline subtracted, 300 ms average), with lines indicating a fitted von Mises function. (**b**) Precision-response functions (PRFs), assessing the changes of orientation tuning as measured by the Circular Variance (CV) as a function of B_θ, fitted with Naka-Rushton (NKR) functions (dashed curves)

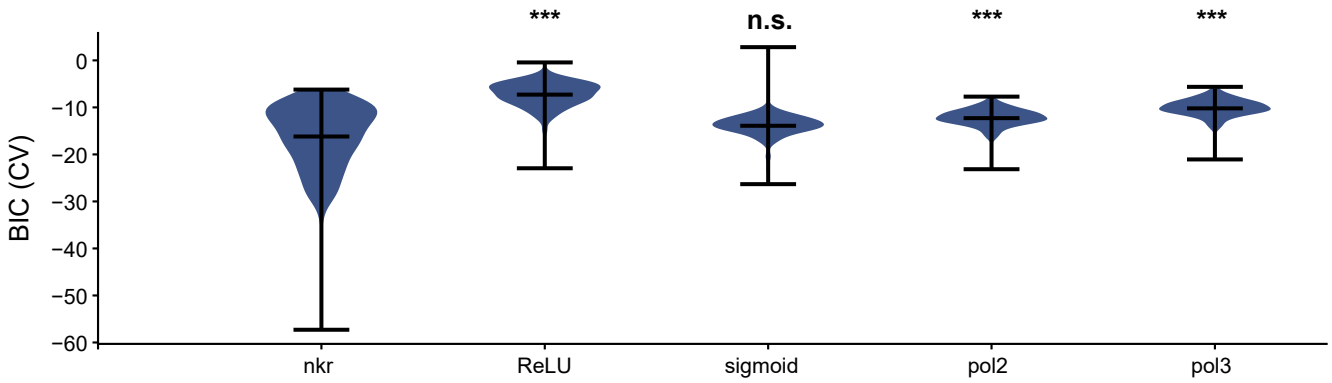

Supplementary Figure 2. Criterion on the model selection for fitting the PRFs. Violin plot of the Bayesian Information Criterion (BIC) of the CV curves of all recorded neurons. Each violin plot represent a different type of fitted equation, respectively: Naka-Rushton (`nkr`, see Materials and Methods) ; Rectified Linear Unit (`ReLU`, $f(x) = max(0, x)$) ; logistic function (`sigmoid`, $f(x) = \frac{e^x}{e^x+1}$) ; second degree polynomial function (`pol2`, $f(x) = ax^2 + bx + c$) and third degree polynomial function (`pol3`, $f(x) = ax^3 + bx^2 + cx + d$). A lower BIC indicates less information lost in the fitting process, hence a better fitting model. Naka-Rushton curves were chosen over sigmoid functions for the skewness of the BIC distributions towards negative values, as well as the explainability of their parameters. `n.s.`, not significant; $\star$, p < 0.05; $\star\star$, p < 0.01; $\star\star\star$, p < 0.001 (Kruskal-Wallis H-test, post-hoc Dunn Pairwise test, Bonferroni corrected).

