A Brain-Wide Map of Neural Activity during Complex Behaviour

International Brain Laboratory*, Brandon Benson[1], Julius Benson[2], Daniel Birman[3], Niccolò Bonacchi[4], Matteo Carandini[5], Joana A Catarino[4], Gaelle A Chapuis[6], Anne K Churchland[7], Yang Dan[8], Peter Dayan[9], Eric EJ DeWitt[4], Tatiana A Engel[10], Michele Fabbri[4], Mayo Faulkner[5], Ila Rani Fiete[11], Charles Findling[6], Laura Freitas-Silva[4], Berk Gerçek[6], Kenneth D Harris[5], Michael Häusser[5], Sonja B Hofer[12], Fei Hu[8], Félix Hubert[6], Julia M Huntenburg[9], Anup Khanal[7], Christopher Krasniak[10], Christopher Langdon[11], Petrina Y P Lau[5], Zachary F Mainen[4], Guido T Meijer[4], Nathaniel J Miska[13], Thomas D Mrsic-Flogel[13], Jean-Paul Noel[2], Kai Nylund[3], Alejandro Pan-Vazquez[11], Alexandre Pouget[6], Cyrille Rossant[5], Noam Roth[3], Rylan Schaeffer[1], Michael Schartner[4], Yanliang Shi[11], Karolina Z Socha[5], Nicholas A Steinmetz[3], Karel Svoboda[14], Anne E Urai[15], Miles J Wells[5], Steven Jon West[13], Matthew R Whiteway[16], Olivier Winter[4], and Ilana B Witten[11]

*Virtual Entity; [1]Stanford University, USA; [2]New York University, USA; [3]University of Washington, USA; [4]Champalimaud Foundation, Portugal; [5]University College London, UK; [6]University of Geneva, Switzerland; [7]University of California Los Angeles, USA; [8]University of California Berkeley, USA; [9]Max Planck Institute; University of Tübingen, Germany; [10]Cold Spring Harbor Laboratory, NY, USA; [11]Princeton University, USA; [12]Massachusetts Institute of Technology, USA; [13]Sainsbury Wellcome Center, University College London, UK; [14]Allen Institute for Neural Dynamics, USA; [15]Leiden University, The Netherlands; [16]Columbia University, USA

Abstract

A key challenge in neuroscience is understanding how neurons in hundreds of interconnected brain regions integrate sensory inputs with prior expectations to initiate movements. It has proven difficult to meet this challenge when different laboratories apply different analyses to different recordings in different regions during different behaviours. Here, we report a comprehensive set of recordings from 115 mice in 11 labs performing a decision-making task with sensory, motor, and cognitive components, obtained with 547 Neuropixels probe insertions covering 267 brain areas in the left forebrain and midbrain and the right hindbrain and cerebellum. We provide an initial appraisal of this brain-wide map, assessing how neural activity encodes key task variables. Representations of visual stimuli appeared transiently in classical visual areas after stimulus onset and then spread to ramp-like activity in a collection of mid- and hindbrain regions that also encoded choices. Neural responses correlated with motor action almost everywhere in the brain. Responses to reward delivery and consumption versus reward omission were also widespread. Representations of objective prior expectations were weaker, found in sparse sets of neurons from restricted regions. This publicly available dataset represents an unprecedented resource for understanding how computations distributed across and within brain areas drive behaviour.

Introduction

How does coherent and effective output emerge from hundreds of interconnected brain areas processing information related to sensation, decisions, action, and behaviour[1–3]? To answer this question we need to know how the activity of individual neurons and populations of neurons reflect variables such as stimuli, expectations, choices, actions, rewards, and punishments[4]. This has duly been studied over decades using electrophysiological recordings. Until recently, however, technical limitations restricted these recordings to small invertebrates[5] or to a small number of brain areas, leaving much of the mammalian brain uncharted or described by fragmentary maps. For instance, the mouse brain comprises over 300 identified regions[6], of which only a minority has been recorded systematically in behaviourally equivalent settings. The regions studied were typically chosen based on *a priori* hypotheses derived from previous recordings and from anatomical connectivity. This approach can suggest a localisation of function, revealing brain regions as being engaged in computations such as the accumulation of sensory evidence in favour of a decision[7]. However, in some circumstances[8–12], a region that shows this encoding can be silenced without substantial behavioural consequences, suggesting the involvement of other regions. Overall, it has proven difficult to obtain a comprehensive picture of neural processing based on different reports from different laboratories recording in different regions during different behaviours and analysing the data with different methods.

A broader search for the neuronal correlates of variables such as sensation and decision-making thus requires systematic recording at a wider scale, using a single task with sufficient behavioural complexity, and using the same analysis methods on all the data. Obtaining such a comprehensive dataset has recently become possible with advances in recording technology. In a species with a small brain such as the mouse, Neuropixels probes[13] have enabled larger-scale recordings, such as sampling activity across eight visual areas[14], or tens of brain regions in mice performing behavioural tasks[15–17], or experiencing changes in physiological state[18]. Modern imaging techniques also provide a wider view over activity across dorsal cortical regions[19–21]. The results emerging from these broad surveys suggest that the encoding of task variables varies greatly: some neural correlates are found only in few brain areas, others in sparse sets of cells, while other variables appear to be distributed much more broadly. Thus, it is critical to record even more fully, because past recordings may have missed essential regions with focused coding of certain variables, and have not fully characterised the nature of distributed coding.

Here we present a publicly available dataset[22] of recordings from 547 electrodes spaced across the entire brain in mice performing a behavioural task that requires sensory, cognitive, and motor processing[23]. This approach enables detection of brain-wide correlates of sensation, choice, action, and reward, as well as internal cognitive states including stimulus expectation or priors. We also describe initial analyses of these data, which indicate that the neural correlates of some variables, such as reward and action, can be found in many neurons across essentially the whole brain, while those of other variables such as prior expectations can be decoded from a narrower range of regions, and significantly influence the activity of many fewer individual neurons. These data are intended to be the starting point for a detailed examination of the algorithmic realisation of decision-making across the brain, and are a unique resource enabling the community to perform a broad range of further analyses at a brain-wide scale.

Results

We first describe the task and recording strategy; further details of how we ensured reproducible behaviour, electrophysiology, and videography are available separately[23,24] and are summarised in the Methods. Next, we describe the set of analysis methods we used to provide different views of a rich and complex dataset. Finally, we report the neural correlates of the main events and variables in the task: the visual stimulus, choice, feedback, the expectation of the upcoming stimulus or choice, and finally movement.

Behavioural task and recording. We trained 115 mice (80 males, 35 females) on the International Brain Laboratory (IBL) decision-making task[23] (Fig. 1a;b). On each trial, a visual stimulus appeared to the left or right, and the mouse had to move it to the centre by turning a wheel with its front paws within 60 seconds (Fig. 1c). The prior probability for the stimulus to appear on the left/right side was constant over a block of trials, at 20/80% (right block) or 80/20% (left block). Blocks lasted for between 20 and 100 trials, drawn from a truncated geometric distribution (mean 51 trials). Stimulus contrast was sampled uniformly from five possible values (100, 25, 12.5, 6.25, and 0%). The 0% contrast trials, when no stimulus was presented, were assigned to a left or right side following the probability distribution of the block, allowing mice to perform above chance by incorporating this prior in their choices. Following a wheel turn, mice received positive feedback in the form of a water reward, or negative feedback in the form of a white noise pulse and a 2 s time-out. The next trial began after a delay, followed by a quiescence period during which the mice had to hold the wheel still.

In these mice, we performed 547 Neuropixels probe insertions (see an example of one recording of three trials in Fig. 1d), following a grid covering the left hemisphere of the forebrain and midbrain, which typically represent stimuli or actions on the contralateral side, and the right hemisphere of the cerebellum and hindbrain, which typically represent the ipsilateral side (Fig. 2a). These recordings were collected by 11 labs in Europe and the USA, with most recordings using two simultaneous probe insertions. To ensure reproducibility, one brain location was targeted in every mouse in every laboratory, as described elsewhere[24]. Only sessions with at least 400 trials and behavioural performance of 90% on the 100% contrast trials were retained for further analysis. Data were uploaded to a central server, preprocessed, and shared through a standardised interface[25]. To perform spike-sorting on the recordings, we used a version of Kilosort[26], with custom additions[27]. This yielded 295,501 neurons (including multi-neuron activity), averaging 540 per probe. To separate individual neurons from clusters of multi-neuron activity, we then applied stringent quality control metrics (based on those described in Ref.[14]), which identified 32,766 'good' neurons, averaging 60 per probe.

Following recordings, electrode tracks were reconstructed using serial-section 2-photon microscopy[28], and each recording site and neuron was assigned a region in the Allen Common Coordinate Framework[6] (CCF; details in online table; statistics in Fig. 2b). Our main analyses are restricted to regions with 10 or more neurons assigned to them in at least two sessions. Due to the grid-based electrode insertion strategy, more recordings were made in larger regions, typically leading to more neurons. Note that it is harder to extract good quality neurons from some regions than others, and so the yield differs substantially. We map our main results into a flatmap of the brain[29] (Fig. 2c); supplementary figures present some of the results

on more conventional 2D sections (Fig. S1). For reference, average activity across all regions aligned to stimulus onset, first movement, and feedback is shown in Fig. S2.

The processed data for each trial consisted of a set of spike trains from multiple brain regions together with continuous behavioural traces and discrete behavioural events (Fig. 1c;d). These were recorded using a variety of sensors including three video cameras and a rotary encoder on the choice wheel. These data were processed with custom scripts and DeepLabCut[30] to yield the times of major events in each trial (stimulus onset, first wheel movement, reward delivery or error tone) along with wheel velocity, whisker motion energy, lick timing, and the positions of body parts. We only analyzed trials in which the first wheel movement time (reaction time) was between 80 ms and 2 s (Fig. 1c).

Instructions for accessing the data[22], together with an online browser, are available at https://data.internationalbrainlab.org.

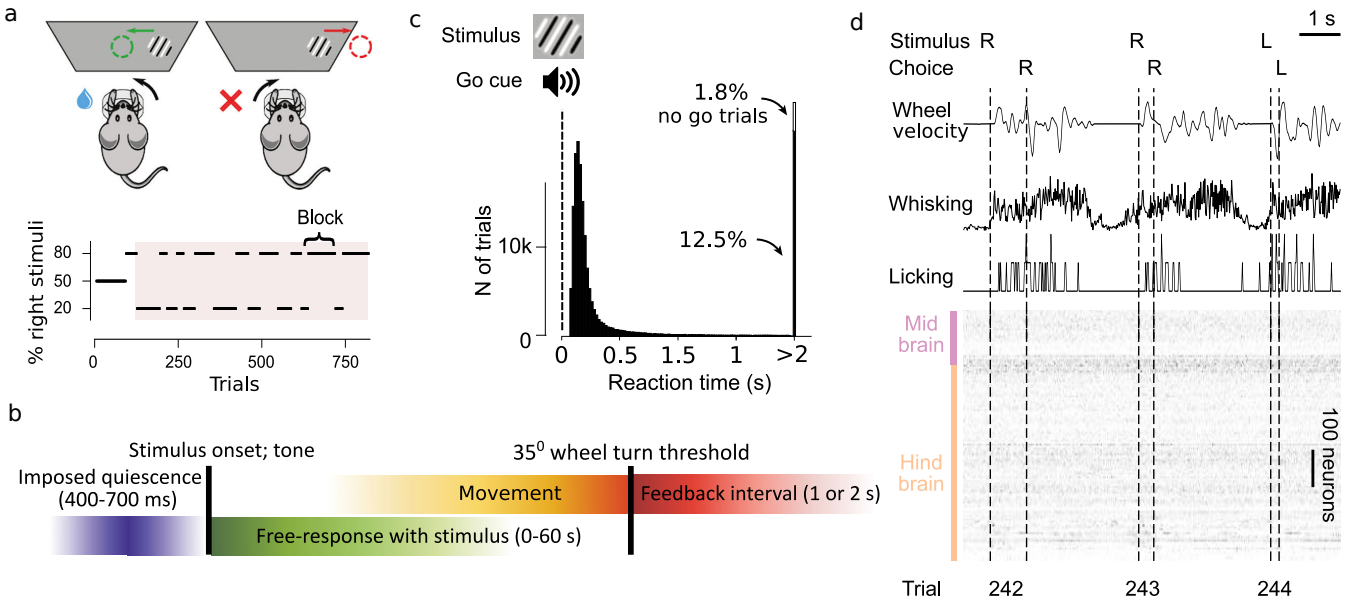

Figure 1. The IBL task and data types. a) Schematic of the IBL task and block structure of an example session. After the first 90 trials in a session, the probability of the stimulus being on the right side was varied in blocks of consecutive trials. **b)** Timeline of the IBL task with main events and the variables on which we concentrate. After a time of imposed stillness, the stimulus appears on the screen, then the first movement is detected by a small threshold of wheel rotation. When the wheel is turned by more than 35°, the animal is considered to have made a choice and feedback is provided. Colours are as in later figures. **c)** Distribution of times between stimulus onset and first wheel movement onset (which can be interpreted as a reaction time), from 354 sessions. The distribution is truncated at 80ms and 2 s, because we only analysed trials with first movement times between these bounds (78.3% of trials, see Methods). **d)** Example time series and key trial information for three trials, including the wheel's rotary encoder output, video analysis and rasters showing the spike times of simultaneously recorded neurons in multiple brain regions.

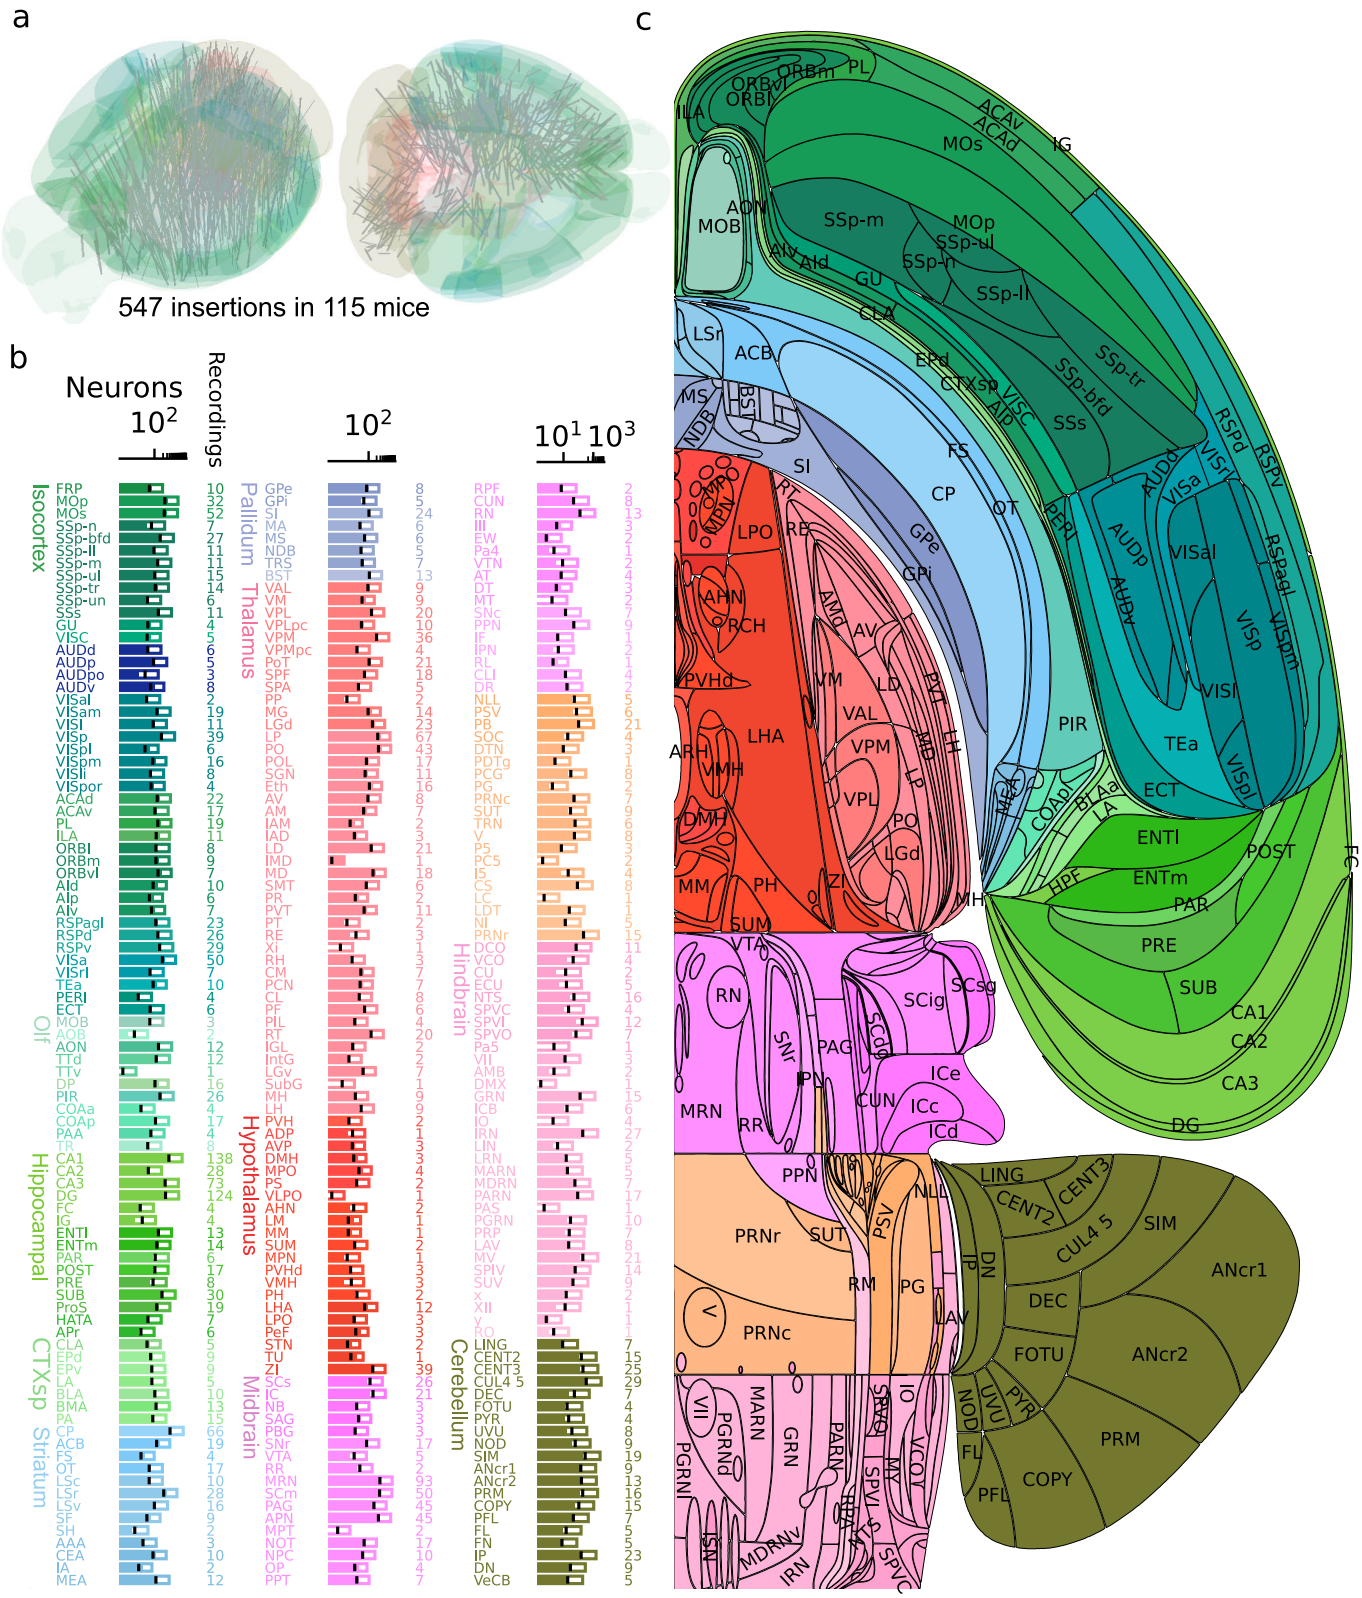

Figure 2. Brain-wide recordings during behaviour. **a)** Neuropixels probe trajectories shown in a 3D brain schematic. **b)** For each region, the number of neurons recorded (full bar length) and the number of good neurons used for analysis (filled portion; for reference, the black line on each bar shows 10% of the number of recorded neurons, which is the average number of neurons that were good). Definitions of the acronyms for brain regions are available online. The same table reports the so-called Cosmos hierarchical grouping of the regions, which distinguishes 'Isocortex', 'OLF' (olfactory areas), 'HPF' (hippocampal formation), 'CTXsp' (cortical subplate), 'CNU' (cerebral nuclei), 'TH' (thalamus), 'HY' (hypothalamus), 'MB' (midbrain), 'HB' (hindbrain), 'CB' (cerebellum). **c)** Flatmap of one hemisphere indicating the region acronyms.

Behavioural performance. As previously shown[23], mice learned both to indicate the position of the stimulus and to exploit the block structure of the task. After training, they made correct choices on 81.7% $\pm$ 0.4% (mean $\pm$ s.d.) of the trials, performing better and faster on trials with high visual stimulus contrast (Fig. 3a). Sessions lasted on average 650 trials (median 605, range [401, 1,525]). Towards the end of the sessions, performance decreased and reaction times increased (Fig. 3b;c). On 0% contrast trials, where no visual information was provided, mice gained reward on 59% $\pm$ 0.4% (mean $\pm$ s.d.) of trials, significantly better than chance (t-test $t_{114} = 21.9, p < 0.001$). After a block switch, mice took around 5 to 10 trials to adjust their behaviour to the new block, as revealed by the fraction of correct choices made on 0% contrast trials after the switch (Fig. 3d). Mice were influenced by their prior estimate also in the presence of visual stimuli: for all contrast values, mice tended to answer right more often on right blocks than left blocks (Fig. 3e). In a minority of sessions, mice embodied these choices in overt untrained behaviour such as whisking (Fig. S3).

Neural analyses. To obtain an initial appraisal of the brain-wide map we performed single-cell and population analyses, assessing how neural activity encodes task variables and how it can be analysed to decode these variables. We considered five key task variables (Fig. 4a): visual stimulus, choice (turning the wheel clockwise or counterclockwise), feedback (reward or time-out), the block, and physical movement (wheel speed). The main figures show the results of these analyses in a canonical dataset of 123 regions for which we had at least two sessions with 10 good neurons each (for a total of 22,113 neurons). Supplementary figures show results for the wider range of neurons and regions appropriate for each analysis.

To provide complementary views on how these task variables are represented in each brain region, we used four analysis techniques (Fig. 4b-e; see also Fig. S4 for a fuller picture). The details of each technique are provided in the Methods, along with a discussion of the corresponding null distributions, permutation tests and false discovery rate controls (FDR_q at a level q; using the Benjamini-Hochberg procedure) that we used to limit statistical artefacts.

First, we used a *decoding model* that attempts to predict the value of the task variable on each trial from neural population activity using regularized logistic or linear regression (Fig. 4d; Fig. S4d). This analysis can detect situations in which a variable is encoded robustly, but only in a sparse subset of neurons.

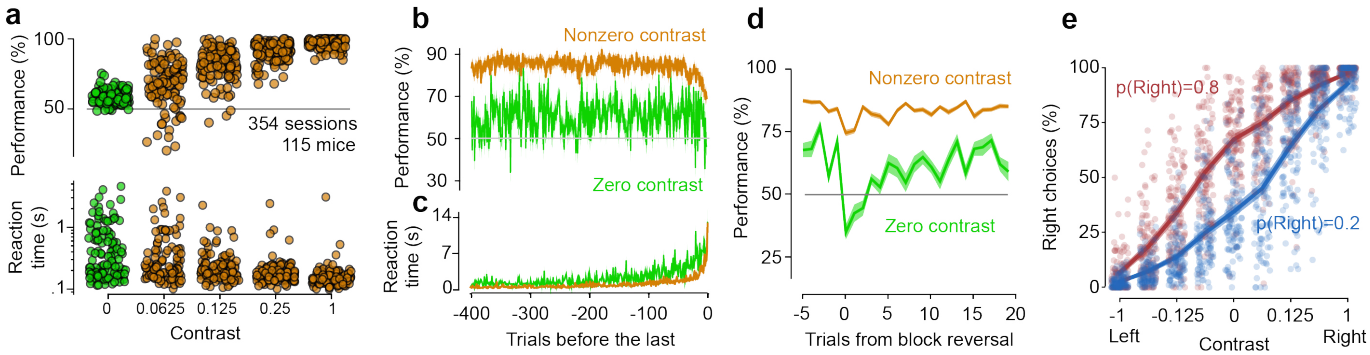

Figure 3. Behaviour during the task. a) Performance (top) and first wheel movement time (reaction time) (bottom) as a function of stimulus contrast (one point per mouse per contrast; green points for 0% contrast; orange points for non-zero contrast). Performance on 0% contrast trials (green) was 59% $\pm$ 0.4% (mean $\pm$ s.e. across mice) correct. **b)** Performance as a function of the number of trials before the end of the session for 0% (green) and non-0% (orange) contrast trials (mean $\pm$ s.e. across mice). **c)** The same analysis for reaction times. **d)** Reversal curves: proportion correct around a block change for trials with 0% contrast and $>$0% contrast (mean +/- s.e. across mice). The first 90 trials (when the stimulus was not biased to appear more frequently on a side) were ignored for this analysis. **e)** Psychometric curves showing the fraction of right choices as a function of signed contrast (positive for right stimuli, negative for left stimuli), for all mice (one dot per contrast per mouse). Right choices were more common in right blocks (red), when the stimulus appeared on the right with probability p(right)=0.8, than on left blocks (blue) when p(right)=0.2.

We did not correct for the correlations between the task variables, and rather assessed decoding for each variable separately. This quantifies what downstream neurons would be able to determine from the activity, but does not disentangle related factors such as stimulus side and block.

As a form of omnibus test, we performed decoding for all the neurons recorded in a session, determining the significance for a session by comparing the R^2 of the fit to a variable to the R^2 of the fit to a suitable null distribution. To combine decoding results across sessions we used Fisher's combined probability test[31,32]. We then performed decoding within each region. Again, we corrected the R^2 of the fit to a variable by the R^2 of the fit to a suitable null distribution and used Fisher's combined probability test[31] to combine decoding results across sessions. To correct for the comparisons over multiple regions we chose a level of 0.01 for the false discovery rate ($\text{FDR}_{0.01}$).

Second, we computed *single-cell* statistics, testing whether the firing rates of single neurons correlated with four task variables (visual stimulus side, choice side, block side, and feedback) in the appropriate epochs of each trial (Fig. 4b; Fig. S4b). Since the task variables of interest are themselves correlated, we used a condition-combined Mann-Whitney U-test[16] which compares spike counts between trials differing in just one discrete task variable, with all others held constant. Using a permutation test, we determined the fraction of individual neurons in a region that are significantly selective to a variable, using a threshold specific to each variable. For each session with recordings in a specific region, we computed the significance score of the proportion of significant neurons by using the binomial distribution to estimate false positive

events. We then combined significance scores across sessions with Fisher's combined probability test[31] to obtain a combined p-value for each region. We report a region as being responsive to a variable if this combined p-value was below the chance level, correcting for multiple comparisons using $\text{FDR}_{0.01}$. This method has lower statistical power than decoding since it only examines noisy single neurons, and it may therefore miss areas that have weak but distributed correlates of a variable. Moreover, it controls for correlated variables in a way that the decoding method does not, so it is selectively able to find neurons that do not only represent a variable by virtue of that variable's correlation with a confound.

Third, we performed a *manifold* analysis (Fig. 4e; Fig. S4e), averaging firing rates of single neurons in a session across all trials in 20 ms bins and then aggregating all neurons across sessions and mice per brain region. We examined how trajectories in the high dimensional neural spaces (manifolds) reflected task variables. We did this by measuring the time-varying Euclidean distance between trajectories $d(\Delta t)$ in the interval of interest, normalised by the square root of the number of recorded cells in the given region to obtain a distance in units of Hz/cell. From this time-varying distance, we extracted response amplitude and latency statistics. For significance testing we permuted trials for the variable of interest while keeping the other variables fixed to minimise the effect of correlations among the variables (as in the single-cell statistics), using $\text{FDR}_{0.01}$ to control for multiple comparisons. For visualisation only, we projected the trajectories into a 3-dimensional principal components space. We performed this analysis on the canonical set of cells and time windows.

Finally, we used an *encoding model*[33] that fits the activity of each cell on each trial as a linear combination of a set of temporal kernels locked to each task event (Fig. 4c; Fig. S4c; Fig. S5a). This general linear model (GLM) provides a picture at a temporally fine scale, at the cost of a potentially lower signal-noise ratio. We measure the impact of a variable by removing its temporal kernels and quantifying the reduction in the fit of the activity of a neuron (typically assessed by ΔR^2). This method lacks a convenient null distribution, and so we report effect sizes rather than significance.

We next describe the results of these four analyses applied to the coding of each of the five task variables.

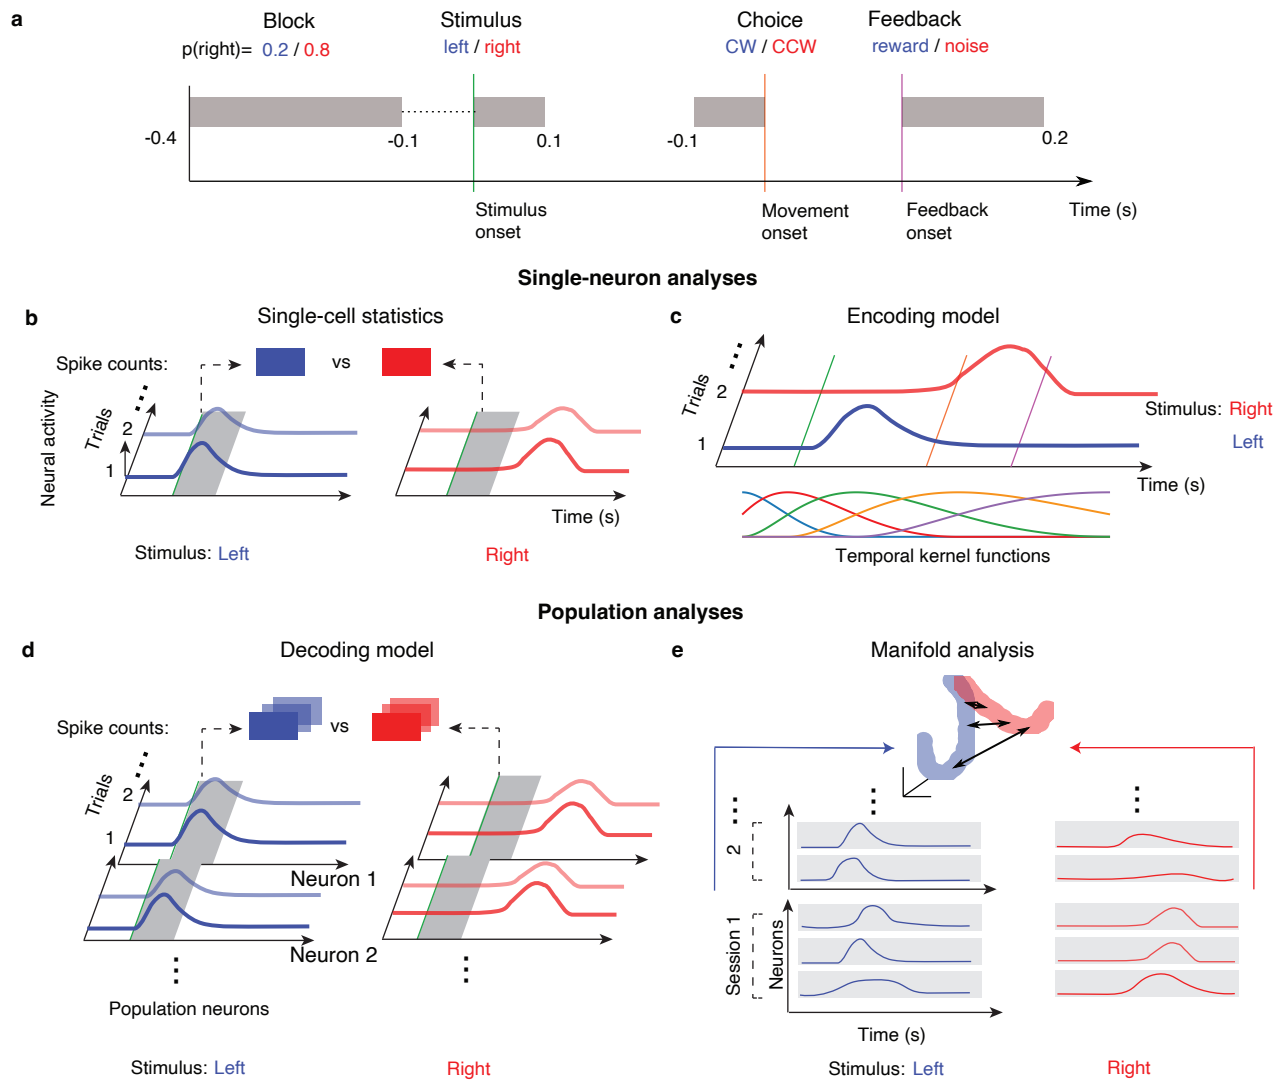

Figure 4. Illustration of neural analyses. See also Fig. S4. **a)** Task structure with time windows used for analysis in grey. **b)** Single-cell analysis quantifies single-cell neural correlates with task variables. **c)** The encoding model uses multiple linear regression of task- and behaviourally-defined temporal kernels on the activity. **d)** The decoding model quantifies neural population correlates with task variables. **e)** Manifold analysis describes the time evolution of the across-session neural population response, pooling cells across all recordings per region.

