A perceptual decision requires sensory but not action coding in mouse cortex

Peter Zatka-Haas[1,3,4,*], Nicholas A. Steinmetz[1,4], Matteo Carandini[2,5], Kenneth D. Harris[1,5]
[1]UCL Queen Square Institute of Neurology, University College London, London WC1E 6BT, UK
[2]UCL Institute of Ophthalmology, University College London, London WC1E 6BT, UK
[3]Department of Physiology, Anatomy & Genetics, University of Oxford, Oxford OX1 3PT, UK
[4]These authors contributed equally
[5]Co-senior authors
*Correspondence: p.zatka@ucl.ac.uk

Abstract

Sensory decisions involve multiple cortical areas, but it is unclear to what extent these areas carry distinct signals and play distinct causal roles. We trained head-fixed mice to discriminate visual contrast and report their decision by turning a wheel. Widefield calcium imaging and Neuropixels recordings revealed stimulus-related activity in visual (VIS) and secondary motor (MOs) areas, and widespread movement-related activity across the dorsal cortex. Optogenetic inactivation biased choices only when it was targeted at VIS and MOs, at times corresponding to peak stimulus decoding. A neurometric model based on summing and subtracting activity in VIS and MOs successfully described performance and predicted the effect of optogenetic inactivation. Thus, sensory signals in VIS and MOs are causally necessary for task performance, while diffuse dorsal cortical signals correlating with movement do not play a causal role.

Introduction

When presented with multiple sensory stimuli and behavioral options, an animal must make a choice and execute the associated action. These processes of sensation, choice, and execution may be served by the multiple brain regions where neural activity correlates with upcoming actions[1–11]. Distinguishing neural activity related to sensation, choice, or action, however, is challenging.

A key confound when interpreting neural responses is the widespread activity that pervades the brain before movements[6,12–16]. The most common neural correlates of behavior appear to be with whether the animal will move at all, rather than which particular action will be chosen. For example, in a multi-alternative visual discrimination task, neurons in nearly all brain structures are modulated prior to upcoming actions, whereas neurons predicting which action will be chosen are rare and confined to specific regions such as basal ganglia, midbrain, and frontal cortex[6].

It is unclear whether this widespread, non-specific pre-movement activity is necessary for performance of a task. Traditional theories would not predict that brain-wide activity is required for simple sensory-motor tasks. Indeed, in some well-trained behaviors, cortical inactivation has little effect[17,18]. However, in other simple tasks such as licking following a whisker deflection, performance is impaired by inactivation of many structures, including medial prefrontal cortex and hippocampus[19].

Inactivating a brain region during a choice task could have three potential consequences. First, inactivation could bias the choices, implicating the inactivated neurons in a sensory or choice circuit[10,20–24]. Second, inactivation might disrupt the execution of the chosen action, for example delaying it or reducing its accuracy[25], but not affecting the probability of the choices. This would implicate the inactivated neurons in circuits for executing a chosen action, but not for the choice itself. Finally,

inactivation might have no effects on the current trial, suggesting that the region's activity in the task reflects corollary discharge (feedback from other brain areas involved in action selection and execution[26]), or cognitive processes such as memory, which only affect subsequent choices[27,28] among other possibilities.

Here we use imaging, electrophysiology, and optogenetic inactivation to reveal the roles of dorsal cortical regions in vision, action and choice. We trained mice in a multi-alternative visual discrimination task[29] and found correlates of vision (regardless of behavioral output) in visual and frontal cortex, and correlates of action (initiation of any movement, regardless of stimulus and choice) everywhere, but strongest in primary motor and somatosensory cortex. Correlates of choice (regardless of stimuli) were revealed electrophysiologically only in rare neurons in frontal areas. The effect of optogenetically inactivating a region closely mirrored its encoding of sensory stimuli rather than its encoding of action or choice. For instance, primary motor cortex had the strongest action correlate but inactivating it had no effect beyond a slight decrease in peak movement velocity. We constructed a model of choice based on cortical activity, and found that it provided a parameter-free prediction of the effects of optogenetic inactivation. These results account for how cortical activity contributes to choice formation, and reveal that the causal role of dorsal cortical regions reflects their representation of sensation rather than action.

Results

A visual discrimination task affording multiple types of error

We trained mice to perform a two-alternative unforced-choice visual discrimination task[29] (Fig. 1). Mice were head-fixed with their forepaws controlling a steering wheel surrounded by three screens (Fig. 1a). Each trial began after the wheel was held still for a minimum duration. Grating stimuli appeared in the left and right screens together with an auditory Go cue (Fig. 1b-d). Mice were rewarded with water for rotating the wheel to bring the higher-contrast stimulus into the center (termed Left or Right choice trials, according to which side's stimulus is driven to the center), or for holding the wheel still for 1.5 s if no stimulus was present (termed NoGo trials). Unrewarded trials ended with a 1 s white noise burst. Mice became proficient in the task, achieving 86 ± 9% correct choices (mean ± s.d.) on trials with single high contrast gratings, and the probability of their choices varied with the contrasts of the two gratings (Fig. 1e-h). In the absence of visual stimuli, moreover, mice correctly held the wheel still (NoGo) in 54 ± 16% of the trials.

In this task design, mice could make multiple types of error (Fig. 1h). When a stimulus was present, mice could err with a NoGo (a "Miss" trial, 16 ± 13% of all trials where a stimulus was present on one side, mean ± s.d. across sessions) or with a wheel turn in the incorrect direction ("Incorrect choice", 7 ± 4%). The probability of both types of error increased with decreasing stimulus contrast, reaching a maximum of 27 ± 21% Miss rate and 9 ± 8% Incorrect rate for single low-contrast stimuli on one side (Fig. 1h). These two types of errors permit us to separate the overall tendency to respond (Choice versus Miss) from the accuracy of perceptual discrimination (Correct versus Incorrect choices).

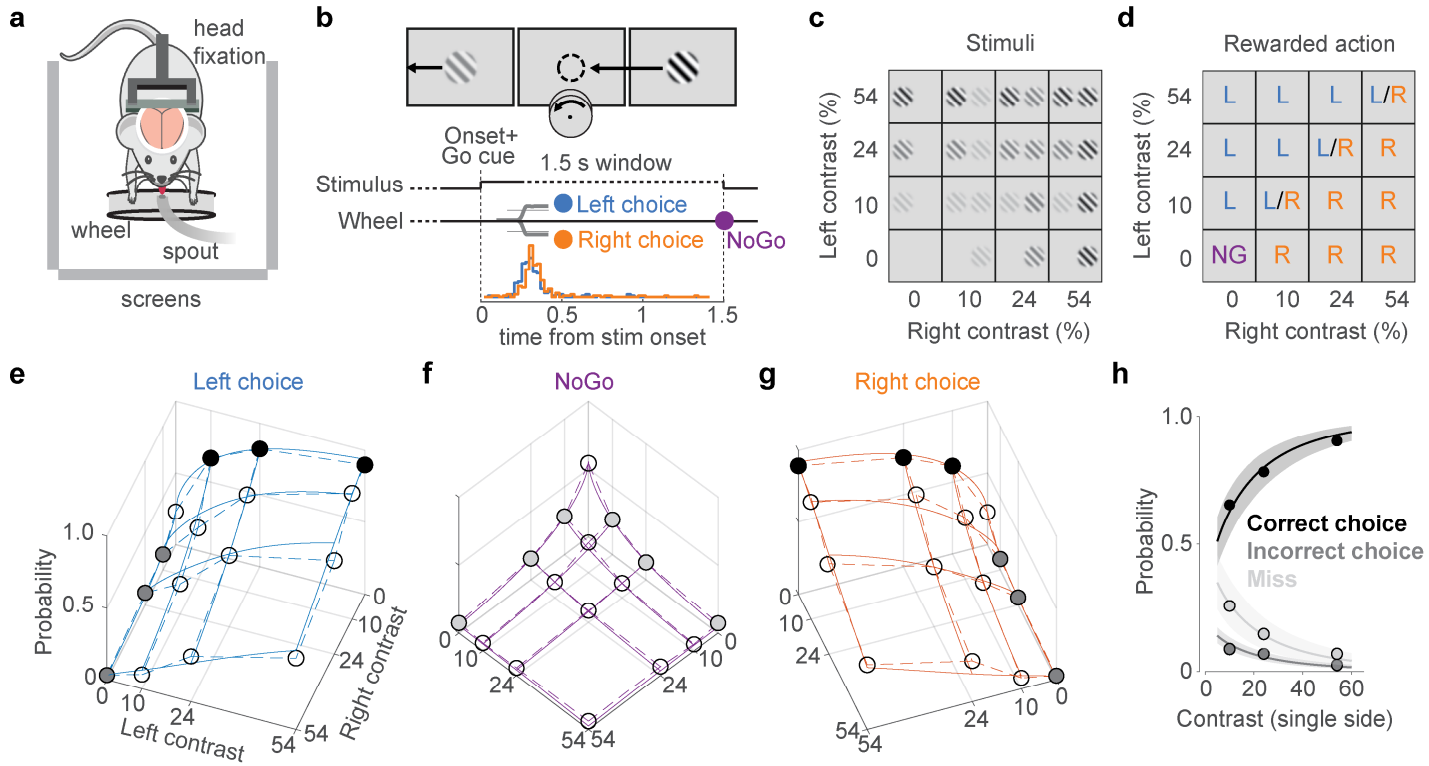

Figure 1. A visual discrimination task affording multiple types of error.
(a) Behavioral setup, with the mouse surrounded by 3 screens. A spout delivers water rewards. **(b)** Example stimulus, with higher contrast on the right than left. The correct movement is to turn the wheel anticlockwise to bring the right stimulus to the middle (arrow and circle, not visible to the mouse). Middle: Task timeline. Trials began after the mice held the wheel still for 0.2-0.6 s. Grating onset was accompanied by an auditory Go cue. Mice could then make a choice (Left/Right) or hold the wheel for 1.5 s (NoGo). Trials were separated by a 1 s interval. Bottom: Distribution of movement onset times ('reaction time', see Methods) for correct Left (blue) and Right (orange) choices from one example mouse. **(c)** The 16 stimulus conditions. Gratings on each side could have one of four contrasts. **(d)** Rewarded actions depended on stimulus condition: Left (L), Right (R), NoGo (NG). When the contrasts were equal, Left and Right choices were rewarded randomly (L/R). **(e)** Probability of Left choices as a function of left and right contrast, averaged over 34 sessions in 6 mice (circles and dashed lines). Colored circles highlight unilateral stimulus conditions summarized in (h), where choices are Correct (black), or Incorrect (dark gray). Blue curves indicate the mean fit of a psychometric model[29] fitted with hierarchical priors (see Methods). **(f,g)** Same format as in (e), showing NoGo trials and Right choice trials. Light gray dots indicate unilateral stimulus conditions where NoGo counts as a Miss. **(h)** Summary of these choices for unilateral stimulus conditions. Circles (with color fills as in e-g) show the probability of a Correct choice (black), Incorrect choice (dark gray) and incorrect NoGo (Miss; light gray). Curves and shaded regions are the mean and 95% credible intervals of the mean fit from the psychometric model.

