nature neuroscience

Article

https://doi.org/10.1038/s41593-025-02071-5

Facial expressions in mice reveal latent cognitive variables and their neural correlates

Received: 10 July 2024

Accepted: 22 August 2025

Published online: 30 September 2025

Check for updates

Fanny Cazettes[1,2,4] ✉, Davide Reato[1,2,3,4], Elisabete Augusto[1], Raphael Steinfeld[1], Alfonso Renart[1,5] & Zachary F. Mainen[1,5]

Brain activity controls adaptive behavior but also drives unintentional incidental movements. Such movements could potentially be used to read out internal cognitive variables that are also neurally computed. Establishing this would require ruling out that incidental movements reflect cognition merely because they are coupled with task-related responses through the biomechanics of the body. Here we addressed this issue in a foraging task for mice, where multiple decision variables are simultaneously encoded even if, at any given time, only one of them is used. We found that characteristic features of the face simultaneously encode not only the currently used decision variables but also independent and unexpressed ones, and we show that these features partially originate from neural activity in the secondary motor cortex. Our results suggest that facial movements reflect ongoing computations above and beyond those related to task demands and demonstrate the ability of noninvasive monitoring to expose otherwise latent cognitive states.

The brain is tasked with controlling the body through extensive nerve pathways that activate muscles. Patterns of motor activity are informed by brain computations that depend on both the current states of the environment and the organism, as well as the organism's past experience. In attempting to unravel the mechanisms of behavior, it is common to assume that movements are adaptive output meant to act upon the world to reach desired goals. However, this premise has an important caveat. It is long known that the body expresses a variety of movements that are neither intentional nor obviously adaptive. Rather, they can be seen as an unintentional and/or unconscious 'leakage' of brain activity into the body. The importance of 'incidental' facial expressions and 'body language' has long been highlighted in popular psychology. For instance, video of the face can be used to detect a large array of nuanced emotional expressions[1], and incidental features of mobile phone interaction can be used to diagnose psychiatric conditions[2]. Although these kinds of studies clearly attest to the potential information latent in bodily expressions, concomitant measurements of neural variables are requisite to better understand how bodily expressions reveal internal cognitive processes.

Leakage of cognitive operations into the body was documented in studies showing that posture reflects the contents of working memory in rats[3,4] as well as in experiments showing that reflex gains in the arm reflect accumulated evidence during perceptual decision-making in humans[5]. In these examples, the influence of cognitive variables on observable behavior is related to the action being used to execute the task at hand. However, not all internal states are directly related to task execution and may serve various other purposes[6–8]. Distinct multidimensional facial expressions can be evoked by salient, behaviorally relevant stimuli in mice[9]. Similarly, pupil diameter is linked not only to changes in luminance, directly relevant for seeing, but also to more abstract cognitive variables related to arousal and uncertainty, even when they have no direct relationship to perception[10–12]. In line with this,

[1]Champalimaud Foundation, Lisbon, Portugal. [2]Institut de Neurosciences de la Timone, CNRS and Aix Marseille Université, Marseille, France. [3]Departement BEL, Centre CMP, Mines Saint-Etienne, Gardanne, France. [4]These authors contributed equally: Fanny Cazettes, Davide Reato. [5]These authors jointly supervised this work: Alfonso Renart, Zachary F. Mainen. ✉e-mail: fanny.cazettes@univ-amu.fr

in mice, even neutral sounds of different spectral or temporal content evoke stimulus-specific facial movements[13,14].

To make the best case for the link between internal state and bodily expressions, rich behavioral observations using high-resolution video recordings during task execution can be combined with large-scale neural recordings, as pioneered by recent work in mice and nonhuman primates[15–21]. These studies have shown that incidental movements can account for a larger fraction of the variance of neural activity than task-related movements (but see ref. 20) and are informative about the internal state of engagement (see also refs. 22,23) and that the extent to which incidental movements shape neural activity depends on the brain area[18,20].

Experiments of this sort, conducted in a task together with neural recordings, can reveal precise relationships between computational variables and both neural and bodily expressions. However, they face a critical challenge: it is difficult to rule out that incidental movements become meaningful, that is, associated with specific cognitive variables, simply because physical constraints of the body link them to other movements that are necessary to report these variables in the first place[18]. To address whether incidental movements are only linked to ongoing computations in this 'trivial' manner, it would be necessary to have a way of defining meaning independently of the current behavioral strategy used to solve the task. Recently, Cazettes et al.[24] showed that during the performance of a foraging task, an entire family of different foraging-related decision variables can be decoded from the brain in a sustained fashion throughout a behavioral session. This family of decision variables is a collection of complex and dynamic internal state markers computed by applying different algorithms to information being accumulated over relatively long time scales (several seconds). During task performance, one of these decision variables primarily manifests itself in the mouse's foraging choices. The rest of the variables in the family, which are equally complex, are computed neurally and are decodable throughout, but are not expressed in the task performance of the mouse at that particular time. Instead, they may be expressed at other times during the session. This provides an ideal situation to ask whether meaningful but currently task-unrelated internal states (decision variables) can be expressed incidentally in the face or body. If existing, this relationship between cognitive variables and movement would, by definition, avoid being trivially explainable through the demands of the task report.

Here, we monitored the facial movements of mice while simultaneously recording or manipulating the activity of large ensembles of neurons in frontal cortical regions during the foraging task of Cazettes et al.[24]. This task requires the integration of streams of binary outcomes (rewarded and unrewarded licks from a spout, representing successful and unsuccessful foraging attempts) to decide whether to stay or leave a foraging site. Multiple strategies, each associated with distinct and quantifiable decision variables, can be used to process the outcome sequence. Use of this task allowed us to directly compare the ability of facial expressions to that of large-scale neural recordings to reveal computationally well-defined but task-irrelevant internal state variables. We found that diverse decision variables with varying dynamics (even currently unused ones) could be linearly decoded from facial movements with a similar accuracy as neural recordings. Each decision variable correlated with characteristic facial features consistently across mice, suggesting a stereotyped link between decision variables and facial expressions. Through correlational (recording) and causal (optogenetics) experiments, we demonstrated that the facial expressions reflecting the decision variables partially stemmed from neural activity in the secondary motor cortex (M2). These results illustrate approaches for decoding seemingly meaningless facial expressions to infer meaningful neural states, demonstrating the potential of noninvasive monitoring to reveal otherwise hidden cognitive activity.