Representation of Visual Stimulus. We first consider neural activity related to the visual stimulus. Classical brain regions in which visual responses are expected include superior colliculus[34,35], visual thalamus[36–38], and visual cortical areas[39–42], with latencies reflecting successive stages in the visual pathway[43,44]. Correlates of visual stimuli have also been observed in other regions implicated in visual performance, such as parietal[45] and frontal[46–49] cortical areas as well as the striatum[50,51]. Substantial encoding of visual stimuli may also be present beyond these classic pathways, since the retina sends outputs to a large number of brain regions[52]. Indeed, an initial survey[16] of regions involved in a similar task uncovered visual responses in

areas such as the midbrain reticular nucleus (MRN). Thus we hypothesised that visual coding would extend to diverse regions beyond those classically described.

Consistent with this hypothesis, a decoding analysis based on the first 100 ms following stimulus onset revealed correlates of the visual stimulus side in many cortical and subcortical regions, with strong signatures in visual cortex, thalamus, midbrain, and hindbrain (Fig. 5a; the top five regions were VISpm, VISam, PRNc, GRN, and SCm). For instance, the activity of neurons in primary visual cortex (VISp) could be used to predict the stimulus side (Fig. 5i). Note that, uniquely amongst our analyses, the decoding analysis does not control for variables correlated with the visual stimulus, such as choice and block, so some of the regions with statistically significant results on this analysis might instead encode these variables.

Decoding performance varied across sessions, and therefore, in Fig. S6, we show performance across sessions for all regions, even those that are not significant after the $\text{FDR}_{0.01}$ correction for multiple comparisons. Given that some regions may well represent visual information in localised sites that were only occasionally targeted by our probes, we also report the fraction of sessions in which we were able to decode a variable from a region (to assess spatial distribution; Fig. S8a).

To distinguish the possible contributions of variables correlated with the visual stimulus, we next analysed responses in the same 100 ms window using a single-cell analysis, which controlled for other variables by holding them constant in each comparison of stimulus side. This analysis yielded a consistent picture but found fewer significant areas, with 0.5% of all neurons correlated with stimulus side (Fig. 5b). The significant regions included visual cortical areas (VISp, VISpm, VISam, VISli, VISl) and visual thalamus (LGd, LP), but also other structures such as retrosplenial cortex (RSPv), ventral thalamus (VM), zona incerta (ZI) and parts of the midbrain (SCm, APN). However, even in those regions containing the largest fractions of responsive neurons (such as visual cortex), this fraction did not exceed $\sim 10\%$. Note that this low percentage of neurons could be the result of neurons not having receptive fields overlapping the stimulus position, given our grid-based approach to targeting recordings.

To provide an overview of the variability across sessions, Fig. S7 shows the fraction of significant neurons broken down by sessions, without applying $\text{FDR}_{0.01}$ correction.

The results of the manifold analysis were consistent with the decoding analysis (Fig. 5c), and provided further information about the time course with which visual signals were encoded (Fig. 5d). For instance, the responses in visual area VISp to right vs. left visual stimuli show early divergence shortly after stimulus onset, followed by rapid convergence (Fig. 5j;k). The null trajectories (shown in grey) are close to the true trajectories because this analysis controls for choice, which is tightly correlated with the stimulus. Altogether, this analysis indicated that distance was significant in 79 regions (Fig. 5c;f).

The evolution of trajectories over time could be distilled into two numbers (Fig. 5l;m): the maximal response (maximum distance or d_{max}) and the response latency (first time to reach 70% of d_{max}; mapped across the brain in Fig. 5d). This characterization of the dynamics of visual representations revealed that some areas had short latencies and early peaks, including classical areas (LGd, LP, VISp, VISam, VIpm),

with a spatiotemporally finely resolved view of latency differences such that LGd $<$ VISp $\approx$ LP $<$ VISpm $<$ VISam (latencies $\approx 34, 42, 42, 57, 78$ ms; Fig. 5d;l;m). This early wave of activity was followed substantially later by significant visual encoding in other areas, including MRN, SCm, PRNr, IRN, and GRN (latencies $\approx 100 - 120$ ms; Fig. 5d;l;m).

Finally, the encoding analysis characterised visual encoding in individual neurons across the brain (Fig. 5e). This analysis asks whether a prediction of single-trial activity can be improved by adding a temporally structured kernel that unfolds over 400 ms after stimulus onset, on top of activity related to feedback, wheel movement speed, and block identity. Since there is no convenient null distribution which could be used to test the significance of this improvement, we only report effect sizes for the encoding analysis.

For instance, as expected, an example VISp neuron showed large differences between stimuli on the left and right (Fig. 5g) such that removing the visual kernel resulted in a poor fit of the neuron's firing rate (Fig. 5h). This analysis indicates that the visual stimulus variable improved fits of encoding models for neurons across a wide range of brain regions (Fig. 5e,f).

Taken together, the decoding, single-cell statistics, and manifold analyses reveal a largely consistent picture of visual responsiveness that includes large and short-latency responses in classical areas but also extends to diverse other regions, particularly at later times relative to stimulus onset.

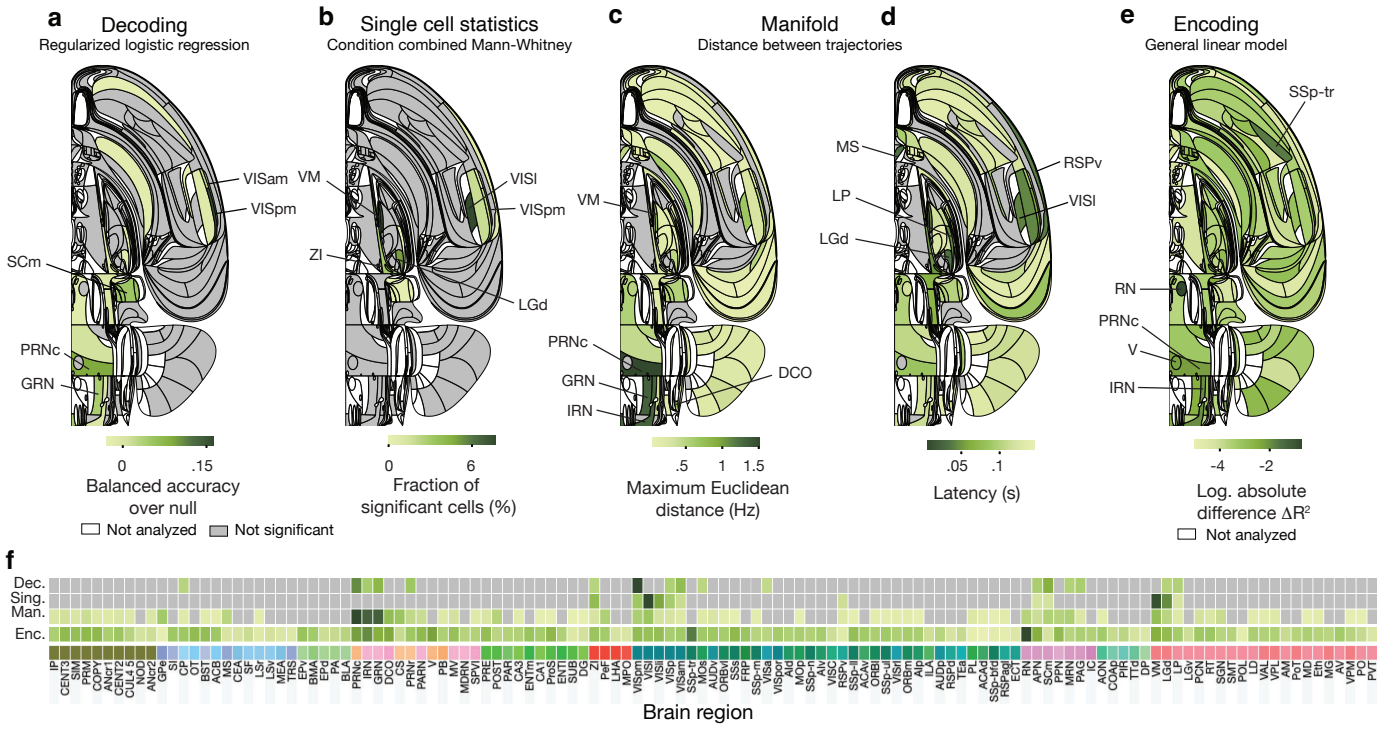

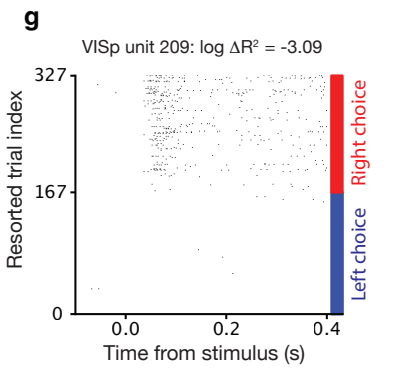

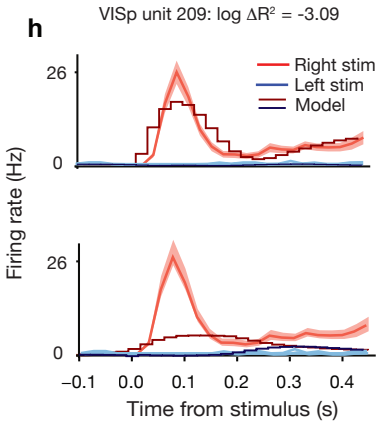

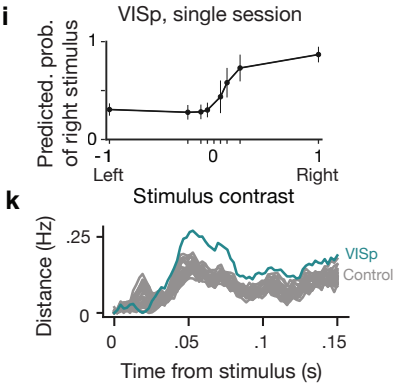

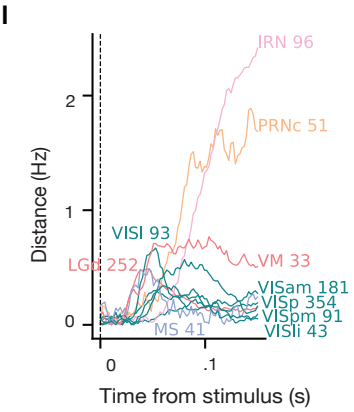

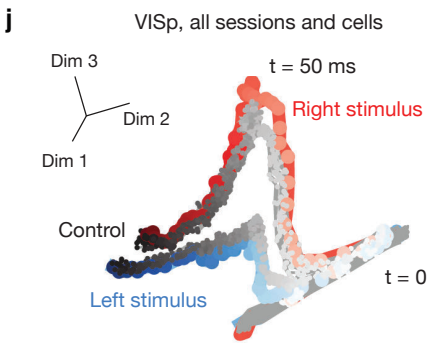

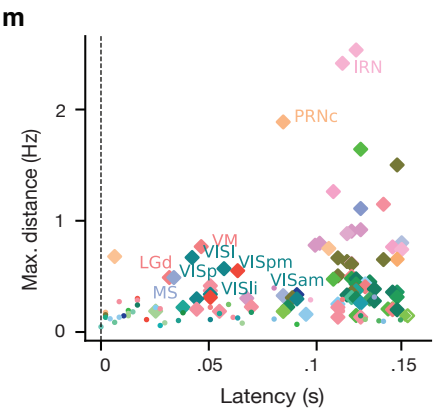

Figure 5. Representation of the Visual Stimulus. See also Fig. S6; S7; S8. An interactive version of this figure can be found on our data website. **a)** Decoding analysis: Flat brain map of decoding score (balanced accuracy) across sessions. The values have been corrected by subtracting the median of the decoding score in the null distribution. Colour represents effect size. Grey: regions in which decoding was not statistically significant, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. White: regions that were not analysed due to insufficient data. The omnibus test of decoding from all neurons in each session was highly significant $p \ll 10^{-10}$. **b)** Single cell statistics: Flat brain map of the fraction of neurons in a region for which firing rates were significantly modulated by stimulus side (left vs right) in the period $[0, 100]$ ms aligned to the stimulus onset. The significance score of single cells were obtained by conjoining two tests: a Mann-Whitney test ($p < 0.001$) and a condition combined Mann-Whitney test ($p < 0.05$). The omnibus test of the fraction of significant neurons from all neurons in each session was highly significant $p \ll 10^{-10}$. **c)** Manifold distance: Flat brain map of the time-resolved maximum distance between trajectories following stimuli on the left vs. right side, based on Euclidean distance (in Hz/cell) in the full-dimensional space (dimension = number of cells) for each brain region. Significance was assessed relative to a shuffle control, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. The omnibus test based on treating all cells from all sessions as one region resulted in highly significant modulation, $p \ll 10^{-3}$. **d)** Manifold latency: The earliest time after stimulus onset at which 70% of the maximum of the trajectory distance was reached, for significant regions only. **e)** Encoding analysis: Flat brain map of mean model improvement differences per region, across all neurons in that region computed as the absolute difference between the improvements ($|\Delta R^2|$) from the right stimulus kernel and left stimulus kernel regressors (400 ms causal kernels aligned to the stimulus onset time). **f)** Table of effect significance (grey - not significant; a-c) and effect size (by darkness; a-c; e), grouped by region. Regions are sorted within each Cosmos group by the sum of effects across analyses (see Methods). **g)** Spike raster of example neuron in VISp identified by the encoding model as being sensitive to stimulus side (see Table 3 for session and neuron details). **h)** Upper panel: Comparison of peri-event time histogram (PETH) of spiking activity for left and right stimuli for the example neuron in panel g conditioned on stimulus onset, along with associated predictions of the full encoding model. The width of the PETH traces reflects standard error of the mean. Lower panel: The same PETHs, but with predictions produced by an encoding model in which the stimulus onset-aligned regressors were omitted. Error bars represent 1 SEM about the mean rate at each time point. **i)** Predicted probability from decoding analysis of stimulus side ('neurometric curve') from VISp as a function of contrast. Note that the left-side stimulus contrasts are shown with < 0 values and the right-side stimulus contrasts are shown with > 0 values. Balanced accuracy for this region session was 0.749; see Table 3 for session and neuron details. **j)** Trajectories obtained from left (right) stimulus side trial-averaged activity of all VISp neurons, visualised by projection via PCA onto 3 dimensions, Each dot corresponds to a single time bin of the trial-averaged population activity, darker colours indicate later times. Grey pseudo-trajectories were obtained by averaging randomly selected trials matched for block and choice but with different visual stimuli presented. (A clear split of pseudo-trajectories indicates correlation of block, choice and stimulus side, as stimulus sides were shuffled within block/choice classes, see methods.) **k)** Trajectory distance as a function of time for example region VISp, showing an early response; pseudo-trajectory distances in grey (control). **l)** Trajectory distances for more example regions showing early response in most visual areas and ramping stimulus side modulation with time in others. The number of neurons is given alongside the acronym. **m)** Maximal manifold distance and modulation latency for all analysed regions (diamonds - significant regions, dots - not significant regions).

Representation of Choice. Next, we examined which regions of the brain represented the mouse's choice, and in which order. Choice-related activity has been observed in parietal, frontal and premotor regions of the primate cortex[7,53,54] where many neurons show ramping activity consistent with evidence accumulation[7,55]. These choice signals develop across frontoparietal regions and appear later in frontal eye fields[56]. Similar responses were found in rodent parietal[57], frontal[58,59], and premotor[60,61] cortical regions. However, in both rodents and primates, choice-related activity has also been found in hippocampal formation[62] and subcortical areas, in particular in striatum[16,63], superior colliculus[16,64,65] and other midbrain structures[16]. Subcortical regions show choice signals with similar timing as cortex[16,66] and play a causal role in the choice[65]. This evidence indicates that decision formation engages a distributed network of cortical and subcortical brain regions. Our recordings allowed us to determine in detail where and when choice-related activity emerges across the brain.

The decoding analysis suggested a representation of choice (left versus right upcoming action) in a larger number of brain regions than the representation of the visual stimulus (Fig. 6a;f). The animal's choice could be decoded from neural population activity during a 100 ms time window prior to movement onset in many analysed regions, with the strongest effect sizes in hindbrain (GRN, PRNc), midbrain (APN, SCm, MRN, PPN), cerebellum (IP), striatum (GPe), and thalamus (VM). For instance, the activity of neurons in the Gigantocellular Reticular Nucleus (GRN) of the medulla could be readily decoded to predict choice in an example session (Fig. 6g;i). Choice was also significantly decodable from somatosensory (SSp-ul), prelimbic (PL), motor (MOs), orbital (ORBvl), and visual (VISpm) cortical areas.

Some of the decodable choice information, however, could be due to responses to the visual stimulus or block, which correlate with choice. We thus performed single-cell analyses that control for correlations between all these task variables. More single neurons responded significantly to choice than to the visual stimulus (Fig. 6b;f): firing rates of 3% of all neurons recorded across all brain regions correlated with choice direction during the 100 ms prior to the movement onset, controlling for correlations with the stimulus and block. The largest fractions of neurons responding significantly to choice were in the hindbrain, cerebellar, midbrain, and thalamic regions, consistent with the results of decoding analysis. Neurons with significant responses to choice were highly prevalent in pons, medulla, and cerebellar nuclei (PRNc, IRN, GRN, DCO, PARN, MV, COPY), most of which did not show visual responses. These results obtained from statistical analyses were confirmed by the single-cell encoding model (Fig. 6e). For instance, an example neuron in GRN showed stronger responses for right choices than left choices (Fig. 6g). The encoding model captured this preference but only in the presence of the kernel associated with choice (Fig. 6h).

The manifold analysis allowed us to compare across brain regions the magnitude of choice representation on the population level, measured as the distance between trajectories in the neural population state space on left versus right choice trials (Fig. 6c;f). The magnitude of the population-level choice representation was highly non-uniform across the brain, with the strongest separation of neural trajectories in hindbrain (IRN, GRN, PRNc) and midbrain (APN, MRN, SCm), dwarfing the population-level choice en-

coding in most other areas (Fig. 6c;f). In our example region GRN, the trajectories for left and right choice trials separated significantly more than in shuffled control (Fig. 6j;k, controlling for correlations with stimulus and block), and the magnitude of this separation was greatest across all brain regions (Fig. 6l). Thus, all our analyses consistently point to a distributed choice representation, with the strongest choice signals in midbrain, hindbrain and cerebellum, and relatively weaker choice encoding across many cortical areas.

Next we analysed when the choice signals emerged across the brain by measuring the latency with which neural population trajectories separated on the left versus right choice trials during the time preceding movement onset (Methods). Some of the earliest choice signals developed nearly simultaneously in pons (PRNr), thalamus (LGd), striatum (LSr), and cortex (VISl, ILA), and later appeared in a larger distributed set of brain areas (Fig. 6d;l;m). The pontine region PRNr and medullary regions GRN and IRN showed, respectively, some of the lowest choice latencies and the strongest magnitude of population level choice representations (Fig. 6m), suggesting a role for these brainstem structures—classically associated with automatic motor processing—in decision formation or movement preparation.

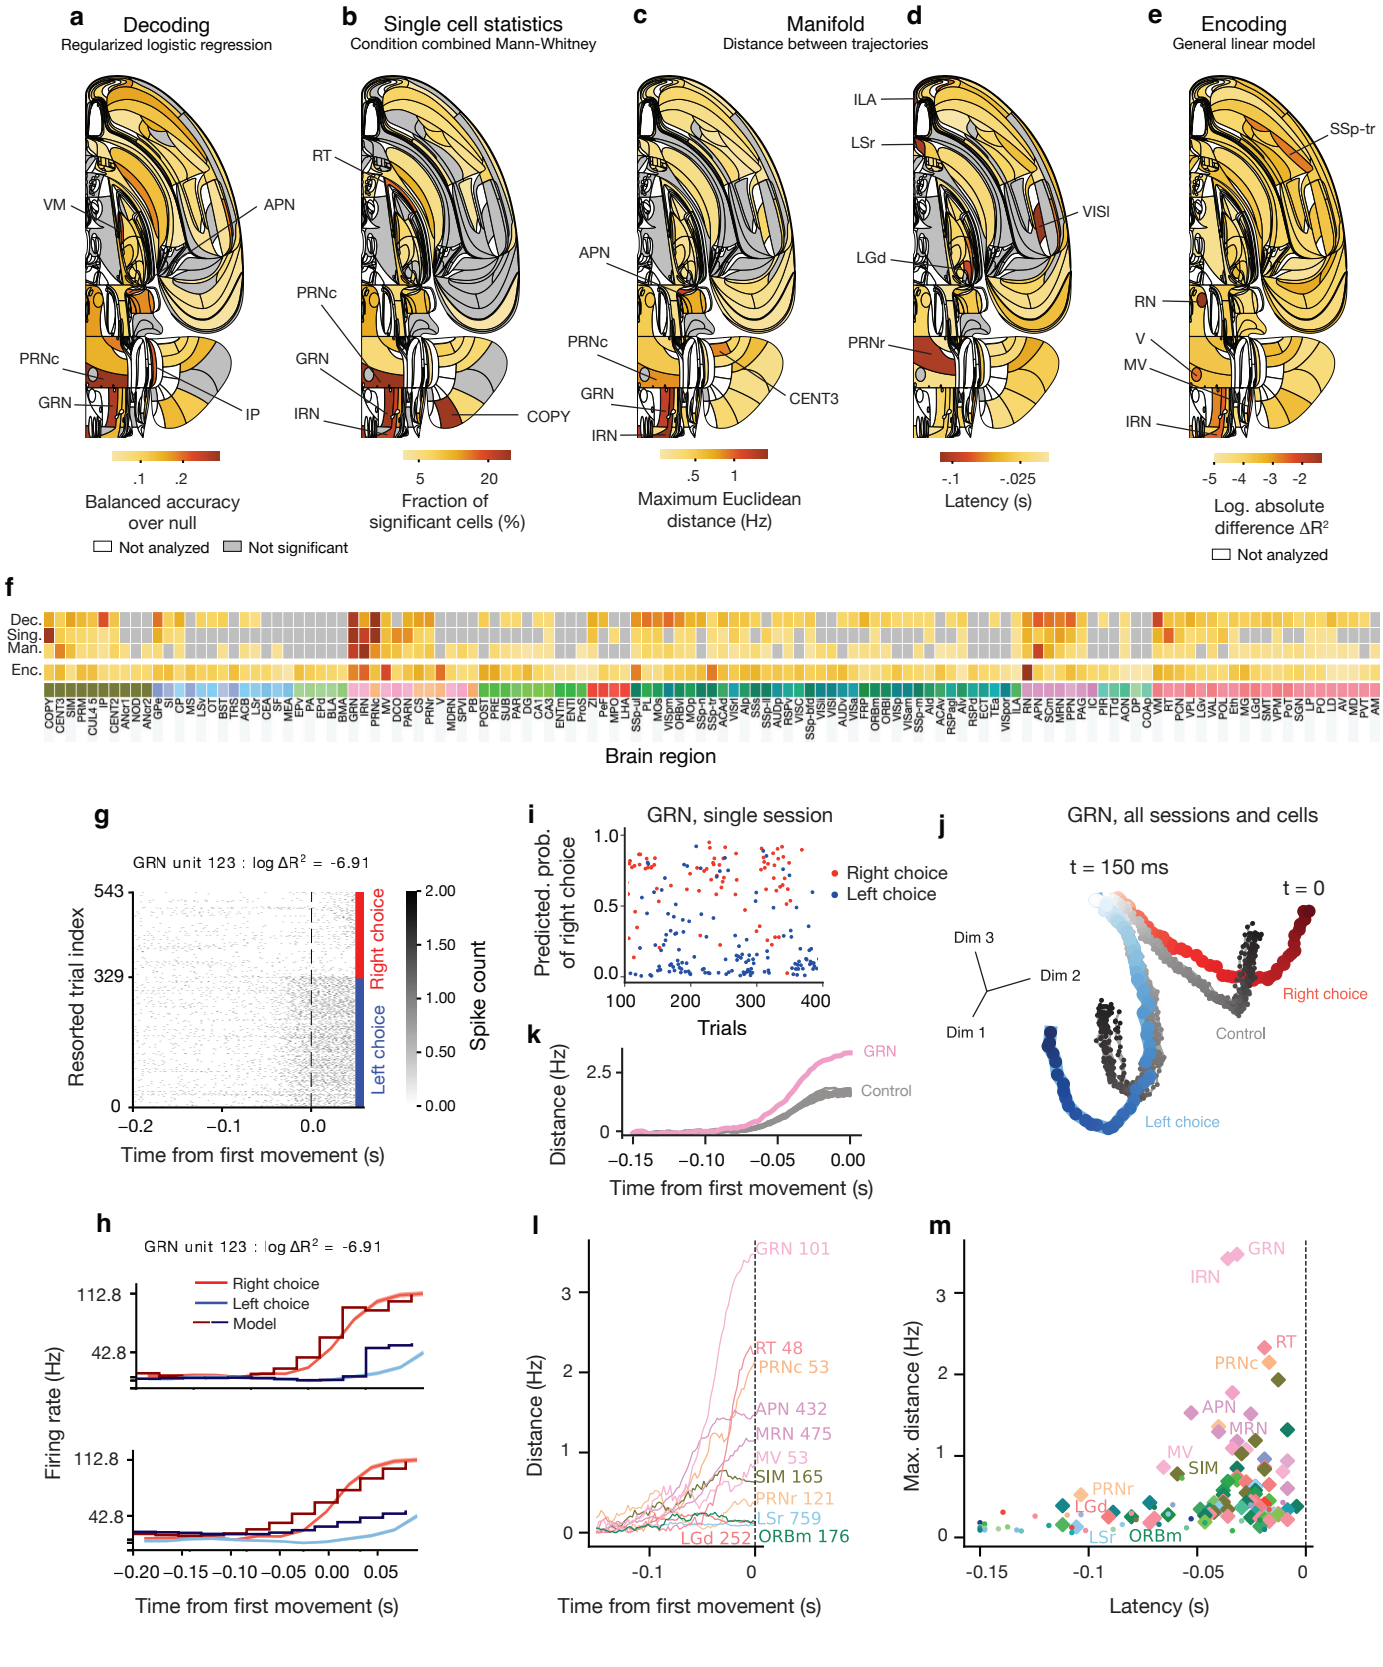

Figure 6. Representation of Choice. See also Fig. S6; S7; S9. An interactive version of this figure can be found on our data website. **a)** Decoding analysis: Flat brain map of median corrected decoding score (balanced accuracy) across sessions. The values have been corrected by subtracting the median of the decoding score in the null distribution. Colour represents effect size. Grey: regions in which decoding was found not to be statistically significant, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. White: regions that were not analysed due to insufficient data. The omnibus test of decoding from all neurons in each session was highly significant ($p \ll 10^{-10}$). **b)** Single-cell statistics: Flat brain map of the fraction of neurons in a region for which firing rates were significantly modulated by choice side (left vs. right) in the period $[-100, 0]$ ms aligned to the first-movement onset. The significance scores of the single cells were obtained by conjoining two tests: a Mann-Whitney test & a condition combined Mann-Whitney test (Methods), using thresholds of $p < 0.001 \& p < 0.05$, respectively. The significance score of each region was computed by assuming a binomial distribution of false positive events, and $\text{FDR}_{0.01}$ was used to correct for multiple comparisons. The omnibus test of the fraction of significant neurons from all neurons in each session was highly significant ($p \ll 10^{-10}$). **c)** Manifold distance: Flat brain map of the maximum distance between neural trajectories for left and right choice, based on time-resolved Euclidean distance (in Hz/cell) in the full-dimensional space (the number of dimensions equals to the number of cells) for each brain region. Significance was assessed relative to a shuffle control, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. Canonical window lengths were employed, as shown in Fig. 4a. The omnibus test based on treating all cells in each session as one population was significant ($p \ll 10^{-3}$). **d)** Manifold latency: The earliest times before movement onset at which 70% of the maximum of the trajectory distance (see part c) was reached, for significant regions only. **e)** Encoding: Flat brain map of mean model improvement differences per region, across all neurons in that region computed as the absolute difference between the improvements ($|\Delta R^2|$) from the right first movement kernel and left first movement kernel regressors. These regressors for choices in each direction were 200 ms anti-causal kernels aligned to the movement onset time. **f)** Table of effect significance (grey - not significant; a-c) and effect size (by darkness; a-c; e), grouped by region. Regions are sorted within each Cosmos group by the sum of effects across analyses (see Methods). **g)** Spike raster of an example neuron in GRN identified by encoding model as sensitive to choice (see Table 3 for session and neuron details). **h)** Upper panel: Comparison of PETHs aligned to movement onset on left (blue) and right (red) choice trials for the example neuron in panel g, along with the encoding model predictions for each condition (dark colours). This neuron was selected for a high difference in ΔR^2_{choice}, computed as the absolute difference of additional variance explained by the left/right first movement regressors. Lower panel: The same PETHs, but with predictions produced using a model lacking the left and right first movement regressors. Error bars represent 1 SEM about the mean rate at each time point. **i)** Decoding to predict choice for part of an example session in GRN (balanced accuracy for this region-session was 0.840; see Table 3 for session and neuron details). Decoder predictions separate in correlation with the mouse' choices (blue - left choice, red - right choice). **j)** Trajectories obtained from trial-averaged activity of all GRN neurons on left (blue) and right (red) choice trials, visualised by projecting onto the first 3 PCA dimensions. Each dot corresponds to a time bin of the population activity, darker colours indicate times closer to the movement onset. Grey pseudo-trajectories were obtained from trial-averaged activity on trials with randomised choice, controlling for correlations with stimulus and block. **k)** Trajectory distance as a function of time for GRN, showing ramping activity; pseudo-trajectory distances in grey (control). **l)** Trajectory distances for more example regions showing ramping choice-modulation with time. **m)** Maximal manifold distance and modulation latency for all analyzed regions (diamonds - significant regions, dots - not significant regions).

Representation of Feedback. At the end of each trial, the mouse received feedback for correct or incorrect responses: a liquid reward at the lick port or a noise burst stimulus with a time-out period. These positive or negative reinforcers influence the learning of the task [67–71]. If received, the liquid is consumed through licking, which may itself have a prominent neural representation which, in this study, we will not be able to distinguish from the more abstract qualities of reward. Feedback also activates neuromodulatory systems such as dopamine[72], which have widespread connections throughout cortical and subcortical regions. However, it is unclear how widespread the feedback signals themselves are in the brain.

The decoding analysis revealed nearly ubiquitous neural responses associated with the delivery of reward on correct versus incorrect trials, and likely the motor responses associated with its consumption (Fig. 7a). Using the neural responses in the 200 ms following feedback onset, we were able to decode whether or not the trial was correct from nearly all recorded brain regions (Fig. 7a;f). In many regions, decoding was practically perfect, including, for instance, the activity of the intermediate reticular nucleus in the hindbrain (IRN) in an example session (Fig. 7i).

Our single-cell statistics applied to the same trial interval confirmed the decoding results. We found that neurons with significant response changes to correct versus incorrect feedback or reward consumption were extremely widespread (Fig. 7b;f), with only a small handful of regions not significant for feedback type. The same was true for feedback versus the inter-trial interval baseline (Fig. S10).

Manifold analysis also found statistically significant response differences for correct versus incorrect responses across every recorded brain region, predominantly consistent with the other analyses (Fig. 7c;f). It confirmed the relative strength of hindbrain, midbrain, and thalamic responses to feedback seen across analyses. Manifold analysis also revealed asymmetries in response to negative versus positive feedback: for positive feedback the response was overall stronger, and multiple brain areas exhibited a coherent $\sim$10 Hz oscillatory dynamics during reward delivery that was phase-locked across brain areas (Fig. 7j;k;l) and sessions (Fig. S11). Across-session coherence is visible as a large oscillatory signal in example area IRN (Fig. 7j;k). These oscillatory dynamics were missing during negative feedback, and were closely related to licking behaviour (Fig. S11)[73–76], pointing to consumption-related activity being the dominant factor over influences of reward on neural activity.