Activity related to stimuli and movement in the dorsal cortex

Widefield calcium imaging revealed that the visual stimulus elicited a sequence of activations in the contralateral cortex, followed by widespread bilateral activity before a movement (Fig. 2a-e). Activity began in the primary visual cortex (VISp) 47 ± 4 ms after stimulus onset (time to 30% of peak; median ± m.a.d. across 6 mice), spreading to secondary visual areas (e.g. VISal at 55 ± 8 ms), and then to frontal cortex (secondary motor area MOs at 98 ± 5 ms; Supplementary Figures S1, S2). The pattern of activity during these first 100 ms appeared to be independent of the trial's outcome of Go (Left or Right) or NoGo (Fig. 2b,c). Later activity, however (150-200 ms following stimulus), strongly de-

pended on whether the stimulus evoked a movement, with widespread bilateral activity seen on Go trials but absent on NoGo trials (Fig. 2b,c; Supplementary Figure S1; Supplementary Movies 1 and 2). This widespread activity preceded the mouse's choice, which did not occur until 229 ± 28 ms after stimulus onset (median ± m.a.d. across 39 sessions; reaction time defined as the earliest detected wheel movement after stimulus onset, see Methods). However, while cortical activity was very different between Go and NoGo trials evoked by the same stimulus, cortical activity preceding Left vs. Right responses to the same stimulus appeared to be indistinguishable (Fig. 2d,e).

The widefield signal was correlated with sensory and action variables, but not with choice (Fig. 2f-h). To show this, we used the Combined Conditions Choice Probability (CCCP) measure, which tests whether one can decode a behavioral or stimulus variable from neural activity observed at a single place and time, while holding all other behavioral variables constant[6]. We first asked which regions' activity could be used to decode the presence of a sensory stimulus on one screen, amongst trials with the same behavioral outcome and the same contrast on the other screen ("Vision" correlates). These Vision signals appeared in contralateral visual and frontal cortices (Fig. 2f), starting in VISp at 50 ms and MOs at 100 ms (defined as the first time bin at which decoding performance significantly exceeded 50% across sessions; $p<0.01$, nested ANOVA). Further decoding performance was possible in ipsilateral MOs at 125 ms, but not elsewhere. We next asked which regions' activity could decode Go vs NoGo choices, amongst trials with the same stimulus conditions ("Action" correlates). This showed a different pattern: Go vs. NoGo could be decoded from cortical activity from 125 ms prior to movement onset, starting in primary motor and frontal cortex, and expanding to most imaged regions by 25 ms before movement (Fig. 2g). Finally we asked which regions could decode upcoming Left vs Right choices amongst Go trials ("Choice" correlates); this was not possible from any recorded region (Fig. 2h).

These imaging results were consistent with recordings from individual neurons (Fig. 2i-l). We analyzed Neuropixels electrode recordings in the same task[6] and we fitted single-neuron firing rates with an encoding model: a sum of kernel functions time-locked to stimulus presentation and to movement onset[3,6] (Fig. 2i). Neurons with significant stimulus encoding (see Methods) were localized to contralateral VISp (25.2%, 255/1010 neurons; Fig. 2j), contralateral MOs (5.6%, 59/1062) and rarely in ipsilateral MOs (1.4%, 15/1062). By contrast, neurons with significant pre-movement action encoding were observed broadly in VISp (2.8%, 28/1010; Fig. 2k), MOs (12.0%, 127/1062), primary motor cortex (MOp; 15.3%, 85/556) and primary somatosensory cortex (SSp; 20.1%, 62/308). Neurons encoding choice were rare, and were found above chance levels only in MOs (2.3%, 24/1062), MOp (1.8%, 10/556) and SSp (3.0%, 9/308); within the recorded hemisphere, neurons with preference for Left and Right choices were found approximately equally (Fig. 2l, Supplementary Figure S3).

To determine the timing of vision, action, and choice coding in the electrophysiologically recorded populations, we used the CCCP decoding strategy (Fig. 2m-p, Supplementary Figure S3). Stimulus decoding began at ~60 ms in contralateral VISp, ~110 ms in contralateral MOs and ~180 ms in ipsilateral MOs (Fig. 2m). Go vs. NoGo decoding was first observed in each region 50-125 ms prior to movement onset (Fig. 2n), while Choice decoding was rare prior to movement (Fig. 2o) but increased substantially afterwards, reaching a peak at 200-300 ms after movement onset (Fig. 2p).

Widefield calcium imaging and Neuropixels recordings therefore suggested two primary cortical correlates of task variables: vision and action. First, neurons in visual cortex VIS and frontal cortex MOs correlated with the sensory visual stimuli. Second, neurons in all recorded regions correlated with upcoming action, with primary motor area MOp showing the strongest action correlates. In

contrast, neurons encoding the direction of the animals' upcoming choice were rare, consistent with the resulting choice signals being absent in the widefield data.

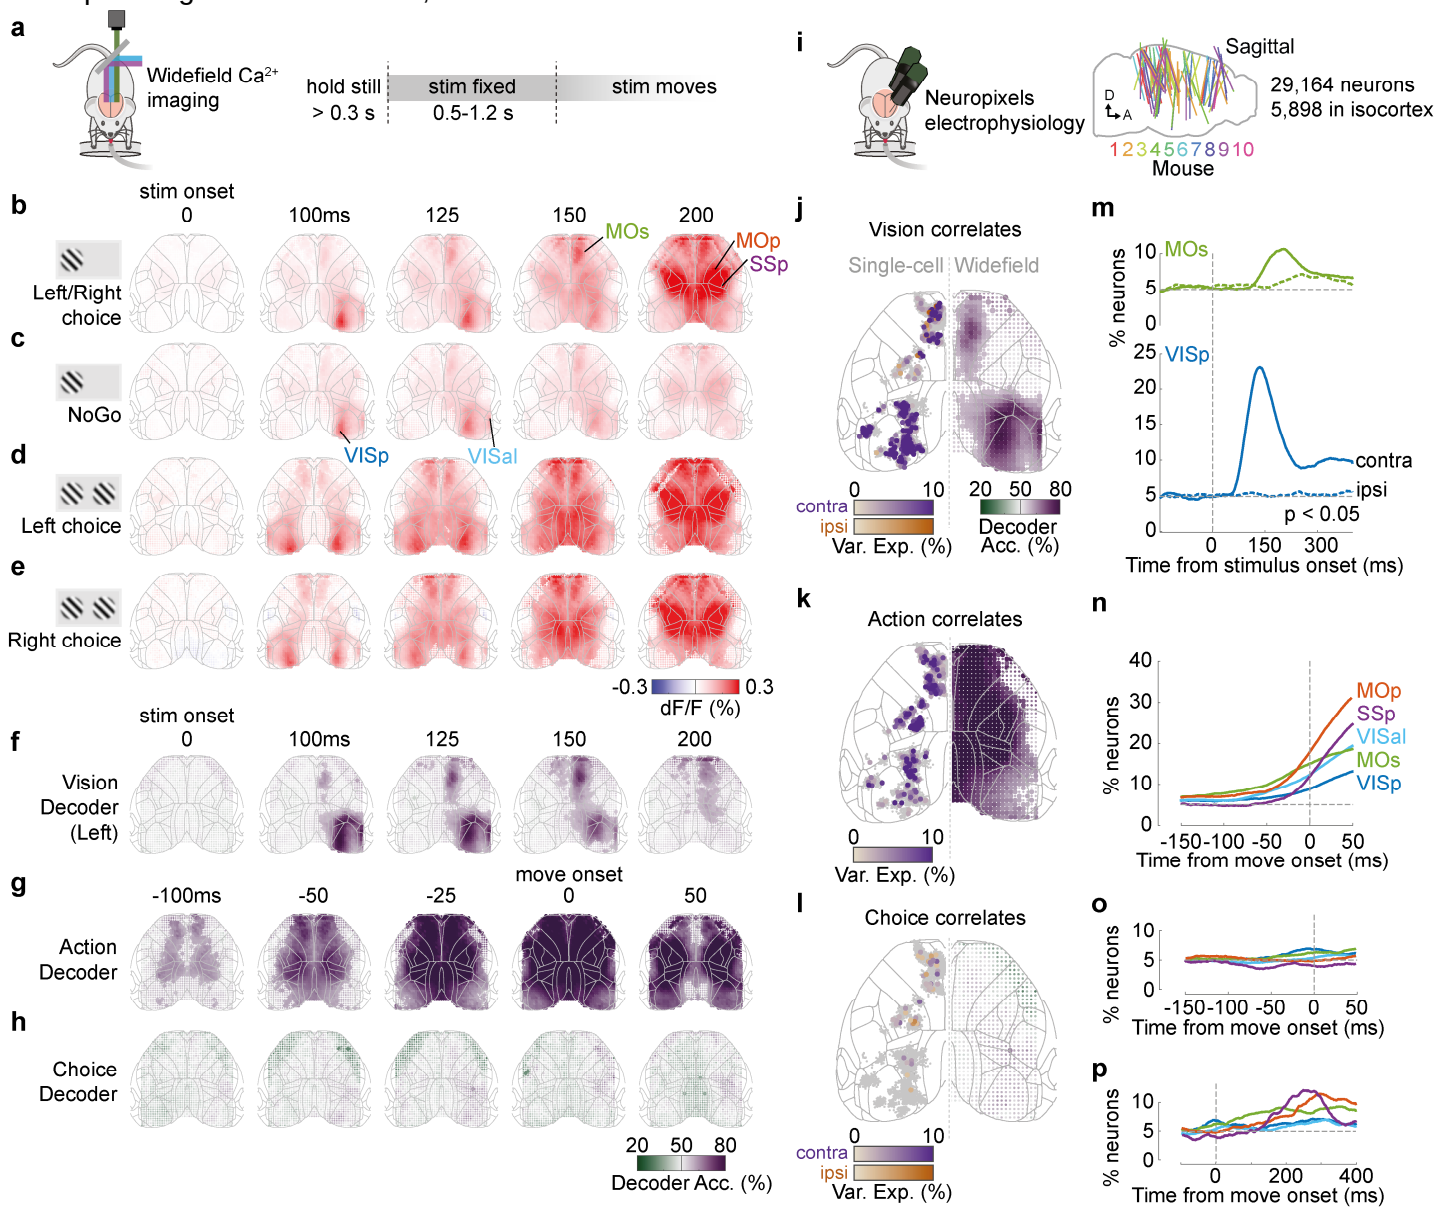

Figure 2. Activity related to stimuli and movement in the dorsal cortex.