Results

Multiple decision variables to time foraging decisions

We analyzed data from a probabilistic foraging task in which mice could navigate between two virtual resource sites and collect drops of sugary water by freely licking at a port (Fig. 1a and Methods and see Cazettes et al.[24] for details on the behavioral approach). By design, only one of the resource sites delivered rewards at any given time. A hidden state determined which of two sites provided rewards with a fixed probability (0.9) after a lick. The hidden state also had a fixed probability (0.3) of transitioning from rewarding to depleted after each lick, but never the other way around. Once a specific resource site was depleted, the only way to receive reward again was to reach the other resource site. The presence of the state transition thus required the animal to decide after every lick whether to stay at or to leave the current resource site.

As described in detail in Vertechi et al.[25] and in Cazettes et al.[24], mice could decide to leave a foraging site based on multiple decision variables that relied on different combinations of observable events (that is, rewarded and failed foraging attempts; Fig. 1b). The optimal solution to time the decision to leave is to infer the hidden state of the resource site by counting the number of consecutive failures, which provides evidence that the site is depleted (with a complete reset after reward delivery, as a single reward signals with certainty that the current site is still active; Fig. 1b, pink). We refer to this reward-independent strategy as inference-based. Alternatively, another common, but suboptimal, strategy is to time the decision to leave based on a running estimate of how much reward has been received at the site, which is equivalent to calculating the site's net negative value (Fig. 1b, blue), staying longer the more rewards are delivered. We refer to this reward-dependent strategy as stimulus-bound.

In previous work, we showed that mice relied on these two strategies interchangeably across different sessions and even within the same sessions across different epochs[24]. Using a framework based on hidden Markov models (HMMs) combined with linear regression models (LMs; Fig. 1c, left), we fit the number of consecutive failures that the animal is willing to accept before switching sites. This was achieved using two inputs: (1) the total number of rewards, which allows a distinction between inference-based (that is, reward independent) and stimulus-bound (that is, reward dependent) strategies (as in Fig. 1b, left), and (2) a constant bias, which reflects the level of impulsivity of the animal (notice that high bias means a large number of failures accepted before switching sites, that is, low impulsivity). Each hidden state in the model captures a specific dependence of consecutive failures on the total rewards and the bias, characterizing a particular decision-making strategy. Here, a model with three states best described the foraging decision and yielded interpretable and persistent strategies (Fig. 1c and Cazettes et al.[24]). One of the strategies had large weight on the number of rewards, indicative of a stimulus-bound strategy, whereas the other two had small weights on rewards, consistent with inference (Fig. 1c, right).

To estimate how the decision variables explained the behavior of the mice in each strategy, we used regularized logistic regressions to model the probability that each lick (n = 5,877 licks in stimulus-bound, n = 6,896 licks in persistent inference and n = 1,593 licks in impulsive inference, across ten sessions) was the last one in the bout, simultaneously considering the consecutive failure and the negative value decision variables as predictors (Fig. 1d, top). This multivariate approach confirmed that the consecutive failures better explained the inference-based strategies, whereas the negative value better explained the stimulus-bound strategy (Fig. 1d, bottom).

Decision variables are reflected in facial movements

To test whether fine movement of the face reflected latent decision variables, we extracted high-dimensional temporal representations of the movement signals by applying singular value decomposition (SVD) to the motion energy movies[26] (Fig. 2a, left). We extracted the principal