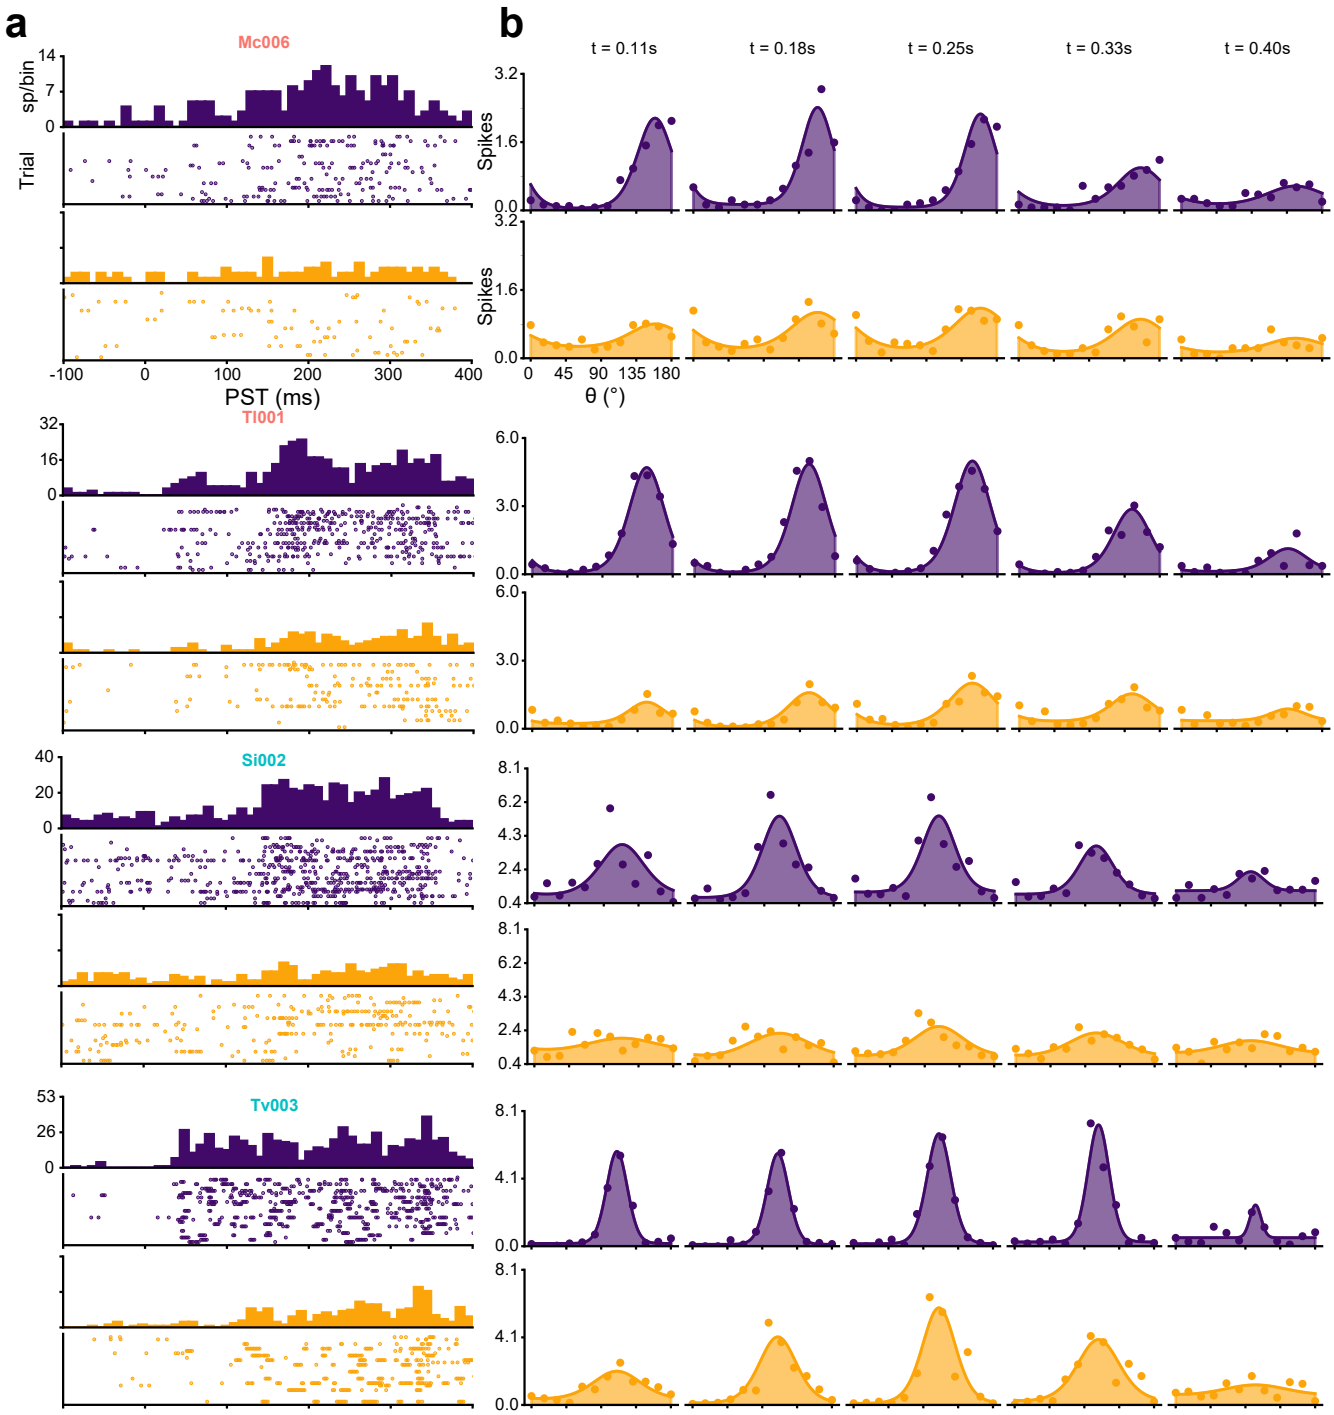

Supplementary Figure 3. Additional examples of dynamical responses of neurons, also shown in Supplementary Figure 1. Neuron labels are colored in red for neurons untuned to stimuli with $B_\theta = 36.0°$ ("vulnerable neurons"), and conversely in teal when still significantly tuned to the same low-precision stimuli ("resilient neurons") (**a**) Peri-stimulus time (PST) histogram and rasterplots of the example neurons, for trials with $B_\theta = 0.0°$ (purple) and $B_\theta = 36.0°$ (yellow). (**b**) Temporal evolution of the tuning curves of the same neurons, taken in a 100 ms window after the reported time.