(a) Schematic of widefield calcium imaging. Mice expressed GCaMP6 in excitatory neurons (n=7 mice) or all neurons (n=2). Right: task timeline. Stimuli were fixed in place for 0.5-1.2 s to dissociate wheel turns from visual motion in recorded activity. **(b)** Cortical fluorescence (dF/F) for successive times relative to stimulus onset, averaged across 39 sessions in 9 mice. Fluorescence is shown for trials with stimuli on the left (avg. 66% contrast) and a Left or Right choice. Gray lines: cortical areas defined from the Allen Common Coordinate Framework[30] (CCF; image cropped to this region). The image is built of dots (one for each image pixel), whose shade reflects dF/F and whose size reflects statistical significance (larger dot: $p<0.001$, nested ANOVA). **(c)** Average dF/F for the same stimulus conditions as in (b) but for Miss (NoGo) trials. **(d,e)** Average dF/F for trials with equal non-zero contrast on both sides and a Left (d) or Right (e) choice. **(f)** Average Vision decoder accuracy, decoding whether a stimulus was present on the left screen from each pixel of the widefield image (see Methods). Dot size reflects statistical significance as in (b). **(g,h)** Average Action decoder accuracy (g; Go vs NoGo) and Choice decoder accuracy (h; Left vs Right), aligned to movement onset. **(i)** Schematic of Neuropixels recordings. Right: sagittal map of probe insertion sites over 39 sessions in 10 mice (5,898 cortical neurons). Panels adapted with permission from Steinmetz et al., 2019. **(j)** Vision correlates. Comparison between contralateral Visual decoding maps obtained from widefield data (right hemisphere; 125 ms after stimulus onset) and a map of single neurons (left hemisphere jittered dots) that significantly encoded contralateral (purple) and ipsilateral (orange) stimuli based on kernel analysis. Dot

brightness for neurons scales with the cross-validated variance explained of the stimulus kernel. Gray dots indicate neurons with no significant stimulus kernel. **(k)** Action correlates. Same plotting scheme as in (j) but comparing widefield Action decoding (Go vs NoGo) at 25 ms before movement onset with the variance explained of a movement kernel. **(l)** Choice correlates. Same plotting scheme as in (j) but comparing widefield choice decoding at 25 ms before movement onset with the variance explained of a choice kernel for contralateral (purple) and ipsilateral (orange) choices. **(m)** Proportion of VISp (blue) and MOs (green) neurons with significant contralateral (solid line) and ipsilateral (dashed colored line) stimulus decoding (p<0.05; CCCP analysis). Dashed gray lines: stimulus onset (vertical) and false-positive rate (horizontal). **(n)** Proportion of neurons which show significant Action decoding (p<0.05; CCCP analysis) relative to movement onset. Colored lines indicate neurons grouped by CCF area. Vertical dashed line: movement onset. **(o)** Proportion of neurons with significant choice decoding (p<0.05; CCCP analysis). **(p)** Same as (o) but showing later times after movement onset.

Optogenetic inactivation of sensory-coding regions biases choice

To assess whether these cortical correlates of stimulus and action are necessary for task performance, we used optogenetic inactivation. We optogenetically inactivated each of 52 sites across the dorsal cortex by shining a blue laser (1.5 mW) in transgenic mice expressing Channelrhodopsin (ChR2) in Parvalbumin-expressing inhibitory interneurons[31], using transcranial laser scanning[21]. Each randomly-selected site was inactivated on ~1.4% of trials. Laser illumination started at the onset of visual stimuli and lasted until a choice (or NoGo) was registered. If the cortical correlates of sensory stimuli in VISp and MOs are necessary for performance, then inactivating these regions should bias the subject's choices away from contralateral stimuli. If the non-specific bilateral correlates of action in other regions such as MOp are necessary for performance, then inactivating these regions should cause an increase in NoGo trials.

Inactivation of each cortical region affected Left/Right choices in direct proportion to the region's encoding of sensory stimuli (Fig. 3a-c). We first focused on trials with equal contrast on both sides, to which animals responded with Left and Right choices (rewarded with 50% probability) and NoGo responses (Miss trials, unrewarded). Optogenetic inactivation affected the animals' choices when targeted at the visual cortex (VIS; light diffusion made it impossible to distinguish effects in individual visual areas[21,32]) and at frontal cortex (MOs; Fig. 3a), biasing choices away from stimuli contralateral to the inactivated hemisphere (-35%, p<0.0002 in VIS, -22%, p<0.0002 in MOs; permutation test; Supplementary Figure S4). The strength of the bias produced by inactivating a cortical location closely matched the accuracy with which that location's activity correlated with the presence of a contralateral stimulus (Fig. 3b,c).

In contrast, no matter how strongly a cortical region encoded upcoming action, inactivation caused no changes in NoGo responses (Fig 3d-f). Inactivation of primary motor and primary somatosensory areas - whose activity most strongly differed between Go and NoGo trials - had no effect on the proportion of Go vs. NoGo choices (Fig. 3d). Similar results were seen across the cortex, and we observed no relationship between the strength with which a region's activity measured by widefield imaging predicted action (Go vs NoGo choices), and the effect of inactivating that region on Go vs. NoGo behavior (Fig. 3e,f). Inactivation of primary motor cortex MOp did, however, reduce wheel peak velocity by a small amount for both Left and Right choices non-specifically (-21.2%, p<0.01, t-test; Supplementary Figure S5).

Consistent with their causal role in sensory encoding, visual and frontal cortex were most needed precisely at the time when they encoded sensory stimuli (Fig. 3g-i). To determine the times at which activity in different regions played a causal role, we inactivated VIS and MOs with a brief laser pulse (25 ms, 15 mW) at different delays from stimulus onset (Fig. 3g). Inactivation of VIS significantly impaired task performance around the time of stimu-

lus onset (-110 to +130 ms; Fig. 3h; Supplementary Figure S6). Inactivation of MOs impaired performance at a later time window (+52 to +174 ms). These windows preceded the time of these mice's choices at 273 ± 21 ms (median ± m.a.d. across 65 sessions, on laser-off trials). Consistent with the scanning inactivation results, pulsed inactivation in MOp produced no significant behavioral impairment at any time (Supplementary Figure S6a). It might seem paradoxical that inactivation in VIS prior to stimulus onset could affect behavior; however, pulse activation of inhibitory cells suppresses cortical activity for over 100 ms (Supplementary Figure S6c,f)[31]. The critical time for inactivation is thus the last moment at which inactivation affects behavior. Notably, this critical time for VIS and MOs mirrored the time of peak contralateral stimulus decoding in these regions as previously identified in electrophysiological recordings (Fig. 3i).

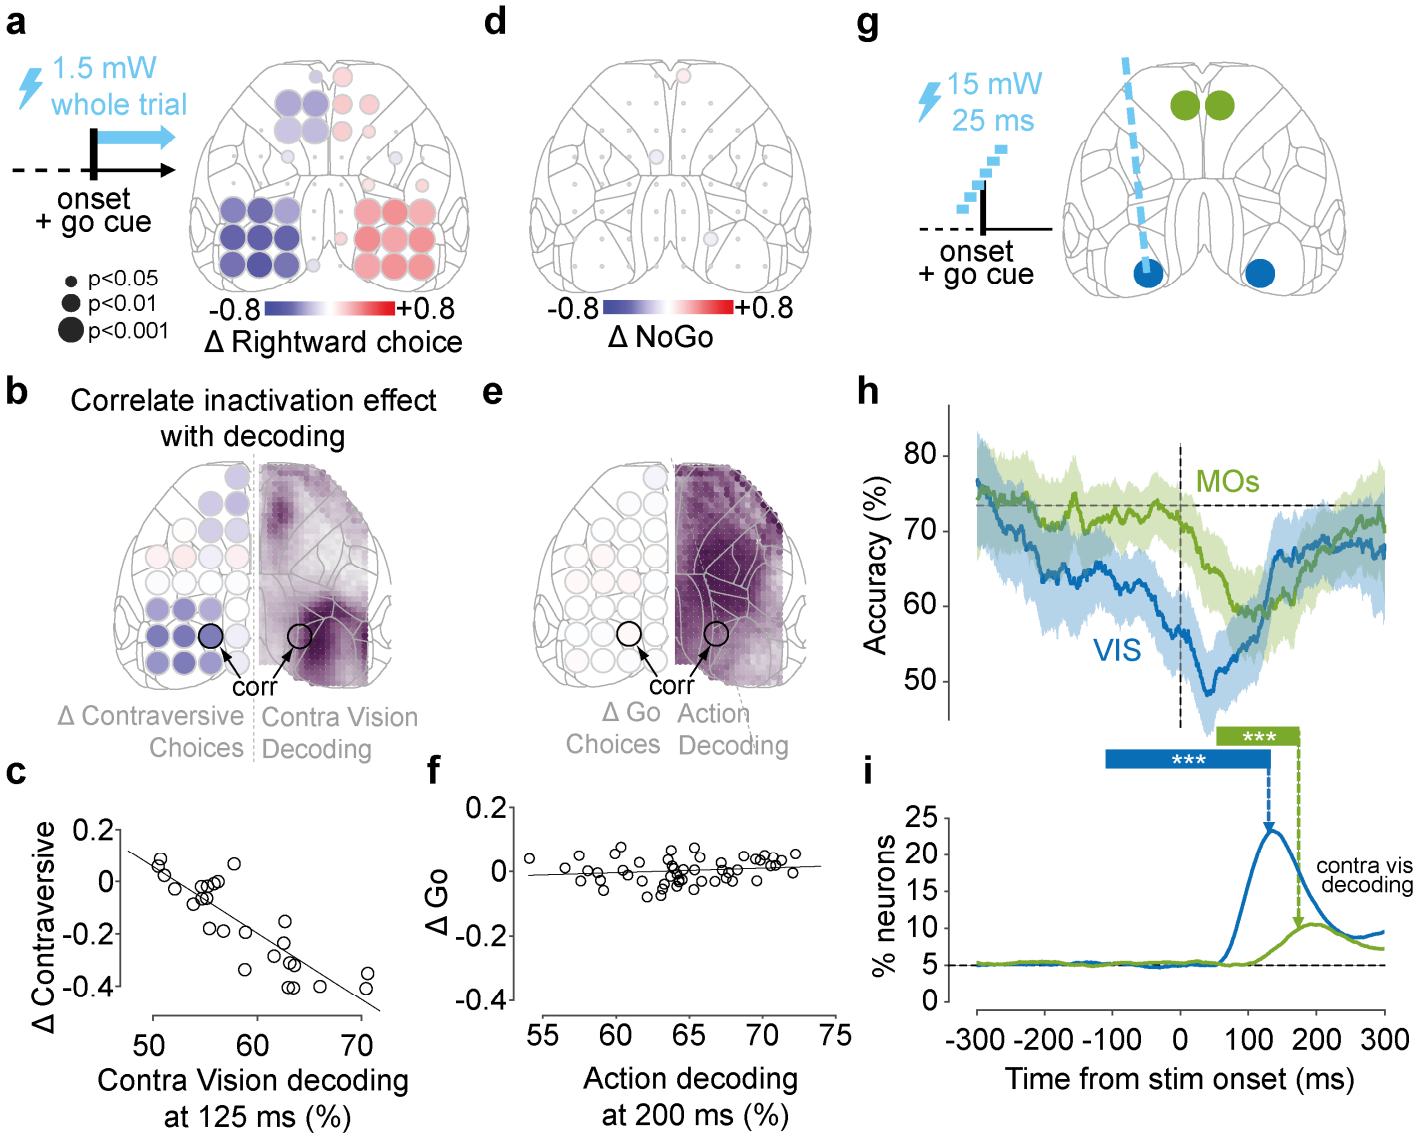

Figure 3. Optogenetic inactivation of sensory-coding regions biases choices.
(a) Scanning inactivation in mice expressing ChR2 in parvalbumin-expressing inhibitory neurons (91 sessions in 5 mice; see Methods). On ~75% of trials, a blue laser (1.5 mW, 40 Hz sine wave) illuminated one of 52 locations, from stimulus onset until a choice or NoGo was registered. Right**:** change in probability of rightward choices for 52 stimulation sites for trials with equal non-zero contrast on each side. Dot size indicates statistical significance (permutation test, see Methods). **(b)** Comparison of the maps for contralateral stimulus decoding (Right hemisphere, from Fig. 2f at 125 ms, dot size not scaled by significance) and for effect of inactivation (Left hemisphere; from panel (a) combining across hemispheres). **(c)**