https://doi.org/10.1038/s41593-025-02071-5

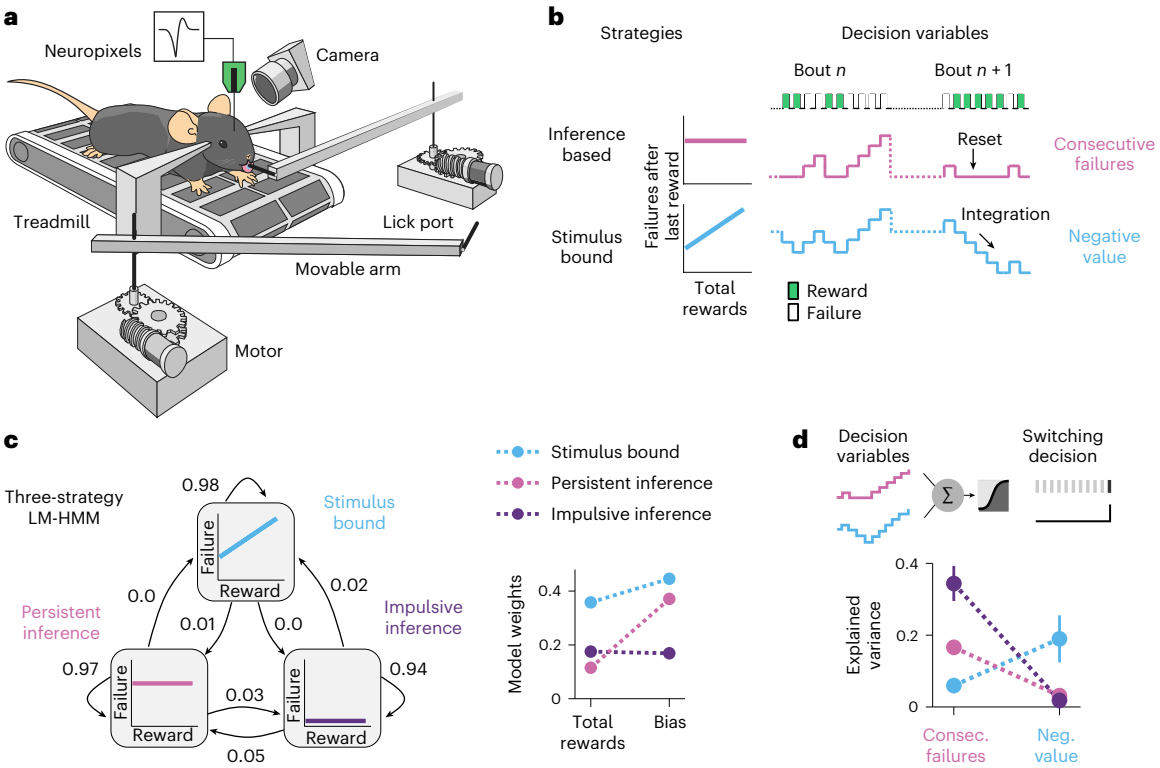

Fig. 1 | Task, strategies and decision variables. a, A head-fixed mouse is positioned on a treadmill and can forage at one of two sites (the two movable arms). Head fixation enables the monitoring of facial movements using high-speed cameras (60 fps) and provides better stability for high-throughput electrophysiology (Neuropixels). The mouse can choose to switch between sites at any time by running a set distance on the treadmill. During site switching, the front arm moves away, and the distal arm moves into place. **b**, Different strategies for timing the decision to leave the site are based on distinct decision variables. The inference-based strategy (pink) is reward independent and resets with each reward, whereas the stimulus-bound strategy (blue) depends on reward integration. **c**, Left: illustration of the LM-HMM with three different strategies (labeled 'Stimulus bound', 'Persistent inference' and 'Impulsive inference'). High self-transition probabilities of 0.98, 0.97 and 0.94 indicate that strategies typically persist for many consecutive bouts. Transition probabilities are shown by the arrows between states. Right: LM weights for the three-strategy model fitted to all sessions simultaneously ($n = 10$). **d**, Variance explained by the two different decision variables during each strategy (data are shown as mean ± s.d. across ten iterations of the cross-validated logistic regression). Consec., consecutive; Neg., negative.

components (PCs) of the movements (movement PCs, $n = 100$) from different regions of interest: (1) the entire face, (2) the upper face, excluding the tongue area where most of the movement occurred during licking, and (3) the body parts visible in the field of view. The movement PCs were binned in 200-ms windows around each licking event (Fig. 2a, middle), which directly correspond to discrete changes in the decision variables (Extended Data Fig. 1a) and capture the period immediately preceding and leading up to the peak decoding accuracy for decision variables (Extended Data Fig. 1b,c, but also see Extended Data Fig. 6d in Cazettes et al.[24]). We then predicted different decision variables on a trial-by-trial basis using cross-validated and regularized generalized LMs (GLMs; Fig. 2a, right, Supplementary Video 1 and Methods).

To test whether facial movements reflect latent computations rather than solely responding to instantaneous stimuli, and to make sure that the encoding of multiple decision variables is not a trivial consequence of potential correlations between the decision variables themselves, we orthogonalized the decision variables both to the observable action outcomes and among themselves (Extended Data Fig. 1d) and attempted to decode the 'unique' aspect of each decision variable (that is, the residual after regressing it to the others and to the action outcomes). As a control, we also tested the decoding of arbitrary signals possessing the same power spectrum as the decision variables. Although we failed to decode the arbitrary signals from any movement PCs, we could decode action outcomes and decision variables with relatively high accuracy (Fig. 2b), especially using the PCs from facial movement rather than body movement.

In principle, it is possible that the link between decision variables and movement PCs might be an indirect consequence of nonspecific changes in facial movement throughout a licking bout, for instance, changes in licking rate. To test this possibility, we attempted to decode the portion of the decision variables that is uncorrelated with licking rate. We found that, although the decision variables can correlate with lick rate changes during a bout (Extended Data Fig. 2a,b), this correlation only accounts for a small fraction of the ability to decode decision variables from facial movement (Extended Data Fig. 2c). Furthermore, we found that we could also decode latent variables with slower time constants, such as the bout numbers (Extended Data Fig. 2d,e), which are not directly linked to the immediate motor actions within a single bout. We also verified that this decoding ability is not task or movement specific: we observed similar decoding of action outcome and slow latent variables in a separate two-alternative forced choice task from Reato et al.[22] (Extended Data Fig. 3). These results show that facial movements encode latent cognitive variables in a robust and generic fashion.

How exactly are the decision variables encoded in facial movements? To answer this question, we examined the averaged spatial distribution of facial motion weighted by the models' weights (Fig. 2c) on co-registered images of the face (Extended Data Fig. 4 and Supplementary Video 2). Noticeable differences in facial movement were associated with different decision variables consistently across mice (Fig. 2d and also see Extended Data Fig. 5a,b for further evidence on an independent cohort of mice). For instance, the unique negative value corresponded to more movement around the nose than around