Assessing response latencies by the divergence of the trajectories over time, we found that the primary auditory region AUDp and the saccade- and gaze-reorienting brainstem region PRNr (rostral pontine reticular nucleus) exhibited the earliest and strongest responses (Fig. 7l;m). Many of these likely reflect residual responses from choice-related activity because the latencies are very short and several identified areas exhibit high choice responsivity. The responses from auditory areas likely reflect responses to the error tone and the click from the reward delivery valve. After these initial responses, the latencies across other brain regions appeared roughly similar, suggesting a common signal broadcast across the brain (Fig. 7d). More detailed manifold distance and latency scatterplots can be found in Fig. S12.

We applied the encoding model to the responses measured in the 400 ms after stimulus onset, and found that the kernel for correct feedback was the largest single contributor to neural response variance (Mean ΔR^2 of 7×10^{-3} averaged across all neurons; Fig. S5b), surpassing all other kernels (left or right stimulus, left or right movement onset, incorrect feedback, block probability, and wheel speed). This high variance-explaining response to reward delivery or consumption held across both wide regions of cortex and subcortical areas. Mid- and hindbrain areas exhibited particularly strong responses to reward delivery, with many additional regions including thalamus and sensory (e.g. AUDp, SSs) and motor (MOp) cortex showing sensitivity as well (Fig. 7e). Removing the regression kernel for correct feedback then refitting the encoding model of an IRN neuron illustrates the large influence of correct feedback on activity (Fig. 7g;h).

In sum, we found feedback signals to be present across nearly all recorded brain regions, with a stronger response to positive than negative feedback (i.e., to reward delivery and consumption) and with particularly strong responses in thalamus, midbrain, and hindbrain. Further research will be needed to distinguish between responses for an internal expectation of feedback or the initiation of choice-related action versus responses to external feedback.

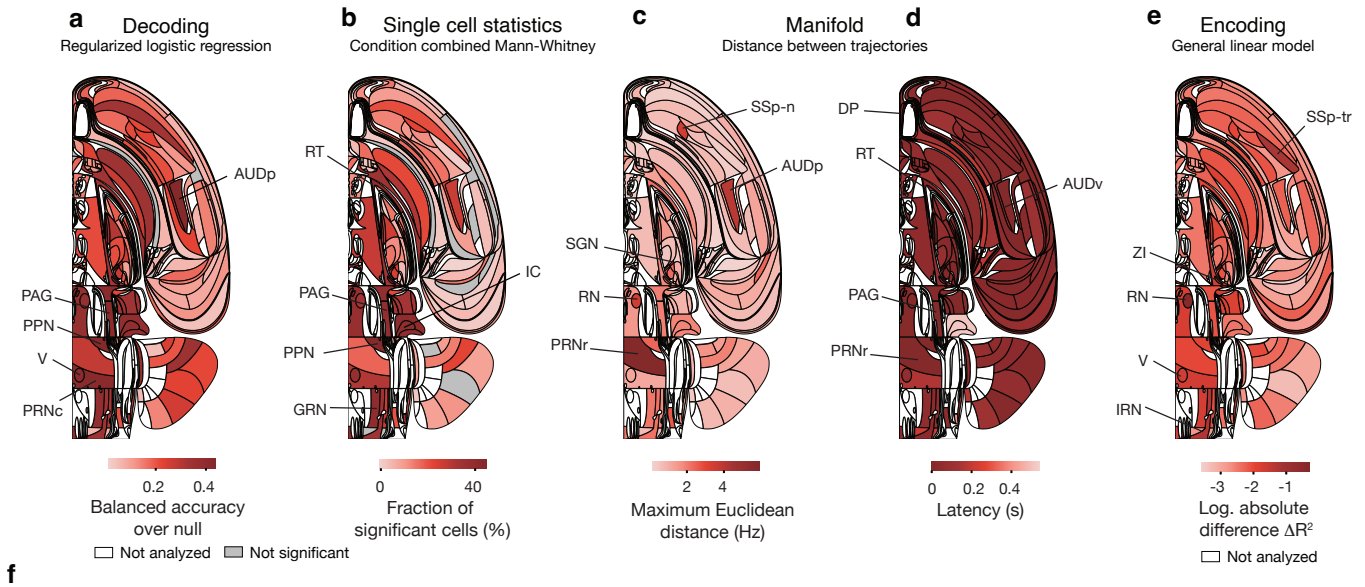

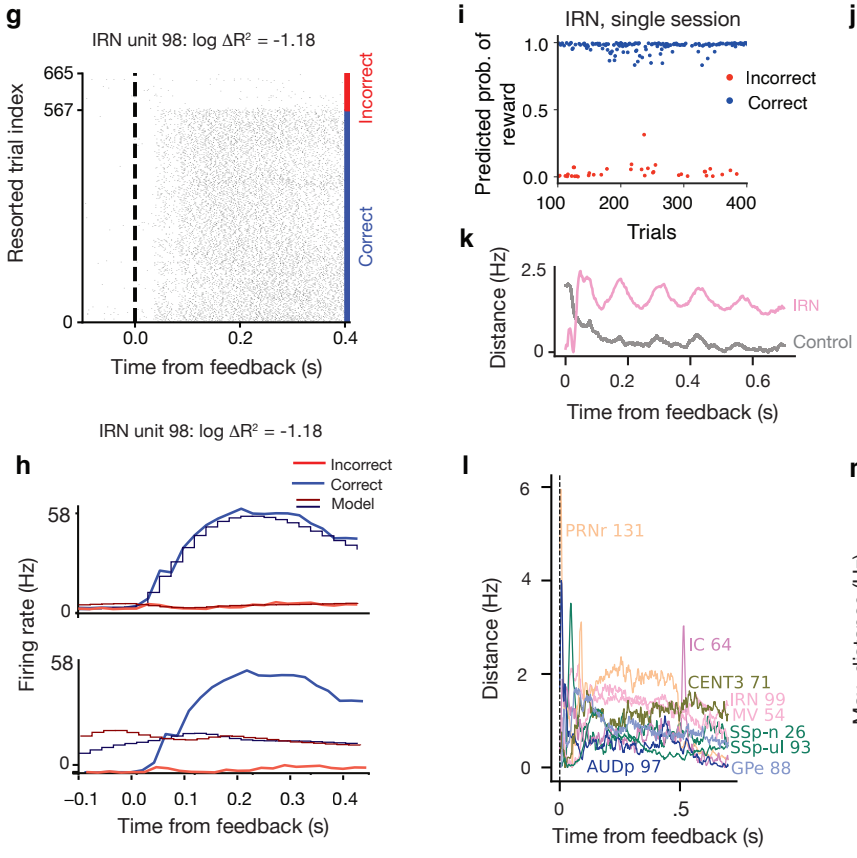

Figure 7. Representation of Feedback. See also Fig. S6; S7; S13. An interactive version of this figure can be found on our data website. **a)** Decoding: Flat brain map of median corrected decoding score (balanced accuracy) across sessions. The values have been corrected by subtracting the median of the decoding score in the null distribution. Colour represents effect size. Grey: regions in which decoding was found not to be statistically significant, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. White: regions that were not analysed due to insufficient data. The omnibus test of decoding from all neurons in each session was highly significant $p \ll 10^{-10}$. **b)** Single-cell analysis: Flat brain map of the fraction of neurons for which firing rates were significantly modulated by feedback compared to baseline activity ($[-200, 0]$ ms aligned to the stimulus onset) in the period following feedback onset ($[0, 200]$ ms). The significance scores of the single cells were obtained by conjoning two tests: a Mann-Whitney test & a condition combined Mann-Whitney test (Methods), using thresholds of $p < 0.001 \& p < 0.05$, respectively. The significance score of each region was computed by assuming a binomial distribution of false positive events, and $\text{FDR}_{0.01}$ was used to correct for multiple comparisons. The omnibus test of the fraction of significant neurons from all neurons in each session was highly significant $p \ll 10^{-10}$. **c)** Manifold distance: Flat brain map of the time-resolved maximum distance between correct and incorrect choice trajectories, based on Euclidean distance (in Hz/cell) in the full-dimensional space (dimension = number of cells across all sessions) for each brain region. Significance was assessed relative to a shuffle control, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. Canonical window lengths were employed, as shown in Fig. 4a. The omnibus test based on treating all cells from all sessions as one region resulted in highly significant modulation, $p \ll 10^{-3}$. **d)** Manifold latency: The earliest time after feedback at which 70% of the maximum trajectory distance (see part c) was reached, for significant regions only. **e)** Encoding: Flat brain map of the mean model differences per region, across all neurons in that region, computed as the log of the absolute difference between the improvements (ΔR^2) from the correct feedback kernel and incorrect feedback kernel regressors. These regressors are 400 ms causal kernels aligned to the feedback time. **f)** Tabular form of effect significance (grey - not significant; a-c) and effect size (by darkness; a-c; e), grouped by region. Regions are sorted within each Cosmos group by the sum of effects across analyses (see Methods).
g) Spike raster of example neuron in IRN identified by encoding model as sensitive to feedback (see Table 3 for session and neuron details). **h)** Upper panel: Comparison of PETHs aligned to movement onset on correct (blue) and incorrect (red) trials for the example IRN neuron in panel g, along with the encoding model predictions for each condition (grey). This neuron was selected for a high difference in $\Delta R^2_{feedback}$, computed as the absolute difference of additional variance explained by the feedback regressors. Lower panel: The same PETHs but with predictions produced using a model lacking the feedback regressors. Error bars represent 1 SEM about the mean rate at each time point. **i)** Predicted probability from decoding analysis about whether reward was received, coloured by true feedback, from region IRN (balanced accuracy for this region-session was 1.000; see Table 3 for session and neuron details). **j)** Trajectories obtained from incorrect (correct) trial-averaged activity of all neurons in IRN vizualised by projection via PCA to 3 dimensions. Each dot corresponds to a time bin of the population activity, darker colours indicate later times. The oscillation of the blue trajectory correlated with licking, and was much stronger in correct trials when the animals received water. Grey pseudo-trajectories were obtained by averaging randomised trials, shuffling choice types within classes of stimulus side and block. **k)** Trajectory distance between correct and incorrect trials, as a function of time for example region IRN, showing oscillatory (licking) activity; pseudo-trajectory distances in grey (control). **l)** Trajectory distances for more example regions showing early response e.g. in auditory areas and prolonged feedback type modulation with time in others. **m)** Maximal manifold distance and modulation latency for all analyzed regions (diamonds - significant regions, dots - not significant regions). Most regions' activity was modulated within 100 ms after feedback.

Representation of Block Prior. The IBL task is a perceptual decision-making task that requires combining sensory evidence with the prior probability over the stimulus side, which itself can be inferred from past trials. Mice perform this computation competently (Fig. 3). We may thus expect neural correlates of the block structure in multiple brain regions, particularly regions in which stimulus value has been reported, as prior information is often associated with value. In primates and humans, these include posterior parietal cortex, orbital cortex, dorsolateral prefrontal cortex, premotor cortex, and striatum, among others[77–82]. In rodents, similarly, they include striatum[83–85], medial prefrontal cortex[84–87], orbitofrontal cortex[84–86,88,89], secondary motor cortex[84,85], dorsal hippocampus[84], and frontal pole[85]. Moreover, one would expect block signals in midbrain dopamine neurons, because changes in block strongly modulate dopamine signals in ventral striatum[90]. Considering that the prior probability may exert its influence through additional means, e.g. through visual attention[91] and motor preparation, one may expect correlates of block structure in an even broader swath of brain regions, potentially impacting not only firing rates but also the strength of synapses[92].

To search for neural correlates of the block structure, we applied the decoding analysis to the activity measured during the inter-trial interval (400 to 100ms before stimulus onset), a time during which the mice must keep the wheel still. After correcting for multiple comparison with a false discovery rate of 0.01, we found that block identity could be decoded significantly ($p < 0.05$) from only the primary motor cortex, MOp (see the example session in Fig. 8g). The single-cell statistics analysis found the block prior to be encoded is a small number of regions, namely, CENT3, CEA, CP, ORBV1 and SCm (Fig. 8b;f).

The encoding analysis revealed that block signals were $\sim$10 times smaller than feedback signals, and much smaller than other signals such as movement (Fig. 8e, Fig. S5b) which possibly explains why only a small number of regions were found to encode the prior significantly. The weak effects of block on the activity of single cells can be observed in an example neuron in MOp: the firing rates measured during the right and left blocks were significantly different according to the model but the difference is modest (4.76 vs 0.68 spikes/s) (Fig. 8h). Finally, the manifold analysis identified no region as reflecting the prior significantly (Fig. 8c). Fig. 8i shows an example of block-conditioned trajectories in MOp. While the right and left block trajectories are distinct, they do not differ from the null trajectories shown in grey (Fig. 8j). As to the latency analysis, it is not informative since it is based on the inconclusive manifold analysis (Fig. 8d).

In sum, while block identity influenced the behaviour of the mice, our standard analyses failed to reveal its neural representation beyond a handful of regions. This result might reflect an extremely sparse encoding of the prior, but our inability to decode the prior from prefrontal, frontal and parietal areas, where the prior has been reported previously[78–81,89], hints at other factors. One is that the active representation of the prior during the inter-trial interval might be very modest. A second is that our analyses were confined to the true block identity, a variable that is not directly available to the animal. In a companion paper [93], we examine in detail the representation of the aspects of the block that we could show actually influenced the behaviour of the animals (a form of subjective prior). By combining Neuropixels recordings with wide field

calcium imaging of cortical regions, we show there a rather wide encoding of the subjective prior across all levels of brain processing.

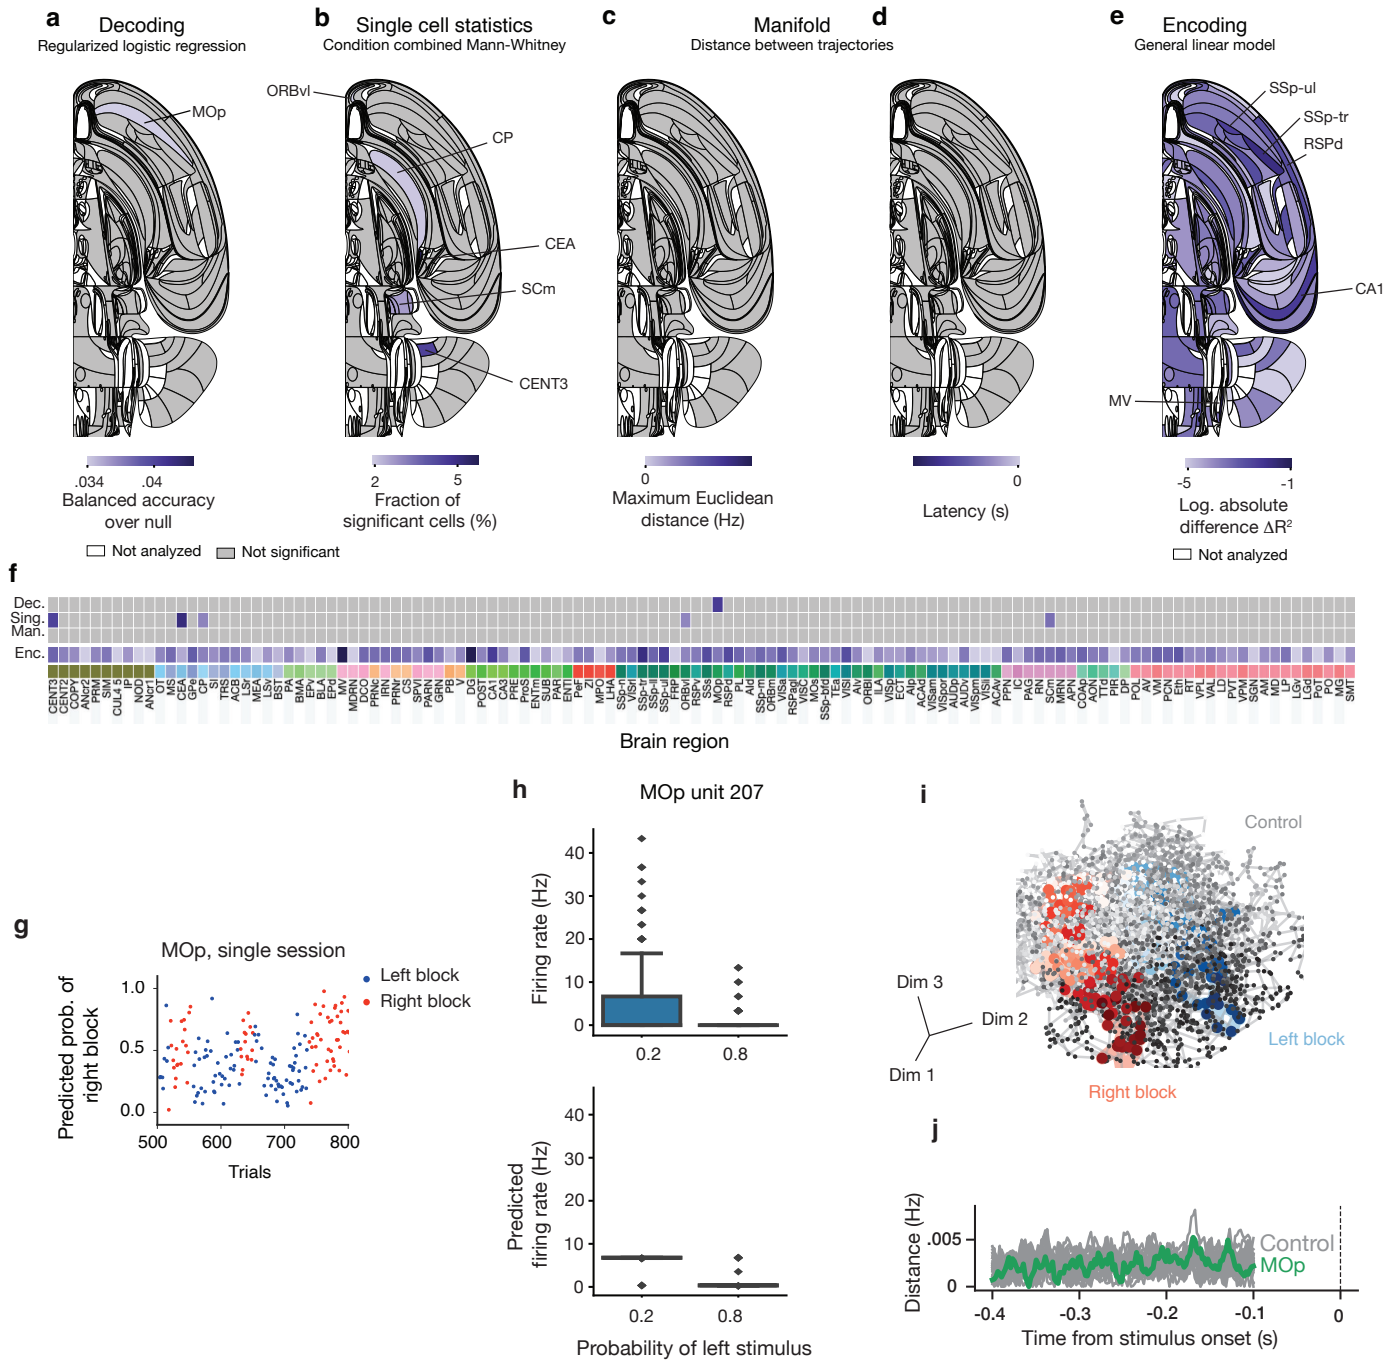

Figure 8. Representation of Block Prior. See also Fig. S6; S7; S14. An interactive version of this figure can be found on our data website. **a)** Decoding: The flat brain map of median corrected decoding score (balanced accuracy) across sessions. The values have been corrected by subtracting the median of the decoding score in the null distribution. Colour represents effect size. Grey: regions in which decoding was found not to be statistically significant, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. White: regions that were not analysed due to insufficient data. The omnibus test of decoding from all neurons in each session was significant $p \ll 10^{-10}$. **b)** Single-cell analysis: The flat brain map of the fraction of neurons for which firing rates are significantly modulated by block prior (left vs right) in the period $[-400, -100]$ ms aligned to the stimulus onset. The significance scores of the single cells were obtained by a Mann-Whitney test, using a threshold of $p < 0.001$. The null distribution was generated via the pseudo-block method (Methods). The significance score of each region was computed by assuming a binomial distribution of false positive events, and $\text{FDR}_{0.01}$ was used to correct for multiple comparisons. The labels of regions are shown in f. The omnibus test of the fraction of significant neurons from all neurons in each session was highly significant $p \ll 10^{-10}$. **c)** Manifold distance: Quantification of the time-resolved maximum distance between left (right) block trajectories, based on Euclidean distance (in Hz/cell) in the full-dimensional space (dimension = number of cells) for each brain region. The omnibus test based on treating all cells from all sessions as one region resulted in insignificant modulation, $p = 0.9$. **d)** Manifold latency: Times at which 70% of the maximum of the trajectory distance was reached (see part c), for significant regions only, of which there was none for block. **e)** Encoding: The flat brain map of mean model improvements per-region, across all neurons in that region, when the ITI P(left) regressor is added to the model. The regressor is a step function bounded to be the value of the block (0.5, 0.2, 0.8) P(left) in the period $[-400, -100]$ ms aligned to stimulus onset. **f)** Tabular form of effect significance (grey - not significant; a-c) and effect size (by darkness; a-c; e), grouped by region. **g)** Decoding analysis uses logistic regression to identify blocks. We show the predicted probability across a window of trials (100 to 400) in region MOp (balanced accuracy in this region-session was 0.666; see Table 3 for session and neuron details). **h)** Distribution of firing rates for an example neuron in region MOp, in the window used in c), along with the distribution of GLM encoding model predicted firing rates in the same window (see Table 3 for session and neuron details). **i)** Example trajectories for region MOp, obtained from left (right) choice trial-averaged activity, reduced via PCA to 3 dimensions, one dot being a time bin of the population activity, darker is later. Grey pseudo-trajectories were obtained by averaging randomised trials, here using the pseudo-session method only. **j)** Trajectory distance as a function of time for example region PL, showing low amplitude and random fluctuations, indistinguishable from pseudo-trajectory distances in grey (control).

Representation of Wheel Movement. A consistent finding from previous large-scale recordings in mice has been the macroscopic impact of movement on neural activity, with both task-related and task-unrelated movements influencing activity well outside premotor, motor, and somatosensory cortical areas [19,20,94,95]. Here, we start from the task-dependent component of movement, namely the movement of the wheel to register a response. Unlike the other task variables, movement trajectories change relatively quickly, necessitating different analysis and null control strategies. Accordingly, we only report simple encoding and decoding analyses.

We observed that different mice (and potentially the same mouse on different sessions) adopted different strategies for moving the wheel – for instance, some used both front paws; others only one. Turning the wheel is also a relatively complex operation, rather than being just a simple, ballistic movement. Thus, one should not expect a simple relationship between these movements and activity in unilateral motor systems. Again, for simplicity, we restrict our analyses to the activity associated with both wheel velocity and its absolute value, wheel speed.

Wheel speed was decodable from a wealth of areas in the brain. Top regions include sensory-motor related regions, both cortical (SSp-ll, VISC, ORBvl, POST) and subcortical (MS, VAL, VPL, GRN, PARN, IRN, RT, DCO) (plus there are cerebellar regions SIM, COPY, IP) (Fig. 9a;e). For example, we could readily decode wheel speed from single trials of activity in region GRN (Fig. 9g).

The encoding analysis confirmed that many regions across the brain were sensitive to the wheel speed during the task, with ΔR^2 taking values orders of magnitude larger than for the other variables considered besides feedback (Fig. S5). The pontine reticular nucleus (PRNc) and gigantocellular reticular nucleus (GRN), in particular, stood out in our analysis for the mean ΔR^2 for neurons within those regions (mean $\Delta R^2 = 9.9 \times 10^{-3}$ in PRNc, and $\Delta R^2 = 24.0 \times 10^{-3}$ in GRN). Many other cortical (e.g. MOs) and subcortical (e.g. GPe, GPi, CP) regions had less substantial, but still above-average, correlations with the wheel speed relative to other regressors (Fig. 9b;e).

Wheel velocity was also significantly decodable from a rather similar collection of areas as wheel speed (Fig. 9c;f;h; S15a), and was also duly encoded (Fig. 9d;f) albeit with generally smaller values of ΔR^2 (Fig. 9i). The apparently high decodability of velocity was unexpected given the complexities of wheel movement mentioned above, and indeed the uncorrected values of R^2 for decoding speed were substantially larger than those for velocity in most regions (Fig. S15b). However, the null distribution based on imposter sessions (i.e., wheel movements from other sessions, including from other mice) can be decoded much more accurately for speed than for velocity (Fig. S15c), reducing the statistical significance of the decoding of speed. We attributed this excess decodability of the null distribution to the more stereotyped, i.e., less variable, trajectory of speed (Fig. S15d.

We also correlated neural activity with behavioural movement traces extracted from video (nose, paw, pupil, and tongue). To test for significance we used the linear-shift method, comparing the correlation of spiking activity with behavioural movement variables against a null ensemble in which the movement variables were shifted in time[96] (Methods, Fig. S16a). More than half of the neurons in most brain regions were significantly correlated with at least one behavioural variable (Fig. S16b).

The widespread relationship between neural activity and motion has various potential sources. These include the specific details of motor planning and execution, efference copy[97], somatosensory feedback, the suppression of input associated with self-motion[98], and more subtle effects such as the change in other sensory inputs caused by the movement[19], or even prediction errors associated with incompetent execution that can fine-tune future performance[99]. Others are more general, including arousal and the calculation and processing of the costs of movement (which would then be balanced against future gain)[100]. More

generally, of the components that are indeed specific, only a fraction are likely to be associated with the wheel movement that we monitored, as opposed to other task-related motor actions. This is especially true given the results of previous studies such as Refs[19,20] showing just how important uninstructed movements are in modulating a wide swathe of neural activity.

Two facts suggest that there might be an even richer coupling of action and activity in the task. One is that, in at least a minority of cases, the mice embody their block-based expectations about the upcoming stimulus in a form of externally observable pose or movement that we can extract with DeepLabCut[30] (Fig. S3). This implies that some of what we would report as an association with the block could really be an association with action. Second, and more subtly, in more detailed analyses of how their choices are influenced by the block (see the companion paper [93]), we found that the mice exhibited a form of action perseveration (i.e., the law of exercise[101]). That is, they had a tendency to repeat recent previous actions, in way that apparently covaries with their estimate of the block. Thus, again, some of what we would report as an association with the block could really be an association with action.

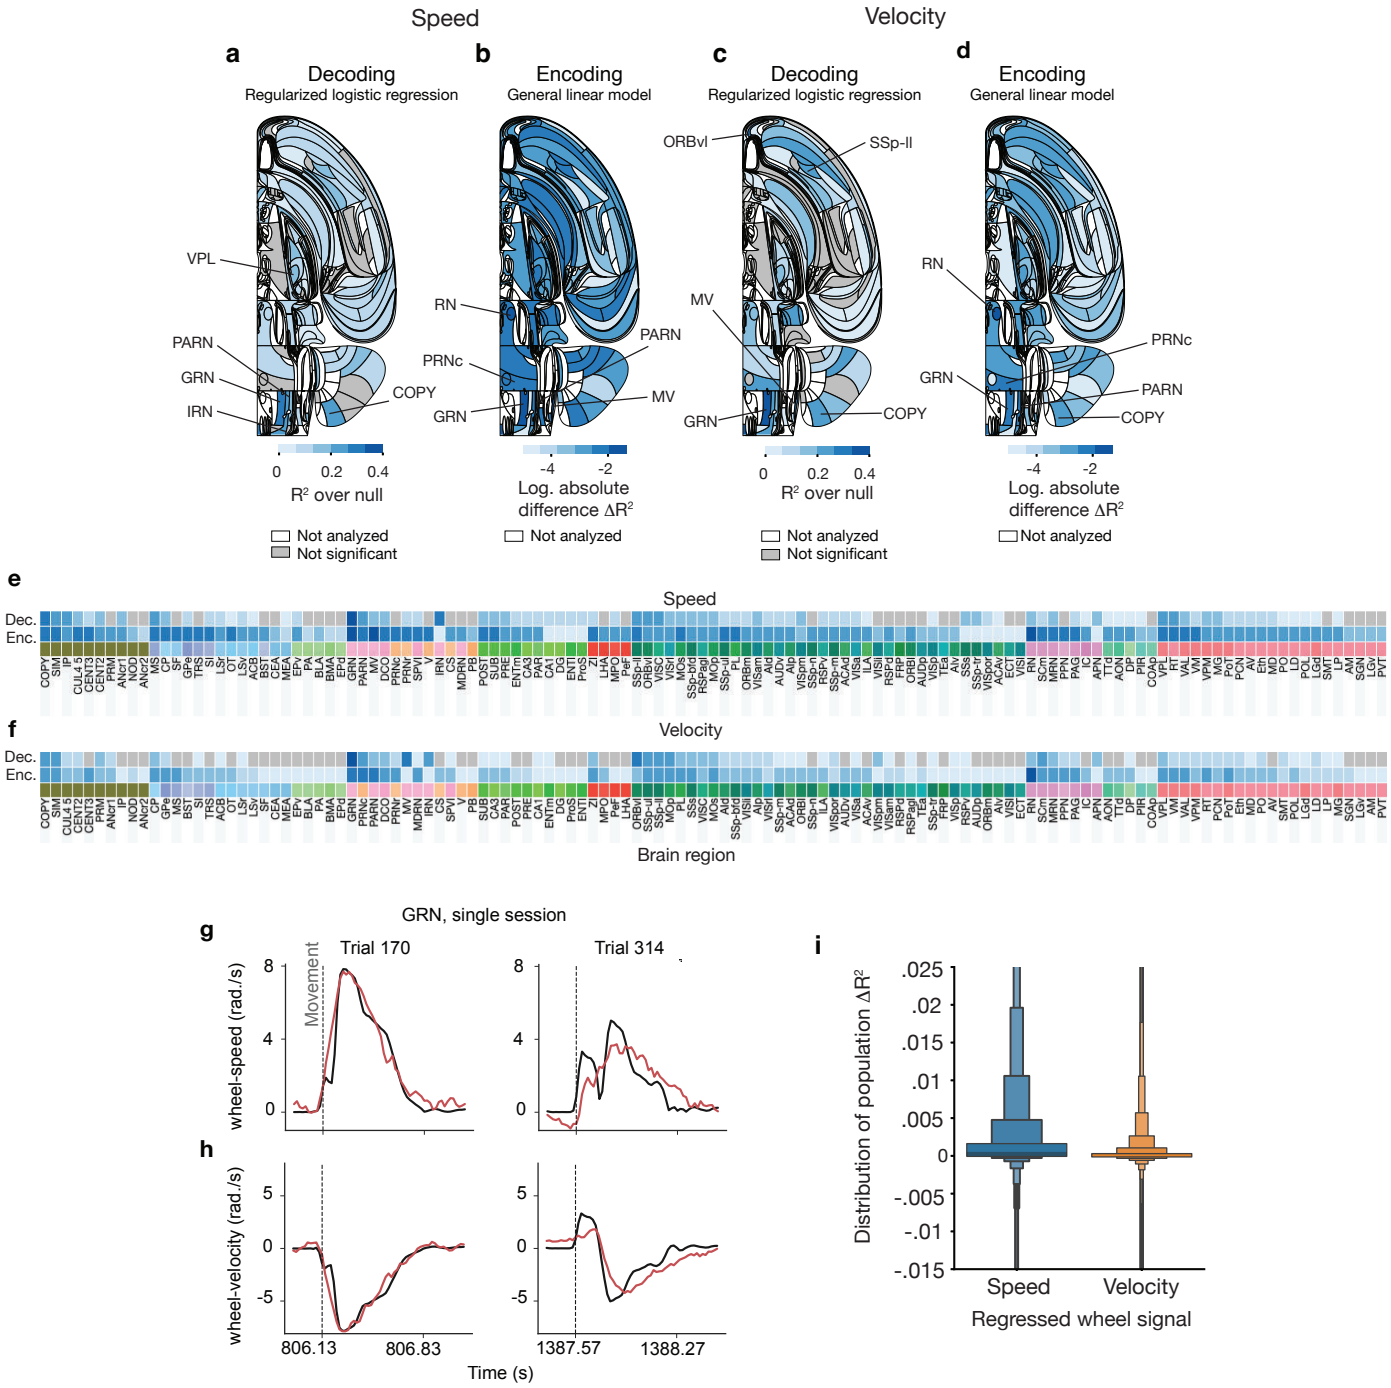

Figure 9. Representation of Wheel Movement. See also Fig. S6; S15. An interactive version of this figure can be found on our data website for speed and velocity. **a)** Decoding: Flat brain map of median corrected decoding score (R^2) across sessions for wheel speed (the absolute value of the wheel velocity). The values have been corrected by subtracting the median of the decoding score in the null distribution. Colour represents effect size. Grey: regions in which decoding was found not to be statistically significant, using $\text{FDR}_{0.01}$ to correct for multiple comparisons. White: regions that were not analysed due to insufficient data. The omnibus tests of both wheel-speed decoding and wheel-velocity decoding from all neurons in each session were significant $p \ll 10^{-10}$ **b)** Encoding analysis: Flat brain map of mean model improvement differences per-region, across all neurons in that region computed as the improvements (ΔR^2) from the wheel speed regressor. These regressors are 200 ms anti-causal kernels convolved with the trace of wheel speed. **c;d)** Same as a;b but for wheel velocity rather than wheel speed. Note that the encoding results involve a completely separate model fit using the velocity rather than speed. **e)** Tabular form of effect significance for wheel speed (grey - not significant; a) and effect size (by darkness; a;b), grouped by region. Regions are sorted within each Cosmos group by the sum of effects across analyses. **f)** Tabular form of effect significance for wheel velocity (grey - not significant; c) and effect size (by darkness; c;d), grouped by region. Regions are sorted within each Cosmos group by the sum of effects across analyses, for significant regions only. **g)** Comparison of actual and predicted wheel speed for an example trial from region GRN (R^2 is 0.662 uncorrected and 0.432 null-corrected for this region-session; see Table 3 for session and neuron details). **h)** Same as g, for a target signal of velocity rather than speed (R^2 is 0.648 uncorrected and 0.623 null-corrected for this region-session). **i)** A comparison of the distributions of additional variance ΔR^2_{wheel} explained, across the whole population, when using either wheel speed or velocity as the base signal. Distributions truncated for clarity.