Relationship between the choice bias caused by optogenetic inactivation across 26 sites, and the strength of contralateral stimulus decoding from widefield imaging at 125 ms post-stimulus, averaged over a region around the inactivated coordinate ('corr' in panel b). **(d)** Same as in (a) for the change in proportion of NoGo choices for equal non-zero contrast trials. **(e,f)** Same as in (b,c), for the relationship between Action (Go vs NoGo) decoding at 200 ms after stimulus onset and the change in Go (Left or Right) choices induced from inactivation. **(g)** Pulse inactivation experiments (65 sessions in 6 mice). On ~66% of trials, a 25 ms laser pulse was presented randomly between -300 and +300 ms relative to the stimulus onset (vertical dashed line), in one of four locations, chosen randomly (map on right). **(h)** Trial accuracy as a function of pulse time, averaged over trials where contralateral contrast > ipsilateral contrast. Blue curve: VIS inactivation; green curve: MOs inactivation; black horizontal dashed line: mean accuracy in no-laser trials. Curves show the average pooled over all sessions and mice, smoothed with a 100 ms boxcar window. Shaded regions: 95% binomial confidence intervals. *** indicate the intervals in which accuracy differs significantly from control trials (p<0.0001; χ2 test). **(i)** Time-course of contralateral Vision decoding (solid lines from Fig. 2m). Vertical dashed arrows indicate the end of the critical intervals in (h). Black horizontal dashed line indicates chance level.

A lawful relationship between cortical activity and decisions predicts the effects of inactivation

To formally relate neural activity to the mouse's behavior, we constructed a neurometric model, which predicts actions and choices based on neural activity. In contrast to the analyses above (Fig. 3), which focused on specific stimulus conditions (equal contrast), this model summarizes how cortical activity shapes the animals' decisions for all contrast combinations.

In the model, a hypothesized subcortical circuit performs weighted sums of population activity of visual and frontal cortex, to compute the evidence for a choice on each side (Fig. 4a). Because only the visual and frontal cortex had a causal role in choice (Fig. 3), we defined the decision variables as a weighted sum of activity in VISp and MOs, obtained from widefield imaging at the time this activity is causally relevant. The inputs to the model are thus the activities V_L and V_R of the left and right visual cortex (VISp) at time 75-125 ms after stimulus onset, and the activities M_L and M_R of the left and right frontal cortex (MOs) at time 125-175 ms. A weighted sum of these 4 variables yields the 'decision variables' representing the evidence for choices on the left and on the right:

$$Z_L = \alpha_L + v_c V_R + v_i V_L + m_c M_R + m_i M_L$$
$$Z_R = \alpha_R + v_i V_R + v_c V_L + m_i M_R + m_c M_L$$

where v_c and m_c represent weights from contralateral visual and frontal cortex, v_i and m_i represent weights from ipsilateral visual and frontal cortex, and α_L and α_R are constant intercept terms. The two decision variables determine the animal's behavior probabilistically, using a 3-way logistic ("softmax") function, where the probabilities of Left, Right, and NoGo responses are:

$$P(Left)/P(NoGo) = exp(Z_L)$$
$$P(Right)/P(NoGo) = exp(Z_R)$$
$$P(NoGo) = 1 - P(Left) - P(Right)$$

We fitted the parameters of this model using the data from widefield imaging, keeping the data from inactivations aside. To determine the activity variables, we calibrated widefield recordings against electrophysiologically-measured firing rates to provide an estimate of firing rate from the widefield fluorescence signal (see Methods; Supplementary Figure S7). To fit the weights we used a hierarchical Bayesian approach that allowed for variability in weights between subjects and recording sessions (Methods). The inactivation data was not used to constrain these weights, allowing us to use these data to later validate the model.

The estimated weights showed a fundamentally different structure for visual and frontal cortex (Fig. 4b). The weights on VISp activity were subtractive (v_c positive and v_i negative): activity on one side of the visual cortex promoted contraversive choices while suppressing ipsiversive choices. By

contrast, the weights on MOs activity were additive ($m_c > m_i$, but both positive): activity on one side of the frontal cortex promoted both contraversive and ipsiversive choices (albeit slightly favoring the contraversive side). Thus, the decision variables represented a subtraction of visual cortical activity between the two hemispheres, and an addition of frontal activity from both hemispheres.

The model captured the behavior of the mice not only during imaging sessions (which were used to fit the model) but also during inactivation sessions, predicting the effects of inactivation even though inactivation data were not used in fitting model parameters (Fig. 4c,d). During inactivation sessions, the model was able to capture the overall dependence on stimulus contrast of each type of choice: Correct choices, Incorrect choices, and Misses (non-laser trials; Fig. 4c; Supplementary Figure S8). To predict the results of optogenetic inactivation, we computed the model's prediction after setting the activity in a selected cortical region to zero. The model predicted a reduction in Correct choices and increase in Miss errors on inactivation of either VISp or MOs contralateral to a stimulus (Fig. 4d).

Intriguingly however, the model's predictions for VISp and MOs inactivation were not identical: due to the subtractive nature of the VISp weights, the model predicted that VISp inactivation should increase Incorrect choices, but that MOs inactivation should not. To test these predictions, we performed an additional optogenetic inactivation experiment in VIS and MOs (1.5 s duration, multiple laser powers; see Methods). The empirical observations in the proportion of Correct, Incorrect and Miss rates resembled the predictions from the neurometric model for VIS and MOs inactivation (Fig. 4d; Supplementary Figure S8). The neurometric model therefore provides a quantitative link between the activity observed in visual and frontal regions, the choices of the animal, and the effects of inactivation of those regions.

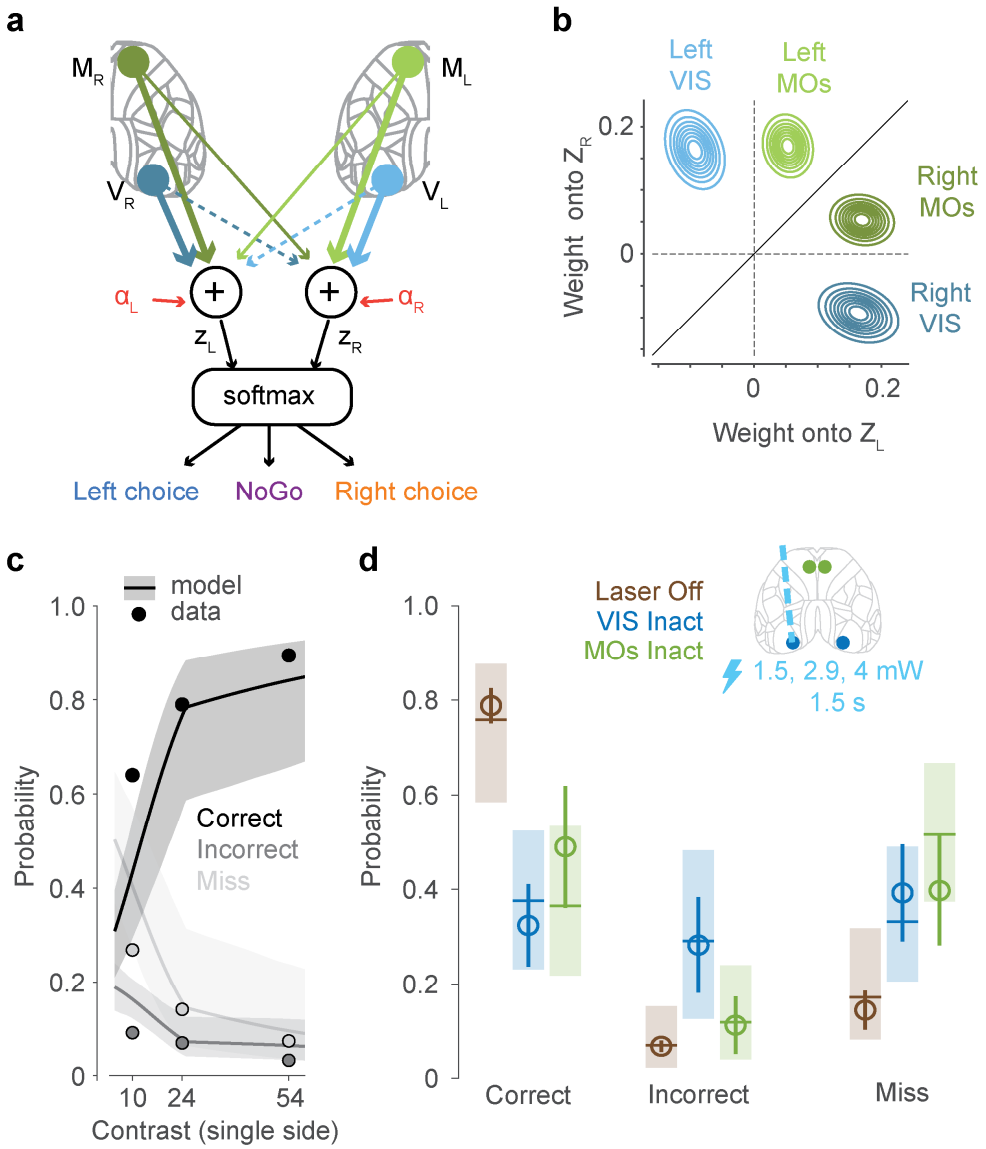

Figure 4. A lawful relationship between cortical activity and decisions predicts the effects of inactivation. **(a)** Schematic of the neurometric model. Activity in VISp and MOs is measured by widefield calcium imaging and used to estimate population firing rates on each trial (Supplementary Figure S7). A weighted sum of activity in both hemispheres determines decision variables Z_L and Z_R , and a softmax function generates the probability of each behavioral choice. **(b)** Posterior distributions of model weights from the four regions to the two decision variables. Contours show Gaussian fit to the posterior distribution over all experiments. Diagonal line illustrates the model weight symmetry. **(c)** Fit of the neurometric model to mouse choices (34 sessions in 5 mice, non-laser trials), using the same plotting scheme as in Fig 1h. **(d)** Model predicts effects of inactivation on Correct, Incorrect and Miss trials, averaging over unilateral stimulus conditions. Open circles show the empirical probability for trials with laser off (brown), and for trials with optogenetic inactivation (1.5 s duration, multiple laser powers; see Methods) of VIS (blue) and MOs (green), averaged across sessions. Error bars indicate the 95% confidence interval of the average estimate. Brown horizontal lines indicate the neurometric model fit to the average contrast value (~30%). Blue and green horizontal lines indicate mean predictions from the neurometric model, obtained by setting the activity of VISp (blue) and MOs (green) in the model to zero. Shaded regions indicate the 95% credible intervals of the model predictions.