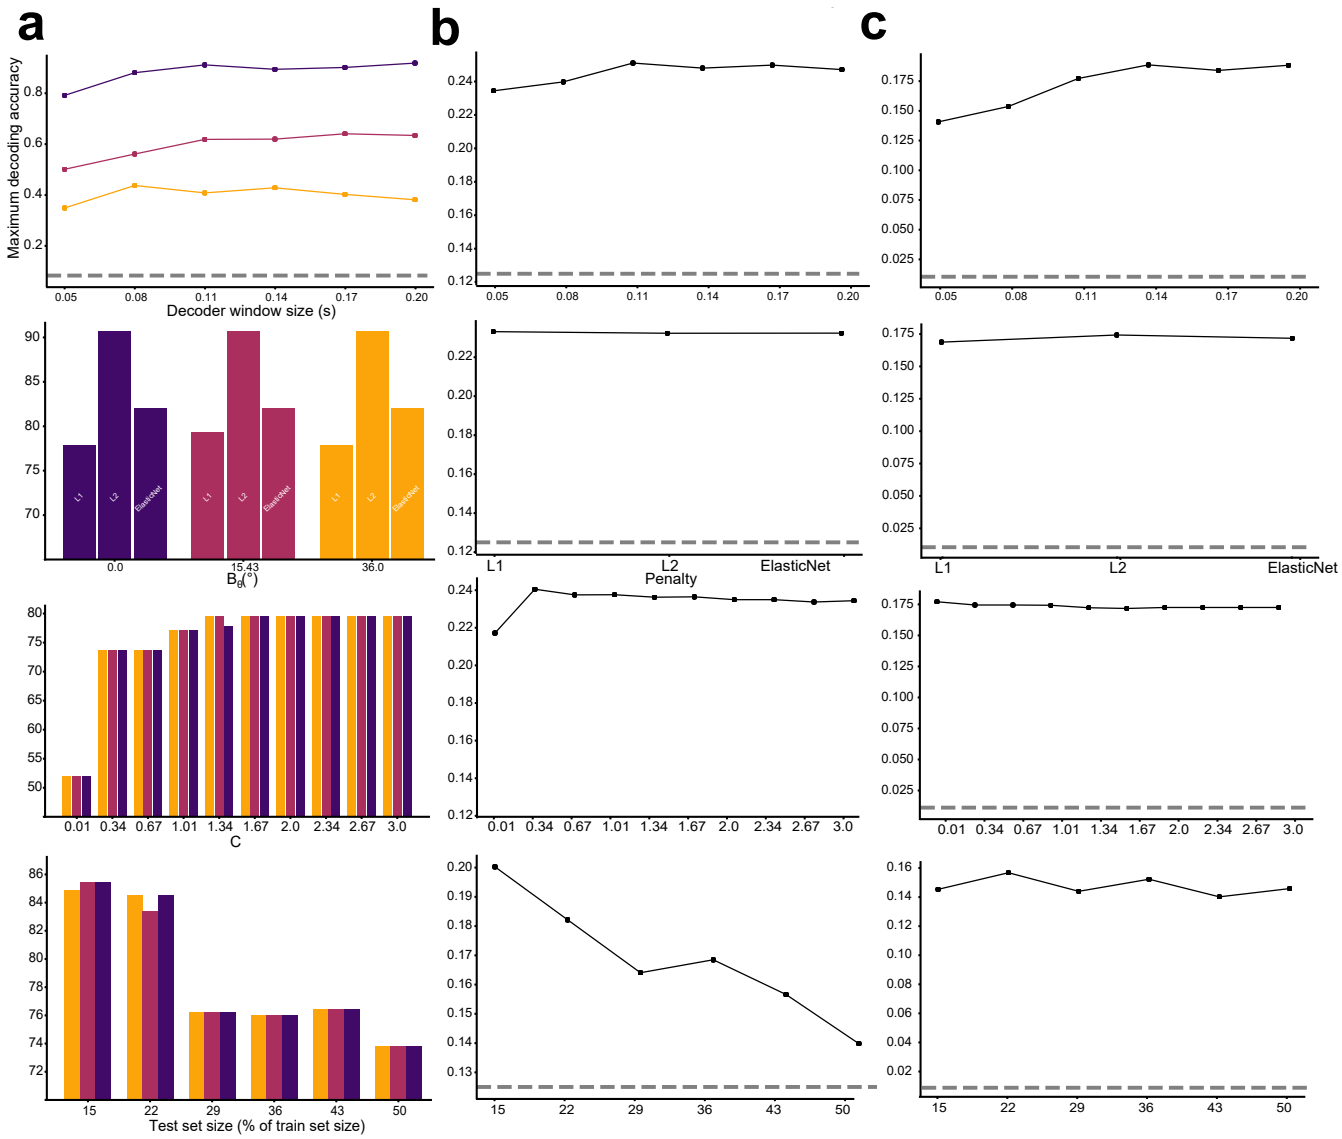

Supplementary Figure 4. Parameters evaluation of the decoders. (**a**) Parameter scanning of the orientation θ decoders for three B_θ. Maximum accuracy reached during the time course of each decoder was reported for each parameter. From top to bottom, the parameters optimized are: the length of the time window ΔT, the type of penalization norm applied to the decoder, the regularization strength parameter and the percentage of data kept out of the training set to evaluate the decoder's accuracy. (**b**) Parameters