Discussion

Building on previous efforts to build brain-wide maps of activity in the mouse at neuronal resolution using Neuropixels probes[15–17] or in other species using imaging (e.g., Refs.[5,102,103]), we have developed a new strategy for assembling large-scale electrophysiological maps by pooling data from numerous laboratories that employed the same standardised and reproducible task[23]. This approach places replication at the very foundation of the data set, because each Neuropixels insertion was repeated in at least two laboratories, with reproducibility of outcomes across laboratories verified with extensive analyses that we have previously reported[24]. The result is an unprecedented brain-wide consensus map of neural activity of the mouse brain. This map reveals the neural activity that underlies performance in a rich task that requires processing of visual stimuli, motor responses, rewards, and stimulus expectations.

In addition to releasing the data – which are available to download via an API and to view interactively at data.internationalbrainlab.org – we also performed a battery of standard analyses coupled with rigorous statistical methods, which suggest that neural activity throughout the entire brain correlates with some aspects of the task, but with major differences in the ubiquity of representation of different task variables (Fig. 5- 9), see also Fig. S6; S7; S17 for side by side comparisons). Importantly, while our coverage is

extensive, we do not have uniform coverage across all regions, particularly when considering the neurons that pass our quality control metrics. This could bias our results against discovering task correlates in regions with certain anatomical arrangements (e.g., where cell bodies are densely packed).

The neural representations of movement[16,19,20] (Fig. 9) and feedback (Fig. 7) were particularly widespread, with the former potentially reflecting a brain-wide change in the state of neural processing during movement periods, along with specific encoding of motor features. The hypothesis of a brain-wide state change is consistent with findings that the neural representation of upcoming movements in cortex is extremely widespread, though not all of this activity is causally related to the performance of the movements[47]. As for the brain-wide correlates of feedback, they may also partly or primarily reflect the licking movements required for reward consumption, rather than the hedonic aspects of reward as such. Distinguishing these possibilities would require experiments recording activity during the presentation of reinforcement with no motor correlates, such as optogenetic stimulation of dopamine systems[104,105], or by comparing correlates of rewards with a motor correlate to the same movements when their hedonic reward was devalued, for example by satiation.

The subject's upcoming choice was represented in the activity of neurons across brain systems including cortex, basal ganglia, thalamus, midbrain, hindbrain, and cerebellum (Fig. 6). These representations cannot reflect sensory reafference (i.e. responses related to sensory stimuli that occur as part of the movements, such as pressure on the paws and movements of the visual stimuli across the screen) since we only analysed the time period prior to the earliest detectable movement onset. Moreover, our carefully controlled task design and pseudo-session statistical methods mean that choice coding reported in the single-cell and manifold analyses cannot reflect processing of the visual stimulus nor non-specific brain states such as arousal. Instead, these responses reflect aspects of decision formation or motor preparation, potentially including corollary discharge specific to the chosen action[106,107]. This study is the first report at single-cell resolution of choice coding across these brain systems in a single task, and while many past studies have focused on the role of cortex, basal ganglia, or midbrain in visual perceptual decisions [7,16,55,56,63–65,108,109], here we discovered that parts of the medulla, pons, and cerebellum are all selectively responsive with similar timing to those areas. Our data are not able to determine whether these different systems make unique contributions to decision formation and execution. However, they rule out a model in which only a limited set of systems subserve a given behaviour according to specific task demands.

Other task variables such as visual stimuli (Fig. 5) and the identity of the bias block (Fig. 8) were represented (prior to movement) in more restricted manners. The processing of visual stimuli followed a temporal sequence from traditional visual areas such as the visual thalamus and cortex to midbrain and hindbrain regions whose activity also correlated with choices. Importantly, the temporal structure of activity in these two groups of regions differed, with visual representations in classical visual regions showing a transient representation of the stimuli, and activity in midbrain and hindbrain showing later, ramping activity, consistent with a role of this activity in decision-making. The representation of stimulus expectation (i.e. the bias block) was more subtle still: it could be significantly decoded from the activity of a small number of neurons in restricted regions but could not be found in a much larger number of areas. The representation of this variable is addressed in more detail in a companion paper [93] by analysing the neural correlates of the

mouse's behavioural estimate of the block prior (rather than the veridical block identity known only to the experimenter, which is what we analysed here), and with additional recording modalities.

Although more than half of the recorded neurons in most brain regions were significantly modulated by at least some aspect of the task (Fig. S18), our ability to explain the total variance of single neurons was limited (the mean R^2 across all neurons was 0.0155 with a median of 0.0027 and standard deviation of $\sigma = 0.044$). This indicates that the vast bulk of activity in the brain is not modulated by the task. It may instead be related to noise, to uninstructed movements[19,20], or to other processes that are not timed to the task events. These processes often result in activity that is strongly symmetrical[110] and are thus best captured through bilateral measurements.[111] Our brain-wide map, however, is unilateral, so it is not ideally suited for this purpose. Even for the activity that is modulated by the task, external cue-driven responses were consistently smaller than internally generated signals such as those arising in relation to integration of the stimulus and movement planning. As ever, absence of evidence for a neural representation of a task variable in a given region cannot be taken to indicate evidence of absence.

In summary, we have provided a database of recordings spanning the entire mouse brain as subjects performed a complex behavioural task. Our analyses have begun to elucidate the mechanisms by which the mouse brain solves this task. Further examination beyond the scope of the present work, including detailed analyses at the level of sub-regions (such as cortical layers or functional zones of the striatum) and cell types (as identifiable from extracellular waveforms, such as broad versus narrow spike shapes in cortex), can be addressed in this rich dataset that is freely available to the community.

Resources

Data access Please follow these instructions to download the data used in this article. You can also browse the data at `http://viz.internationalbrainlab.org`.

Data quality Please read our white paper on the released data for additional details about quality control and metrics. For electrophysiology data, please also see the spreadsheet listing the quality of the raw data and spike sorting output.

Code repository Please visit our Github repository to access the code used to produce the results and figures presented in this article.

Protocols and pipelines Please visit our two Figshare repositories to access the protocols and pipelines used to train mice and perform the electrophysiology recording and histology validation.

Methods

All procedures and experiments were carried out in accordance with the local laws and following approval by the relevant institutions: the Animal Welfare Ethical Review Body of University College London; the Institutional Animal Care and Use Committees of Cold Spring Harbor Laboratory, Princeton University, University of Washington, University of California at Berkeley and University of California at Los Angeles; the University Animal Welfare Committee of New York University; and the Portuguese Veterinary General Board.

Animals Mice were housed under a 12/12 h light/dark cycle (normal or inverted depending on the laboratory) with food and water available ad libitum, except during behavioural training days. Electrophysiological recordings and behavioural training were performed during either the dark or light phase of the cycle depending on the laboratory. The data from N=115 adult mice (C57BL/6; 80 male and 35 female, obtained from either Jackson Laboratory or Charles River) is used in this study. Mice were aged 13-122 weeks (mean 34.43 weeks, median 26.0 weeks) and weighed 16.1-36.2 g (mean 24.17 g, median 24.0 g) on the day of electrophysiological recording.

Headbar implant surgery A detailed account of the surgical methods for the headbar implant is in Appendix 1 of Ref.[23]. Briefly, mice were anesthetised with isoflurane and head-fixed in a stereotaxic frame. The hair was then removed from their scalp, which was subsequently removed along with the underlying periosteum. was Once the skull was exposed, Bregma and Lambda were marked. The head was positioned along the anterior-posterior and left-right axes using stereotaxic coordinates. The head bar was then placed in one of three stereotactically defined locations and cemented (Super-Bond C&B) in place. Future craniotomy positions were marked on the skull relative to Bregma. The exposed skull was then covered with cement and clear UV curing glue (Norland Optical Adhesives).

Materials and apparatus For detailed parts lists and installation instructions for the training rigs, see Appendix 3 of Ref.[23]; for the electrophysiology rigs, see Appendix 1 of Ref.[24].

Each lab installed a standardised electrophysiological rig, which differed slightly from the apparatus used during behavioural training [23]. The structure of the rig was constructed from Thorlabs parts and was placed on an air table (Newport, M-VIS3036-SG2-325A) surrounded by a custom acoustical cabinet. A static head bar fixation clamp and a 3D-printed mouse holder were used to hold a mouse such that its forepaws rested on the steering wheel (86652 and 32019, LEGO)[23]. Silicone tubing controlled by a pinch valve (225P011-21, NResearch) was used to deliver water rewards to the mouse. Visual stimuli were displayed on an LCD screen (LP097Q × 1, LG). To measure the timing of changes in the visual stimulus, a patch of pixels on the LCD screen flipped between white and black at every stimulus change, and this flip was captured with a photodiode (Bpod Frame2TTL, Sanworks). Ambient temperature, humidity, and barometric air pressure were measured with the Bpod Ambient module (Sanworks) and wheel position was monitored with a rotary encoder (05.2400.1122.1024, Kubler). Videos of the mouse were recorded from 3 angles (left, right and body) with USB cameras (CM3-U3-13Y3M-CS, Point Grey) sampling at 60, 150, 30 Hz respectively (for details see Appendix 1 of Ref.[24]). A custom speaker (Hardware Team of the Champalimaud Foundation for the Unknown, V1.1) was used to play task-related sounds, and an ultrasonic microphone (Ultramic UM200K, Dodotronic) was used to record ambient noise from the rig. All task-related data was coordinated by a Bpod State Machine (Sanworks). The task logic was programmed in Python and the visual stimulus presentation and video capture was handled by Bonsai[112] and the BonVision package[113].

Neural recordings were made using Neuropixels 1.0 (3A or 3B) probes (Imec[13]), advanced in the brain using a micromanipulator (Sensapex, uMp-4). Typically, the probes were tilted at a 15 degree angle from the vertical line. Data were acquired via an FPGA (for 3A probes) or PXI (for 3B and 1.0 probes, National Instrument) system using SpikeGLX, and stored on a PC.

Habituation, training, and experimental protocol For a detailed protocol on animal training, see Methods in Refs.[23,24]. Briefly, at the beginning of each trial, the mouse was required to not move the wheel for a quiescence period of 400–700 ms. After the quiescence period, a visual stimulus (Gabor patch) appeared on either the left or right ($\pm 35^\circ$ azimuth) of the screen, with a contrast randomly sampled from a predefined set (100, 25, 12.5, 6, 0%) . A 100 ms tone (5 kHz sine wave) was played at stimulus onset. Mice had 60 s to move the wheel and make a response. Stimuli were yoked to the rotation of the response wheel, such that a 1 millimetre movement of the wheel moved the stimulus by 4 visual degrees. A response was registered if the centre of the stimulus crossed the $\pm 35^\circ$ azimuth line from its original position. If the mouse correctly moved the stimulus 35° to the centre of the screen, it immediately received a 3 μL reward; if it incorrectly moved the stimulus 35° away from the centre, it received a timeout. If the mouse responded incorrectly or failed to reach either threshold within the 60 s window, a white noise burst was played for 500 ms and the inter-trial interval was set to 2 s. In trials where the visual stimulus contrast was set to 0%, the mouse had to respond as for any other trial by turning the wheel in the correct direction (assigned according to the statistics of the prevailing block) to receive a reward, but the mouse was not able to perceive whether the stimulus was presented on the left or right side of the screen. The mouse also received feedback (noise

burst or reward) on 0% contrast trials.

Each session started with 90 trials in which the probability of a visual stimulus appearing on the left or right side was equal. Specifically, the 100%, 25%, 12.5%, and 6% contrast trials were each presented 10 times on each side, and the 0% contrast was presented 10 times in total (i.e. the ratio of the [100 : 25 : 12.5 : 6 : 0] % contrasts is set at [2 : 2 : 2 : 2 : 1]). The side (and thus correct movement) for the 0% contrast trials was chosen randomly between the right and left with equal probability. This initial block of 90 trials is referred to as the unbiased block (50/50).

After the unbiased block, trials were presented in biased blocks: in right bias blocks, stimuli appeared on the right on 80% of the trials, while in left bias blocks, stimuli appeared on the right on 20% of the trials. The ratio of the contrasts remained as above ([2 : 2 : 2 : 2 : 1]). Whether the first biased block in a session was left or right was chosen randomly, and blocks then alternated.

Mice were trained on this task following a stepwise shaping strategy in which progressively lower contrast stimuli were introduced as performance increased[23].

Electrophysiological recording using Neuropixels probes For details on the craniotomy surgery, see Appendix 3 of Ref.[24].

Briefly, upon the first day of electrophysiological recording, the animal was anaesthetised using isoflurane and surgically prepared. The mouse was administered with analgesics (typically Carprofen) subcutaneously. The glue was removed (typically using a biopsy punch (Kai Disposable Biopsy Punches (1mm)) or a drill), exposing the skull over the planned craniotomy site(s). A test was made to check whether the implant could hold liquid; the bath was then grounded either via a lose or implanted pin. One or two craniotomies (approximately 1 × 1 mm) were made over the marked locations. The dura was left intact, and the brain was lubricated with ACSF. A moisturising sealant was applied over the dura (typically DuraGel (Cambridge NeuroTech) covered with a layer of Kwikcast (World precision instruments)). The mouse was left to recover in a heating chamber until locomotor and grooming activity were fully recovered.

Subjects were head-fixed for recording after a minimum recovery period of 2 hours. Once a craniotomy was made, up to 4 subsequent recording sessions were made in that same craniotomy. Once the first set of craniotomy was fully recorded from, a subject could undergo another craniotomy surgery in accordance with the institutional licence. Up to two probes were implanted in the brain on a given session. CM-Dil (V22888 Thermofisher) was used to label probes for subsequent histology.

Serial section two-photon imaging Mice were given a terminal dose of pentobarbital and perfuse-fixed with PBS followed by 4% formaldehyde solution (Thermofisher 28908) in 0.1M PB pH 7.4. The whole mouse brain was dissected, and post-fixed in the same fixative for a minimum of 24 hours at room temperature. Tissues were washed and stored for up to 2-3 weeks in PBS at 4C, prior to shipment to the Sainsbury Wellcome Centre for image acquisition. For full details, see Appendix 5 of Ref.[24].

section	filter	remaining		
		sessions	probes	neurons
5.9.1	session and insertion QC	354	547	295,501
5.9.2	reaction time & missing events	353	545	293,566
5.9.2	minimum 3 error trials	352	543	292,165
5.9.3	single unit QC	352	543	32,525
5.9.3	gray matter regions	351	542	28,972
5.9.3	minimum 10 units/region	312	450	22,818
5.9.3	minimum 2 sessions/region	308	439	22,113

Table 1. Session and probe filtering. The table indicates the progressive filtering of the sessions and probes based on the various inclusion criteria described in the text.

For imaging, brains were equilibrated with 50mM PB solution and embedded into 5% agarose gel blocks. The brains were imaged using serial section two-photon microscopy[114,115]. The microscope was controlled with ScanImage Basic (Vidrio Technologies, USA), and BakingTray, a custom software wrapper for setting up the imaging parameters[116]. Image tiles were assembled into 2D planes using StitchIt[117]. Whole brain coronal image stacks were acquired at a resolution of 4.4 x 4.4 x 25.0 μm in XYZ, with a two-photon laser wavelength of 920 nm, and approximately 150 mW at the sample. The microscope cut 50 µm sections but imaged two optical planes within each slice at depths of about 30 µm and 55 µm from the tissue surface. Two channels of image data were acquired simultaneously using multialkali PMTs ('Green' at 525 nm $\pm$25 nm; 'Red' at 570 nm low pass).

Whole brain images were downsampled to 25 μm XYZ pixels and registered to the adult mouse Allen common coordinate framework[6] using BrainRegister[118], an elastix-based[119] registration pipeline with optimised parameters for mouse brain registration. For full details, see Appendix 7 of Ref.[24].

Probe track tracing and alignment Neuropixels probe tracks were manually traced to yield a probe trajectory using Lasagna[120], a Python-based image viewer equipped with a plugin tailored for this task. Traced probe track data was uploaded to an Alyx server[121]; a database designed for experimental neuroscience laboratories. Neuropixels channels were then manually aligned to anatomical features along the trajectory using electrophysiological landmarks with a custom electrophysiology alignment tool[122,123]. For full details, see Appendix 6 of Ref.[24].

Spike sorting The spike sorting pipeline used at IBL is described in detail in Ref.[27]. Briefly, spike sorting was performed using a modified version of the Kilosort 2.5 algorithm[124]. We found it necessary to improve the original code in several aspects (scalability, reproducibility, and stability, as discussed in Ref.[24]), and developed an open-source Python port; the code repository is in Ref.[125].

Inclusion criteria We applied a set of inclusion criteria to sessions, probes and neurons to ensure data quality. Table 1 lists the consequences of these criteria for the number of sessions and probes that survived. They are described below.

Sessions and insertions Sessions were included in the data release if the mice performed at least 400 trials, with a performance of at least 90% correct on 100% contrast trials for both left and right blocks, and if there were at least three trials with incorrect choices (after applying the trial exclusions below). Furthermore, sessions were included in the release only if reached threshold on a collection of hardware tests (definitions can be found at `https://int-brain-lab.github.io/iblenv/_autosummary/ibllib.qc.task_metrics.html`)

Insertions were excluded if the neural data failed the whole recording visually assessed criteria of the 'Recording Inclusion metrics and Guidelines for Optimal Reproducibility' (RIGOR) from Ref.[24], by presenting major artefacts (see examples in Ref.[27]), or if the probe tract could not be recovered during the histology procedure. Furthermore, only insertions whose alignments had been resolved (see Appendix 6 of Ref.[24] for definition) were used in this study.

After applying these criteria, a total of 354 sessions, 547 insertions and 295,501 neurons remained, constituting the publicly released data set.

Trials For the analyses presented here, trials were excluded if one of the following trial events could not be detected: `choice`, `probabilityLeft`, `feedbackType`, `feedback times`, `stimOn times`, `firstMovement times`. Trials were further excluded if the time between stimulus onset and first movement of the wheel was outside the range of [0.08, 2.00] seconds.

Neurons and brain regions Neurons output by the spike sorting pipeline were excluded from the analyses presented here if they failed one of the three criteria described in Ref.[27] (the single unit computed metrics of RIGOR[24]): amplitude >50 μV ; noise cut-off < 20 μV; refractory period violation. Neurons that passed these criteria are termed 'good' neurons (or often just 'neurons') in this study. Out of the 295,501 neurons collected, 32,766 were considered good neurons. Final analyses were additionally restricted to regions that a) are designated gray matter in the adult mouse Allen common coordinate framework[6], b) contained at least 10 good neurons per session, and c) were recorded from in at least 2 such sessions.

Video analysis We briefly describe the video analysis pipeline; full details can be found in Ref.[126]. The recording rigs contain three cameras, one called 'left' at full resolution (1280x1024) and 60 Hz filming the mouse from one side; one called 'right' filming the mouse symmetrically from the other side at half resolution (640x512) and 150 Hz; and one called 'body' at half resolution and 30 Hz filming the body of the mouse from above. We developed several quality control metrics to detect raw video issues such as poor illumination (as infrared light bulbs broke) or accidental misplacement of the cameras[126].

We computed the motion energy (the mean across pixels of the absolute value of the difference between adjacent frames) of the whisker pad areas in the 'left' and 'right' videos. The whisker pad area was defined empirically using a rectangular bounding box anchored between the nose tip and the eye, both found using DeepLabCut[127] (DLC; see more below). This metric quantifies motion in the whisker pad area and has a temporal resolution of the respective camera.

We also performed markerless pose estimation of body parts using DLC[30], which is used within a fully automated pipeline in IBL (version 2.1) to track various body parts such as the paws, nose, tongue, and pupil. In all analyses using DLC estimates, we drop predictions with likelihood < 0.9. Furthermore, we developed several quality control metrics for the DLC traces[126].

Assessing statistical significance In this work, we studied the neural correlates of task and behavioural variables. In order to assess the statistical significance of these analyses, we need to account properly for spurious correlations. Spurious correlations can be induced in particular by slow continuous drift in the neurophysiological recordings, due to various factors including movement of the Neuropixels probes in the brain. Such slow drifts can create temporal correlations across trials. Because standard correlation analyses assume that all samples are independent, they can yield apparently significant nonsense-correlations even for signals that are completely unrelated[128,129].

Null distributions were generated against which we tested the statistical significance of our results. More specifically, we used distinct null distributions for each of the three types of variables we considered: discrete behaviour-independent variables (e.g., the block prior and the stimulus side), discrete behaviour-dependent variables (e.g., reward and choice), and continuous behaviour-dependent variables (e.g., wheel speed and wheel velocity). For the rest of the section, we will denote the aggregated neural activity across L trials and N neurons by $S \in \mathbb{R}^{L\times N}$, and denote the vector of scalar targets across all trials by $C \in \mathbb{R}^{L}$.

For discrete behaviour-independent variables, we generated the null distribution from so-called "pseudo-sessions", i.e., sessions generated from the same generative process as the one used for the mice. This ensured that the time series of trials in each pseudo-session shares the same summary statistics as the ones used in the experiment. We generated M (typically $M = 200$) pseudo-targets $\widetilde{C}_i, i \in [1, M]$, and performed the given analysis on the pair $\left(S, \widetilde{C}_i\right)$ and obtained a fit score $\widetilde{F}_i$. In pseudo-sessions, the neural activity S should be independent of $\widetilde{C}_i$ as the mouse did not see $\widetilde{C}_i$ but rather C. Any predictive power from $\widetilde{C}_i$ to S (or from S to $\widetilde{C}_i$) would arise, for instance, from slow drift in S unrelated to the task itself. These pseudo-scores $\widetilde{F}_i$ can be compared to the actual score F obtained from the neural analysis on (S, C) to assess statistical significance.

For discrete behaviour-dependent variables (such as choice or reward), we could not use the pseudo-session procedure above as we did not know the underlying generative process in the mouse. We therefore used "synthetic" sessions to create a null distribution. These depend on a generative model of the process governing the animals' choices. In turn, this requires a model of how the animals estimated the prior probability that the stimulus appears on the right or left side of the screen, along with a model of its response to different contrasts given this estimated prior. We found (see our companion paper on the subjective prior [93]) that the best model of the prior across all animals uses a running average of the past actions as a subjective prior of the side of the next stimulus, which we refer to as the 'action-kernel' model. The subjective prior π_t follows the update rule:

$$\pi_{t+1}|\pi_t, a_t, \alpha = (1-\alpha)\cdot\pi_t + \alpha\cdot(a_t > 0)$$

with $a_t \in \{-1, 1\}$ the action performed by the mouse on trial t and α the learning rate. This effectively

models how mice use information from previous trials to build a subjective prior of where the stimulus is going to appear at the next trial. The details of how this prior is integrated with the stimulus to produce a decision policy is described in the companion paper[93].

With this model of the mouse's decision-making behaviour in hand, we fit an action-kernel model to the mouse's behaviour on each session. Then, we generated stimuli from pseudo-sessions and simulate the fitted model on the pseudo-sessions to obtain times series of choice and reward. This leads to the "synthetic" targets $\widetilde{C}_i$. Then, as we did with pseudo-sessions, we obtained pseudo-scores $\widetilde{F}_i$ and assessed statistical significance by comparing the distribution of pseudo-scores to the actual score F obtained from the neural analysis on (S, C).

For the third type of variable – continuous behaviour-dependent variables such as wheel speed – generating synthetic sessions is harder, as we do not have access to a reasonable generative model of these quantities. We instead used what we call "imposter" sessions, generated from the continuous behaviour-dependent variable from another mouse on another session. More precisely, an imposter session for an original session of L trials is generated by performing the following steps:

1. concatenate trials across all sessions analysed in this paper (leaving out the session under consideration)
2. randomly select a chunk of L consecutive trials from these concatenated sessions
3. return the selected chunk, the imposter session

The continuous behaviour-dependent variable can then be extracted from the imposter session. As with the pseudo-sessions and the synthetic sessions, we obtained pseudo-scores $\widetilde{F}_i$ from a collection of imposter sessions and assess statistical significance by comparing the distribution of pseudo-scores to the actual score F obtained from the neural analysis on (S, C).

Additional information about assessing statistical significance for individual analyses are detailed in the analysis-specific sessions below. For decoding, single-cell and manifold analyses, the results come in the form of per-region p-values. We used the false discovery rate to correct for comparisons across all the regions involved in each analysis (123 for the main figures), at a level of $q = 0.01$. We employed the Benjamini-Hochberg procedure[130] as we expected substantial independence among the tests. As noted, we were not able to assess significance for the encoding analysis because of a lack of a convenient null distribution.

Motor correlates of block The mice in this study were able to leverage block information to guide their choice behaviour (Fig. 3). We described our analyses of the neural representation of this information in Sec. ; but we must also consider the possibility that this information was represented in the non-choice *behaviour* of the mouse; for example, a mouse might have rested its paws on the right side of the wheel during

periods where the stimulus appears on the right more often. To search for such "embodied" strategies, video and rotary encoder (wheel) data were used to extract 7 behavioural variables as being potentially modulated according to the block structure: licking, left/right whisker pad motion, wheel movement, nose position, and left/right paw positions (all described below). For each behavioural variable we assessed whether the variable differed across blocks during the time window -400 ms to 0 ms relative to stimulus onset. We chose to analyse the behaviour before stimulus onset in order not to confound the behaviour with other trial properties such as the stimulus side and action taken. We summarised the amplitude of each of these behaviours by taking the mean across time bins in the window, and then compared the average amplitude across left block trials to that across right block trials (i.e. taking the absolute value of their difference). We then assessed statistical significance using a null distribution obtained by randomizing the block labels of the trials with 100 pseudo-sessions. A one-sided p-value was defined with the percentile of our statistic in this null distribution, say q, as $1-(q/100)$. We repeated this analysis using another time window, -600 ms to -200 ms relative to stimulus onset, in order to test how sensitive the metrics are to the exact position of the window. Below we provide more details on how the individual behaviours are extracted:

- "Licking" is obtained from videos from both side cameras of the animal using DLC to obtain two points on the tongue (left and right edge of the tongue). A lick is defined to have occurred in a frame if the difference for either coordinate to the subsequent frame is larger than 0.25 times the standard deviation of the difference of this coordinate across the whole session. This measure was empirically confirmed as detecting licks, via manual inspection of sample videos.
- "Whisking left/right" is defined by the motion energy in a bounding box around the whisker pad of the corresponding camera (see Sec.).
- "Wheeling" is obtained by computing the speed (rather than velocity) of the rotary encoder.
- "Nose position" is defined as the x coordinate of the DLC-tracked nose in the left camera video.
- "Paw position left/right" is the Euclidean distance between the left/right paw (tracked by DLC in the left/right camera) and the upper left frame corner of the frame.

Single-cell correlates of sensory, cognitive, and motor variables We quantified the sensitivity of single neurons to four task variables: the visual stimulus (left versus right location of the visual stimulus), choice (left versus right direction of wheel turning), prior expectation (left versus right block side), and feedback (reward versus non-reward). We computed the sensitivity metric for each task variable using the condition combined Mann-Whitney U statistic[16,131,132] (Fig. S4a,b). Specifically, we compared the firing rates from those trials with one task-variable value V_1 (e.g., trials with stimulus on the left side) to those with the other value V_2 (e.g., with stimulus on the right side), while holding the values of all other task variables fixed. In this way, we could isolate the influence of individual task variables on neural activity. To compute the U statistic, we first assigned numeric ranks to the firing rate observations in each trial. We then computed the sum of ranks R_1 and R_2 for the observations coming from n_1 and n_2 trials associated with the task-variable

values V_1 and V_2, respectively. The U statistic is defined as:

$$U = \min\left[R_1 - \frac{n_1(n_1+1)}{2}, R_2 - \frac{n_2(n_2+1)}{2}\right] . \tag{1}$$

The probability that the firing rate on V_1 trials is different (greater or smaller) from the firing rate on V_2 trials is computed as $1 - P$, where P is given by

$$P = \frac{U}{n_1 n_2} , \tag{2}$$

which is equivalent to the area under the receiver operating characteristic curve[133,134]. The null hypothesis is that the distributions of firing rates on V_1 and V_2 trials are identical.

To obtain a single probability across conditions, we combined observations across different trial conditions j by a sum of U statistic in these conditions[16]:

$$P = \frac{\sum_j U_j}{\sum_j n_{1,j} n_{2,j}} . \tag{3}$$

Here $n_{1,j}$ and $n_{2,j}$ are the numbers of V_1 and V_2 trials, respectively, within the condition j.

For the visual stimulus, we compared firing rate on trials with the stimulus on the left versus stimulus on the right during time-window $[0, 100]$ ms aligned to the stimulus onset time. For choice, we compared firing rate on trials with the left versus right choice during time-window $[-100, 0]$ ms aligned to the movement onset time. For the block side, we compared firing rate on trials with the left versus right block prior during time-window $[-400, -100]$ ms aligned to the stimulus onset. For the feedback, we compared firing rate on trials with reward versus non-reward during time-window $[0, 200]$ ms aligned to the feedback onset time.

To estimate statistical significance, we used a permutation test in which trial labels for one task variable were randomly permuted 3,000 times within each subset of trials with fixed values of all other task variables, and the Mann-Whitney U statistic was computed for each permutation. We computed the p-value for each task variable as the fraction of permutations with the statistic P greater than in the data. This approach controls for correlations among task variables and allows us to isolate the neuron's sensitivity to stimulus that is not due to sensitivity to block and choice, and vice versa. Random permutations, however, do not control for spurious correlations that can arise due to autocorrelations in the time series of the firing rate and task variable[128]. To control for spurious correlations, we used the pseudo-session approach[128] to generate a null distribution of block-side labels. The pseudo-session approach controls for temporal correlations but destroys correlations among task variables, which is complementary to the random permutations. For the single neuron correlates with block-side, we tested the statistical significance with pseudo-session null distribution. For the single neuron correlates with visual stimulus, choice and feedback, we cannot use the pseudo-session approach to generate null distributions, because it cannot preserve the original statistical distribution of stimulus or choice in permuted blocks. Instead, for the single neuron correlates with these task variables (visual stimulus, choice and feedback), we used a within-block permutation test to simultaneously control for both temporal correlations and correlations among task variables. Specifically, we

generated the null distribution by randomly permuting trial labels with fixed values of all other task variables within each individual block, which effectively reduces the serial dependencies of task variables at the timescale of block duration.

For visual stimulus, choice, and feedback, the combined condition Mann-Whitney U statistic can have a relatively high false positive rate due to the limited number of trials in each condition. To obtain a sufficient number of trials, we also performed simple Mann-Whitney U statistic without separating different conditions. We defined $p < 0.001$ ($\alpha_{MW} = 0.001$) as the criterion of significance for the simple Mann-Whitney U statistic, and $p < 0.05$ ($\alpha_{CCMW} = 0.05$) for the combined condition Mann-Whitney U statistic. We defined neurons that were statistically significant in both tests to be sensitive neurons for a specific task variable. For block variable, we performed Mann-Whitney U test and defined $p < 0.001$ ($\alpha_{MW} = 0.001$) as the criterion of significance.

To quantify the overall responsiveness of single neurons to the behavioural task, we used the Wilcoxon rank-sum test to compare firing rates between the baseline ($[-200, 0]$ ms window aligned to the stimulus onset) and different task periods: $[50, 150]$ ms and $[0, 400]$ ms aligned to stimulus onset, $[-100, 50]$ ms and $[-50, 200]$ ms aligned to the first-movement onset, and $[0, 150]$ ms aligned to the reward delivery.