Discussion

Our results reveal a tight link between the behavior of a mouse in a visual choice task, and the sensory, but not the motor, activity of the dorsal cortex. First, the visual stimuli elicited activity in the visual cortex (VIS) followed by frontal cortex (MOs). Second, activity all over the dorsal cortex – and particularly in primary motor cortex, MOp – correlated with action before movements in either direction. Only a few neurons in frontal cortex showed correlates of choice between the two directions. Optogenetic inactivation showed that the localized cortical correlates of vision were causal, but the widespread correlates of action were not. Inactivation of VIS and MOs biased choices away from the contralateral stimulus, but inactivating MOp barely had any effect. The behavioral effects of inactivating a cortical region scaled linearly with the degree to which that region's activity correlated with sensory information, and not at all with the degree to which it correlated with motor information. Moreover, the behavioral effects of inactivation matched the predictions of a neurometric model trained to compute choice based on cortical activity. These results suggest that the main contribution of dorsal cortex in this task is that VIS and MOs transfer sensory information to downstream circuits that select the animal's choice. Other areas such as MOp make a minor contribution to movement execution.

Our data confirm the critical role of rodent frontal cortex (MOs) in sensory decisions but indicate that at least in this task, MOs' role is primarily to route sensory information to other circuits, which feed it back a decision only once the movement is underway. MOs is thought to be critical for decisions based on sensory input[33], from diverse modalities, such as somatosensation[21,34], audition[20,35], olfaction[13,36], and vision[37]. Our data are consistent with this hypothesis, and suggest a primarily sensory role in this task. Widefield imaging revealed no difference in MOs activity between trials of identical stimulus conditions but different choices (Left or Right). Electrophysiology indicated that single neurons which could predict Left vs Right choices prior to movement onset were much rarer than neurons carrying a sensory signal, although choice direction could be decoded from more MOs neurons after the action had started. Inactivation of MOs impaired behavior in a manner predicted by a neurometric model in which MOs contributed weighted sensory information to a downstream decision circuit. While inactivation of MOs could in principle affect behavior through off-target effects in VISp[38], this is unlikely because MOs inactivation maximally impairs behavior ~50 ms after the time when inactivation of VISp ceases to have any effect.

The rarity of pre-movement choice coding in dorsal cortex suggests that at least in this task, the subject's choice is largely determined by subcortical circuits. Our neurometric model accurately predicted the subject's average choices based on cortical activity, but it assumed that the choice on each trial was determined stochastically by a circuit downstream of dorsal cortex. Indeed, previous electrophysiological recordings in the same task identified neurons with strong choice signals in a variety of subcortical structures, including basal ganglia, superior colliculus and zona incerta[6]. (Although cortical areas not accessible to widefield imaging could also in principle carry strong choice signals, previous electrophysiology did not find choice signals in any recorded non-dorsal cortical areas[6].) The subcortical circuits carrying choice signals feed back to cortex, so it is certainly possible that a small number of cortical neurons participate in a distributed recurrent circuit that determines the animal's choice. However, even in MOs, encoding of the subject's choice was sparse prior to the choice and only became widespread after the action had already started. Thus, if MOs does play a role in selecting the subject's choice, the neurons required for this would be substantially fewer in number than the neurons encoding the sensory stimulus. Moreover, while our data show that VIS plays an important role in routing sensory information to decision structures, they do not constrain whether this information flows directly from VIS to subcortical targets, or indirectly via MOs.

The limited choice signals observable in MOs in this task stand in contrast to several previous studies[10,21,34,36]. We suggest that the reason for this difference is that these tasks all involved a delay period in between presentation of the sensory stimulus and execution of the action. The more widespread, sustained choice activity seen in MOs during these delay tasks might therefore reflect storage of a choice in memory, rather than formation of the choice itself.

The strongest and most widespread cortical correlate of behavior was with the upcoming action (Go vs. NoGo), a correlate that may reflect a corollary discharge or the generation of movements irrelevant to ask performance. For instance, primary motor cortex MOp showed the strongest correlation with upcoming action, but this correlation did not correspond to causation: inactivating MOp had no effect on the animals' choices, causing only a slight decrease in wheel velocity. Therefore, these strong action correlates are not involved in the decision to choose Go over NoGo. Our present data cannot rule out that unilateral MOp/SSp inactivation engages mechanisms that compensate for its effect on choices[39], or that a distributed cortical representation of action is necessary for movement in this task, and our localized inactivations did not sufficiently disrupt this distributed representation. Further studies incorporating simultaneous multi-region inactivation would be required to test this. Nevertheless the most parsimonious explanation for these results is that MOp does not causally contribute to Go/NoGo choices. Our data are thus consistent with previous studies showing that neural circuits exhibiting correlates of choice or action are not always necessary for it[18,20,40].

What then might be the function of the strong MOp/SSp activity observed prior to action execution? Several possibilities remain. First, this activity might contribute to finely-detailed execution of the ongoing movement. Indeed, some studies have demonstrated a causal role for MOp areas in dexterous fine forepaw movements[25,41,42]. It is possible that inactivation of MOp/SSp would perturb these fine details of paw configuration in a manner that our current videographic methods cannot detect, while sparing the control of more proximal muscles that suffice to make a wheel turn. Second, these regions might contribute to movements irrelevant to task performance, such as postural adjustments, fidgeting or whisking. Third, they may reflect efference copy or corollary discharge[26,43] from other circuits involved in producing the choice. The function of such corollary discharge is unclear, however it is not restricted to choice tasks, as even spontaneous movements increase neuronal activity across the brain[16]. We speculate that this increased activity might serve to improve neural coding capacity throughout the brain by boosting spike counts and lowering time constants[44].

The fact that the non-specific action correlates are so large and widespread highlights the difficulties of experimentally identifying neural correlates of choice in Go/NoGo tasks. Several studies have shown that in Go/NoGo tasks, neurons of many regions, including sensory cortex, change their activity prior to the movement[8,13,45–47]. Given the present results, and our recordings across the brain in this task[6], it is not surprising that correlates of action are found in sensory cortex: they are found everywhere. However, as the inactivation of MOp in the present task makes clear, we should not generally expect brain activity that predicts an upcoming Go/NoGo choice to play a causal role in that choice.

In conclusion, the dorsal cortex's involvement in a sensory discrimination task appears primarily to relay sensory evidence to a downstream circuit that determines behavior based on this evidence. While activity all over cortex reflects the occurrence of an impending action, this activity is not locally causal. Instead, the causal effect of a cortical region depends on whether that region's activity encodes sensory stimuli. Modeling suggests that a stochastic circuit downstream of cortex might have also learned to optimally weigh cortical activity to produce the required behavior. The location of this downstream circuit is currently unconstrained, but may include subcortical structures with choice correlates (such as basal ganglia and midbrain), non-

dorsal cortical structures, as well as the rare neurons in frontal cortex that encode choice. Testing the causal relevance of these circuits will require further targeted inactivation experiments.

Acknowledgements

We thank Michael Krumin for assistance with the experimental setup; Charu Reddy, Miles Wells, Laura Funnell and Hamish Forrest for help with mouse husbandry and training; Pip Coen, Kevin Miller, Hamish Forrest and Andy Peters for feedback on earlier forms of the manuscript. This work was supported by the Wellcome Trust (grants 095668 and 095669 to M.C. and K.D.H.), the Engineering and Physical Sciences Research Council (CoMPLEX PhD studentship to P.Z-H.), the Human Frontiers Science Program (Fellowship LT001071, to N.A.S.), and the European Union's Horizon 2020 research and innovation programme (Marie Sklodowska-Curie fellowship 656528 to N.A.S.). M.C. holds the GlaxoSmithKline/Fight for Sight Chair in Visual Neuroscience.

Author contributions

P.Z-H., N.A.S., M.C. and K.D.H. conceived of and designed the study. P.Z-H. and N.A.S. collected and analyzed data. P.Z-H., N.A.S., M.C. and K.D.H. wrote the manuscript. Correspondence and material requests should be directed to Kenneth D. Harris, kenneth.harris@ucl.ac.uk.

Competing Interests

The authors declare no competing interests.

Methods

All experimental procedures were conducted at UCL according to the UK Animals Scientific Procedures Act (1986) and under personal and project licenses granted by the Home Office following appropriate ethics review.

Mouse transgenic lines

For the widefield calcium imaging, we used transgenic mice expressing GCaMP6s in excitatory neurons (tetO-G6s[Jax #024742, RRID:IMSR_JAX: 024742] x CaMK2a-tTA[Jax #007004, RRID:IMSR_JAX: 007004]), GCaMP6f in excitatory neurons (Ai95[Jax #024105, RRID:IMSR_JAX: 024105] x VGlut1-Cre[Jax#023527, RRID:IMSR_JAX: 023527]), or GCaMP6s in all neurons (Snap25-GCaMP6s[Jax #025111, RRID:IMSR_JAX: 025111]). For the optogenetic inactivation experiments, we used transgenic mice expressing ChR2 in Parvalbumin-positive inhibitory interneurons (Ai32[Jax #012569, RRID:IMSR_JAX: 012569] x PV-Cre[Jax #008069, RRID:IMSR_JAX: 008069]). For the Neuropixels electrophysiology experiments, mice of multiple genotypes were used[6]. All mice were 10-73 weeks of age at the time of data collection. Mouse genotype and session counts for each Figure are detailed in Supplementary Table 1.

Surgery

For widefield imaging, optogenetic inactivation and electrophysiological recording experiments, mice were prepared with a clear skull cap similar to that of Guo et al. (2014) and described previously[6,29]. The implantation surgery proceeded as follows. The dorsal surface of the skull was cleared of skin and periosteum, and the junction between cut skin and skull was sealed with cyanoacrylate. The exposed skull was prepared with a brief application of green activator to ensure strong connection between cement and bone (Super-Bond C&B, Sun Medical Co, Ltd, Japan). The junction between skin and skull was again covered, using dental cement (Super-Bond C&B). In most cases, a 3D printed 'cone' was attached to the head with cyanoacrylate and dental cement at this stage, surrounding the exposed skull and providing light isolation. A thin layer of cyanoacrylate was applied to the skull and allowed to dry. Two to four thin layers of UV-curing optical glue (Norland Optical Adhesives #81, Norland Products Inc., Cranbury, NJ; from ThorLabs) were applied to the skull and cured (~10 s per layer) until the exposed skull was covered (thin layers were used to prevent excessive heat production). A head-plate was attached to the skull over the interparietal bone with SuperBond polymer.

Behavioral task

Apparatus: The two-alternative unforced choice task design was described previously[29]. In this task, mice were seated on a plastic apparatus with forepaws on a rotating wheel, and were surrounded by three computer screens (Adafruit, LP097QX1) at right angles covering 270 x 70 degrees of visual angle (d.v.a.). Each screen was ~11 cm from the mouse's eyes at its nearest point and refreshed at 60 Hz. The screens were fitted with Fresnel lenses (Wuxi Bohai Optics, BHPA220-2-5) to ameliorate reductions in luminance and contrast at larger viewing angles near their edges, and these lenses were coated with scattering window film ("frostbite", The Window Film Company) to reduce reflections. The wheel was a ridged rubber Lego wheel affixed to a rotary encoder (Kubler 05.2400.1122.0360). A plastic tube for delivery of water rewards was placed near the subject's mouth. For full details of the experimental apparatus including detailed parts list see www.ucl.ac.uk/cortexlab/tools/wheel. All behavioral experiments were run using MATLAB Rigbox[48].