https://doi.org/10.1038/s41593-025-02071-5

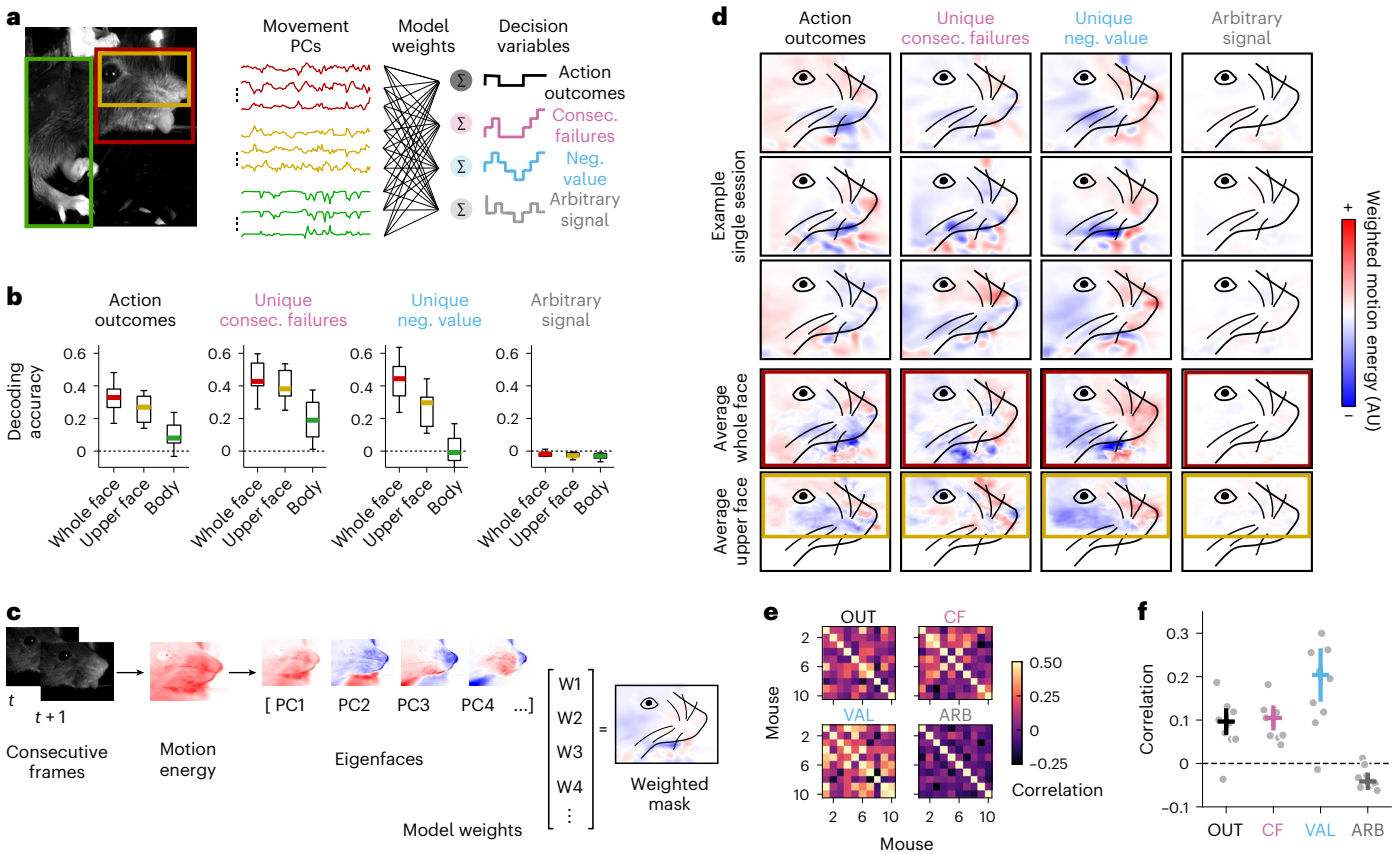

Fig. 2 | Stereotyped facial expressions of decision variables. a, Left: example PCs extracted using SVD on the motion energy movie (movement PCs) from three regions of interest. Right: the movement PCs are used as predictors in multivariate regression models to predict a set of decision variables and arbitrary signals. **b**, Decoding accuracy (cross-validated R^2) in each region for different decision variables (median across ten sessions and 25th and 75th percentiles; whiskers represent minimum and maximum values; dashed line indicates chance level). **c**, After computing the motion energy as the absolute value of the difference of consecutive frames, eigenfaces corresponding to the facial movement PCs were weighted by the models' weights. The weighted sum of the facial masks represents the spatial pattern that is most strongly associated with changes in a given decision variable. In this example, we show the facial mask for the negative value during a single session. **d**, Weighted masks for single-example sessions (top) and averaged across sessions ($n = 10$; middle for the whole face and bottom for the upper face). Red represents more movement than average, whereas blue indicates less movement than average. AU, arbitrary units. **e**, Interanimal similarity of facial expression of decision variables. OUT, action outcome; CF, unique consecutive failure; VAL, unique negative value; ARB, arbitrary signal. Colors represent the normalized two-dimensional (2D) cross-correlation at zero lag between the weighted masks of two mice ($n = 10$, 8 sessions from distinct animals, 2 from the same animal). **f**, Mean weighted mask similarity for each mouse ($n = 10$, 8 sessions from distinct animals, 2 from the same animal). Each gray dot represents the average pairwise correlation of the weighted mask of a mouse for a given decision variable with the weighted masks of all the other mice for the same decision variable. Color error bars represent median and m.a.d. across mice.

the cheek, whereas the unique consecutive failures corresponded to a subtler pattern of facial movement. We found substantial positive correlations across mice between facial expressions for the same decision variable that were higher than those obtained with an arbitrary control signal (Fig. 2e,f and Extended Data Fig. 5c–e). This between-mouse consistency was substantially greater than the correlation between the facial expression of different decision variables within a single mouse (Extended Data Fig. 5d). Furthermore, within individual mice, expressions remained highly stable across multiple sessions (Extended Data Fig. 5f). Together, these results provide strong evidence that individual decision variables are reflected in stereotypical facial expressions.

Facial expressions reflect a reservoir of decision variables

Because the mice transition between periods where they use one of the three identified strategies (Fig. 1c), and because behavior during each strategy is best predicted by different decision variables (Fig. 1d), the results in Fig. 2 are consistent with two scenarios. First, it could be that the characteristic patterns of facial motion become associated with each of the decision variables only because the mice move slightly differently during each of the behavioral strategies. As noted in the introduction, this corresponds to the 'trivial' case where cognitive variables leak into incidental movements only because there are different patterns of task-related movement associated with the report of those cognitive variables. Alternatively, the link between facial movement and decision variables could exist above and beyond the strategy currently used by the mouse. For instance, it could be that the face reflects the stimulus-bound decision variable even during periods where the mouse is solving the task using an inference-based decision variable. Although this might seem surprising, in previous work we have shown that neuronal activity in the frontal cortex of mice retains information about multiple decision variables simultaneously, even when the variables do not explain the foraging decision of mice[24]. Thus, we attempted to test if the face, like the neurons, contains information about particular decision variables during periods where those decision variables are not guiding behavior.