To measure the behavioural movement correlates of single neurons in the entire recording sessions, we computed zero time-lag Pearson correlation coefficient between time-series of spike counts in 50ms bins and time-series of four behavioural variables (nose, pupil, paw, and tongue) each extracted from videos of the subject using DeepLabCut software[30]. To assess the significance of these correlations, we applied a time-shift test[135] and computed $2K = 40$ time-shifted correlations varying the offset between time-series of spiking activity and behavioural variables from 50 to 1,000 ms (both positive and negative offsets). We then counted the number of times m where the absolute value of time-shifted correlation exceeds that of zero time-lag correlation and assigned the p-value as the fraction of the absolute value of permuted correlations greater than in the data $p = m/(K + 1)$. We can then assign each neuron as being significantly responsive relative to a particular threshold on this p-value.

We then computed the fraction of neurons in each brain region that were significantly responsive to the behavioural task, movement, visual stimulus, choice and prior expectation, and identified brain regions that were most responsive to these conditions. Specifically, for each region, we computed p-value of the fraction of neurons (f_i) in i-th session by comparing the fraction to a binomial distribution of fractions due to false positive events: Binomial(N_i, α), where N_i is the number of neurons in i-th session, and α is the false positive rate:

$$\alpha = \begin{cases} \alpha_{MW} \times \alpha_{CCMW} = 0.001 \times 0.05 \,, & \text{for stimulus, choice, feedback} \\ \alpha_{MW} = 0.001 \,, & \text{for block} \end{cases} \tag{4}$$

We defined the p-value p_i as the probability of the fraction f_i that is larger than the distribution Binomial(N_i, α). Next, we used Fisher's method to compute a combined p-value of each brain region by combining the p-

values of all sessions ($i = 1, 2, ...m$):

$$\text{combined p-value} : \chi^2_{2m} = -2\sum_{i}^{m} \log(p_i) \ , \tag{5}$$

where χ^2_{2m} has a chi-squared distribution with $2m$ degrees of freedom. After computing combined p-values of each brain region, these p-values are then entered into the false discovery rate (FDR) procedure (Benjamini–Hochberg procedure) at level $q = 0.01$ to correct for multiple comparisons. We declare a list of regions to be significant based on this FDR procedure.

Multiple linear regression model of single neuron activity We fit linear regression models to single neuron activity, measured as spikes binned into 20 ms intervals. These models aim to express $\{s_{lt}\}$, the neural activity in time bin $t \in [1, T]$ on trial $l \in [1, L]$ based on D time-varying task-related regressors $X \in \mathbb{R}^{L,T,D}$. We first represented the regressors across time using a basis of raised cosine "bump" functions in log space[136]. Each basis function was associated with a weight in the regression model, with the value of the basis function at time t described by $\cos(\frac{2(t-\tau)\pi}{2w} + \frac{1}{2})$. The basis functions were computed in log space and then mapped into linear time to more efficiently capture both fast neuronal responses in the $< 100\,\text{ms}$ range and slow changes beyond that time (Fig. 4c). The width w and centre τ of each basis were chosen to ensure even coverage of the total duration of the kernel. In an example kernel with 3 bases, 3 separate weights would be fit to the event in question with weights describing early, middle, and late activity predicted by the event. These bases were convolved with a vector describing the effects of each regressor. In the case of timing events, the bases were convolved with a Kronecker delta function, resulting in a copy of the kernel at each time where the event occurred. We describe the simple case that each regressor enjoys the same number B of basis functions. This produced a new regression tensor $\hat{X} \in \mathbb{R}^{L,T,D,B}$

We then sought regression weights $\boldsymbol{\beta} \in \mathbb{R}^{D,B}$ such that, as closely as possible, $s_{lt} = \beta_0 + \sum_{d,b} \beta_{db} \hat{x}_{ltdb}$, where $\{\beta_{db}\}$ are linear regression weights (Fig. S4c). Each single-neuron model used regressors for stimulus onset (left and right separately), first movement onset (L/R), correct feedback, incorrect feedback, value of the block probability, movement initiation, and wheel speed. Fitting was performed using an L2-penalised objective function (as implemented in the scikit-learn python ecosystem as $||\mathbf{s} - \boldsymbol{\beta}_0 - \hat{X} \cdot \boldsymbol{\beta}||_2^2 + \alpha * ||\boldsymbol{\beta}||_2^2$), with the weight of the regularization α determined via cross-validation. Note that the intercept of the model is not included in the regularization, in order to capture fully the mean of the distribution of $\mathbf{s}$.

We used a kernel composed of 5 basis functions to parameterise left and right stimulus onset, and correct and incorrect feedback. These bases spanned 400 ms, and corresponded to 5 weights per regressor for each of these 4 regressors in the model.

Previous work has shown that difficulty in perceptual decision-making tasks[137], along with neural responses, do not change linearly with contrast. To account for this we modulated the height of the stimulus onset kernels as a function of contrast c with height $h = \frac{\tanh 5c}{\tanh 5}$. The resulting kernels would produce a response that was lower at low contrasts for the same set of weights $\{\beta_{db}\}$.

To capture statistical dependencies between wheel movements and spiking, we used anti-causal ker-

nels (in which the convolution of signal and kernel produces a kernel peak before peaks in the signal) describing the effect of movement onset for leftward and rightward movements. These kernels described 200 ms of activity preceding first movement using 3 basis functions. We also used an additional anti-causal kernel of 3 bases covering 300 ms describing the effect of wheel speed, and was convolved with the trace of wheel speed for each trial. With these regressors we aimed to capture preparatory signals preceding movements related to the wheel.

Finally, two terms describing the modulation of firing by the block probability were used. Neither term used the basis function parameterization, instead using a scalar value. One term (a single slice of the regression tensor X) which was the value of the prior (0.8, 0.5 or 0.2 depending on the block) in the previous trial from 400 ms pre-stimulus to 100 ms pre-stimulus, and another which was the same value from stimulus onset until 600 ms post-feedback. The latter term changed value from the previous trial's value of the prior to the current trial at feedback time.

Models were fit on a per-neuron basis with a ridge objective function using 5-fold cross-validation. Trials for cross-validation were chosen from a uniform distribution, and not in contiguous blocks. Models were then fit again using a leave-one-out paradigm, with each set of regressor weights $\beta_{d1} \dots \beta_{dB}$ being removed as a group and the resulting model fit and scored again on the same folds. The change between the base model score R^2_{full} and the omission model $R^2_{-regressor}$ was computed as $\Delta R^2_{regressor} = R^2_{full} - R^2_{-regressor}$. Additionally, the sensitivity for several pairs of associated regressors, such as left/right stimulus onset and correct/incorrect feedback, were defined as $log|\Delta R^2_A - \Delta R^2_B|$. This computation was applied to: right/left stimulus, right/left movement onset, and correct/incorrect feedback.

Decoding

Overview We performed a decoding analysis to measure how much information the activity of populations of neurons contained about task variables such as stimulus side and choice. To do this we used cross-validated, maximum likelihood regression with L1 regularization. The neural regressors are defined by binning the spike counts from each neuron in each session in a given region within a specific time window on each trial. The duration of the time window, the number of bins in that time window (i.e. bin size), and the trial event to which it is aligned depend on the variable that is the target of our regression (Table 2). These are discussed further below and include a variety of behavioural and task variables: stimulus side, block side, choice, feedback, and wheel speed/velocity. Although a session may include multiple probe insertions, we did not perform decoding on these probes separately because they are not independent. Instead, neurons in the same session and region were combined across probes for our decoding analysis. Decoding was cross validated (see Sec.) and compared to a null distribution for significance (see Sec.). A given region may contain multiple sessions, and thus in the main figures (Fig. 5, Fig. 6, Fig. 7, Fig. 8, and Fig. 9) the region p-value is defined by combining session p-values (see Sec.), and the region effect sizes are defined by subtracting the median of the null distribution from the decoding score and reporting the median of the resulting values across sessions. The p-values for all regions are then subjected to false discovery rate correction for multiple comparisons at level $q = 0.01$.

Decoded variable	Trial-relative time	Window start (ms)	Window end (ms)	Bin size (ms)	Regression type	Score	Null distribution
Block	Stim onset	-400	-100	300	Logistic	bal. acc.	pseudo-session
Stim side	Stim onset	0	100	100	Logistic	bal. acc.	pseudo-session
Wheel speed	Movement onset	-200	1000	20	Linear	R^2	imposter
Wheel velocity	Movement onset	-200	1000	20	Linear	R^2	imposter
Choice	Movement onset	-100	0	100	Logistic	bal. acc.	synthetic
Feedback	Feedback time	0	200	200	Logistic	bal. acc.	synthetic

Table 2. Details of decoding analysis. Decoded variables are the targets for regression. Spike sorted activity is summed within trial-relative bins and used as the regressors (Fig. S4d). Regression is performed with the specified regression type using L1 regularization and the described cross-validation scheme (see text). Performance is reported on held-out trials using the specified score, and this score is compared to a null distribution of the form listed to evaluate statistical significance (Section).

An example of our procedure (Fig. S4d) is as follows: to decode the mouse's choice from neurons within a given region from a single session, we first sum the spikes from each included trial across a single bin spanning 100ms before movement onset to movement onset, separately for each neuron. If there are N neurons and L trials, this binning procedure results in a matrix of size $N \times L$. We then construct the binary target vector that is 1 to indicate a left choice and 0 to indicate a right choice. This results in a binary vector of length L. Finally, we perform L1-regularized logistic regression in order to predict the choice from the binned spike count data.

Target variables Stimulus side, choice, feedback, and block side are treated as binary target variables for logistic regression. For stimulus side, trials which have zero-contrast are excluded. To make the bias block compatible with binary classification, we exclude the first 90 trials of the sessions which are unbiased. We use the `LogisticRegression` module from scikit-learn[138] (version 1.1.2) with 0.001 tolerance, 1000 maximum iterations, "l1" penalty, "liblinear" solver, and "fit_intercept" set to True. We balanced decoder classes by weighting samples by the inverse of the class frequency, $1/(2p_{i,class})$. Decoding performance is evaluated using the balanced accuracy of classification, which is the average of the recall probabilities for the two classes. Fig. S19 shows histograms of the regression coefficients for all the variables.

Wheel values (speed and velocity) change over the course of a trial, unlike the previous decoding targets, and we must therefore treat this target variable differently. We average wheel values in non-overlapping 20ms bins, starting 200ms before movement onset and ending at 1000ms after movement onset. Spike counts are binned similarly. The target value for a given bin (ending at time t) is decoded from spikes in a preceding (causal) window spanning W bins (ending at times t, . . . , t-W). Therefore, if decoding from N neurons, there are $(W + 1)N$ predictors of the target variable in a given bin. In practice we use $W = 10$. To decode these continuous-valued targets we perform linear regression using the `Lasso` module from scikit-learn[138] (version 1.1.2) with 0.001 tolerance, 1000 maximum iterations, and "fit_intercept" set to True. Decoding performance is evaluated using the R^2 metric.

Cross validation We performed all decoding using nested cross-validation. Each of five "outer" folds is based on a training/validation set comprising 80% of the trials and a test set of the remaining 20% of trials. We selected trials at random (in an "interleaved" manner). The training/validation set of an outer fold is itself split into five "inner" folds, again using an interleaved 80%/20% partition. When logistic regression is performed (two classes), the folds must be selected such that the trials used to train the decoder include at least 1 example of each class. Because both outer and inner folds are selected at random, it is possible that this requirement is not met. Therefore, when there is an insufficient number of class examples to complete training, we re-sample the outer or inner folds. Likewise, we disallow null sessions which have an insufficient number of class examples. We trained regression models on the 80% training set of the inner fold using various regularization coefficients: $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^{0}, 10^{1}\}$ for Logistic regression (denoted as input parameter C by sklearn) and $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^{0}, 10^{1}\}$ for Linear regression (denoted as input parameter α by sklearn). We then used each model to predict targets on the remaining 20% of the trials of the inner fold, i.e. the validation set. We repeated this procedure such that each trial in the original training/validation set of the outer fold is used once for the validation set and 4 times for the train set. We then take the regularization coefficient that performed best across all validation folds and retrained a regression model using all trials in the training/validation set of the outer fold. This final model is used to predict the target variable on the 20% of trials in the test set of the outer fold. We repeated the above training/validate/test procedure 5 times, each time holding out a different 20% of test trials such that, after the 5 repetitions, each trial has been included in the test set exactly once, and included in the training/validation set exactly 4 times. The concatenation of all test set predictions, covering 100% of the trials, is used to evaluate the decoding score.

We found for some regions and sessions, the resulting decoding score was sensitive to the precise assignment of trials to different folds. Therefore, to provide additional robustness to this procedure, we repeated the full five-fold cross-validation for multiple separate runs, each of which used a different random seed for selecting the interleaved training/validation/test splits. We then took the average decoding score across all runs as the final reported decoding score. When decoding block, stim side, choice, and feedback we perform ten runs, and while decoding wheel speed and wheel velocity we used two runs due to the added computational burden of decoding the wheel values which include multiple bins per trial.

Significance testing with null distributions We assessed the significance of the decoding score resulting from the multi-run cross-validation procedure described above by comparing it to those of a bespoke null distribution of decoding scores. To construct appropriate null distributions, we fixed the regressor matrices of neural activity and generated new vectors of target values that follow similar statistics. The ways in which these new vectors of target values were constructed depends on the target variable (see Table 2), as described in Section . Stimulus side and block are both behaviour-independent variables controlled by the experimenter, and thus we used the pseudo-session method. Choice and reward are behavioural variables that are not under the control of the experimenter, and thus we used the synthetic method. Finally, wheel speed and velocity are behavioural variables for which we lack a good generative model, and therefore we used the imposter session method. Once the new target values were generated, we carried out the full multi-run cross validation procedure described above to get a new decoding score. This was repeated multiple times to produce a null distribution of decoding scores: stimulus side, choice, and feedback are repeated 200

times; block is repeated 1000 times; and wheel speed and wheel velocity are repeated 100 times because wheel values impose a greater computational burden.

The null distribution is used to define a p-value for each region-session pair, where the p-value is defined as $1 - \rho$ where ρ is the percentile relative to the null distribution. Each brain region is recorded in ≥ 2 sessions, and we employed two different methods for summarizing the decoding scores across sessions: a) the median corrected decoding score amongst sessions which is used as the effect size in the main figures; the values have been corrected by subtracting the median of the decoding score of the null distribution; b) the fraction of sessions in which decoding is significant, that is if the p-value is less than $\alpha = 0.05$, which is shown in the supplement. We combined session-wide p-values using Fisher's method (also known as Fisher's combined probability test [31,32]) when computing a single statistic for a region. Finally, the combined p-value for a region is subjected to a false discovery rate correction for multiple comparisons at level $q = 0.01$. We note the combined p-value may be significant, but the computed effect size may be negative. This is because many sessions used for decoding in that region may be insignificant driving the effect size down, while a small number of sessions may be significant causing Fisher's method to produce a significant combined p-value.

Population manifold analysis methods We examine how responsive different brain regions are to a task variable v of interest. To do so, we construct a pair of variable-specific "supersessions" (s_v, s'_v): We partition all the IBL data into two, corresponding to the opposing pair of conditions for the variable (e.g. for stimulus discrimination, we split the trials into the L and R stimulus conditions) and replace the trial-by-trial responses of each cell within the condition and within each session with one trial-averaged response (Fig. 4e). These trials are aligned to a variable-specific reference time (e.g., stimulus onset time for stimulus discrimination). We used the canonical time windows shown in Fig. 4a around the alignment time for the main figures unless stated otherwise (e.g. for feedback we use a longer time window in the temporal evolution plot to illustrate licking), time bins of length $12.5\,\mathrm{ms}$, and stride $2\,\mathrm{ms}$. The supersessions S_v, S'_v have a number of rows equalling the number of IBL sessions passing quality control for that variable condition times the number of cells per session; columns correspond to time-bins.

We then subdivide the supersessions by brain region r $(S_{v,r}, S'_{v,r})$. These define a pair of across-IBL response trajectories (temporal evolution of the response) to the pair of variable v conditions for each brain region.

We next compute the time-resolved difference in response of brain region r to the opposing conditions of task variable v. We restrict our analyses to regions with ≥ 20 rows in $(S_{v,r}, S'_{v,r})$ for all analyses if not stated to use the canonical set of cells, same across all analyses. Our primary distance metric, which we call $d_{v,r}(t)$, is computed as a simple Euclidean distance in neural space.

Given a time-resolved distance curve, we compute the maximum and minimum distances along the curve to define a variable- and region-specific modulation amplitude:

$$A_{v,r} = \max_t [d_{v,r}(t)] - \min_t [d_{v,r}(t)]. \tag{6}$$

We obtain a variable- and region-specific response latency by defining it as the first time t at which $d_{v,r}(t) = \min_t[d_{v,r}(t)] + 0.7(\max_t[d_{v,r}(t)] - \min_t[d_{v,r}(t)])$. Using modulation amplitude as a measure of effect size, we then quantified the combined modulation amplitude and latency of regions as a function of task variable.

To generate a significance measure for the variable- and region-specific distance measures, we used a pseudo-trial method for generating null distance distributions, as described below. Distances were significant if they were greater in size than the corresponding null distance distribution with $p < 0.01$. Although the statistical significance of regions is therefore controlled for the effects of other task variables, note that the distance amplitudes and latencies are not.

Below we list the 4 task variables examined and the associated null distributions:

- "Stimulus" supersession: S_v, S'_v correspond to trials with the stimulus on the left or right, respectively, aligned by stimulus onset time and including 0 ms before to 150 ms after onset. To generate pseudo-trials, we permuted the stimulus side labels among trials that share the same block and choice side.
- "Choice" supersession: $S_v, S'_{v'}$ correspond to trials with the animal's response (wheel movement) to the left or right, respectively, aligning by wheel movement onset and including 0 ms before to 150 ms after onset. To generate pseudo-trials, we permuted the choice labels among trials with the same block and stimulus side.
- "Feedback" supersession: S_v, S'_v correspond to trials where the animal's response was correct (recall that the feedback was water delivery) or incorrect (recall that the feedback was tone and timeout delivery), respectively, aligning by feedback onset including 0 ms before to 150 ms after onset. To generate pseudo-trials, we permuted the choice labels among trials with the same block and stimulus side and then compared these pseudo-choices with the true stimulus sides to obtain pseudo-feedback types.
- "Block supersession": S_v, S'_v correspond to trials in left or right blocks, respectively, aligning by stimulus onset from 400 ms to 100 ms before stimulus onset. To generate pseudo-trials, we permuted the block side labels using the pseudo-session method, i.e. the same random function that was initially used to create block structure.

For each $S_{v,r}, S'_{v,r}$ pair we repeated the pseudo-trial process M (=1000) times, then followed the same distance computation procedures described above to obtain a null distribution of M modulation amplitude scores. We obtain a p-value by counting n (as the number of pseudo-scores that are greater than the true score for this region) like this: $p = \frac{n+1}{M+1}$.

For regions with significant and large effect sizes to a given variable, we generated visualizations of the population dynamics by projecting the trajectories in $S_{v,r}, S'_{v,r}$ into a low-dimensional subspace defined by the first three principal components of the pair $S_{v,r}, S'_{v,r}$. In addition to the main figure results, manifold results on the maximal dataset are shown in Fig. S12.

Visualization and comparison of results across neural analyses. To facilitate the comparison of neural analyses across brain regions, for each task variable we visualised effect sizes in a table (e.g. Fig. 5f) specifying the effect size for each analysis and brain region. Cells of the table were coloured according to effect size using the same colour map as in the corresponding flatmap. We then sorted brain regions within each Cosmos region according to the sum of (significant) effects across analyses. Before summing, the effect sizes for each analysis were normalised to lie in the interval $[0, 1]$. This method highlights regions with large effects across all analyses, and indicates the extent to which the analyses agree.

Acknowledgements

This work was supported by grants from the Wellcome Trust (216324), the Simons Foundation, The National Institutes of Health (NIH U19NS12371601), the National Science Foundation (NSF 1707398), the Gatsby Charitable Foundation (GAT3708), and by the Max Planck Society and the Humboldt Foundation. Part of the data analysis for this project was performed on Stanford University's Sherlock cluster; we would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results. Another part was performed at University of Geneva on "Baobab" and "Yggdrasil" high performance computing clusters. We thank Peter Latham, Tom Mrsic-Flogel and IBL colleagues for helpful comments on the manuscript. The production of all IBL Platform Papers is led by a Task Force, which defines the scope and composition of the paper, assigns and/or performs the required work for the paper, and ensures that the paper is completed in a timely fashion. The Task Force members for this platform paper include authors AP, BG, BB, CF, CL, DB, FH, GAC, IRF, JMH, KDH, KZS, MRW, MC, MS, NJM, NAS, OW, PD, TAE, YS.

Competing interests

The authors declare no competing interests.

References

1. Broca, P. Remarques sur le siège de la faculté du langage articulé, suivies d'une observation d'aphémie (perte de la parole). *Bulletin et Memoires de la Societe anatomique de Paris* **6**, 330–357 (1861).

2. Lashley, K. S. Brain mechanisms and intelligence: A quantitative study of injuries to the brain. (1929).

3. Tizard, B. Theories of brain localization from flourens to lashley. *Medical history* **3**, 132–145 (1959).

4. Alivisatos, A. P. *et al.* The brain activity map project and the challenge of functional connectomics. *Neuron* **74**, 970–974 (2012).

5. Kato, S. *et al.* Global brain dynamics embed the motor command sequence of caenorhabditis elegans. *Cell* **163**, 656–669 (2015).

6. Wang, Q. *et al.* The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **181**, 936–953.e20 (2020).

7. Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of neurophysiology* **86**, 1916 – 1936 (2001). URL https://www.physiology.org/doi/10.1152/jn.2001.86.4.1916.

8. Katz, L. N., Yates, J. L., Pillow, J. W. & Huk, A. C. Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* **535**, 285–288 (2016).

9. Jeurissen, D., Shushruth, S., El-Shamayleh, Y., Horwitz, G. D. & Shadlen, M. N. Deficits in decision-making induced by parietal cortex inactivation are compensated at two timescales. *Neuron* **110**, 1924–1931 (2022).

10. Yao, J. D., Gimoto, J., Constantinople, C. M. & Sanes, D. H. Parietal cortex is required for the integration of acoustic evidence. *Current Biology* **30**, 3293–3303 (2020).

11. Pisupati, S., Chartarifsky, L. & Churchland, A. K. Decision activity in parietal cortex–leader or follower? *Trends in Cognitive Sciences* **20**, 788–789 (2016).

12. Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D. & Brody, C. D. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* **4**, e05457 (2015).

13. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).

14. Siegle, J. H. *et al.* Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).

15. Allen, W. E. *et al.* Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* **364**, eaav3932 (2019).

16. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).

17. Chen, S. *et al.* Brain-wide neural activity underlying memory-guided movement. *bioRxiv* 2023–03 (2023).

18. Hsueh, B. *et al.* Cardiogenic control of affective behavioural state. *Nature* **615**, 292–299 (2023).

19. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, eaav7893 (2019).

20. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677 – 1686 (2019). URL http://www.nature.com/articles/s41593-019-0502-4.

21. de Vries, S. E. *et al.* A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature neuroscience* **23**, 138–151 (2020).

22. The International Brain Laboratory. Data release - Brainwide map - Q4 2022 (2023). URL https://figshare.com/articles/preprint/Data_release_-_Brainwide_map_-_Q4_2022/21400815.

23. The International Brain Laboratory *et al.* Standardized and reproducible measurement of decision-making in mice. *eLife* **10**, e63711 (2021).

24. The International Brain Laboratory *et al.* Reproducibility of in-vivo electrophysiological measurements in mice. *bioRxiv* 2022–05 (2022).

25. The International Brain Laboratory *et al.* Data architecture for a large-scale neuroscience collaboration. *Nature Methods, in press* (2020). URL https://www.biorxiv.org/content/early/2020/02/23/827873. https://www.biorxiv.org/content/early/2020/02/23/827873.full.pdf.

26. 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 Lee, D., Sugiyama, M., Luxburg, U., Guyon, I. & Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29 (Curran Associates, Inc., 2016). URL https://proceedings.neurips.cc/paper/2016/file/1145a30ff80745b56fb0cecf65305017-Paper.pdf.

27. The International Brain Laboratory *et al.* Spike sorting pipeline for the international brain laboratory. *figshare* (2022).

28. Ragan, T. *et al.* Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nature methods* **9**, 255–258 (2012).

29. Swanson, L. W. & Hahn, J. D. A qualitative solution with quantitative potential for the mouse hippocampal cortex flatmap problem. *Proceedings of the National Academy of Sciences* **117**, 3220–3231 (2020).

30. Mathis, A. *et al.* Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience* **21**, 1281–1289 (2018).

31. Fisher, R. A. *Statistical methods for research workers* (Oliver and Boyd, Edinburgh, 1925).

32. Fisher, R. Questions and answers # 14. *The American Statistician* **2**, 30–31 (1948).

33. Park, I. M., Meister, M. L., Huk, A. C. & Pillow, J. W. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature neuroscience* **17**, 1395–1403 (2014).

34. Wang, L., Sarnaik, R., Rangarajan, K., Liu, X. & Cang, J. Visual receptive field properties of neurons in the superficial superior colliculus of the mouse. *The Journal of Neuroscience* **30**, 16573–16584 (2010).

35. Drager, U. C. & Hubel, D. H. Physiology of visual cells in mouse superior colliculus and correlation with somatosensory and auditory input. *Nature* **253**, 203–204 (1975).

36. Grubb, M. S. & Thompson, I. D. Quantitative characterization of visual response properties in the mouse dorsal lateral geniculate nucleus. *Journal of Neurophysiology* **90**, 3594–3607 (2003).

37. Piscopo, D. M., El-Danaf, R. N., Huberman, A. D. & Niell, C. M. Diverse Visual Features Encoded in Mouse Lateral Geniculate Nucleus. *The Journal of Neuroscience* **33**, 4642–4656 (2013).

38. Roth, M. M. *et al.* Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nature neuroscience* **19**, 299–307 (2016).

39. Han, X., Vermaercke, B. & Bonin, V. Diversity of spatiotemporal coding reveals specialized visual processing streams in the mouse cortex. *Nature Communications* **13**, 3249 (2022).

40. Wang, Q. & Burkhalter, A. Area map of mouse visual cortex. *Journal of Comparative Neurology* **502**, 339–357 (2007).

41. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience* **28**, 7520–7536 (2008).

42. Glickfeld, L. L. & Olsen, S. R. Higher-order areas of the mouse visual cortex. *Annual review of vision science* **3**, 251–273 (2017).

43. Siegle, J. H. *et al.* Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).

44. Schmolesky, M. T. *et al.* Signal timing across the macaque visual system. *Journal of neurophysiology* **79**, 3272–3278 (1998).

45. Licata, A. M. *et al.* Posterior parietal cortex guides visual decisions in rats. *Journal of Neuroscience* **37**, 4954–4966 (2017).

46. Bichot, N. P., Schall, J. D. & Thompson, K. G. Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature* **381**, 697–699 (1996).

47. Zatka-Haas, P., Steinmetz, N. A., Carandini, M. & Harris, K. D. Sensory coding and the causal impact of mouse cortex in a visual decision. *Elife* **10**, e63163 (2021).

48. Wal, A., Klein, F. J., Born, G., Busse, L. & Katzner, S. Evaluating visual cues modulates their representation in mouse visual and cingulate cortex. *Journal of Neuroscience* **41**, 3531–3544 (2021).

49. Orsolic, I., Rio, M., Mrsic-Flogel, T. D. & Znamenskiy, P. Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making. *Neuron* **109**, 1861–1875 (2021).

50. Peters, A. J., Fabre, J. M., Steinmetz, N. A., Harris, K. D. & Carandini, M. Striatal activity topographically reflects cortical activity. *Nature* **591**, 420–425 (2021).

51. Yartsev, M. M., Hanks, T. D., Yoon, A. M. & Brody, C. D. Causal contribution and dynamical encoding in the striatum during evidence accumulation. *eLife* **7**, e34929 (2018).

52. Martersteck, E. M. *et al.* Diverse central projection patterns of retinal ganglion cells. *Cell reports* **18**, 2058–2072 (2017).

53. Ding, L. & Gold, J. I. Neural Correlates of Perceptual Decision Making before, during, and after Decision Commitment in Monkey Frontal Eye Field. *Cerebral Cortex* **22**, 1052 – 1067 (2012). URL http://cercor.oxfordjournals.org/cgi/doi/10.1093/cercor/bhr178.

54. Chandrasekaran, C., Peixoto, D., Newsome, W. T. & Shenoy, K. V. Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nature Communications* **8**, 996 – 16 (2017). URL http://www.nature.com/articles/s41467-017-00715-0.

55. Gold, J. I. & Shadlen, M. N. The Neural Basis of Decision Making. *Annual review of neuroscience* **30**, 535 – 574 (2007). URL http://www.annualreviews.org/doi/10.1146/annurev.neuro.29.051605.113038.

56. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science* **348**, 1352 – 1355 (2015). URL http://www.sciencemag.org/cgi/doi/10.1126/science.aab0551.

57. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784 – 1792 (2014). URL http://www.nature.com/doifinder/10.1038/nn.3865.

58. Hanks, T. D. *et al.* Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220 – 223 (2015). URL http://www.nature.com/doifinder/10.1038/nature14066.

59. Scott, B. B. *et al.* Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron* **95**, 385–398.e5 (2017). URL http://linkinghub.elsevier.com/retrieve/pii/S0896627317305111.

60. Guo, Z. V. *et al.* Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* **81**, 179 – 194 (2014). URL https://linkinghub.elsevier.com/retrieve/pii/S0896627313009240.

61. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459 – 464 (2016). URL http://www.nature.com/doifinder/10.1038/nature17643.

62. Nieh, E. H. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* 1 – 5 (2021). URL https://www.nature.com/articles/s41586-021-03652-7.

63. Ding, L. Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys. *Journal of neurophysiology* **114**, 1850 – 1861 (2015). URL http://jn.physiology.org/lookup/doi/10.1152/jn.00395.2015.

64. Duan, C. A. *et al.* Collicular circuits for flexible sensorimotor routing. *Nature Neuroscience* **24**, 1110–1120 (2021).

65. Jun, E. J. *et al.* Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. *Nature Neuroscience* **24**, 1121–1131 (2021).

66. Chen, S. *et al.* Brain-wide neural activity underlying memory-guided movement (2023).

67. Schultz, W. & Dickinson, A. Neuronal coding of prediction errors. *Annual review of neuroscience* **23**, 473–500 (2000).

68. Niv, Y. Reinforcement learning in the brain. *Journal of Mathematical Psychology* **53**, 139–154 (2009).

69. Kostadinov, D. & Häusser, M. Reward signals in the cerebellum: origins, targets, and functional implications. *Neuron* (2022).

70. Averbeck, B. & O'Doherty, J. P. Reinforcement-learning in fronto-striatal circuits. *Neuropsychopharmacology* **47**, 147–162 (2022).

71. Iordanova, M. D., Yau, J. O.-Y., McDannald, M. A. & Corbit, L. H. Neural substrates of appetitive and aversive prediction error. *Neuroscience & Biobehavioral Reviews* **123**, 337–351 (2021).

72. Montague, P. R., Dayan, P. & Sejnowski, T. J. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of neuroscience* **16**, 1936–1947 (1996).

73. Gutierrez, R., Carmena, J. M., Nicolelis, M. A. & Simon, S. Orbitofrontal ensemble activity monitors licking and distinguishes among natural rewards. *Journal of neurophysiology* **95**, 119–133 (2006).

74. Gutierrez, R., Simon, S. A. & Nicolelis, M. A. Licking-induced synchrony in the taste–reward circuit improves cue discrimination during learning. *Journal of Neuroscience* **30**, 287–303 (2010).

75. Amarante, L. M., Caetano, M. S. & Laubach, M. Medial frontal theta is entrained to rewarded actions. *Journal of Neuroscience* **37**, 10757–10769 (2017).

76. Karakaş, S. A review of theta oscillation and its functional correlates. *International Journal of Psychophysiology* **157**, 82–99 (2020).

77. Hansen, K. A., Hillenbrand, S. F. & Ungerleider, L. G. Human Brain Activity Predicts Individual Differences in Prior Knowledge Use during Decisions. *Journal of Cognitive Neuroscience* **24**, 1462–1475 (2012). URL https://doi.org/10.1162/jocn_a_00224. https://direct.mit.edu/jocn/article-pdf/24/6/1462/1944040/jocn_a_00224.pdf.

78. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience* **31**, 6339–6352 (2011). URL https://www.jneurosci.org/content/31/17/6339. https://www.jneurosci.org/content/31/17/6339.full.pdf.

79. Forstmann, B., Brown, S., Dutilh, G., Neumann, J. & Wagenmakers, E.-J. The neural substrate of prior information in perceptual decision making: a model-based analysis. *Frontiers in Human Neuroscience* **4** (2010). URL https://www.frontiersin.org/articles/10.3389/fnhum.2010.00040.

80. Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience* **32**, 2335–2343 (2012). URL https://www.jneurosci.org/content/32/7/2335. https://www.jneurosci.org/content/32/7/2335.full.pdf.

81. Niv, Y. Learning task-state representations. *Nature Neuroscience* **22**, 1544–1553 (2019).

82. Lee, D., Seo, H. & Jung, M. W. Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience* **35**, 287–308 (2012). URL https://doi.org/10.1146/annurev-neuro-062111-150512. PMID: 22462543, https://doi.org/10.1146/annurev-neuro-062111-150512.

83. Ito, M. & Doya, K. Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *Journal of Neuroscience* **29**, 9861–9874 (2009). URL https://www.jneurosci.org/content/29/31/9861. https://www.jneurosci.org/content/29/31/9861.full.pdf.

84. Shin, E. J. *et al.* Robust and distributed neural representation of action values. *eLife* **10**, e53045 (2021). URL https://doi.org/10.7554/eLife.53045.

85. Ottenheimer, D. J., Hjort, M. M., Bowen, A. J., Steinmetz, N. A. & Stuber, G. D. A stable, distributed code for cue value in mouse cortex during reward learning (2023). URL https://doi.org/10.7554/elife.84604.1.

86. Lak, A. *et al.* Dopaminergic and prefrontal basis of learning from sensory confidence and reward value. *Neuron* **105**, 700–711.e6 (2020). URL https://www.sciencedirect.com/science/article/pii/S0896627319309821.

87. Bari, B. A. *et al.* Stable representations of decision variables for flexible behavior. *Neuron* **103**, 922–933.e7 (2019).

88. Kuwabara, M., Kang, N., Holy, T. E. & Padoa-Schioppa, C. Neural mechanisms of economic choices in mice. *eLife* **9**, e49669 (2020). URL https://doi.org/10.7554/eLife.49669.

89. Nogueira, R. *et al.* Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nature Communications* **8**, 14823 (2017).

90. Moss, M. M., Zatka-Haas, P., Harris, K. D., Carandini, M. & Lak, A. Dopamine axons in dorsal striatum encode contralateral visual stimuli and choices. *Journal of Neuroscience* **41**, 7197–7205 (2021). URL https://www.jneurosci.org/content/41/34/7197. https://www.jneurosci.org/content/41/34/7197.full.pdf.

91. Shipp, S. The brain circuitry of attention. *Trends in Cognitive Sciences* **8**, 223–230 (2004).

92. Mongillo, G., Barak, O. & Tsodyks, M. Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).

93. Findling, C. *et al.* Brain-wide representations of prior information in mouse decision-making. *bioRxiv* (2023). URL https://www.biorxiv.org/content/early/2023/07/04/2023.07.04.547684.

94. Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).

95. Saleem, A. B., Ayaz, A., Jeffery, K. J., Harris, K. D. & Carandini, M. Integration of visual motion and locomotion in mouse visual cortex. *Nature neuroscience* **16**, 1864–1869 (2013).

96. Harris, K. D. A shift test for independence in generic time series (2020). URL https://arxiv.org/abs/2012.06862.

97. Straka, H., Simmers, J. & Chagnaud, B. P. A new perspective on predictive motor signaling. *Current Biology* **28**, R232–R243 (2018).

98. Blakemore, S.-J., Frith, C. D. & Wolpert, D. M. Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of cognitive neuroscience* **11**, 551–559 (1999).

99. Sokolov, A. A., Miall, R. C. & Ivry, R. B. The cerebellum: adaptive prediction for movement and cognition. *Trends in cognitive sciences* **21**, 313–332 (2017).

100. Shadmehr, R. & Ahmed, A. A. *Vigor: Neuroeconomics of movement control* (MIT Press, 2020).

101. Thorndike, E. L. *Animal intelligence: Experimental studies* (Transaction Publishers, 1911).

102. Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M. & Keller, P. J. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature methods* **10**, 413–420 (2013).

103. Lovett-Barron, M. Learning-dependent neuronal activity across the larval zebrafish brain. *Current opinion in neurobiology* **67**, 42–49 (2021).

104. Witten, I. B. *et al.* Recombinase-driver rat lines: tools, techniques, and optogenetic application to dopamine-mediated reinforcement. *Neuron* **72**, 721–733 (2011).

105. Kim, K. M. *et al.* Optogenetic mimicry of the transient activation of dopamine neurons by natural reward is sufficient for operant reinforcement. *PloS one* **7**, e33612 (2012).

106. Crapse, T. B. & Sommer, M. A. Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience* **9**, 587–600 (2008).

107. Wurtz, R. H. & Sommer, M. A. Identifying corollary discharges for movement in the primate brain. *Progress in brain research* **144**, 47–60 (2004).

108. Horwitz, G. D., Batista, A. P. & Newsome, W. T. Representation of an abstract perceptual decision in macaque superior colliculus. *Journal of neurophysiology* **91**, 2281–2296 (2004).

109. Felsen, G. & Mainen, Z. F. Midbrain contributions to sensorimotor decision making. *Journal of neurophysiology* **108**, 135–147 (2012).

110. Mohajerani, M. H., McVea, D. A., Fingas, M. & Murphy, T. H. Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *Journal of Neuroscience* **30**, 3745–3751 (2010).

111. Shimaoka, D., Steinmetz, N. A., Harris, K. D. & Carandini, M. The impact of bilateral ongoing activity on evoked responses in mouse cortex. *elife* **8**, e43533 (2019).

112. Lopes, G. *et al.* Bonsai: an event-based framework for processing and controlling data streams. *Frontiers in neuroinformatics* **9**, 7 (2015).

113. Lopes, G. *et al.* Creating and controlling visual environments using bonvision. *Elife* **10**, e65541 (2021).

114. Ragan, T. *et al.* Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nat. Methods* **9**, 255–8 (2012).

115. Economo, M. N. *et al.* A platform for brain-wide imaging and reconstruction of individual neurons. *eLife* **5** (2016).

116. Campbell, R. Bakingtray (2020). URL https://github.com/SainsburyWellcomeCentre/BakingTray.

117. Campbell, R. Stitchit (2021). URL https://github.com/SainsburyWellcomeCentre/StitchIt.

118. West, S. J. Brainregister (2021). URL https://github.com/stevenjwest/brainregister.

119. Klein, S., Staring, M., Murphy, K., Viergever, M. A. & Pluim, J. P. W. elastix: a toolbox for intensity based medical image registration. *IEEE Transactions on Medical Imaging* **29**, 196–205 (2010).

120. Campbell, R., Blot, A., Rousseau, C. & Winter, O. Lasagna (2020). URL https://github.com/SainsburyWellcomeCentre/lasagna.

121. Rossant, C. *et al.* Alyx (2021). URL https://github.com/cortex-lab/alyx.

122. Faulkner, M. Ephys atlas gui (2020). URL https://github.com/int-brain-lab/iblapps/tree/master/atlaselectrophysiology.

123. Liu, L. D. *et al.* Accurate localization of linear probe electrode arrays across multiple brains. *eNeuro* **8** (2021). https://www.biorxiv.org/content/10.1101/2020.02.25.965210v1.full.pdf.

124. Steinmetz, N. A. *et al.* Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**, eabf4588 (2021).

125. The International Brain Laboratory. pykilosort (2021). URL https://github.com/MouseLand/pykilosort.

126. The International Brain Laboratory *et al.* Video hardware and software for the international brain laboratory. *figshare* (2022).

127. The International Brain Laboratory. iblvideo (2021). URL https://github.com/int-brain-lab/iblvideo.

128. Harris, K. D. Nonsense correlations in neuroscience. *biorxiv preprint doi: https://doi.org/10.1101/2020.11.29.402719* (2021).

129. Elber-Dorozko, L. & Loewenstein, Y. Striatal action-value neurons reconsidered. *ELife* **7**, e34248 (2018).

130. Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* **57**, 289–300 (1995).

131. Britten, K., Shadlen, M., Newsome, W. & Movshon, J. The analysis of visual motion: a comparison of neuronal and psychophysical performance **12**, 4745–4765 (1992).

132. Nienborg, H., R. Cohen, M. & Cumming, B. G. Decision-related activity in sensory neurons: Correlations among neurons and with behavior. *Annual Review of Neuroscience* **35**, 463–483 (2012).

133. Hanley, J. A. & McNeil, B. J. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology* **143**, 29–36 (1982).

134. Mason, S. J. & Graham, N. E. Areas beneath the relative operating characteristics (roc) and relative operating levels (rol) curves: Statistical significance and interpretation. *Quarterly Journal of the Royal Meteorological Society* **128** (2002).

135. Harris, K. D. A shift test for independence in generic time series. *arXiv preprint arXiv:2012.06862* (2020).

136. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).

137. Roy, N. A., Bak, J. H., Akrami, A., Brody, C. D. & Pillow, J. W. Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron* **109**, 597–610.e6 (2021). URL https://www.sciencedirect.com/science/article/pii/S0896627320309636.

138. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).

Variable	Region	Analysis	Session eid	Neuron ID	Number of neurons
Stimulus	VISp	Encoding	e0928e11-2b86-4387-a203-80c77fab5d52	209	-
Stimulus	VISp	Decoding	5d01d14e-aced-4465-8f8e-9a1c674f62ec	-	18
Choice	GRN	Encoding	671c7ea7-6726-4fbe-adeb-f89c2c8e489b	123	-
Choice	GRN	Decoding	671c7ea7-6726-4fbe-adeb-f89c2c8e489b	-	23
Feedback	IRN	Encoding	a7763417-e0d6-4f2a-aa55-e382fd9b5fb8	98	-
Feedback	IRN	Decoding	e012d3e3-fdbc-4661-9ffa-5fa284e4e706	-	11
Block	MOp	Encoding	7bee9f09-a238-42cf-b499-f51f765c6ded	207	-
Block	MOp	Decoding	9e9c6fc0-4769-4d83-9ea4-b59a1230510e	-	30
Wheel speed/velocity	GRN	Decoding	671c7ea7-6726-4fbe-adeb-f89c2c8e489b	-	23
Wheel speed/velocity	All	Encoding	-	-	32,304

Table 3. Session and neurons used for example encoding and decoding analysis. For the stimulus (Fig. 5), choice (Fig. 6), feedback (Fig. 7), block (Fig. 8), and wheel speed/velocity (Fig. 9) we included figure panels showing example encoding and decoding analyses. Here, we list the variable name, region, and analysis type along with the session eid, encoding neuron ID number, and decoding number of neurons.

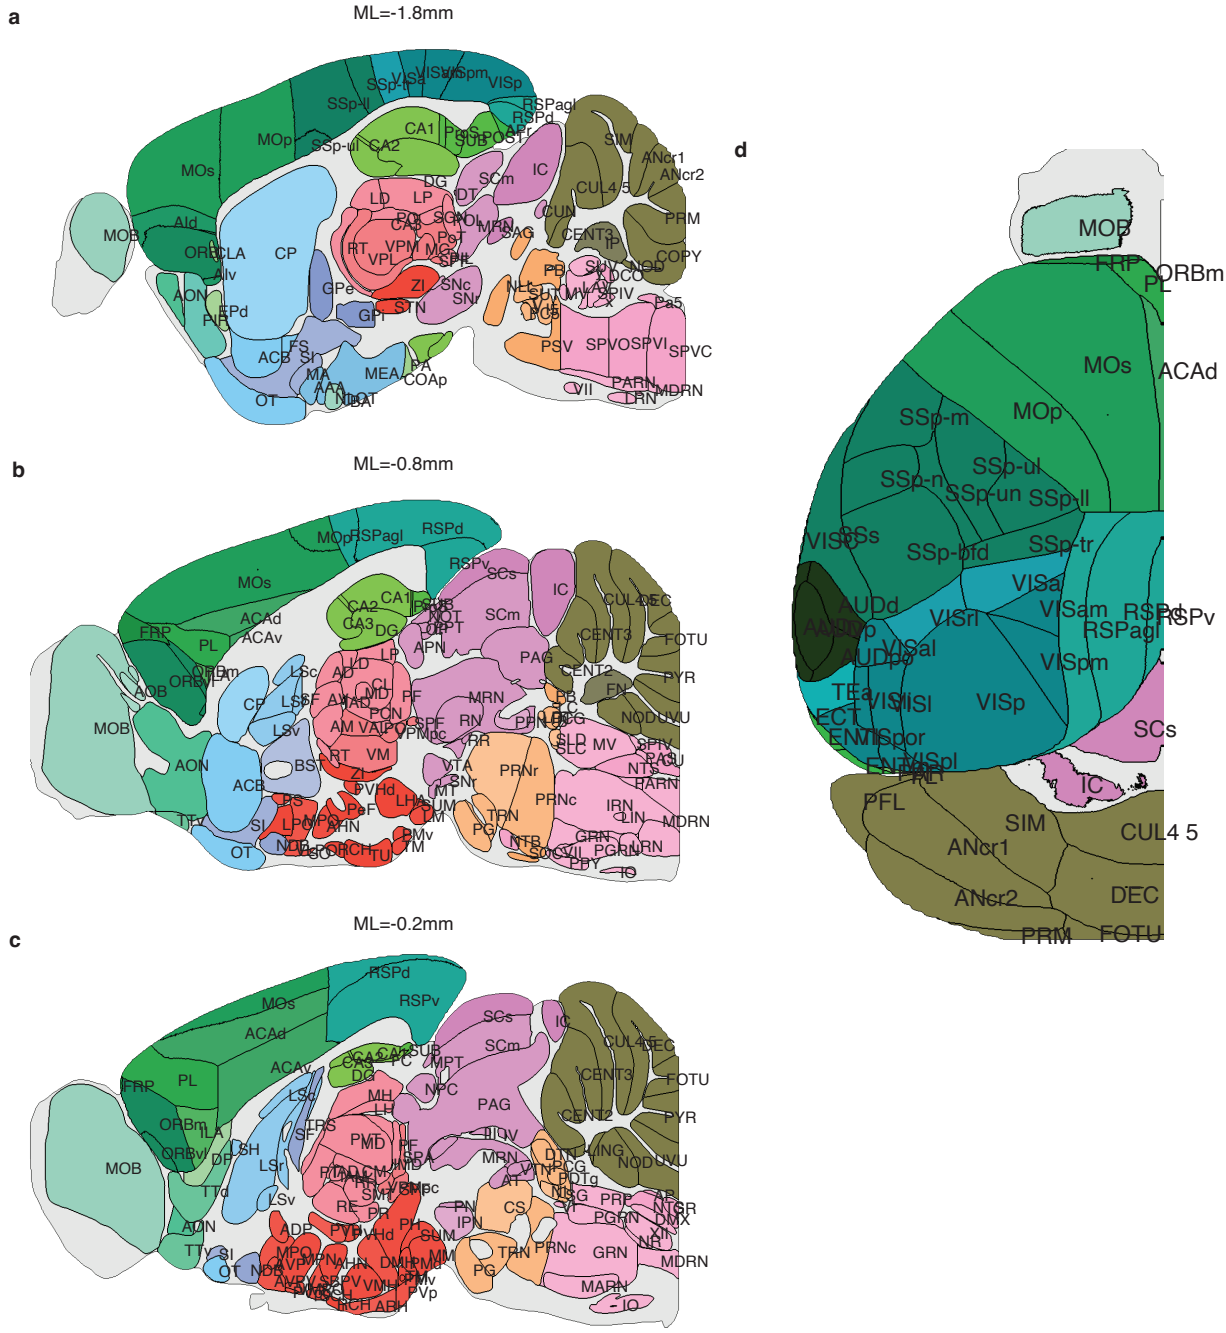

Figure S1. 2d-brain slices maps annotated with region acronyms. **a)** Region acronyms for sagittal slices with coordinates: ML=-1.8mm, **b)** ML=-0.8mm, **c)** ML=-0.2mm. **d)** Region acronyms for the top view of the dorsal cortex.

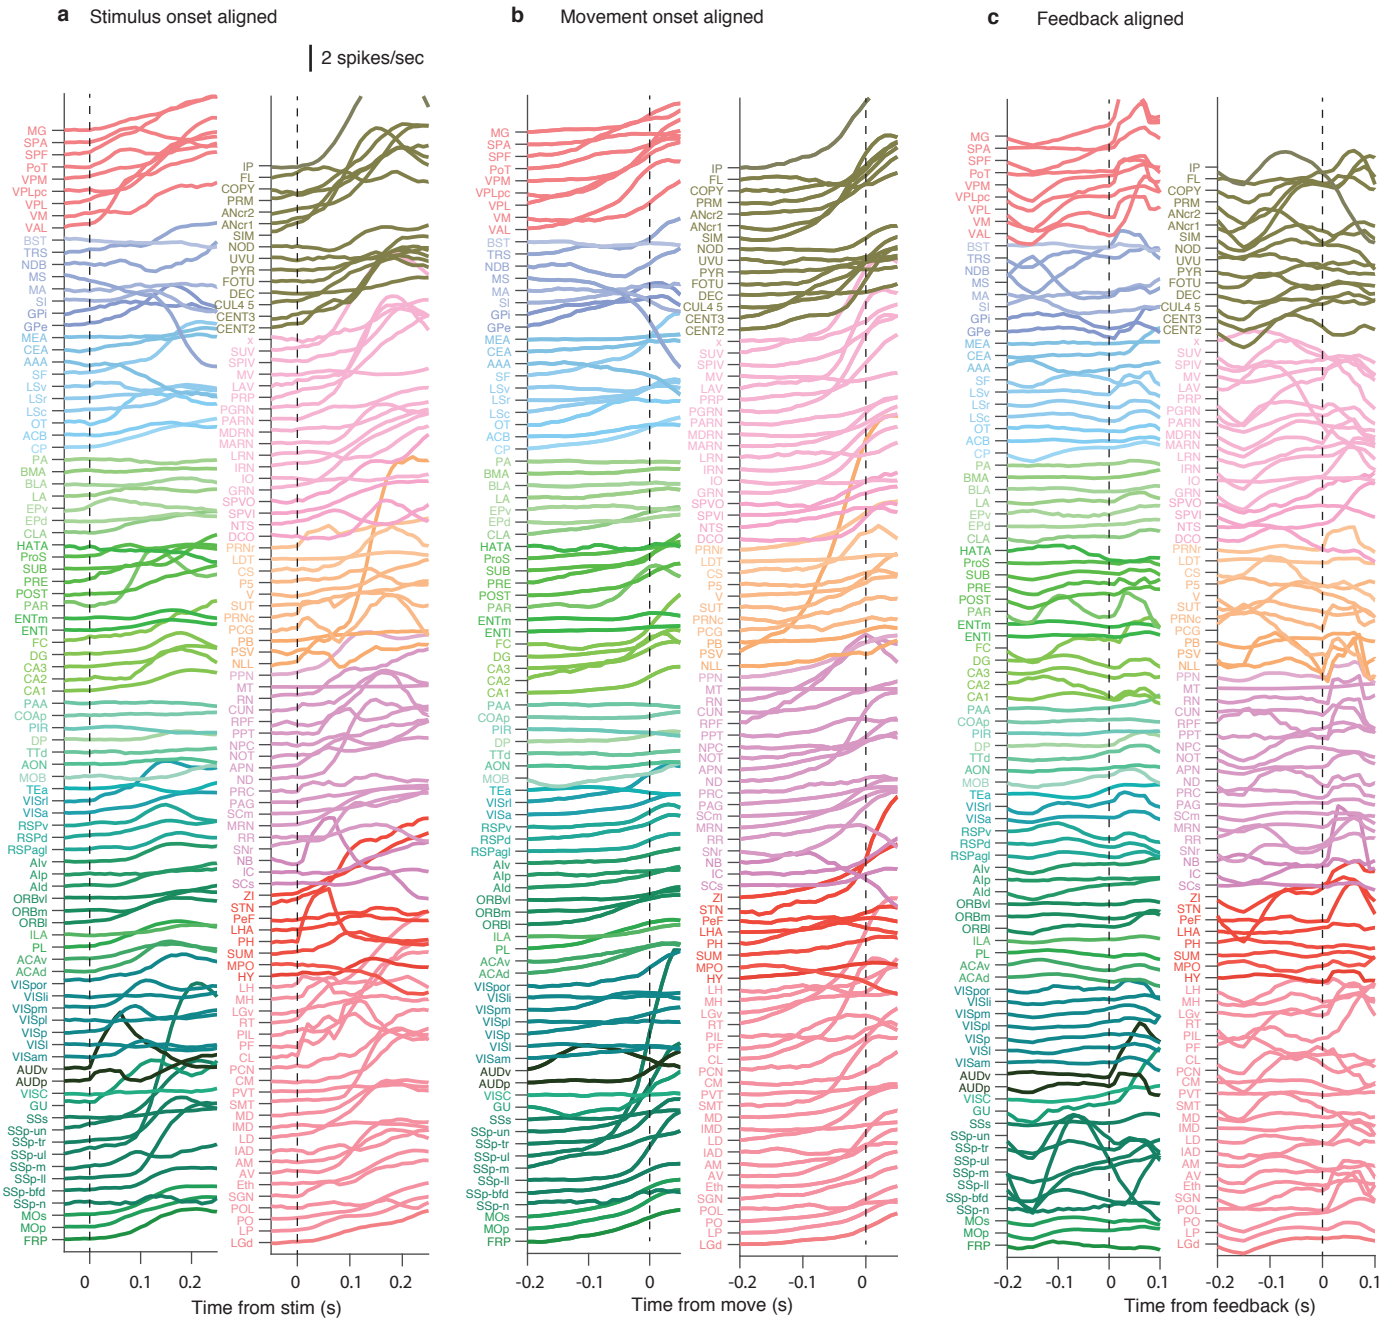

Figure S2. PSTH of average neural activity across the brain. a) PSTH of average neural activity aligned to stimulus onset, **b)** Movement onset, and **c)** Feedback onset.

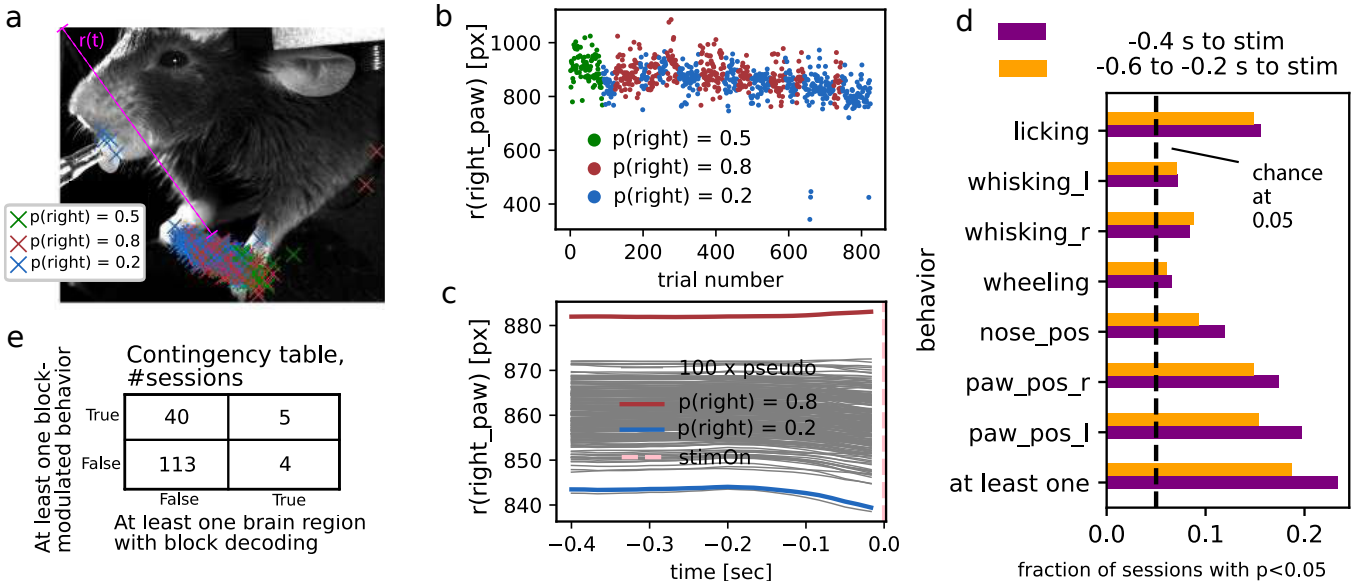

	False	True
True	40	5
False	113	4

Figure S3. Motor correlates of block. **a)** Example frame of 'left' video with DLC-tracked left paw position in the pre-stimulus interval for all trials of this session, coloured by block identity of the trial. $r(t)$ indicates the distance of the left paw to the upper left corner of the frame, here averaged across $[-400, 0]$ms pre-stimulus interval. **b)** For each trial, the average pre-stimulus position of the left paw is shown, displaying some tendency to be higher in right block trials (p(right) = 0.8, red dots), along with a trend to be lower later in the session. **c)** Trial averaged $r(t)$, splitting trials by block type (blue/red) or into a pseudo-block (null-distribution, in grey). The left paw is on average closer to the upper left corner of the frame for trials with p(left) = 0.8 (blue line). **d)** The prevalence of block-modulated behaviours across all sessions is shown for paw position and other similar behaviours in the pre-stimulus interval, such as licking, whisking or the nose position. For both pre-stimulus windows shown, approximately 20% of sessions has at least one of these behaviours being block-modulated, see bar "at least one". To test for the statistical significance of this correlation, we used a "pseudo-session method", comparing the difference in mean paw position between left and right blocks to a null ensemble generated by re-randomizing the trial's block structure using the same distribution that it was generated from (panel c). **e)** Contingency table showing counts of sessions with at least one brain region from which the block prior could be decoded, respectively where there was at least one block-modulated behaviour. A Barnard exact test applied to this table results in $p = 0.06$ and statistic $s = 1.9$, i.e. no strong evidence for a connection between motor correlates of block and decoding of block from neural activity.

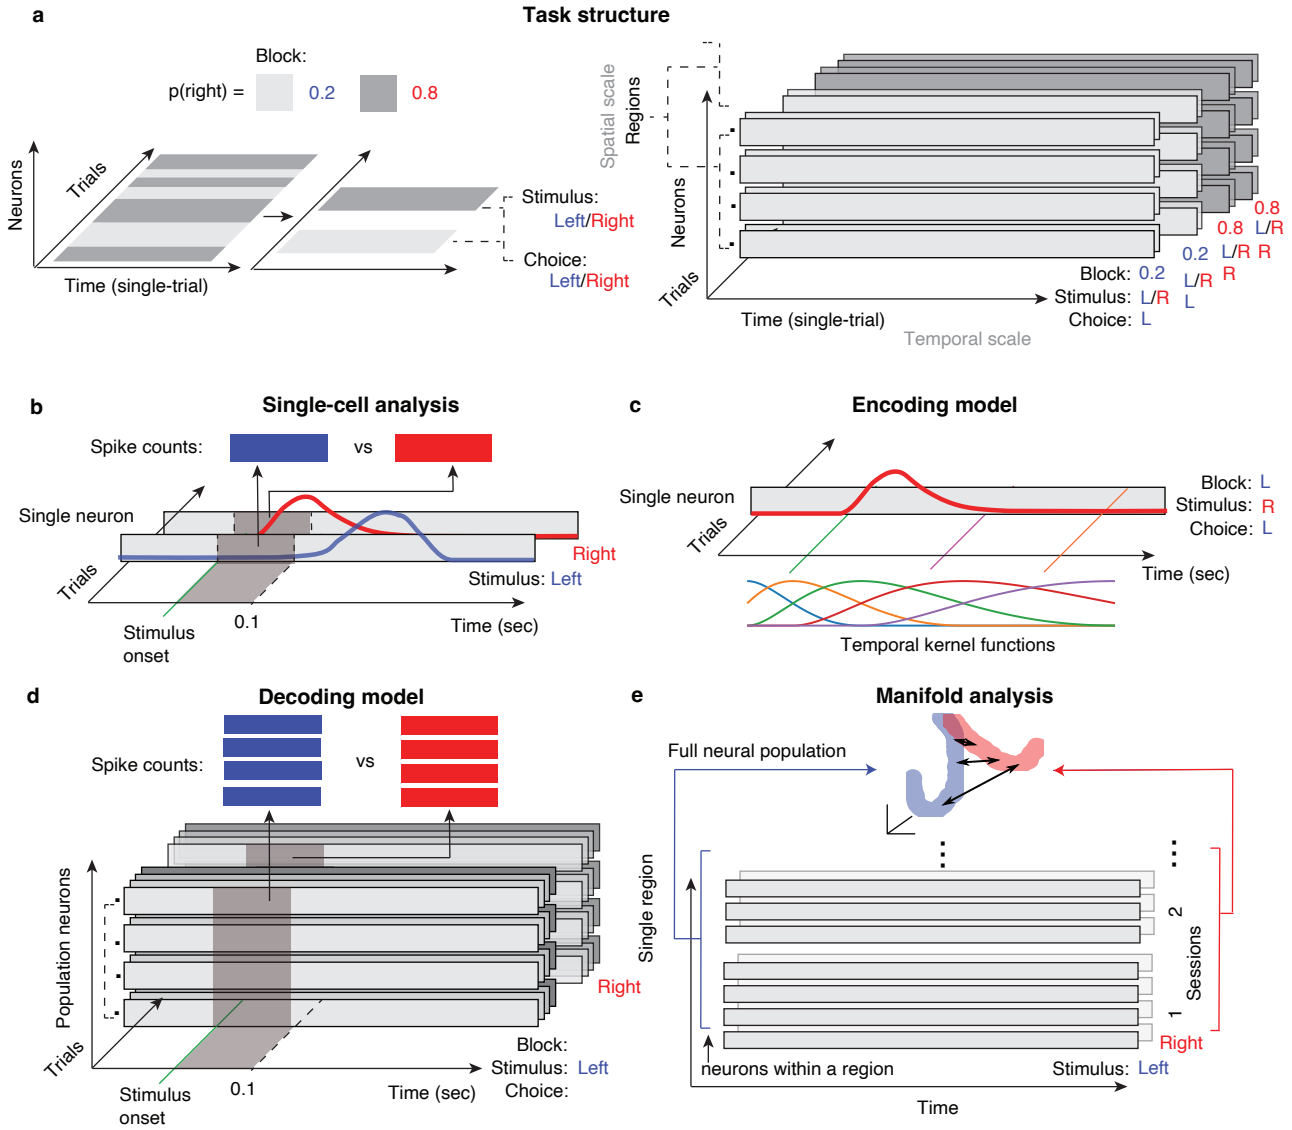

Figure S4. Detail of spatiotemporal structure of neural analyses. a) Task structures. In each session, consecutive trials form a block structure with the probability of a right-side stimulus being 0.8 and 0.2. In each block, there are trials with stimulus and choice side that are left or right. By regrouping trials, we can obtain 8 categories of trials with different combinations of stimulus side, block identify, and choice side. **b)** Single-cell analysis studies the modulation of single-neuron activity by individual task variables within a short time window. **c)** The encoding model uses temporal kernel functions to describe single-neuron activity during the entire trial at high temporal resolution. **d)** The decoding model studies the modulation of population neural activity by individual task variables (without marginalizing other variables) within a short time window. **e)** Manifold analysis combines neural responses across multiple sessions (each neuron averaged across trials within a session) and analyses the trajectory of population neural activity.

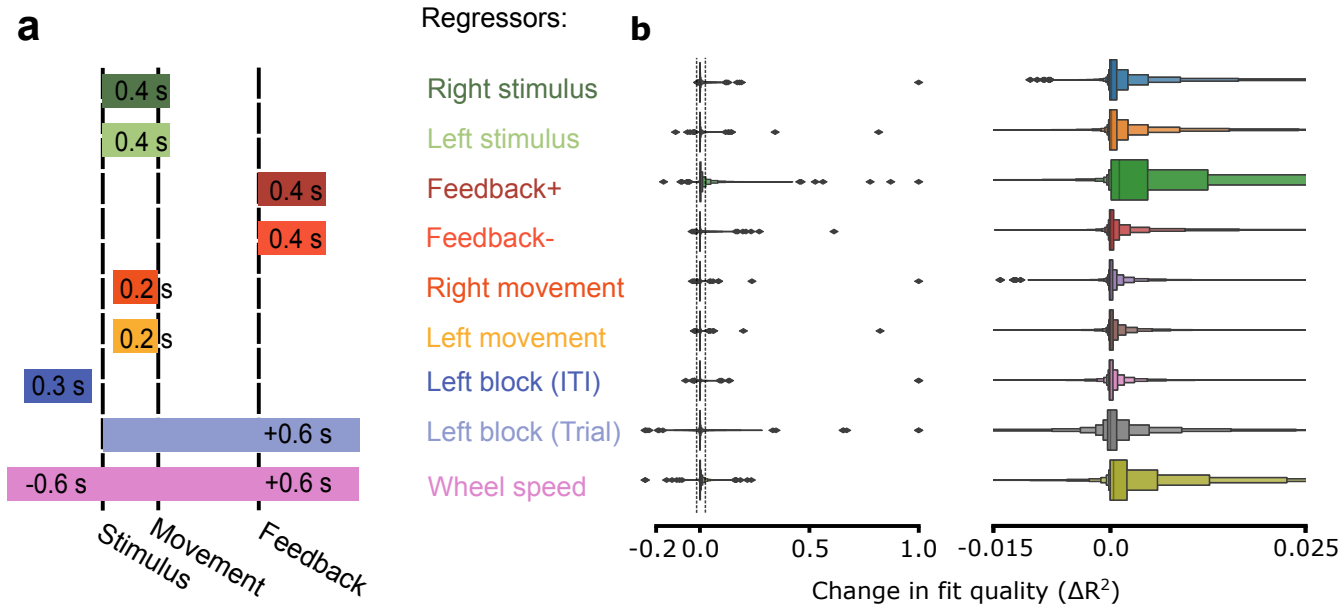

Figure S5. Regressor windows and variance explained in linear encoding model. a) Schematic of within-trial windows in which different regressors in the encoding model apply to firing predictions. **b)** Additional variance explained in a leave-one-out paradigm by each regressor for the full distribution (left) and zoomed-in to the medians of the distributions (right). Note that the range on the right panel is depicted on the left via dotted lines.