Pre-stimulus quiescence: For three experiments, trials began after a period of no wheel movement (widefield imaging: 0.3-0.7 s, 52-coordinate inactivation experiment: 0.2-0.6 s, Neuropixels electrophysiology 0.2-0.5 s). For all

other behavioral sessions, there was no constraint however trials were excluded post-hoc if wheel movement was detected -0.15 to +0.05 s from stimulus onset.

Stimulus onset: At trial initiation, a visual stimulus was presented on the left, right, both, or neither screen. The stimulus was a Gabor patch with orientation 45 degrees, sigma 9 d.v.a., and spatial frequency 0.1 cycles/degree. The grating stimuli on the left and right screens displayed at all combinations of four contrast levels, totaling 16 contrast conditions. The proportion of trials of each stimulus type were weighted towards easy trials (high contrast vs zero, high vs low, medium vs zero, and no-stimulus trials) to encourage high overall reward rates and sustained motivation. Zero contrast trials made up ~31% of trials; High contrast and Medium contrast single-side contrast trials made up ~24%; equal contrast trials made up ~12% of trials; all other comparison contrast trials uniformly comprised the remaining trial types. For all experiments except for widefield imaging and Neuropixels electrophysiology (see 'open-loop period' below), the onset of the visual stimulus also coincides with the onset of an auditory 'go cue' (12 kHz tone, 100 ms duration), marking the time at which the mouse can respond.

Wheel movements: Wheel turns in which the top surface of the wheel was moved to the subject's right led to rightward movements of stimuli on the screen, i.e. a stimulus on the subject's left moved towards the central screen. Put another way, clockwise turns of the wheel, from the perspective of the mouse, led to clockwise movement of the stimuli around the subject. A left or right choice was registered when the wheel was turned by an amount sufficient to move the visual stimuli by 90 d.v.a. in either direction. Movement onset time ('reaction time') is defined as time of the earliest detected wheel movement using the findWheelMoves3 algorithm[6] (https://github.com/cortexlab/wheelAnalysis/blob/master/+wheel/findWheelMoves3.m).

When at least one stimulus was presented, the subject was rewarded for driving the higher contrast visual stimulus to the central screen (if both stimuli had equal contrast, Left/Right choices were rewarded with 50% probability). When no stimuli were presented, the subject was rewarded if no turn (NoGo) was registered during the 1.5 s following the go cue.

Open-loop period: For widefield calcium imaging & Neuropixels electrophysiology sessions, after stimulus onset there was a random delay of 0.5-1.2 sec, during which time the subject could turn the wheel without penalty, but visual stimuli were locked in place and rewards could not be earned. The subjects nevertheless typically responded immediately to the stimulus onset, and trials were excluded if the initial wheel movement onset time was greater than 0.5 s. At the end of the delay interval, an auditory go cue was delivered (8 kHz pure tone for 0.2 sec) after which the visual stimulus position became coupled to movements of the wheel and a choice could be made. Initial wheel turns were nearly always in the same direction as the final choice (96.6% ± 3.4%, mean ± std across 39 sessions) indicating rare changes of mind[49] during the open-loop period. This small task modification was important to ensure that visual-stimulus-related cortical activity was not inter-mixed with activity related to the auditory go cue, and that movement-related activity was not inter-mixed with signals related to visual motion of the stimulus on the screen.

Feedback: Immediately following registration of a choice or expiry of the 1.5 s window, feedback was delivered. If correct, feedback was a water reward (0.7 – 2.5 µL) delivered by the opening of a valve on the water tube for a calibrated duration. If incorrect, feedback was a white noise sound played for 1 s. During the 1 s feedback period, the visual stimulus remained on the screen. After a subsequent inter-trial interval of 1 s (or 2 s for the 52-coordinate inactivation experiment), the mouse could initiate another trial

by again holding the wheel still for the prescribed duration.

Training: Mice were trained on this task with the following shaping protocol. First, high contrast stimuli (50 or 100%) were presented only on the left or the right, with an unlimited choice window, and repeating trial conditions following incorrect choices ('repeat on incorrect'). Once mice achieved high accuracy and initiated movements rapidly – approximately 70 or 80% performance on non-repeat trials, and with reaction times nearly all < 1 second – trials with no stimuli were introduced, again repeating on incorrect. Once subjects responded accurately on these trials (70 or 80% performance, at experimenter's discretion), lower contrast trials were introduced without repeat on incorrect. Finally, contrast comparison trials were introduced, starting with high vs low contrast, then high vs medium and medium vs low, then trials with equal contrast on both sides. The final proportion of trials presented was weighted towards easy trials (high contrast vs zero, high vs low, medium vs zero, and no-stimulus trials) to encourage high overall reward rates and sustained motivation.

Trial-exclusion: Trials were excluded from analyses based on several criteria. For all experiments, error trials were excluded if they represented the second (or more) consecutive error on easy trials (see 'repeat on incorrect' in previous section), as these errors likely reflect mouse disengagement. In addition, the first 5-10 trials of every session were excluded as these trials may include periods when the mouse is not settled into the task. For electrophysiological recording experiments, trials were excluded if they did not exhibit clear Left or Right wheel movements for Left and Right trials within 0.1 to 0.4 s from stimulus onset, or if they exhibited twitch movements for NoGo trials within -0.05 to 0.5 s from stimulus onset.

Widefield calcium imaging

Mice and apparatus: Imaging was performed in transgenic mice expressing GCaMP6 in excitatory neurons (tetO-G6s x CaMK2a-tTA; VGlut1-cre x Ai95) or all neurons (Snap25-GCaMP6s). For all widefield analyses, data was averaged across these mouse genotypes since all groups showed qualitatively similar responses. Aberrant epileptiform activity has not been observed in these mouse lines[50]. Details of the imaging have been described before[51] and are summarized here. We imaged using a macroscope (Scimedia THT-FLSP) with sCMOS camera (PCO Edge 5.5) and dual-wavelength illumination (Cairn OptoLED). The macroscope used a 1.0x condenser lens (Leica 10450028) and 0.63x objective lens (Leica 10450027). Images were acquired from the PCO Edge with ~10 ms exposures and 2 x 2 binning in rolling shutter mode. Images were acquired at 70 Hz, alternating between blue and violet illumination (35 Hz each). The light sources were 470 nm and 405 nm LEDs (Cairn OptoLED, P1110/002/000; P1105/405/LED, P1105/470/LED). Excitation light passed through excitation filters (blue: Semrock FF01-466/40-25; violet: Cairn DC/ET405/20x), and through a dichroic (425 nm; Chroma T425lpxr). Excitation light then went through 3 mm core optical fiber (Cairn P135/015/003) and reflected off another dichroic (495 nm; Semrock FF495- Di03-50x70) to the brain. Emitted light passed through the second dichroic and an emission filter (Edmunds 525/50-55 (86-963)) to the camera. Alternation was controlled with custom code on an Arduino Uno, and illumination was restricted to the 'global' phase of the rolling shutter exposures, i.e. only the times when all pixels of a frame were being exposed together.

Preprocessing: We de-noised the signal with singular value decomposition[52] and normalized the signal to the mean fluorescence at each pixel. The signal from the 405 nm illumination frames was used to correct for parts of the 470 nm signal that were due to changes in blood flow that obstruct the fluorescence signal[53] and the correction was performed with custom Matlab code (www.github.com/cortex-lab/widefield). We then low-pass filtered the signal at 8.5 Hz and applied a derivative filter to the fluorescence

trace to approximate deconvolution of the calcium sensor's time course from the underlying neural activity. Fluorescence was extracted on a grid centered at bregma for each mouse and session, making it possible to average activity across sessions/mice. When computing stimulus-aligned averages of the fluorescence, pre-stimulus baseline activity was removed, removing the impact of long-term trends.

ROI selection: Each mouse's cortical fluorescence map was aligned to the Allen Common Coordinate Framework (CCF) atlas[30]. This alignment was performed manually, by matching primary visual, secondary visual, and secondary motor areas to the corresponding hotspots of fluorescence following presentation of a contralateral stimulus. Single-pixel regions of interest (ROIs) for each cortical area were selected based on this atlas, and manually adjusted to allow for inter-mouse differences. VISp was selected as the peak of the most posterior-medial activated site in the visual cortex in response to a contralateral stimulus. VISal was selected as the center of VISal according to the Allen CCF. VISal was taken as an exemplary secondary visual cortical area because it was furthest from the part of VISp activated by our visual stimuli, ensuring minimal contamination of fluorescence between these two ROIs. The MOs ROI was selected as the most anterior site activated by the contralateral stimulus which was also within the CCF bounds for this region. MOp and SSp ROIs were selected within the cortical region active during wheel movements, positioned equidistant from the MOp-SSp border in the CCF.

Decoding analysis: To decode task variables from neural activity at different cortical regions and time-bins, we used a strategy that we have called 'combined conditions choice probability' (CCCP)[6], which measures how neural activity can predict one behavioral or sensory variable, while holding the others constant. We describe this method by explaining how it is used to predict Action (Go vs NoGo) for constant stimulus conditions. For each of the 16 possible stimulus conditions (contrast pairs), we compute a Mann-Whitney U statistic: the number of Go-NoGo trial pairs of this stimulus condition for which the Go trial had more activity than the NoGo trial. We then sum the 16 U statistics, sum the total number of trial pairs for each condition, and divide to obtain an auROC (area under receiver operating characteristic) value between 0 and 1. This auROC therefore quantifies how well a decoder could distinguish between Go and NoGo trials from neural activity, analogous to Choice Probability analysis[54] but controlling for different stimulus conditions.

For Vision decoding a similar strategy is used: to decode left contrast, we first divide trials into 12 groups corresponding to 3 possibilities for the animal's choice (Left, Right, NoGo) and 4 right stimulus contrasts. Within each group we calculate a U statistic: the number of trials, for which one trial had non-zero contrast on the left and the other had zero contrast, and the non-zero contrast trial had more activity. We sum these U-statistics across groups and divide by the total trial count to obtain an auROC value. For stimulus decoding on the right side, the same analysis is performed but swapping left for right contrast conditions. Contralateral and ipsilateral stimulus decoding is computed by combining the auROC values between hemispheres and hemifields. For choice decoding, NoGo trials are excluded, and a U statistic is computed between fluorescence for Left and Right choices, within each of the 16 stimulus conditions. The auROC value here gives an estimate for how well a decoder can discriminate Left from Right choices from neural activity, controlling for the stimulus contrast.

For the widefield dataset, decoding significance was measured for each coordinate and 25 ms time slice using a nested ANOVA across sessions and subjects (sessions nested within subjects), testing whether decoding performance was significantly different from 0.5. To compute Action decoding at movement-aligned time bins, NoGo trials were assigned a 'movement time' based on randomly sampling from the distribution of reaction times for correct Go trials.

Optogenetic inactivation

While mice performed the task, we optogenetically inactivated several cortical areas through the skull using a blue laser. For these experiments we utilized transgenic mice expressing ChR2 in Parvalbumin-expressing inhibitory interneurons (Ai32 x PV-Cre).