To address this question, we used the registered and concatenated videos to extract facial movement PCs across sessions (Fig. 3a, Extended Data Fig. 4 and Supplementary Video 2). This procedure allowed us to compare the same facial movement PCs in different

https://doi.org/10.1038/s41593-025-02071-5

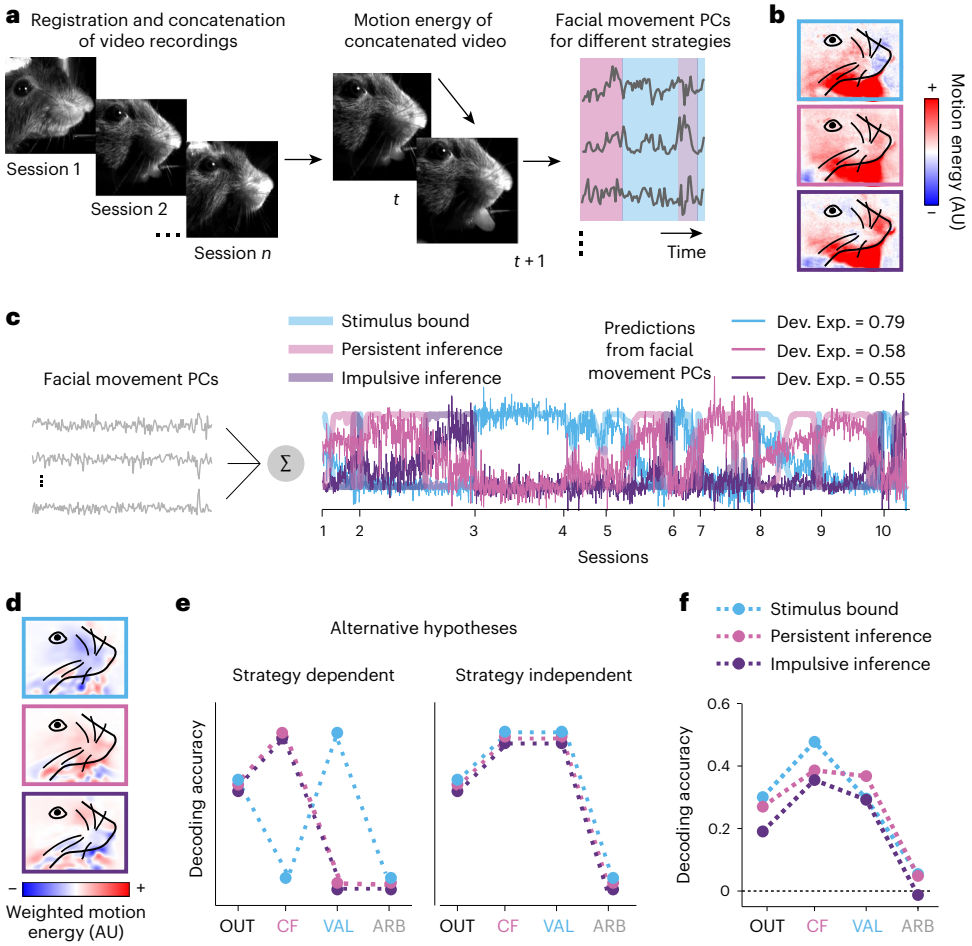

Fig. 3 | Facial expressions of decision variables do not depend on the strategy. **a**, Registration and concatenation of video recordings to extract the movement PCs of the face across different strategies. **b**, Average facial motion (calculated from raw movement PCs) between the three different strategies, across mice and licks (top: stimulus-bound; middle: persistent inference; bottom: impulsive inference). **c**, Illustration of the multivariate regression method for predicting the strategies (thick colored traces) from the facial movement PCs (gray traces). The predictions of the models (thin colored lines, which are the weighted sums of the facial movement PCs) track quite faithfully the probability of being in a given strategy (thick colored traces). The deviance explained (estimated on testing data) is reported in the top right (Dev. Exp.). **d**, Weighted masks (eigenfaces of the concatenated videos weighted by the models' weights from **c**) for the three different strategies (top: stimulus-bound; middle: persistent inference; bottom: impulsive inference). **e**, Schematic representation of hypothetical performance of strategy-dependent (left) versus strategy-independent (right) models. Left: the decoding accuracy of the decision variables depends on the strategy. Right: the decoding accuracy of the decision variables is independent of the strategy. **f**, The decoding accuracy of the decision variables from facial movement PCs is strategy independent (like data on the right in **e**; the dashed line indicates chance level).

strategies. First, we asked whether there was any difference in motion energy between strategies. Although, on average, most of the motion was around the tongue in the three different strategies, there were noticeable differences in motion across the face (Fig. 3b). Because of these differences in motion between strategies, we could predict the probability of being in a given strategy from facial movement PCs using the same regularized and cross-validated multivariate regression approach as previously described (Fig. 3c,d). Yet, to disambiguate the previously described scenarios, the critical test was to decode the different decision variables during each strategy separately. If the decision variables are expressed on the face only when guiding behavior, the portion of the signal that is unique to each variable should only be decoded in the corresponding strategy (Fig. 3e, left). Alternatively, if the decision variables are expressed on the face even when not guiding behavior, the decoding accuracy should be independent of the strategy (Fig. 3e, right). We found that the decoding accuracy of the decision variables was unequivocally strategy independent (Fig. 3f). These results suggest that, at least in the context of our task, when the brain calculates a given decision variable, that variable becomes reflected in the facial expressions of the mouse, even if that decision variable is not currently used to guide behavior (and is orthogonal to the alternative decision variables, which do play this guiding role).