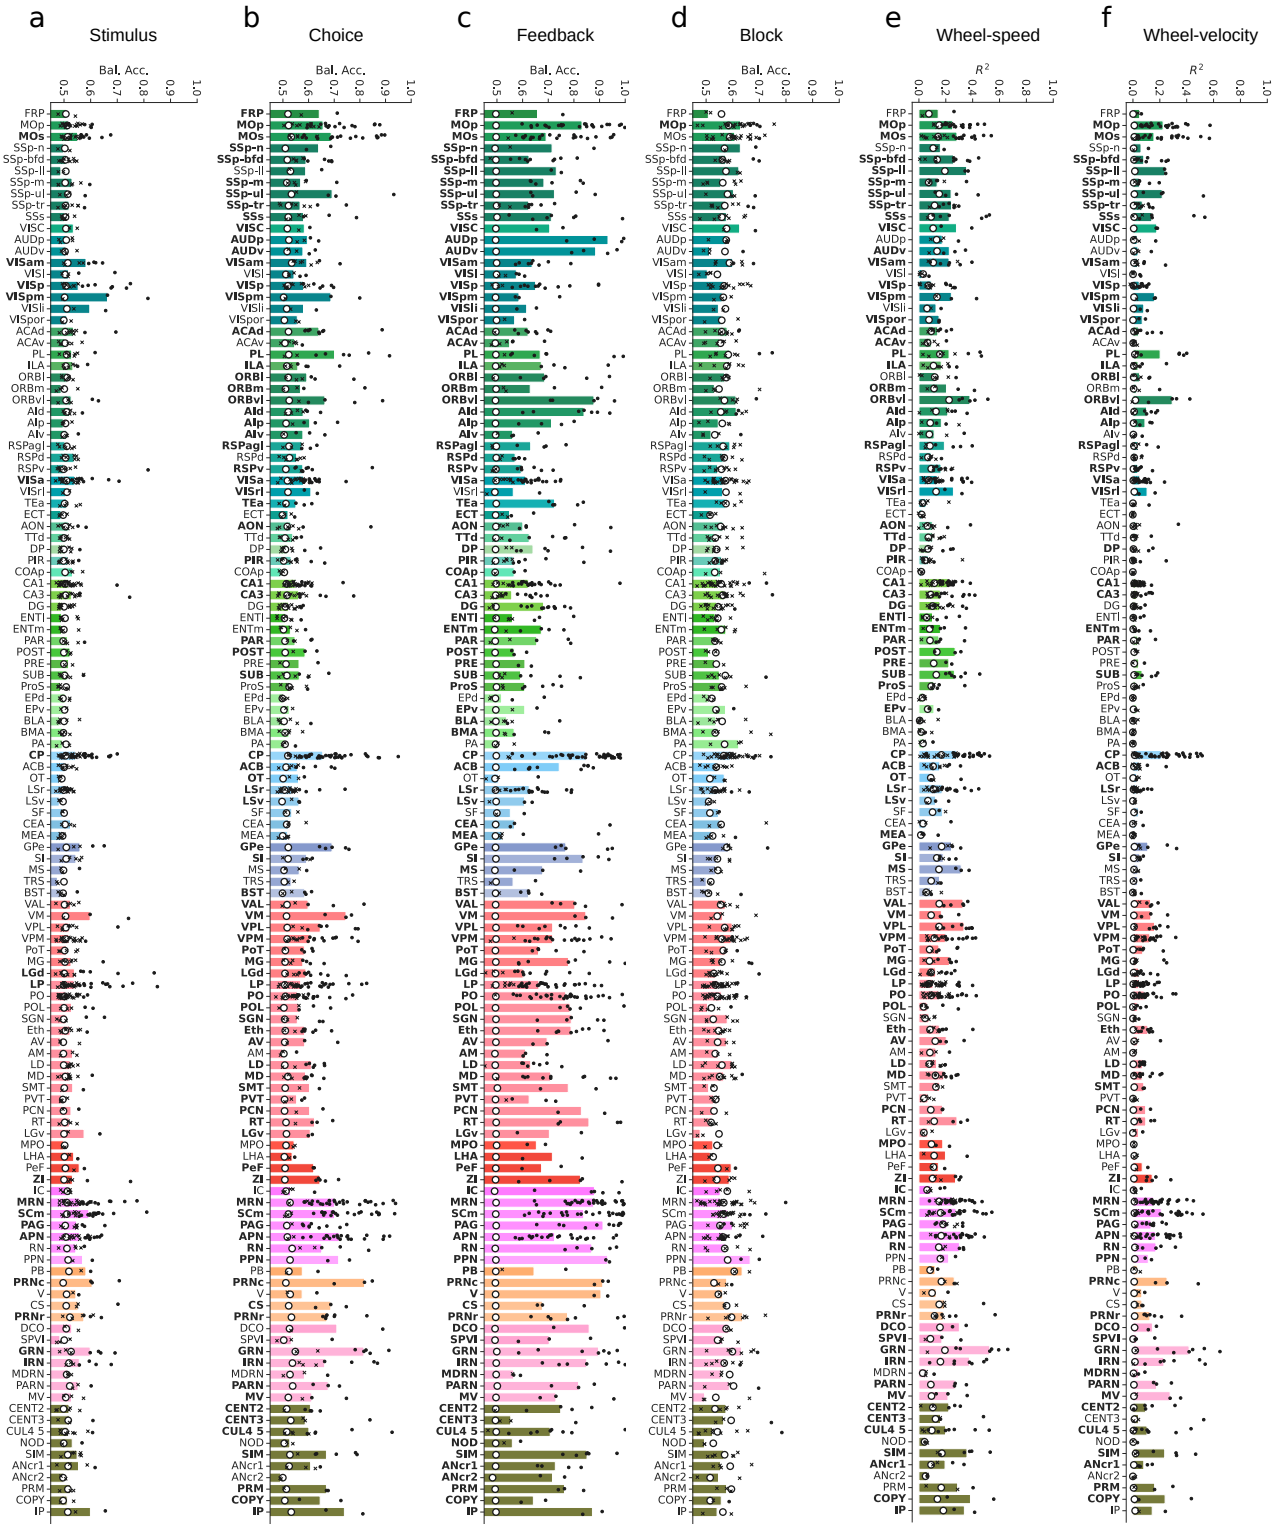

Figure S6. Decoding performance per region with per session results. Decoding analysis as done for **a)** stimulus in Fig. 5, **b)** choice in Fig. 6, **c)** feedback in Fig. 7, **d)** block in Fig. 8, and **e,f)** wheel-speed and wheel-velocity in Fig. 9. No FDR correction has been applied in the bar plots; but the bold labels indicate those regions that survive $\text{FDR}_{0.01}$ (and are shown in the figures in the main paper). Black dots and x's indicate decoding perform on individual sessions where dots are significant at $\alpha = 0.05$ and x's are insignificant. The bar height is the median of all sessions within that region, and the white dot is the across-session median of the null distribution medians.

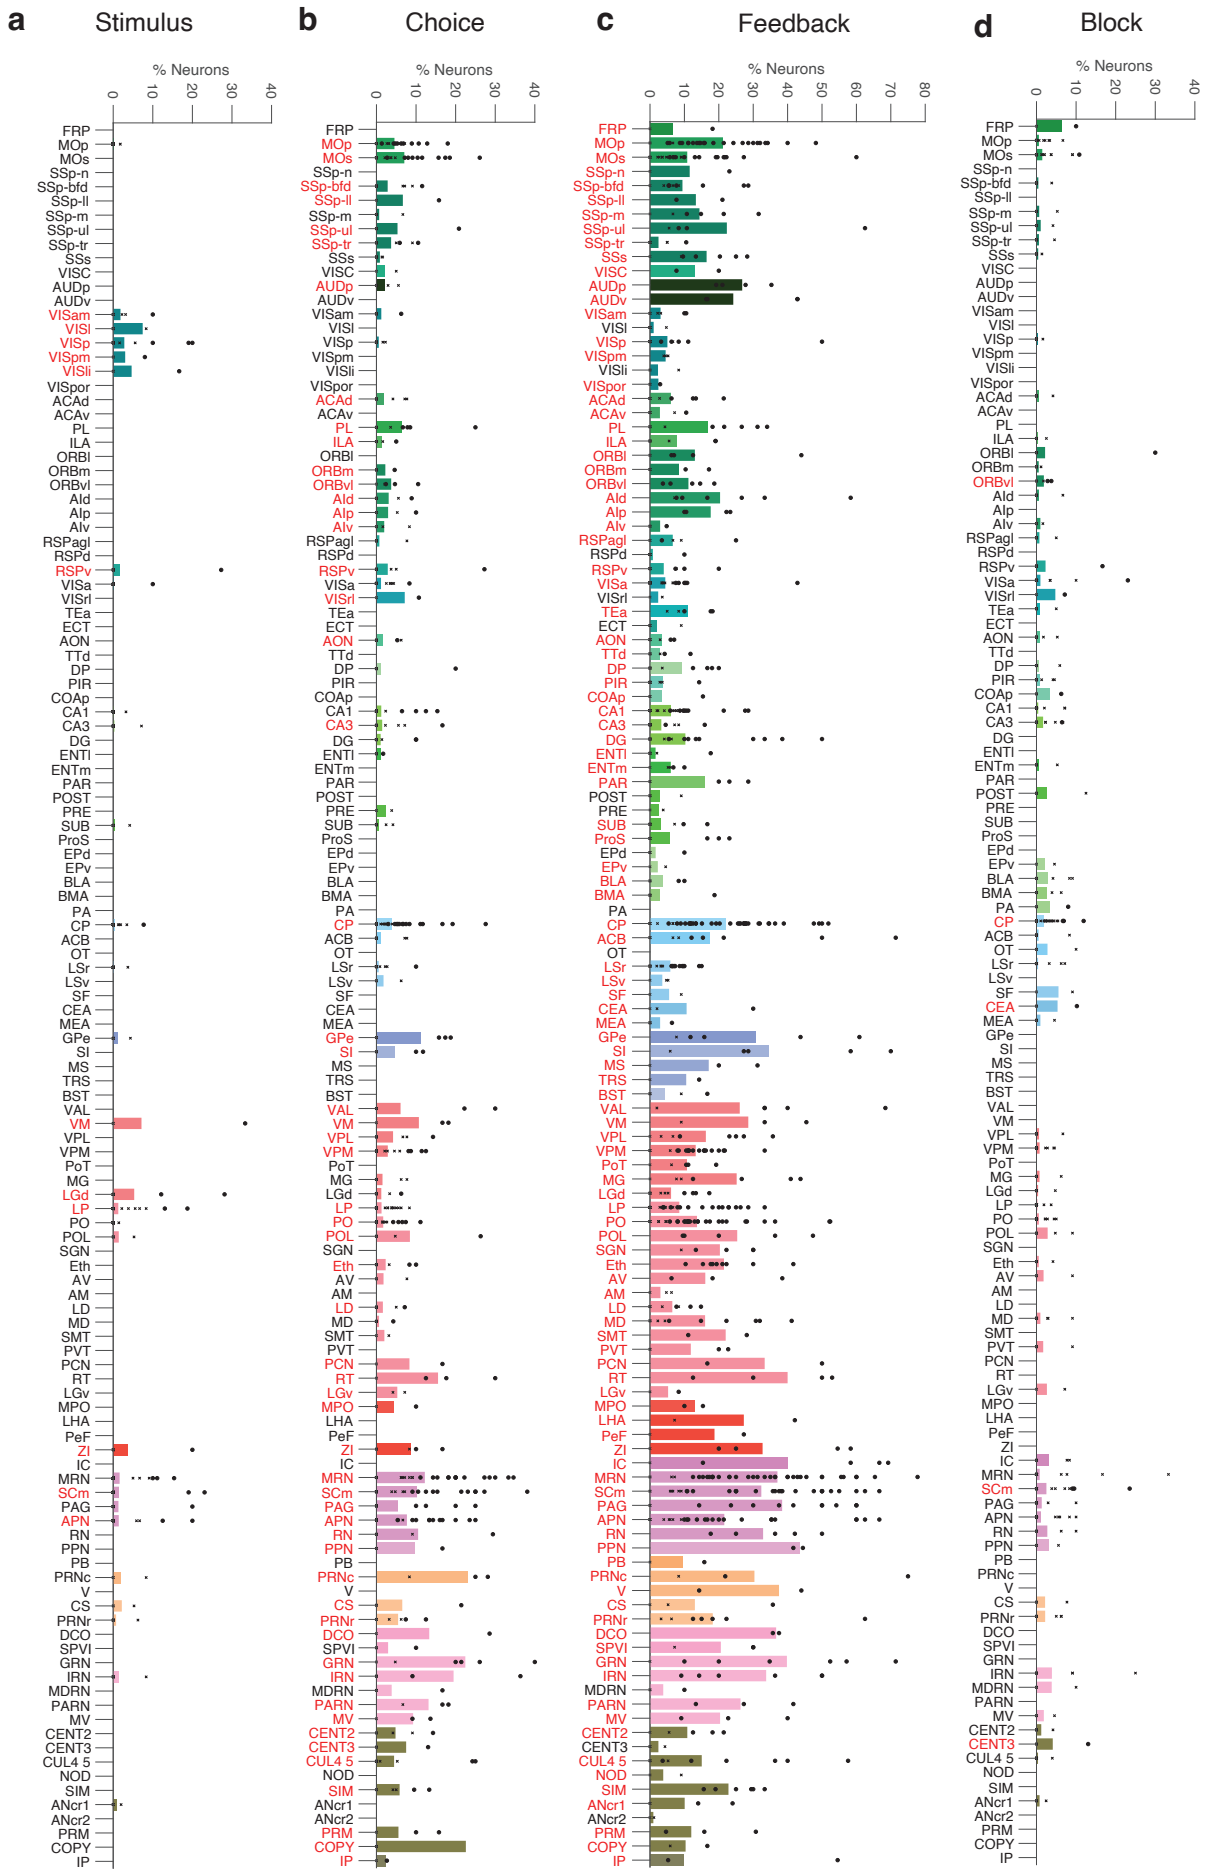

Figure S7. Fraction of significant cells per region in single-cell analysis. Summary of single-cell analysis for stimulus in Fig. 5, **b)** choice in Fig. 6, **c)** feedback in Fig. 7, **d)** block in Fig. 8. No FDR correction has been applied in the bar plots; but the red colour labels indicate those regions that survive $\text{FDR}_{0.01}$ (and are shown in the figures in the main paper). Black dots and x's indicate single-cell analysis is done on individual sessions where dots are significant at $\alpha = 0.05$ and x's are insignificant. The bar height is the mean of all sessions within that region.

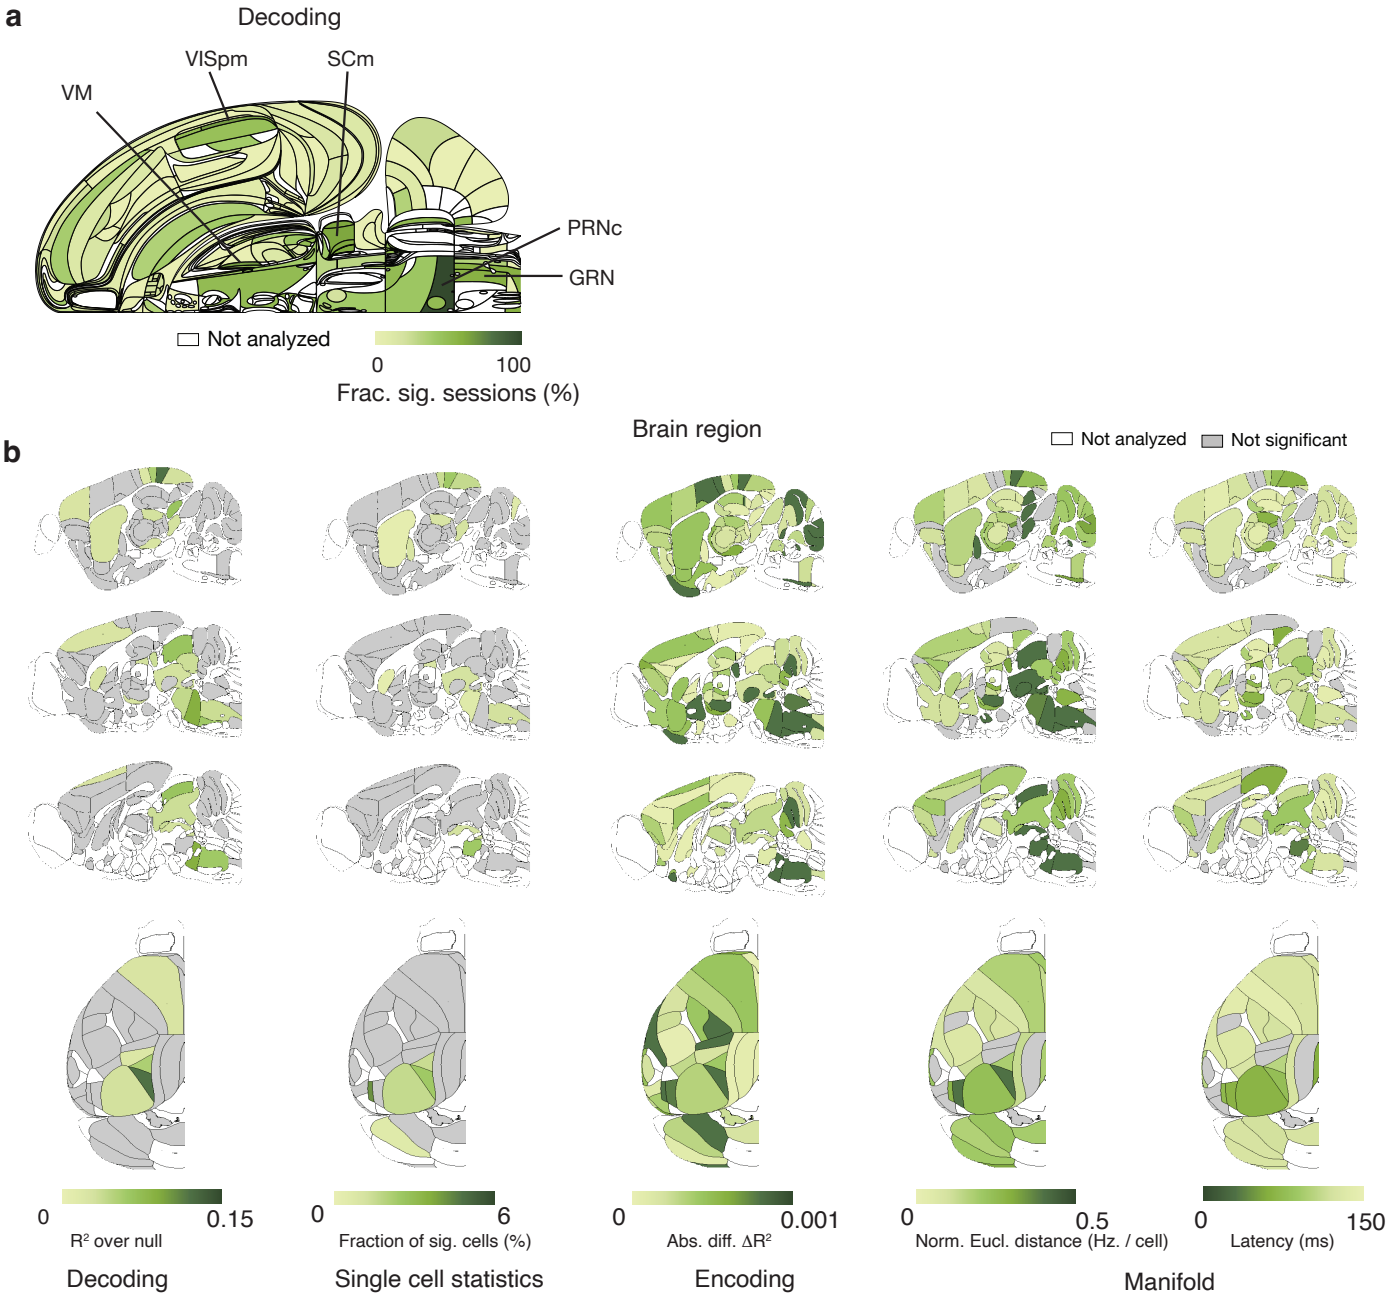

Figure S8. The stimulus variable. **a)** Fraction of sessions with significant decoding performance for the stimulus variable relative to the null. **b)** 2d-brain slices of analysis results for the stimulus variable in Fig. 5a-e. Instead of Swanson flat map, here we use 3 sagittal slices with coordinates ML=-1.8mm, -0.8mm, -0.2mm, and the top view of the dorsal cortex to visualise the representation of task variables across the brain. The locations of sagittal brain slices are optimised to display 252 brain regions. The region acronyms for these slices are listed in Fig. S1.

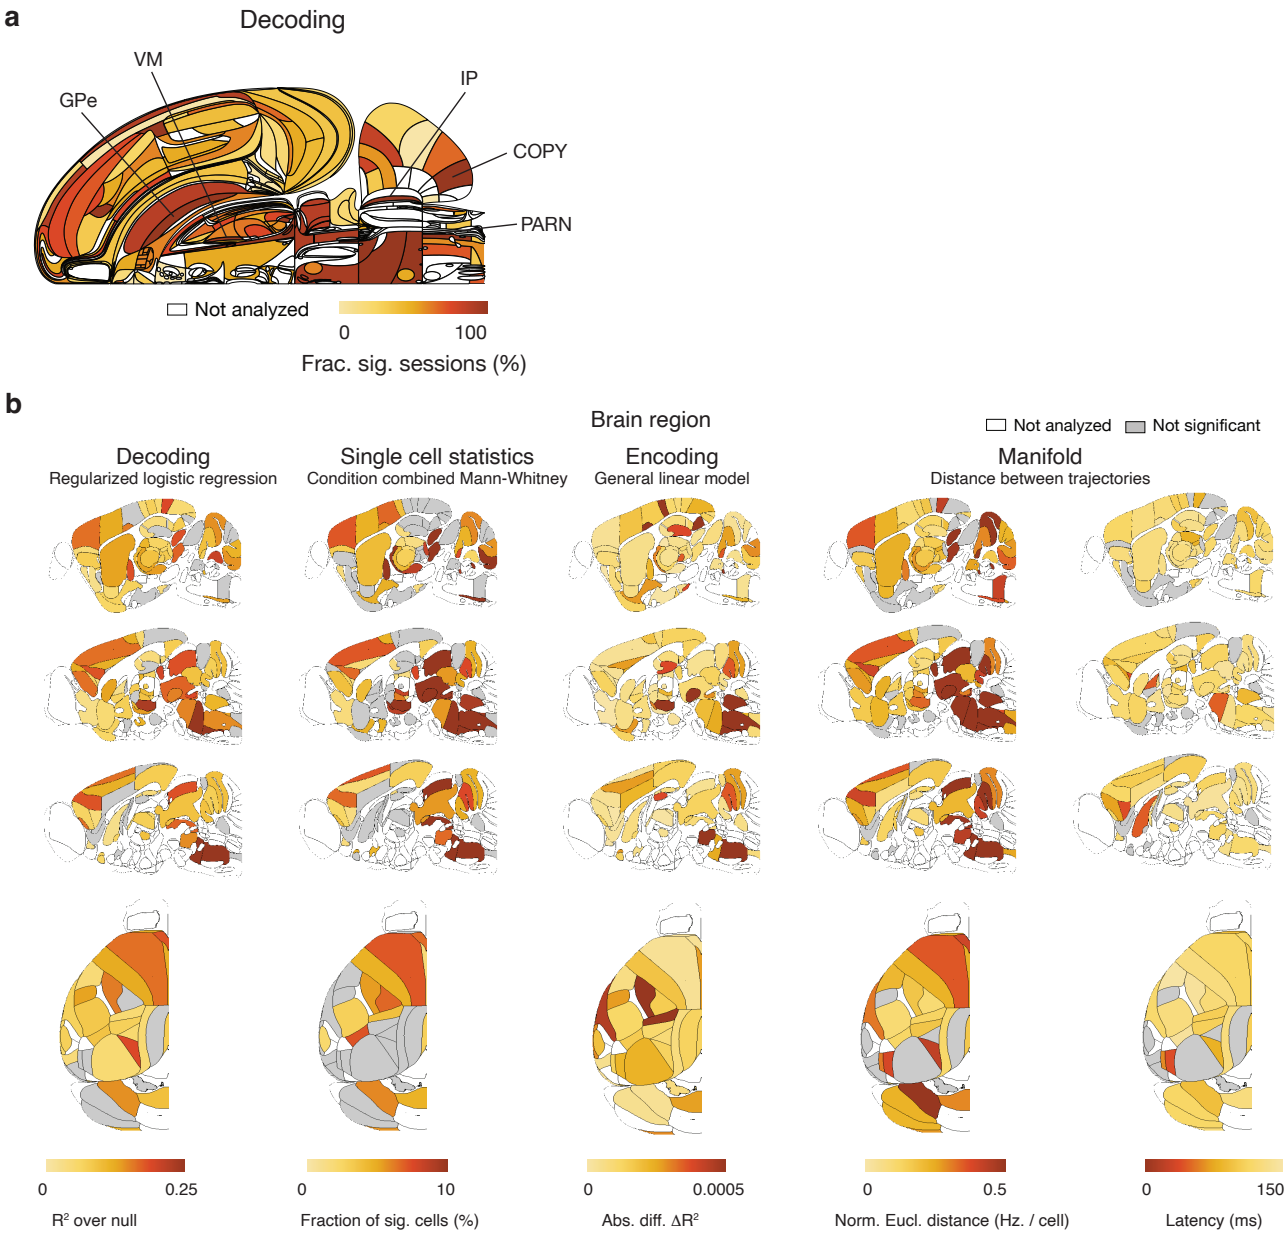

Figure S9. The choice variable. Analysis of the choice variable, with conventions as in Fig. S8.

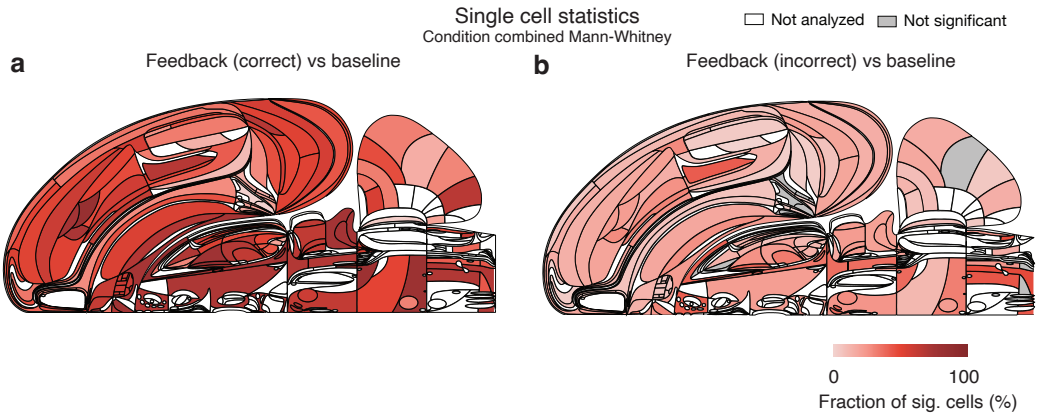

Figure S10. The modulation of neural activity by feedback signal across the brain. a) Fraction of significant neurons per region identified by the condition combined Mann-Whitney test. We compared neural activity after correct feedback ([0, 200] ms) with baseline inter-trial neural activity ([-200, 0] ms aligned to stimulus onset). We deemed a region significant if the number of significant neurons there exceeded the $(1-\alpha)$th percentile of a binomial (N, α) distribution (α=0.001), using $\text{FDR}_{0.01}$ to correct for multiple comparisons. (Methods). **b)** Comparison between neural activity after incorrect feedback ([0, 200]ms) with baseline inter-trial neural activity ([-200, 0]ms aligned to stimulus onset).

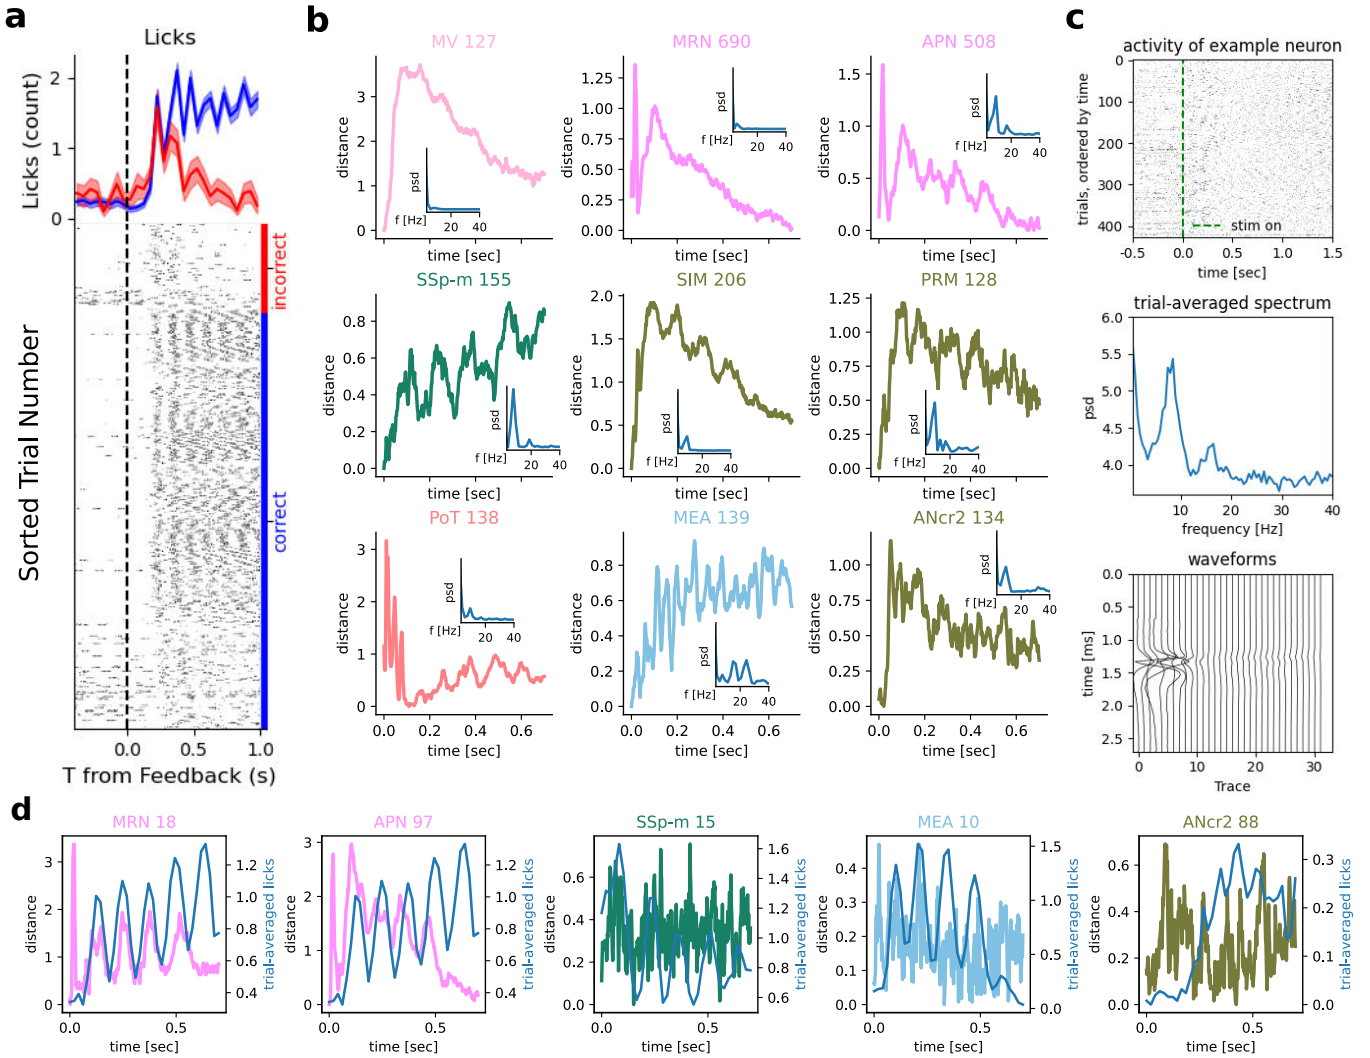

Figure S11. Neural correlates of licking. **a)** Example lick activity for a single session, top trial-averaged, bottom per trial. Animals lick more for correct trials (blue) with a clear rhythm around 10 Hz. Licks were detected using tongue tracking via DLC from side videos. **b)** Neural manifold distance (see manifold analysis) between correct and incorrect trials for example regions selected manually for visible oscillations, with the number of cells (pooled across sessions) next to the region acronym in the title, aligned to feedback. The inset of each panel shows the power spectral density of the distance curve, several having a peak around 10 Hz, correlating with licking. **c)** One example neuron's activity to show activity is physiological and not an artefact. Top panel, raster per trial with rhythmic 10 Hz activity, also shown in the middle panel by the power spectral density of the raster, averaged across trials. Bottom panel, waveforms of this neuron across adjacent traces, illustrating that the spikes we counted are physiological rather than being caused by an electrical artefact. Artefacts could arise, for example, from current flowing through the drinking spout into the Neuropixels probe, which would result in all traces having a strong waveform. We found no evidence for such artefacts when sampling various neurons and inspecting the waveforms. **d)** Single-session manifold distance for select regions with trial-averaged lick activity in blue on top. E.g. in MRN a clear correlation with licking was found when restricting the analysis to a single session, while much less so when considering the session-averaged results, shown in **b)**.