52-coordinate inactivation experiment: Unilateral inactivation was achieved by mounting a fiber-optic cable (50 µm) with collimating and focusing lenses attached, on a moving manipulator (Scientifica, Patch-Star). On every trial, custom code drove the manipulator to set the position of the fiber-optic cable to one of 52 different coordinates distributed across the cortex. Inactivation coordinates were defined stereotaxically from bregma. On ~75% of trials, the laser was switched on (473 nm, 1.5 mW, 40 Hz sine wave) to inactivate the cortical site. The laser dot was collimated and focused on the brain surface to ~100 µm radius, resulting in light power density of ~25 mW/mm^2 at the skull surface for the trial duration (peak power 1.5 mW / (pi * (0.1 mm)^2), delivered in a sine wave such that average power is 50% of maximum, with additional power attenuation through the skull[21]. Laser and non-laser trials, and the location of the cortical inactivation, was randomized. The duration of the laser was from visual stimulus onset, until a behavioral choice was made. The laser positioning was independent of laser power, so auditory noise from the manipulator did not predict inactivation.

For the majority of sessions, laser illumination was targeted uniformly at the 52 coordinates thereby inactivating each coordinate a handful of times (~1.4% of trials). This discouraged any adaptation effects that may occur on more frequent inactivation paradigms. However for 9 sessions, the set of inactivation coordinates was restricted to coordinates within anterior MOs (2 mm AP, ±1.5 mm ML; 2 mm AP, ±0.5 mm ML; 3 mm AP, ±0.5 mm ML). For 14 sessions, the set of coordinates was restricted fall between VISp and SSp (-1 mm AP, ±2.5 mm ML; -2 mm AP, ±2.5 mm ML; -3 mm AP, ±2.5 mm ML). For these restricted coordinate sessions, the proportion of inactivated trials was reduced to 50%.

Pulsed inactivation experiment: Unilateral pulsed inactivation was achieved using a pair of mirrors mounted on galvo motors to orient the laser (462 nm; collimated and focused to ~100um dot size on cortical surface) to different points on the skull. We also introduced improved light isolation to ensure no light could reflect from the skull surface and be seen by the mouse. For ~66% of randomly-interleaved trials, the laser was switched on for 25 ms (DC) at random times relative to stimulus onset (-300 to +300 ms drawn from a uniform distribution). Inactivation was targeted at visual areas (VIS; -4 mm AP, ±2 mm ML), secondary motor area (MOs; +2 mm AP, ±0.5 mm ML), and primary motor area (MOp; -0.5 mm AP, ±1 mm ML). The VIS coordinate was chosen stereotaxically as the medial part of VISp, corresponding to the part of VISp with retinotopy for the stimulus location. The MOs coordinate was chosen in anterior MOs, corresponding to the region which exhibited significant behavioral effects from the 52-coordinate inactivation experiment. The MOp coordinate was selected as the most posterior part of MOp from the CCF atlas, positioned as far as possible from MOs to reduce the possibility that MOp inactivation also silences MOs activity.

Mixed-power inactivation experiment: Unilateral inactivation was achieved using the same laser/galvo hardware as in the pulsed inactivation experiment. The mirrors oriented the laser at visual and secondary motor areas. For ~70% of randomly-interleaved trials, the laser was switched on from stimulus onset for 1.5 s fixed duration. Laser power was chosen randomly between 1.5, 2.9 and 4.2 mW. For subsequent analyses, trials using different laser powers were pooled together.

Electrophysiological recordings

The Neuropixels electrophysiological dataset was described previously[6]. Methods relating to electrophysiological recordings have been detailed in this work, and are reproduced here.

Hardware: Recordings were made using Neuropixels electrode arrays[55]. Probes were mounted to a custom 3D-printed PLA piece and affixed to a steel rod held by a micromanipulator (uMP-4, Sensapex Inc.). To allow later track localization, prior to insertion probes were coated with a solution of DiI (ThermoFisher Vybrant V22888 or V22885) by holding 2 µL in a droplet on the end of a micropipette and touching the droplet to the probe shank, letting it dry, and repeating until the droplet was gone, after which the probe appeared pink.

Procedure: On the day of recording or within two days before, mice were briefly anaesthetized with isoflurane while one or more craniotomies were made, either with a dental drill or a biopsy punch. The craniotomies for VISp were targeted in some cases using measured retinotopic maps in the same mice, and in other cases to the same position stereotaxically (-4 mm AP, 1.7 mm ML, left hemisphere). The craniotomies for MOs were targeted stereotaxically (+2 mm AP, 0.5 mm ML, left hemisphere), to match the coordinates with strong inactivation effects. After at least three hours of recovery, mice were head-fixed in the setup. Probes had a soldered connection to short external reference to ground; the ground connection at the headstage was subsequently connected to an Ag/AgCl wire positioned on the skull. The craniotomies as well as the wire were covered with saline-based agar. The agar was covered with silicone oil to prevent drying. In some experiments a saline bath was used rather than agar. Probes were advanced through the agar and through the dura, then lowered to their final position at ~10 µm/s. Electrodes were allowed to settle for ~15 min before starting recording. Recordings were made in external reference mode with LFP gain = 250 and AP gain = 500. Recordings were repeated on multiple subsequent days. All recordings were made in the left hemisphere.

Preprocessing: The data were automatically spike sorted with Kilosort[56] (www.github.com/cortex-lab/Kilosort) and then manually curated with the 'phy' gui (www.github.com/kwikteam/phy). Extracellular voltage traces were preprocessed using common-average referencing: subtracting each channel's median to remove baseline offsets, then subtracting the median across all channels at each time point to remove artifacts. During manual curation, each set of events ('unit') detected by a particular template was inspected and if the spikes assigned to the unit resembled noise (zero or near-zero amplitude; non-physiological waveform shape or pattern of activity across channels), the unit was discarded. Units containing low-amplitude spikes, spikes with inconsistent waveform shapes, and/or refractory period contamination were labeled as 'multi-unit activity' and not included for further analysis. Finally, each unit was compared to similar, spatially neighboring units to determine whether they should be merged, based on spike waveform similarity, drift patterns, or cross-correlogram features. For calculating event-triggered averages and decoding performance, spike counts were binned with a 1 ms window and then smoothed with a 25 ms causal Gaussian filter.

Decoding analysis: Action, Vision and Choice were decoded from electrophysiological recordings using the same CCCP analysis outlined above for widefield decoding, with one exception: since single-neuron data could not be averaged across sessions, statistical significance could not be assessed by comparing decoding performance across sessions. Instead, decoding performance as measured by auROC was compared to a shuffled null distribution of auROC values obtained by shuffling the trial labels 2000 times (e.g. shuffling Go vs NoGo trial labels for Action decoding).

Kernel analysis: The cross-validated variance of each neuron's firing rate explained by Vision, Action, or Choice was determined using a reduced-rank kernel regression method detailed previously[6]. In brief, a predictor matrix was formed using the times and identities of stimuli and actions. Reduced-rank regression between the predictor matrices and the binned spike counts was performed across all recordings simultaneously. The variance explained by the regression was assessed with cross-validation. Vision/Action/Choice kernels were deemed significant if their cross-validated variance explained exceeded 2%. The false positive rate for this threshold was found with a shuffle test to be 0.33%.

Psychometric model

We modeled probabilistic choice behavior using a multinomial logistic regression described previously[29], but with the addition of a hierarchical Bayesian framework to account for inter-subject and inter-session variability in the model parameters.

For each trial i, the probability ratio of Left vs NoGo and Right vs NoGo choices are set by the exponential function of two decision variables $Z_L^{(i)}$ and $Z_R^{(i)}$:

$$P^{(i)}(Left)/P^{(i)}(NoGo) = exp(Z_L^{(i)})$$
$$P^{(i)}(Right)/P^{(i)}(NoGo) = exp(Z_R^{(i)})$$
$$P^{(i)}(NoGo) = 1 - P^{(i)}(Left) - P^{(i)}(Right)$$

Choices $y^{(i)} \in (Left, Right, NoGo)$ are drawn from a categorical probability distribution with these parameters:

$$y^{(i)} \sim Cat(P^{(i)}(Left), P^{(i)}(Right), P^{(i)}(NoGo))$$

To capture the behavioral dependence on choice bias and stimulus sensitivity, the decision variables are modelled as a saturating non-linear transformation of stimulus contrast. We denote the session number in which trial i occurred as $d[i]$, and the subject (mouse) number that performed the session as $m[d]$. The decision variables depended on parameters which varied between sessions (and thus also between subjects) according to the following formulae:

$$Z_L^{(i)} = b_L^{d[i]} + s_L^{d[i]} . (c_L^{(i)})^{n^{d[i]}}$$
$$Z_R^{(i)} = b_R^{d[i]} + s_R^{d[i]} . (c_R^{(i)})^{n^{d[i]}}$$

Here, $b_L^{d[i]}$ and $b_R^{d[i]}$ are bias parameters, which capture stimulus-independent choice behavior in session d, while $s_L^{d[i]}$ and $s_R^{d[i]}$ are session-dependent sensitivity parameters scaling the visual input on the left and right side. The visual input consists of the contrast presented on the left ($c_L^{(i)}$) and right ($c_R^{(i)}$), raised to a session-dependent exponent parameter $0 \leq n^{d[i]} \leq 1$ to allow for a saturating non-linear contrast transformation.

To capture how these parameters vary across sessions and subjects, we expanded the model to incorporate a hierarchical prior on sessions and subjects. Let θ_d be a 5-element vector containing the 5 session-specific parameters stated above, $\theta_d = [b_L^{(d)}, b_R^{(d)}, s_L^{(d)}, s_R^{(d)}, n^{(d)}]$. We model each session's parameter vector θ_d as drawn from a multivariate Gaussian distribution whose mean $\underline{\theta}_{m[d]}$ depends on the subject, with a common covariance matrix Σ,

$$\theta_d \sim N(\underline{\theta}_{m[d]}, \Sigma)$$

The covariance matrix Σ is given the following prior. It is first converted to a correlation matrix, which is given a $LKJ(2)$ prior[57] to penalize large positive or negative parameter correlations across sessions. The standard deviation terms for each parameter are given a $HalfCauchy(0,1)$ prior to penalize large variability in each parameter across sessions.

The subject-level mean vector $\underline{\theta_m}$ is drawn from a Gaussian grand-average mean $\underline{\theta^*}$, with a covariance matrix Σ^* which quantifies covariation in the parameters across subjects,

$$\underline{\theta_m} \sim N(\underline{\theta^*}, \Sigma^*)$$

The covariance matrix Σ^* is given the same prior as Σ. Finally, the grand-average parameters are given a weakly-informative Gaussian hyperprior with mean $[0,0,5,5,0.5]$, variances $[2,2,2,2,0.25]^2$ and covariances all zero.

The full joint posterior distribution of all parameters was numerically estimated with Hamiltonian Monte Carlo (No-U-Turn) sampling, using the Stan language[58]. Sampling was performed in 4 chains, each with 500 warm up iterations and 500 sampling iterations. The samples were checked manually to ensure convergence within and between chains. The posterior prediction/fit from the model is constructed by computing the model prediction from each of the posterior distribution samples, and then computing the mean and 95% credible intervals on the prediction across samples. All model predictions shown in the figures use the grand-average parameter posterior $\underline{\theta^*}$, unless specified otherwise.