Facial and neural expressions of decision variables

The capability to simultaneously decode multiple variables from facial motion, independent of strategy, mirrors the properties of the M2, where we have previously shown the simultaneous decoding of multiple decision variables from neural activity[24]. This raises the question of which of the two brings about the other. On the one hand, decision variables could manifest in neural activity as efference copies of motor signals. On the other hand, neurons could compute decision variables that are subsequently reflected in facial movements. If the latter is true, there should be 'covert' neural representations of decision variables in neural activity preceding the 'overt' expressions of decision variables in facial movement. To test this idea, we evaluated the temporal lag of the relationship between decision variables and neural populations from different cortical regions (M2: $n = 67 \pm 26$ neurons; orbitofrontal cortex (OFC): $n = 58 \pm 18$ neurons; olfactory cortex (OC): $n = 28 \pm 13$ neurons; data are shown as median ± median absolute deviation (m.a.d.) across ten sessions; Extended Data Fig. 6a–c) and compared them to

https://doi.org/10.1038/s41593-025-02071-5

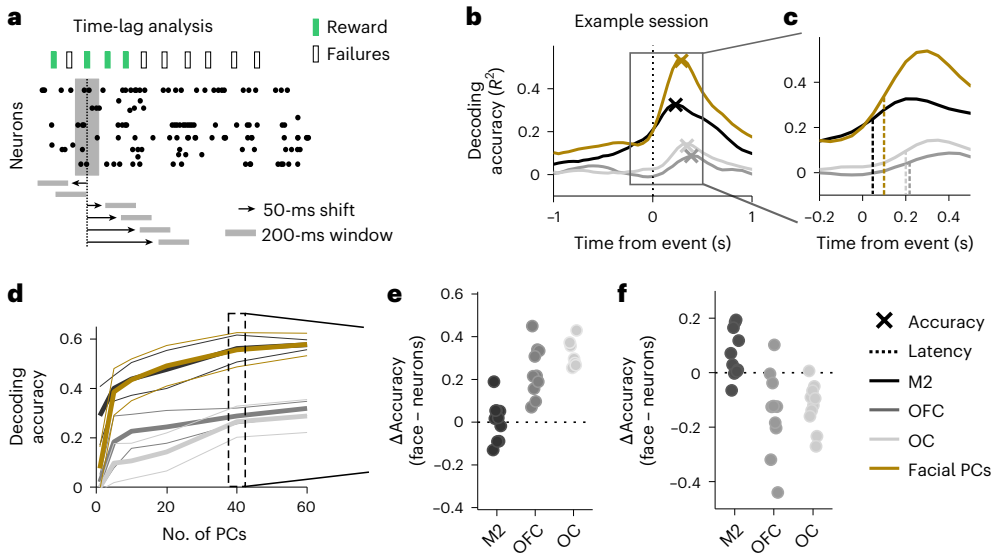

Fig. 4 | Expressions of decision variables in facial movement versus neurons. **a**, A 200-ms sliding window with a 75% overlap was used to generate a time-lag series of neural and movement activity. **b**, Example decoding accuracy (cross-validated R^2) of multivariate regression predicting decision variables from facial movement PCs (gold) or neural activity in different brain regions (grays) as a function of the time delay of the predictors. The crosses indicate the maximum accuracy. **c**, Example estimation of latencies as the times corresponding to the midpoint between the peak and the trough of each curve. **d**, Decoding accuracy of decision variables as a function of the number of PCs used. Accuracy is shown for PCs derived from facial movements (gold) and from neural activity in individual brain regions (gray lines, one line per region). Decoding accuracy is reported as the median ± m.a.d. across ten sessions. **e**, Difference in decoding accuracy between facial movement PCs and neural PCs (using 40 PCs for both) for each individual session. Positive values indicate higher decoding accuracy from facial movements. **f**, Difference in decoding latency between facial movement PCs and neurons for each individual session. Positive values indicate faster decoding (that is, shorter latencies) from neurons.

the temporal lags associated with the encoding of the decision variables by facial movements.

We generated, for each neuron and facial movement PC, time-lagged versions of the binned activity using a sliding window with increasing delays from the time of the lick (Fig. 4a). We then used each lagged series of neural or facial movement activity to independently predict the different decision variables with cross-validated and regularized GLMs. The representations of decision variables by each cortical region and facial movement PCs varied depending on the time delay, peaking at different accuracy levels (Fig. 4b). In particular, the accuracy of models using facial movement regressors was similar to that of models using neural regressors in M2 (Fig. 4d,e; ΔAccuracy M2: 0.02 ± 0.07; data are shown as median ± m.a.d. across ten sessions; $P = 0.92$, Wilcoxon signed-rank test, Holm–Bonferroni corrected) and greater than that of models using neural regressors in the OFC and OC (Fig. 4d,e; ΔAccuracy OFC: 0.22 ± 0.10, $P = 0.006$; ΔLatency OC: 0.30 ± 0.06, $P = 0.03$; data are shown as median ± m.a.d. across ten sessions, Wilcoxon signed-rank test, Holm–Bonferroni corrected).

The latency of representations of decision variables, estimated from the time delay corresponding to the midpoint decoding accuracy, also varied across brain regions and facial movement PCs (Fig. 4c). We found that decision variables were decodable from the facial movements before being decodable from neural activity in the OFC and OC (Fig. 4f; ΔLatency OFC: −0.125 ± 0.112 s; data are shown as median ± m.a.d.; Wilcoxon signed-rank test, $P = 0.0195$, Holm–Bonferroni corrected; ΔLatency OC: −0.107 ± 0.063 s; data are shown as median ± m.a.d.; Wilcoxon signed-rank test, $P = 0.0117$, Holm–Bonferroni corrected; Extended Data Fig. 6d). By contrast, predictions of decision variables from facial movement emerged slightly, but significantly, later than predictions using neural activity from M2 (ΔLatency M2: 0.050 ± 0.076 s; data are shown as median ± m.a.d.; Wilcoxon signed-rank test, $P = 0.0273$, Holm–Bonferroni corrected). This finding is consistent with the evidence that M2 neuronal activity preceding movement is the strongest predictor of facial movement PCs (Extended Data Fig. 6e,f). Together, these results suggest that M2 may participate in the generation of the facial expressions of decision variables.