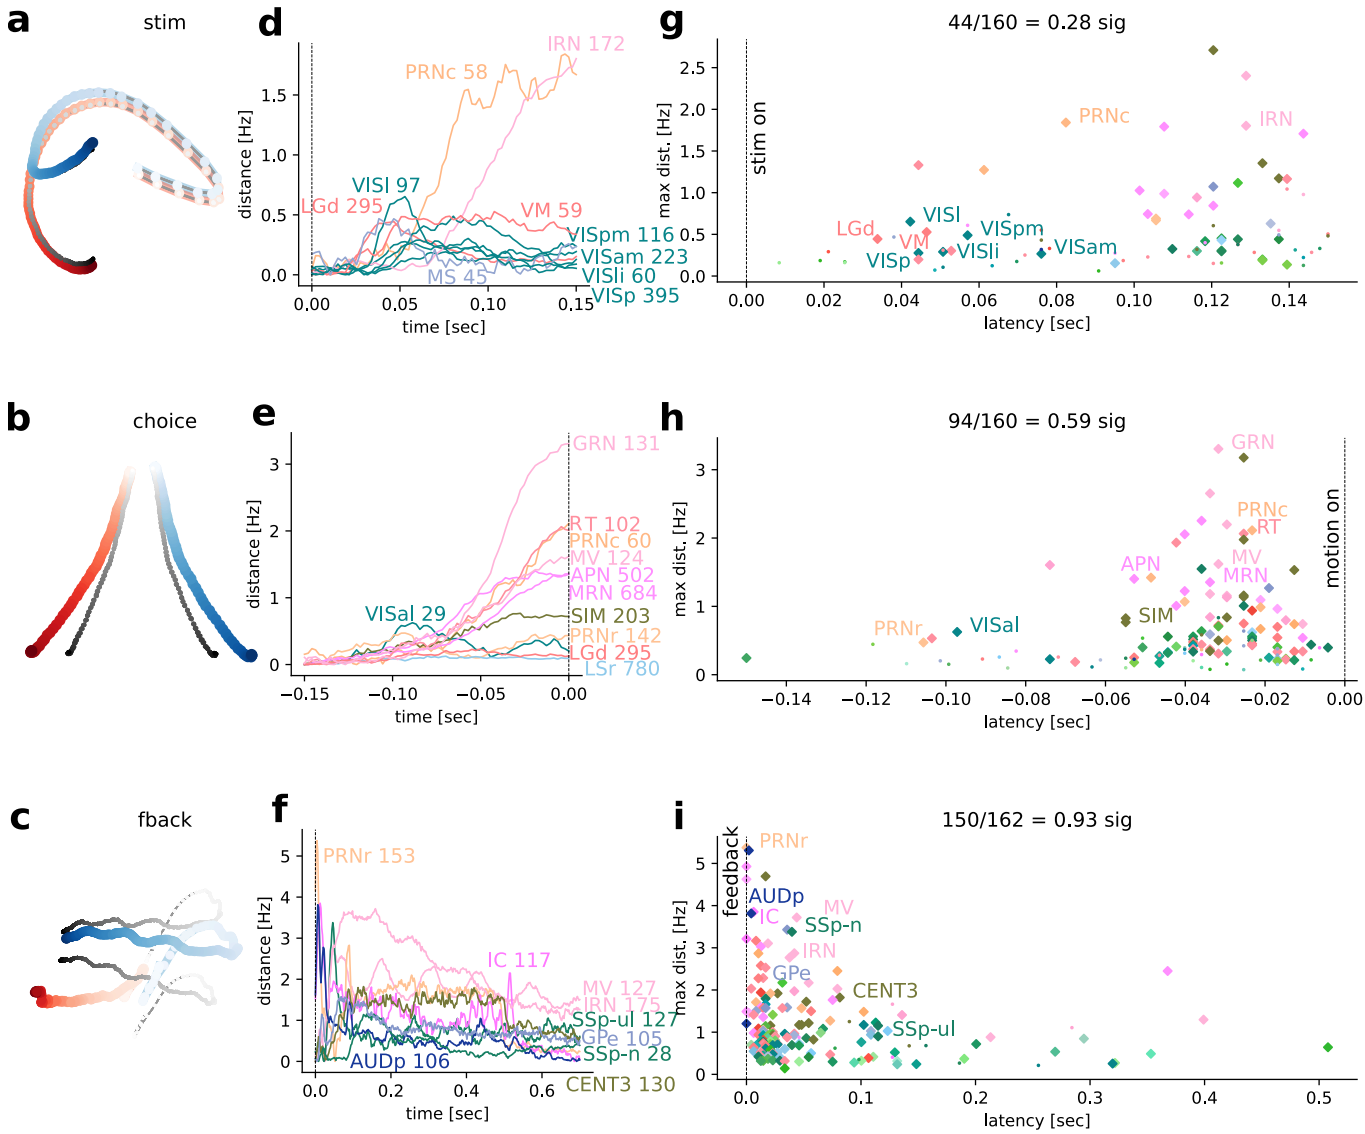

Figure S12. Population dynamics across the brain on the full dataset. Using all good units and considering regions with at least 20 neurons after pooling across sessions, results in about 25% more neurons than in the canonical set of cells that are used across analyses and shown in the main figures. **a-c)** Visualizations (through low-dimensional PCA-embedding) of whole-brain population dynamics (combined across all cells, all sessions, all regions) for three task variables (left versus right stimulus, left versus right choice, correct versus wrong feedback. Blue/red dots represent one time-bin of the population response for left/right (or correct/wrong) trials; colour gradient indicates temporal evolution (darker is later). Grey dots: pseudo-trials. **d-f)** Quantification of the time-resolved distance between opposite trajectories for each variable, based on Euclidean distance (in Hz/cell) in the full-dimensional space (dimension = number of cells) for example brain regions, selected based on response magnitude and to illustrate different response profiles. Curves are annotated by region name and number of cells. Scalebars in all panels represent spikes/s/cell. **g-i)** Summary of variable discriminability for stimulus side, choice side, and feedback type, respectively, by magnitude and latency of response across all recorded brain regions. Diamonds indicate all regions that have statistically significant discrimination ($p < 0.01$ relative to pseudo-trial controls), and line plot examples are labelled by region name. Dots indicate responses of non-significant regions.

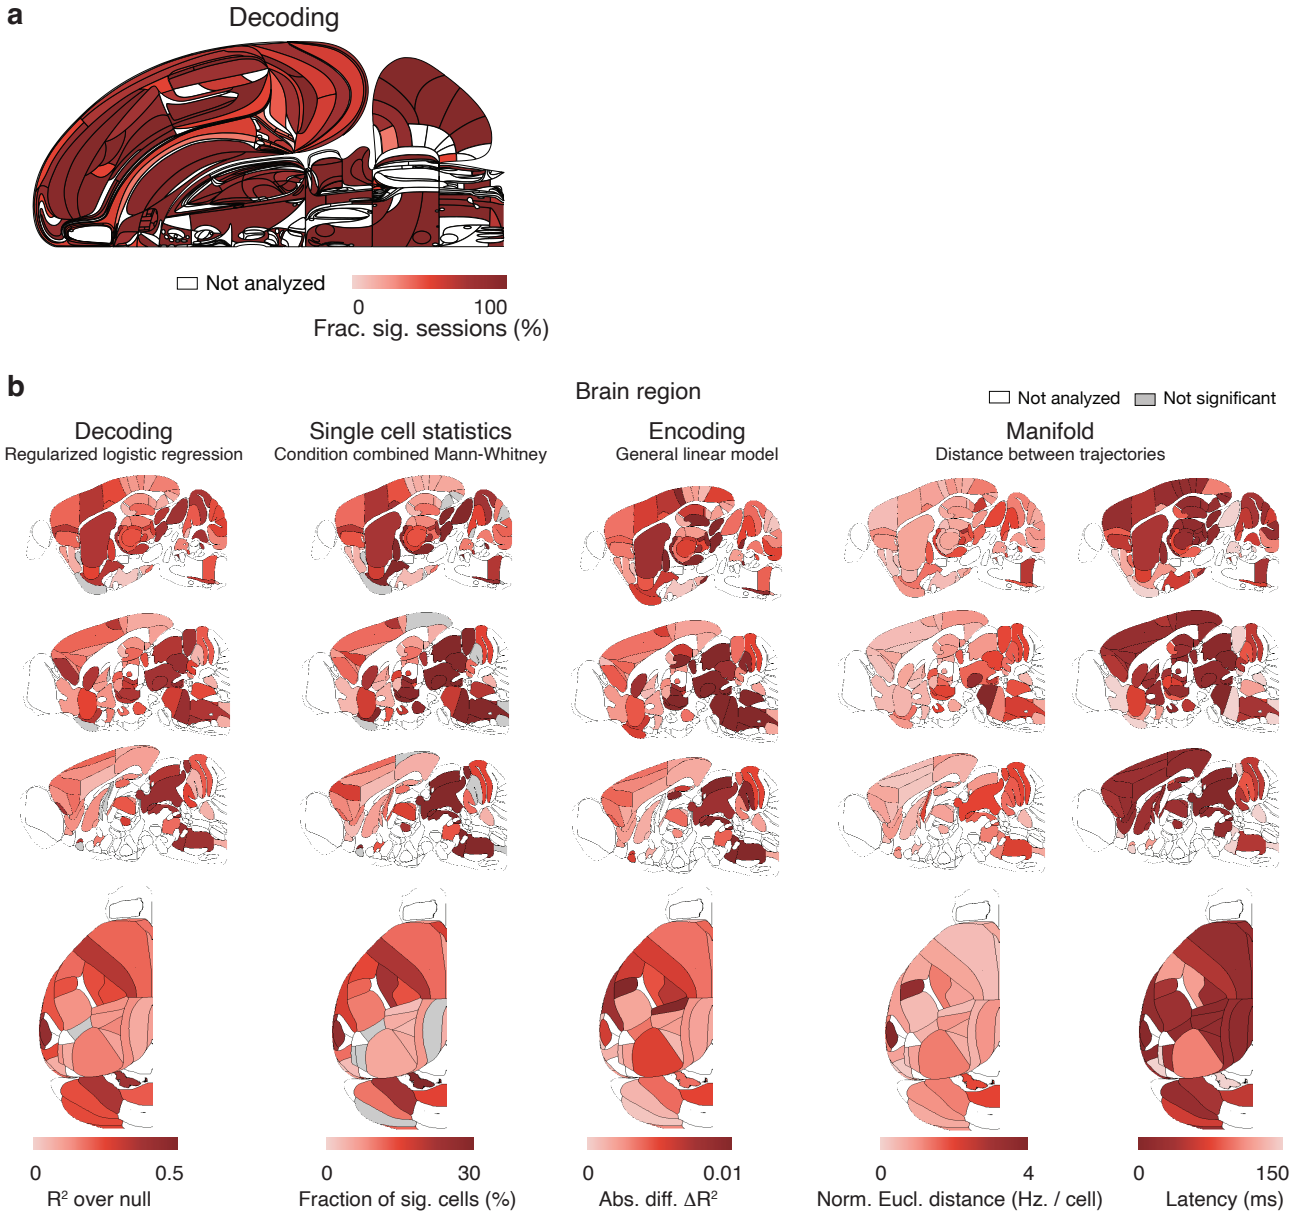

Figure S13. The feedback variable. Analysis of the feedback variable, with conventions as in Fig. S8.

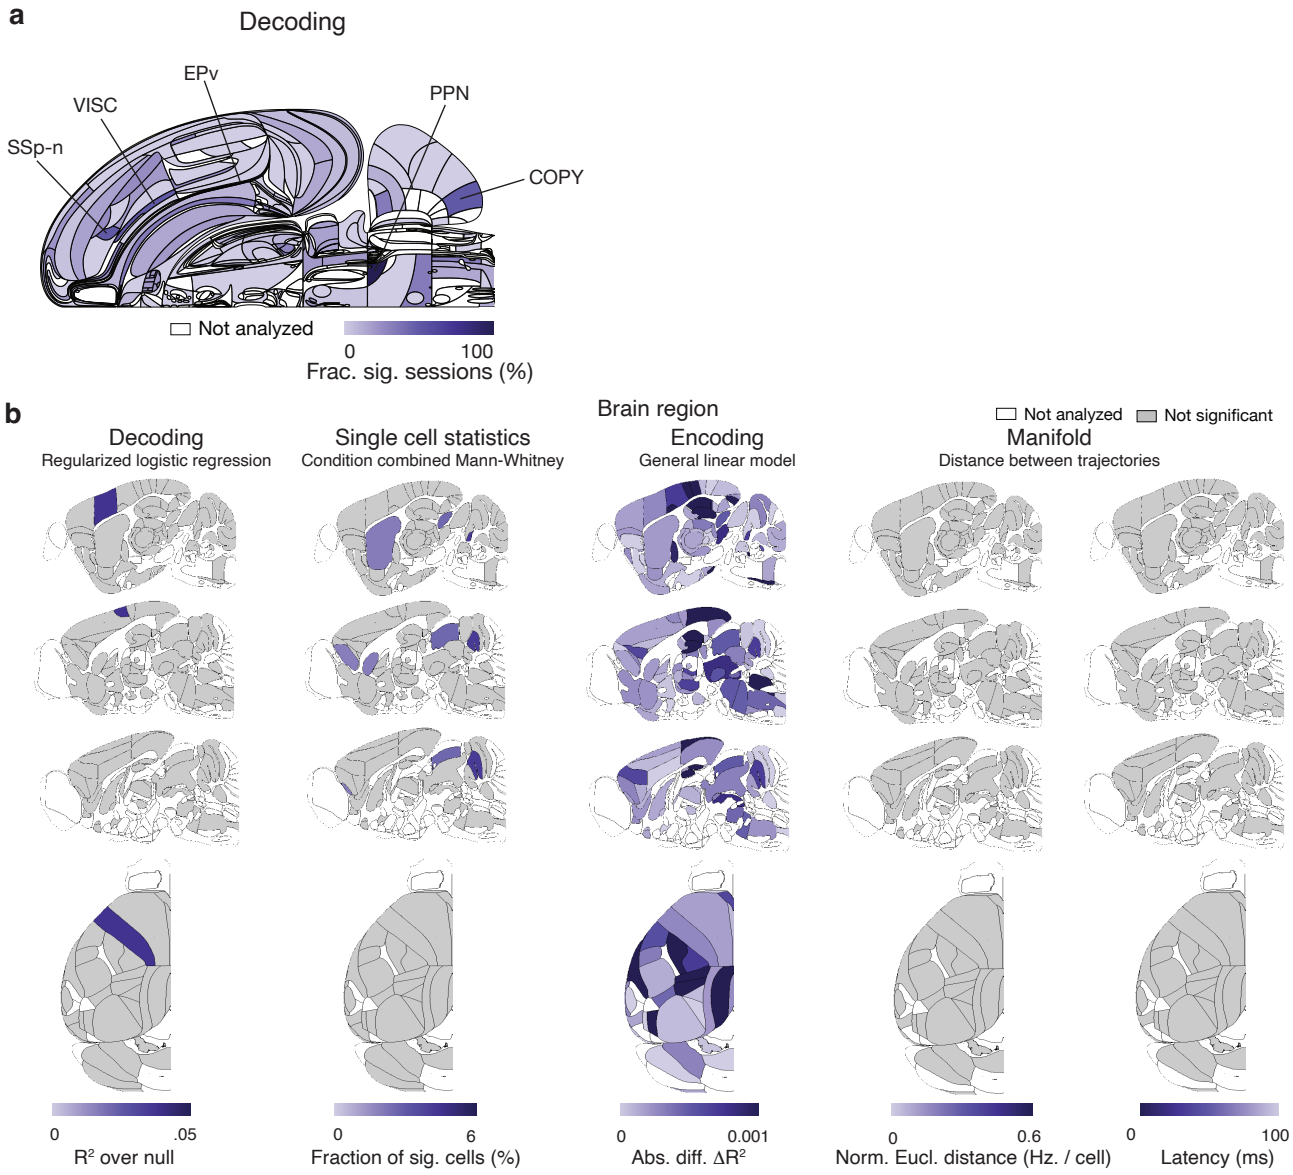

Figure S14. The block variable. Analysis of the block variable, with conventions as in Fig. S8.

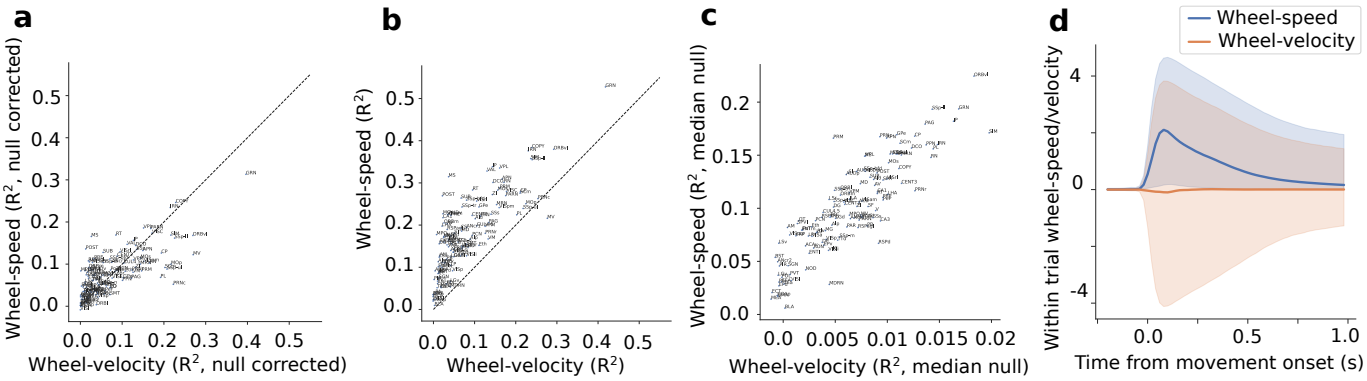

Figure S15. Decoding wheel-velocity versus wheel-speed. Scatter plots comparing per-region wheel-velocity decoding results against wheel-speed for all canonical regions. Decoding is performed for all session-region pairs in the canonical set and the median metric of all such pairs in a given region is plotted. Three such metrics are shown: **(a)** R^2 scores corrected by the median of the null distribution, **(b)** R^2 scores, and **(c)** median of the null distribution. Note the difference in scales for the axes in c. **(d)** The median wheel-speed and wheel-velocity trajectories across all trials are shown and the 5th to 95th percentiles are lightly shaded. The stereotyped shape of wheel-speed produces higher R^2 scores for sessions and null sessions. For example, computing R^2 between 400 randomly chosen trials and 400 repetitions of the median wheel-speed trajectory gives an R^2 of 0.245 (averaged across 1000 repeats). The same computation for wheel-velocity yields $R^2 = 0.000$.

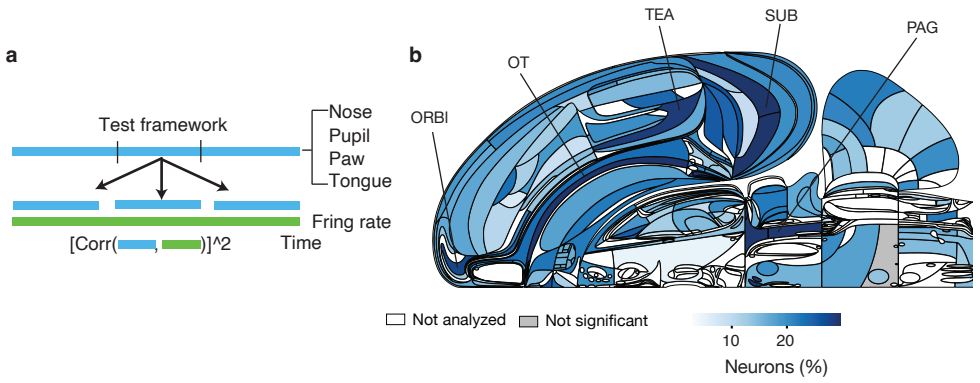

Figure S16. The behavioural correlates of single-neuron activity across the brain. a) Statistical tests to measure the behavioural correlates of single neurons across all sessions. We compute the Pearson correlation coefficient between the time series neural activity and five behavioural variables (nose position, pupil diameter, paw position, and licks, extracted from behaviour video by using DLC; see Sec. S3). The significance of correlation is estimated by a time-shift test[135] (Methods), using $\text{FDR}_{0.01}$ to correct for multiple comparisons. **b)** The flat brain map of the fraction of neurons that significantly correlates with at least one of the movement variables.

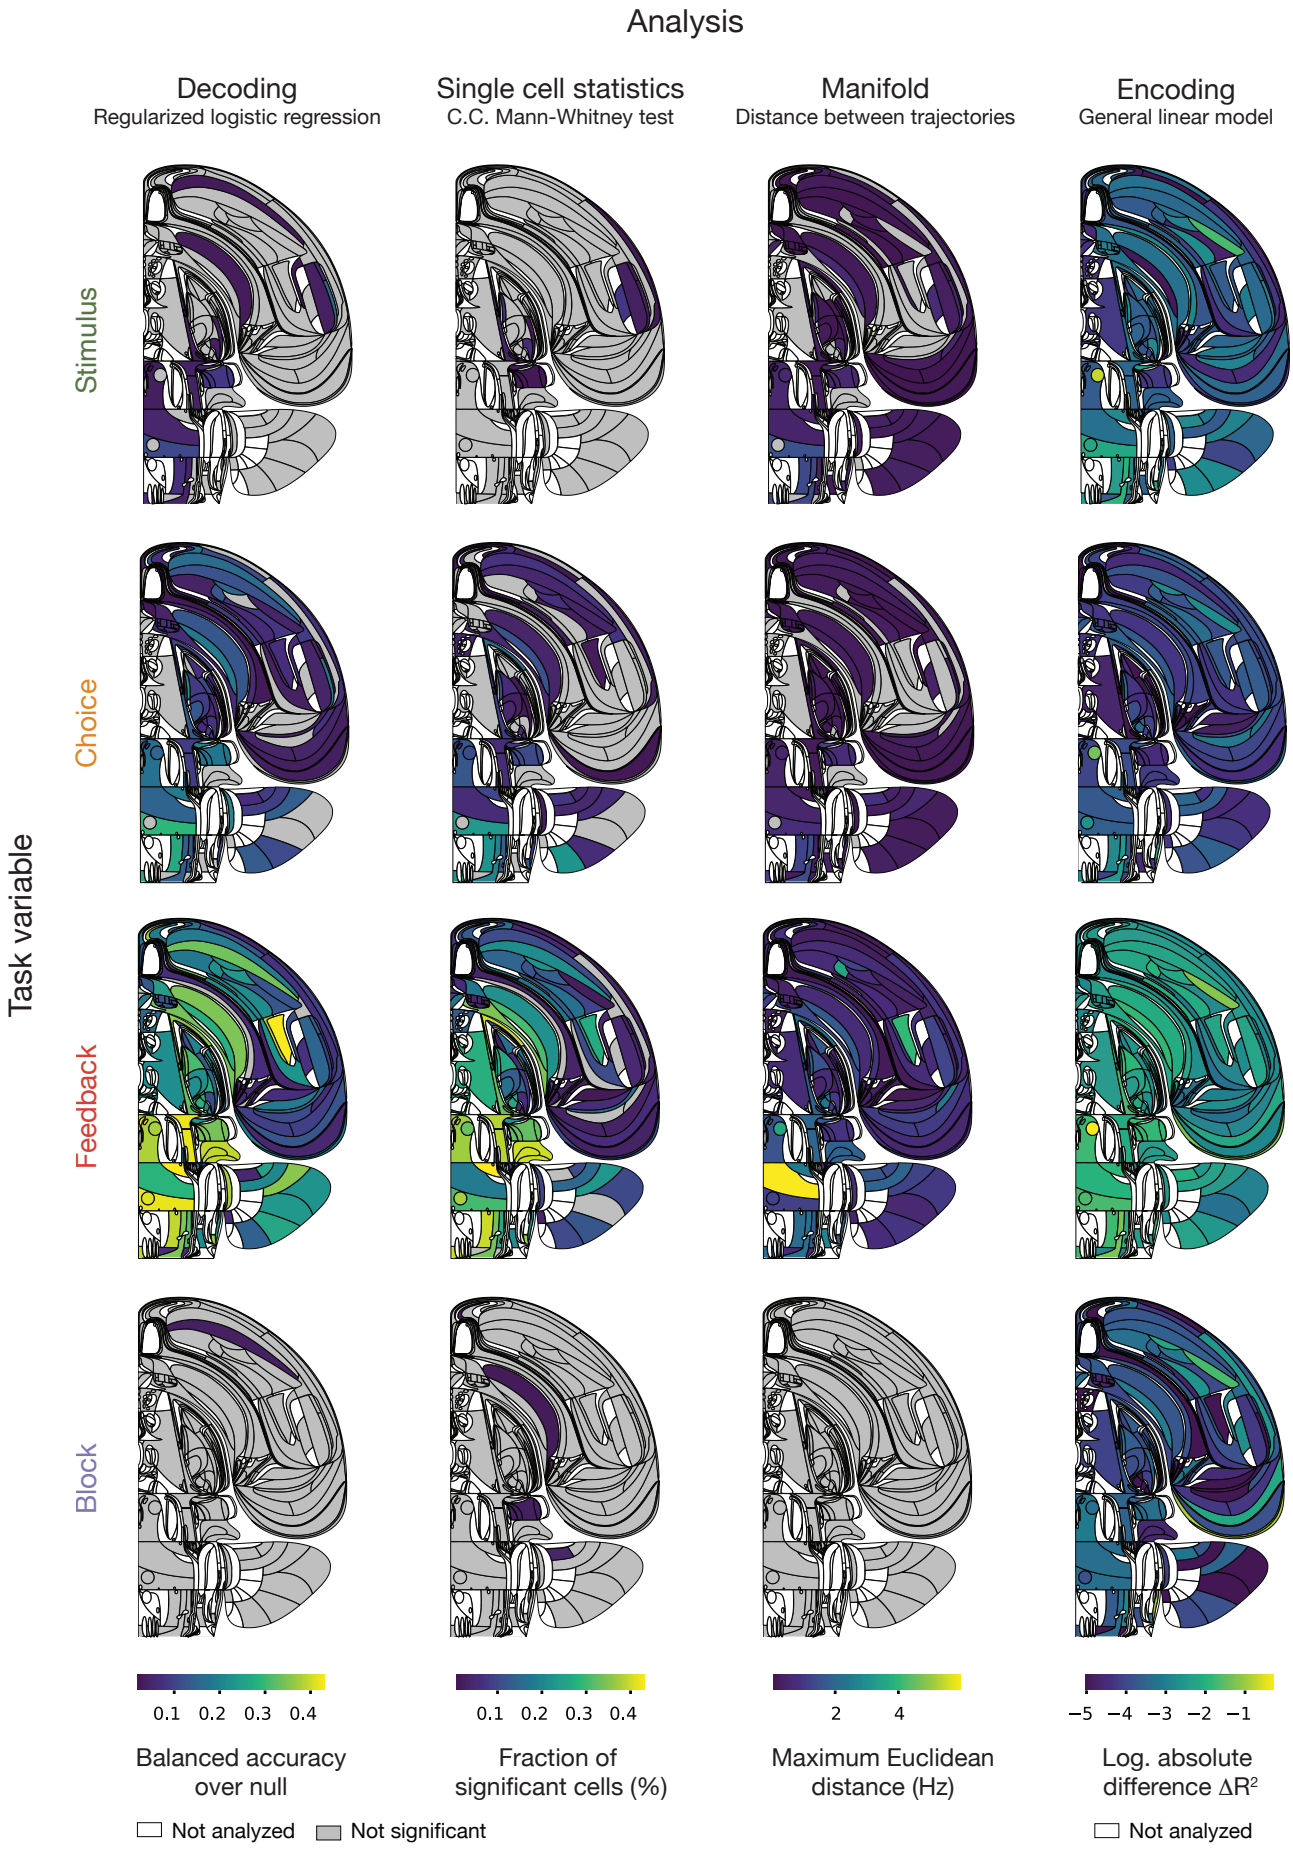

Figure S17. Comparison of effect sizes across task variables. Each column corresponds to a particular neural analysis and each row a task variable. For each analysis, the colour scale is fixed across all variables to enable comparison of effects between variables. For most analyses, the feedback variable has the largest effect amongst all task variables.

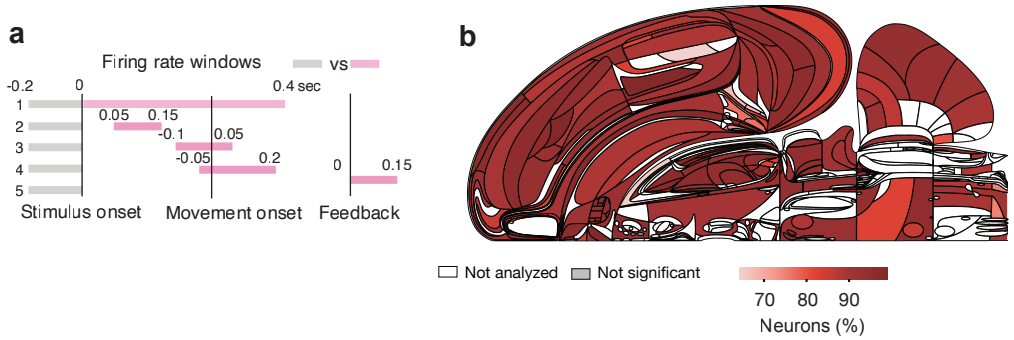

Figure S18. The neural correlates of task across the brain. a) Statistical tests to measure responsiveness in different task windows. The schematics show the summary of all tests, superimposed on the task timeline. Each row represents a separate Wilcoxon rank-sum test comparing firing rates in two different periods over which firing rates were estimated. **b)** The flat brain map of the fraction of neurons that show significant task response during at least one of the task epochs (test of responsiveness: **a**), using $\text{FDR}_{0.01}$ to correct for multiple comparisons.

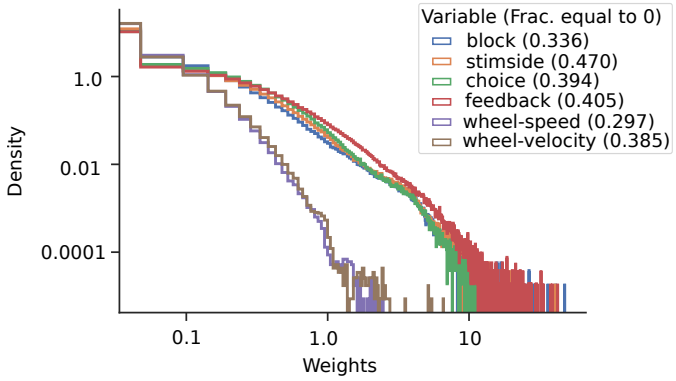

Figure S19. Histogram of regressor weights in decoding analysis. The distributions include decoding weights across all region-session pairs in the canonical set. The distribution combines all weights used on held-out test folds (typically 5 folds) including all temporal bins (for wheel-speed and wheel-velocity) and all repeated decoding runs, but excludes regression intercepts. The legend indicates the decoded variable with, in parentheses, the fraction of weights equal to zero (due to L-1 regularization).