Neurometric model

The neurometric model predicts the animal's choice from the activity of left VISp, right VISp, left MOs and right MOs ROIs. Like the psychometric model, it has two decision variables, which are determined on each trial from a weighted sum of activity in the four cortical areas:

$$Z_L^{(i)} = \alpha_L^{d[i]} + f^{(i)}.w_L^{d[i]}$$
$$Z_R^{(i)} = \alpha_R^{d[i]} + f^{(i)}.w_R^{d[i]}$$

The decision variable are thus a sum of offsets $\alpha_L^{d[i]}$ and $\alpha_R^{d[i]}$, and the inner product between two 4x1 session-dependent weight vectors $w_L^{d[i]}$ and $w_R^{d[i]}$, and a vector $\boldsymbol{f}^{(i)}$ containing the estimated population firing rate in the four regions on trial i. The weights and offset parameters are given a hierarchical prior allowing for variation between sessions and subjects similar to in the psychometric model:

$$\underline{\alpha_L^*}, \underline{\alpha_R^*} \sim N(0, 4^2)$$
$$\underline{w_L^*}, \underline{w_R^*} \sim N(0, I \times 4^2)$$

The estimated population rate $\boldsymbol{f}^{(i)}$ is derived from widefield calcium imaging at the four ROIs (see Widefield Methods for detail). In subsequent analyses, $\boldsymbol{f}^{(i)}$ will be set to zero to generate behavioral predictions for cortical inactivation. However widefield dF/F data has no meaningful baseline due to background fluorescence. Therefore, prior to fitting the model, widefield dF/F data is calibrated to electrophysiologically measured firing rates. This calibration was achieved by recording extracellular spiking activity in VISp and MOs using Neuropixels probes in separate sessions, and computing trial-averaged firing rates for each of the contrast conditions over a time window (Supplementary Figure S7; VISp: 75-125 ms, MOs: 125-175 ms). Calcium fluorescence was also averaged over the same windows but 30 ms later to allow for slower GCaMP6 kinetics. The transformation from widefield fluorescence to firing rate was computed by simple linear regression over the 16 contrast conditions. This linear transformation was then applied to the fluorescence value for each individual trial, thereby providing a population firing rate estimate for the four cortical regions on every trial. To improve fit stability, parameter symmetry was enforced between the left and right hemispheres (e.g. the weight of left VISp onto Z_L was the same as right VISp onto Z_R).

To generate predictions from the neurometric model for the effects of optogenetic inactivation, the model was modified in two ways. Firstly, since neural activity in VISp and MOs was not measured during optogenetic inactivation ses-

sions, the trial-by-trial activity in $\boldsymbol{f}^{(i)}$ was replaced with the trial-averaged firing rate measured electrophysiologically for each contrast condition. Secondly, since the overall tendency to NoGo differed idiosyncratically between widefield imaging and optogenetic inactivation sessions, the model offset parameters (α_L and α_R) were re-fit to the non-laser trials contained within optogenetic inactivation sessions (Mixed-power inactivation experiment).

To simulate the effect of optogenetic inactivation of a single cortical area, one element of the firing rate vector $\boldsymbol{f}$ was set to zero. This effect propagates forward through the model based on the fixed weights estimated from the widefield sessions, thereby affecting the decision variables and the probability associated with each choice. Importantly, the behavioral prediction obtained from the model when simulating inactivation did not depend on any empirical data involving actual optogenetic inactivation. In this sense, the neurometric model predicts behavior in a new dataset on which it was not fit.

Statistical tests

For the 52-coordinate inactivation experiment (Fig. 3a), statistical significance of the inactivation effect was assessed using a permutation test. The test statistic used was the difference in the proportion of a specific choice type, between laser and non-laser off trials (on trials with equal left and right contrast) for each of the inactivated coordinates, averaged across sessions. The null distribution of the test statistic was computed by repeated shuffling of laser and non-laser trial identities within each session. All other statistical tests are specified in the main text.

Code & Data availability

The code used in the current study is available from the corresponding authors on reasonable request. The datasets generated and/or analyzed during the current study are available from the corresponding authors on reasonable request.

References

1. Hernández, A. *et al.* Decoding a Perceptual Decision Process across Cortex. *Neuron* **66**, 300–314 (2010).
2. Hirokawa, J., Vaughan, A., Masset, P., Ott, T. & Kepecs, A. Frontal cortex neuron types categorically encode single decision variables. *Nature* 1–6 (2019) doi:10.1038/s41586-019-1816-9.
3. Park, I. M., Meister, M. L. R., Huk, A. C. & Pillow, J. W. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat. Neurosci.* **17**, 1395–1403 (2014).
4. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
5. 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).
6. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* 1–8 (2019) doi:10.1038/s41586-019-1787-x.

7. Wei, Z., Inagaki, H., Li, N., Svoboda, K. & Druckmann, S. An orderly single-trial organization of population dynamics in premotor cortex predicts behavioral variability. *Nat. Commun.* **10**, 1–14 (2019).

8. Yang, H., Kwon, S. E., Severson, K. S. & O'Connor, D. H. Origins of choice-related activity in mouse somatosensory cortex. *Nat. Neurosci.* **19**, (2016).

9. Pho, G. N., Goard, M. J., Woodson, J., Crawford, B. & Sur, M. Task-dependent representations of stimulus and choice in mouse parietal cortex. *Nat. Commun.* **9**, 2596 (2018).

10. Goard, M. J., Pho, G. N., Woodson, J. & Sur, M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife* **5**, e13764 (2016).

11. Freedman, D. J. & Assad, J. A. Experience-dependent representation of visual categories in parietal cortex. *Nature* **443**, 85–88 (2006).

12. Ahrens, M. B. *et al.* Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471–477 (2012).

13. Allen, W. E. *et al.* Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* **94**, 891-907.e6 (2017).

14. 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).

15. Salkoff, D. B., Zagha, E., McCarthy, E. & McCormick, D. A. Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task. *Cereb. Cortex* **30**, 421–437 (2020).

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

17. Pinto, L. *et al.* Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* (2019) doi:10.1016/j.neuron.2019.08.025.

18. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).

19. Le Merre, P. *et al.* Reward-Based Learning Drives Rapid Sensory Signals in Medial Prefrontal Cortex and Dorsal Hippocampus Necessary for Goal-Directed Behavior. *Neuron* **97**, 83-91.e5 (2018).

20. 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).

21. Guo, Z. V. *et al.* Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* **81**, 179–194 (2014).

22. Licata, A. M. *et al.* Posterior parietal cortex guides visual decisions in rats. *J. Neurosci.* 0105–17 (2017) doi:10.1523/JNEUROSCI.0105-17.2017.

23. Znamenskiy, P. & Zador, A. M. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* **497**, 482–485 (2013).

24. Seidemann, E., Zohary, E. & Newsome, W. T. Temporal gating of neural signals during performance of a visual discrimination task. *Nature* **394**, 72–75 (1998).

25. Guo, J.-Z. *et al.* Cortex commands the performance of skilled movement. *eLife* **4**, e10774 (2015).

26. Crapse, T. B. & Sommer, M. A. Corollary discharge across the animal kingdom. *Nat. Rev. Neurosci.* **9**, 587–600 (2008).

27. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).

28. Lak, A. *et al.* Dopaminergic and Prefrontal Basis of Learning from Sensory Confidence and Reward Value. *Neuron* **0**, (2019).

29. Burgess, C. P. *et al.* High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Rep.* **20**, 2513–2524 (2017).

30. Wang, Q. *et al.* The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **0**, (2020).

31. Olsen, S. R., Bortone, D. S., Adesnik, H. & Scanziani, M. Gain control by layer six in cortical circuits of vision. *Nature* **483**, 47 (2012).

32. Li, N. *et al.* Spatiotemporal limits of optogenetic manipulations in cortical circuits. *bioRxiv* 642215 (2019) doi:10.1101/642215.

33. Barthas, F. & Kwan, A. C. Secondary Motor Cortex: Where ‘Sensory’ Meets ‘Motor’ in the Rodent Frontal Cortex. *Trends Neurosci.* (2017) doi:10.1016/j.tins.2016.11.006.

34. Li, N., Chen, T.-W., Guo, Z. V., Gerfen, C. R. & Svoboda, K. A motor cortex circuit for motor planning and movement. *Nature* **519**, 51–56 (2015).

35. Siniscalchi, M. J., Phoumthipphavong, V., Ali, F., Lozano, M. & Kwan, A. C. Fast and slow transitions in frontal ensemble activity during flexible sensorimotor behavior. *Nat. Neurosci.* **19**, 1234–1242 (2016).

36. Wu, Z. *et al.* Context-Dependent Decision Making in a Premotor Circuit. *Neuron* **106**, 316-328.e6 (2020).

37. 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. *bioRxiv* 552026 (2019) doi:10.1101/552026.

38. Otchy, T. M. *et al.* Acute off-target effects of neural circuit manipulations. *Nature* **528**, 358–363 (2015).

39. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).

40. 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).

41. Sauerbrei, B. A. *et al.* Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386–391 (2020).

42. Harrison, T. C., Ayling, O. G. S. & Murphy, T. H. Distinct Cortical Circuit Mechanisms for Complex Forelimb Movement and Motor Map Topography. *Neuron* **74**, 397–409 (2012).

43. Kaplan, H. S. & Zimmer, M. Brain-wide representations of ongoing behavior: a universal principle? *Curr. Opin. Neurobiol.* **64**, 60–69 (2020).

44. Destexhe, A., Rudolph, M. & Paré, D. The high-conductance state of neocortical neurons in vivo. *Nat. Rev. Neurosci.* **4**, 739–751 (2003).

45. Allen, W. E. *et al.* Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* eaav3932 (2019) doi:10.1126/science.aav3932.

46. Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y. & Petersen, C. C. H. Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* **16**, 1671–1677 (2013).

47. Poort, J. *et al.* Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron* **86**, 1478–1490 (2015).

48. Bhagat, J. *et al.* Rigbox: an Open-Source Toolbox for Probing Neurons and Behavior. *bioRxiv* 672204 (2019) doi:10.1101/672204.

49. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).

50. Steinmetz, N. A. *et al.* Aberrant Cortical Activity in Multiple GCaMP6-Expressing Transgenic Mouse Lines. *eNeuro* **4**, ENEURO.0207-17.2017 (2017).

51. Jacobs, E. A. K., Steinmetz, N. A., Carandini, M. & Harris, K. D. Cortical state fluctuations during sensory decision making. *bioRxiv* 348193 (2018) doi:10.1101/348193.

52. Shimaoka, D., Harris, K. D. & Carandini, M. Effects of Arousal on Mouse Sensory Cortex Depend on Modality. *Cell Rep.* **22**, 3160–3167 (2018).

53. Ma, Y. *et al.* Wide-field optical mapping of neural activity and brain haemodynamics: considerations and novel approaches. *Phil Trans R Soc B* **371**, 20150360 (2016).

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

55. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, nature24636 (2017).

56. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M. & Harris, K. D. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* 061481 (2016) doi:10.1101/061481.

57. Lewandowski, D., Kurowicka, D. & Joe, H. Generating random correlation matrices based on vines and extended onion method. *J. Multivar. Anal.* **100**, 1989–2001 (2009).

58. Carpenter, B. *et al.* Stan: A probabilistic programming language. *J. Stat. Softw.* **76**, (2017).