of the orientation precision B_θ decoder. (**c**) Parameters of the orientation and orientation precision $B_\theta \times \theta$ decoder.

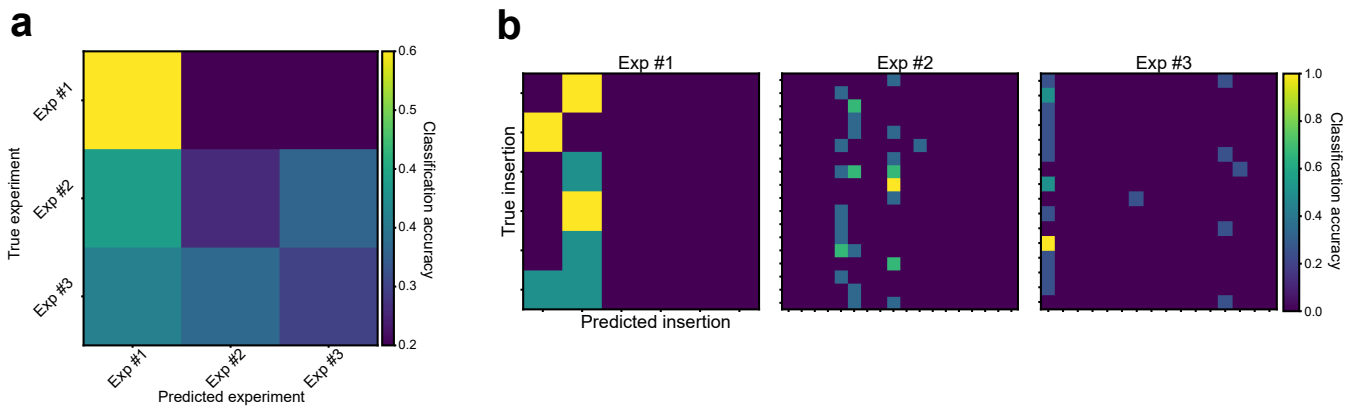

Supplementary Figure 5. Decoding the experiment identity and insertion identity of neurons. (**a**) Confusion matrix of a decoder trained to retrieve the experiment identity using three groups of 30 neurons (bootstrapped 1000 times). (**b**) Confusion matrices of three decoders trained to retrieve the insertion identity of the neurons recorded in each experiments.