Effect of M2 inactivation on facial expressions

We have previously reported in Cazettes et al.[24] that the partial inactivation of M2 significantly decreased the predictive power of the decision variables for explaining the switching behavior, suggesting that M2 is part of the neural pathway through which the decision variables shape behavior during the foraging task. Here, the emergence of decision variable representations in M2 before their facial expression (Fig. 4f and Extended Data Fig. 6f) suggests that M2 activity may also contribute to generating micromovement related to decision variables. To determine whether M2 activity is required for the facial expression of decision variables, we performed bilateral optogenetic inactivation of M2 while video monitoring facial movements during the foraging task. Specifically, we used VGAT–ChR2 mice, which express the excitatory opsin channelrhodopsin-2 in inhibitory GABAergic neurons, to silence M2 in 30% of randomly selected behavioral bouts (Fig. 5a). We examined 47 sessions from 8 mice, 5 of which were ChR2 expressing and 3 of which were control wild-type littermates that expressed no inhibitory opsin but were implanted and stimulated in the same manner.

First, we compared facial movements between laser ON and OFF conditions in inactivated and control groups by estimating the difference in average facial movement (mean of movement PCs) between the two conditions. Inactivated animals showed altered facial movements compared to controls (Fig. 5b,c; Wilcoxon rank-sum test on ΔMovement, $P = 0.036$), starting immediately after laser activation (after first lick) and persisting throughout the activation period (Extended Data Fig. 7a,b). These changes manifested as a change in motion energy of the whole face, but the most substantial differences were observed in mouth and tongue regions.

We then used separate GLMs to decode the decision variables from facial movement PCs in the laser ON condition (that is, M2 silencing)

https://doi.org/10.1038/s41593-025-02071-5

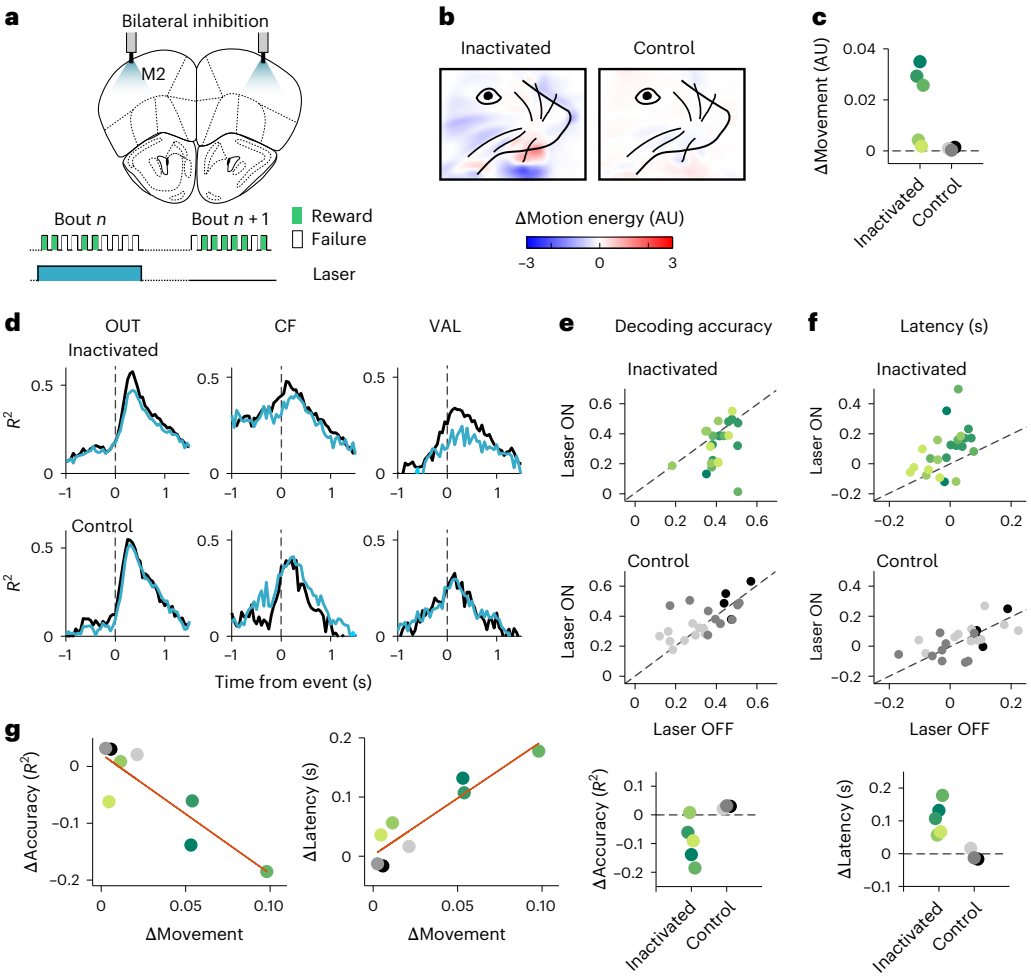

Fig. 5 | Effect of M2 inactivation on facial expressions of decision variables. **a**, Schematic representation of optic fiber placement on top of M2 (+2.5 anterior, ±1.5 lateral of bregma). Bilateral photostimulation (5 mW power per fiber, 10-ms pulses at 75 Hz) was triggered by the first lick in 30% of trials and lasted until the last lick of the bout. **b**, Laser-induced changes in facial movement patterns. Two-dimensional masks show the difference in average facial motion (calculated from movement PCs) between laser ON and OFF conditions across mice and licks. **c**, Magnitude of laser-induced facial motion changes around lick events. The variance of the difference in motion energy between laser ON and OFF (normalized by the variance in laser OFF) for inactivated (green, $n = 5$ mice) and control (black, $n = 3$ mice) groups is shown. **d**, Decoding accuracy of multivariate regression models predicting decision variables from facial movement PCs as a function of the time delay of the predictors during laser OFF (black) and laser ON (cyan) conditions (median across $n = 24$ sessions for both inactivated and control groups). **e**, Decoding accuracy (top and middle): comparison of laser ON versus laser OFF conditions. Dots below the unity line indicate that representations of decision variables derived from facial movement PCs were decoded less accurately during laser ON than during laser OFF. Difference in decoding accuracy (bottom): laser ON minus laser OFF (mean across sessions for each mouse in the inactivated and control groups). Dots represent individual sessions; colors represent individual mice. **f**, Same as in **e** but for decoding latency. Dots above the unity line indicate that representations of decision variables derived from facial movement PCs were decoded later during laser ON than during laser OFF. **g**, Correlations between decoding accuracy (left) and latency (right) with changes in movement. Green dots show inactivated mice, gray dots show control mice, and the red line indicates the linear regression.