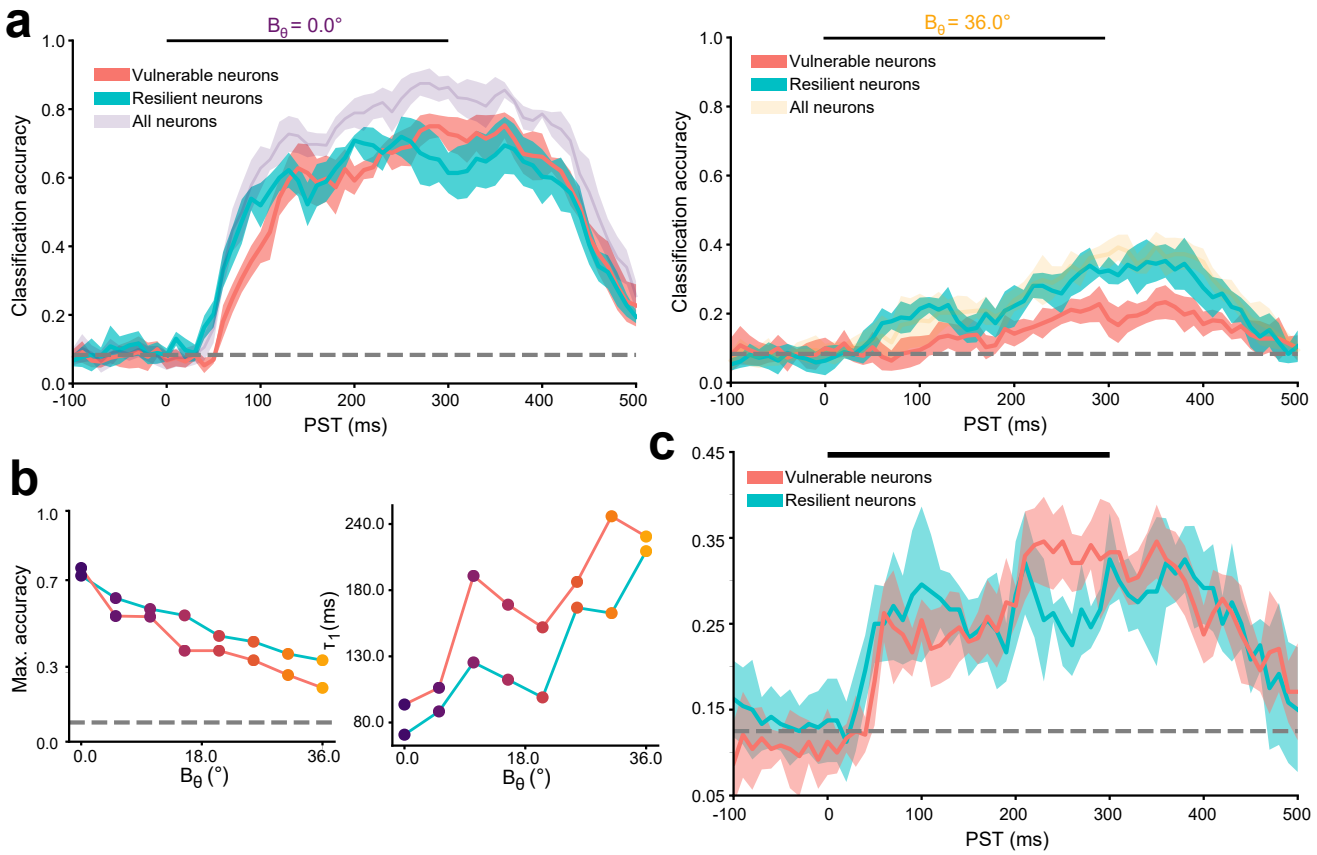

Supplementary Figure 6. Decoding stimulation properties from vulnerable and resilient neurons. (**a**) Left: Time course around the peri-stimulus time (PST) of the accuracy of three independent decoders, trained to predict the orientation θ of maximally precise Motion Clouds from either vulnerable, resilient, or all neurons (identical to Figure 5b). Right: Time course of three independent decoders, trained on the same task but for Motion Clouds of the lowest precision. Solid dark line represent the mean accuracy of a 5-fold cross validation and filled contour the SD. Decoding at chance level (here, $\frac{1}{12}$) is represented by a gray dashed line. (**b**) Parameters of the decoders' time courses trained from vulnerable and resilient neurons data, estimated by fitting a difference of two sigmoids. (**c**) Time course around the PST of the accuracy of two decoders, trained to retrieve the precision of the stimulation with the preferred θ prior (identically to Figure 6b,teal), from either vulnerable or resilient neurons.

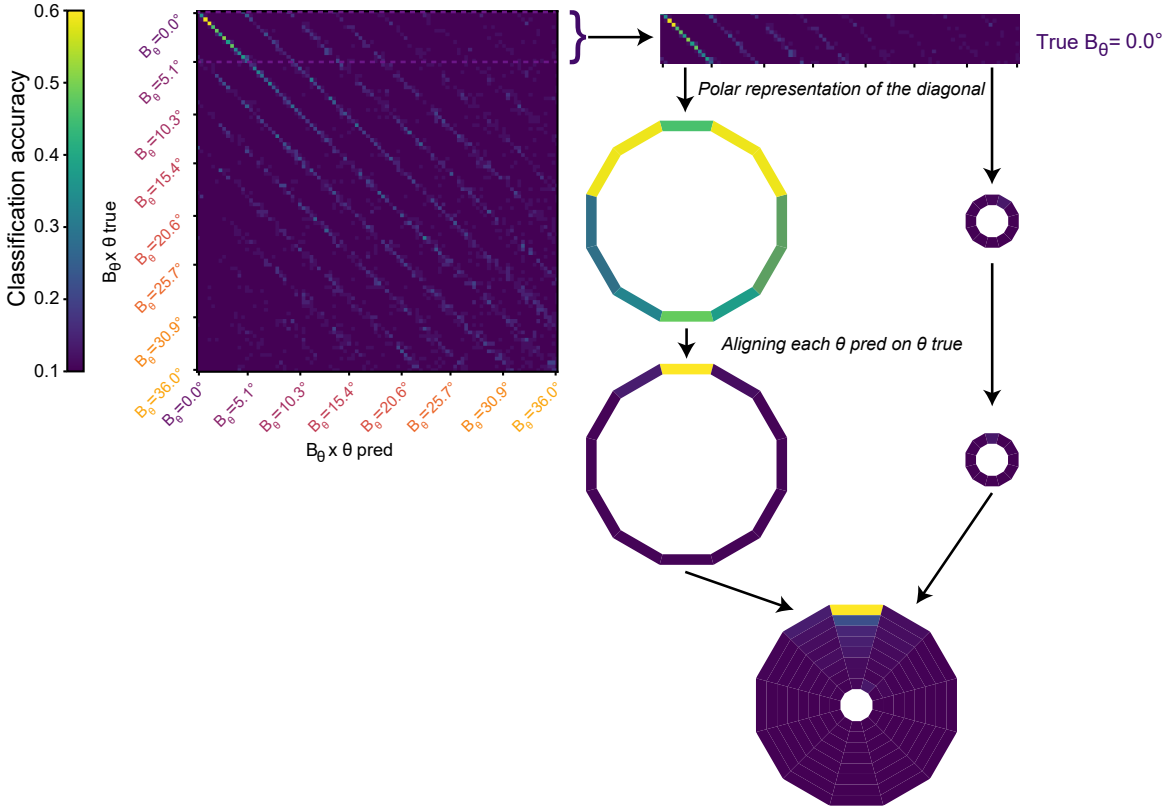

Supplementary Figure 7. Schematic representation of the process used to obtain polar representations from the $\theta \times B_\theta$ decoder's confusion matrices. For each row, in which all stimuli decoded belong to one given precision B_θ, the diagonal from each possible column, in which all stimuli are inferred as belonging to one of the 8 possible B_θ is taken and aligned over the true θ. This process is repeated for each row, which produces a full polar representation of the decoding output. When representing weight matrices, as in Figure 8, the initial data comes from weight vectors β_k (see materials and methods).

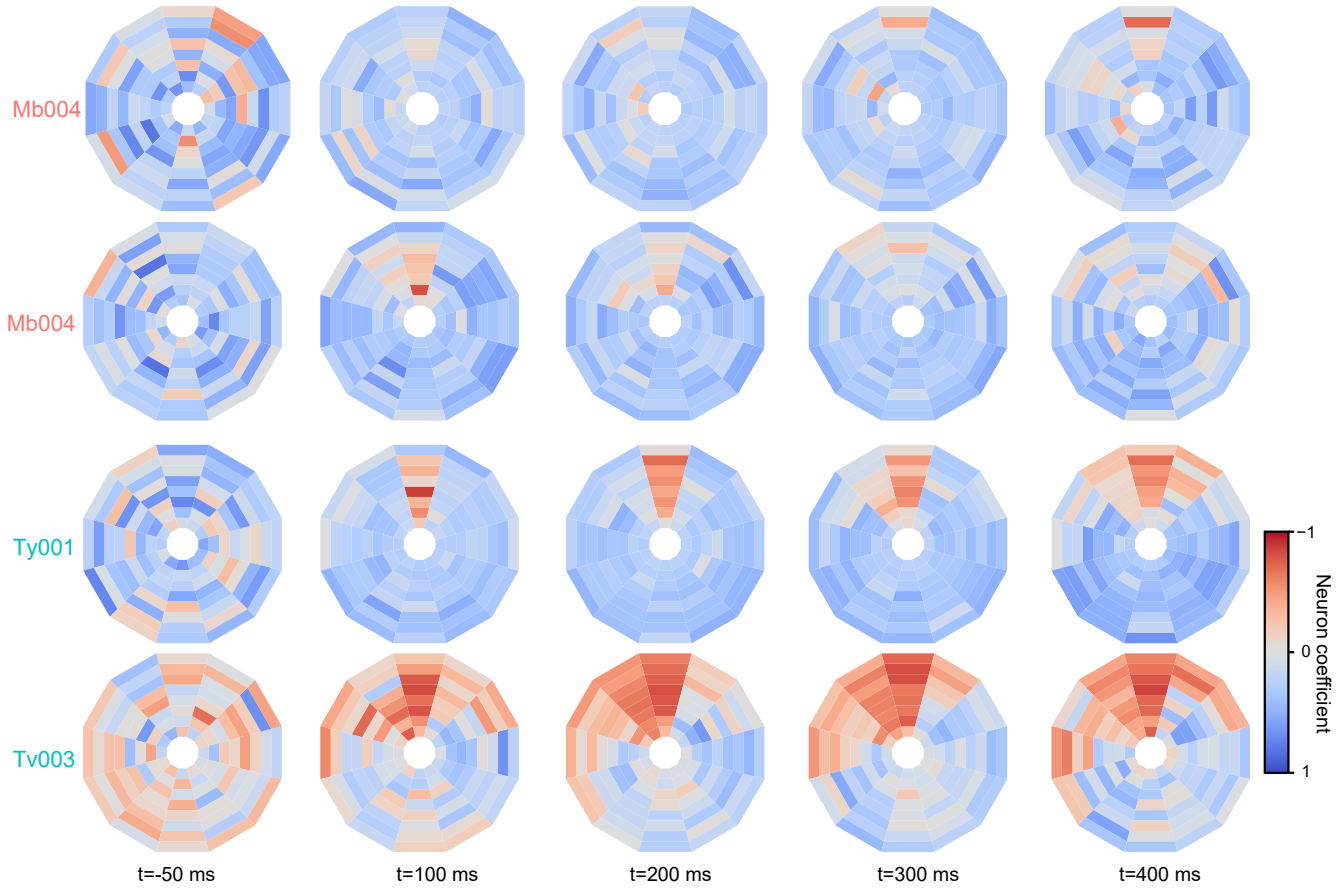

Supplementary Figure 8. Polar plot of weight matrices from additional neurons, coming from two vulnerable (top rows) and two resilient (bottom rows) neurons. As in Figure 7, the angle of each bin represents the error on the θ identity of the stimulus, Δ_θ, and the eccentricity corresponds to the weight for each B_θ (lower precisions towards the center). The temporal evolution of the weights from the decoder are normalized for each neurons by row.