and compared the latency and accuracy of the decodings to those observed in the laser OFF condition (that is, unperturbed M2 activity). We found that partial silencing of M2 decreased the accuracy by which decision variables could be predicted with facial movement PCs (Fig. 5d,e, $P = 0.036$, Wilcoxon signed-rank test between control and inactivated groups, and Extended Data Fig. 7c). Partial silencing of M2 also delayed the representation of decision variables from facial movement PCs (Fig. 5d,f, $P = 0.036$, Wilcoxon signed-rank test between control and inactivated groups, and Extended Data Fig. 7d).

To determine whether the altered facial movements observed during M2 inactivation were responsible for the changes in decoder performance of decision variables, we correlated the difference in decoding accuracy (ΔAccuracy) and the difference in decoding latency (ΔLatency) between laser ON and OFF conditions with changes in movement (ΔMovement) around peak decoding (250 ms after the lick; Fig. 5d). We found that both decrease in decoding accuracy (Fig. 5g, left; $R^2 = 0.76$, $P = 0.0049$) and increase in decoding latency (Fig. 5g, right; $R^2 = 0.87$, $P = 0.0007$) were highly correlated with the change in movement. This suggests that the atypical facial movements generated during M2 inactivation are responsible for the impaired decoder performance.

Discussion

In this study, we investigated how facial movements in mice reflect ongoing brain computations during a foraging task. We found that there is a remarkable richness in the representational capacity of the face, which reflected not only the identity of the current behavioral strategy used by each mouse and the time-varying decision variable that underlies this strategy but also the alternative time-varying decision variables that support strategies used during other epochs.

The musculature of the face is thus a high-capacity information channel that tracks in real time a multiplicity of latent dynamical variables computed by the brain that are related in complex ways (for example, temporal integration or reset) to immediately available external events, such as individual licks or their rewarding consequences. At a mechanistic level, our results are consistent with the hypothesis that activity in the premotor cortex (M2) contributes to the facial expressions associated with these decision variables. These facial expressions, in turn, may influence activity in other cortical regions.

Traditional studies of the neural control of movement have focused on elucidating purposeful and adaptive behaviors, such as those required to perform a task[27]. More recently, studies of seemingly meaningless expressions and movements of the face and body are also being linked to specific neural activity patterns[15,26]. Establishing these relationships is important in at least two ways. The first, which has received more emphasis, relates to the potential ability of facial movements to explain neural activity that would otherwise be considered 'noise' or misattributed. The second, which we have focused on in this report, relates to their potential to be used as visible readouts of latent cognitive processes.

The face is a particularly interesting site of neural control in mammals, being well known as a locus of emotional expression. Emotional expressions, considered the canonical markers of internal states, are still quite mysterious despite their importance to human life. Although traditional accounts identify a handful of emotional states, classification using statistical methods suggests a much richer variety[7,28]. The discovery of facial expressions of emotions in mice[9,29] has opened the door for a better understanding of their neural basis. A key step in this direction has been to link facial expressions to controllable and quantifiable variables. Recent examples of this are facial expressions being linked to auditory-evoked responses[13,14], arousal state[16,22,23] and motor readiness for multiple actions[30]. Clayton et al.[14], in particular, showed not only that facial movement is more sensitive to the intensity of auditory stimuli but also that it reports the spectral and temporal structure of sound. Their results, however, point to a sophisticated but very fast transformation from sound to muscle pattern, leaving open the extent to which the face can accurately report temporally extended cognitive variables associated with high-level cognition.

Decision-making tasks offer a powerful opportunity to link incidental expressions to temporally extended neural variables such as accumulated evidence. Previous studies showed that a running estimate of accumulated sensory evidence is present in the musculature of the arm before a choice is made in a decision-making task where the arm is used to report the decision[5]. Movement vigor has also been shown to report uncertainty in a foraging task in mice (nose-poking vigor[31]) and decision confidence in a perceptual decision-making task in humans (finger speed[32]). The contents of working memory are also reflected in postural adjustment in rats[3,4]. These results are consistent with accounts of computation that emphasize readiness in real time[33], so that the state of the motor system continuously reflects the agent's best guess of what to do at the moment.

In our study, we explored whether the cognitive significance of incidental facial expressions extends beyond (1) fast, reflex-like encoding of the current sensory environment[13,14] and (2) the immediate demands that tasks requiring temporally extended computations place on the motor system. To do so, we took advantage of the recent finding that mice compute several different decision variables simultaneously when performing a foraging task[24]. This allowed us to search for correlates of internal variables that were unrelated (and indeed mathematically orthogonal) to the currently used behavioral strategy. Unexpectedly, we found that these non-task variables are expressed as much as task-related ones and that the face contains as much information about these decision variables as recordings from multiple tens of neurons in the M2.

The decision variables could be decoded from M2 activity earlier (approximately 50 ms) than facial movements. The delay between the representation of decision variables in M2 and in facial movement is consistent with previous work reporting that movement-related signals in the brain occurred at least 50 ms before detectable movement onset[34]. Furthermore, partially inactivating M2 with optogenetics reduced our ability to decode decision variables from facial expressions (it increased the latency and reduced the accuracy of the decoding). Although we observed consistent optogenetic effects, variations in fiber placement, which could not be fully assessed histologically, may have contributed to interanimal variability in responses. Interestingly, Clayton et al.[14] found that inactivating the auditory cortex enhanced, rather than disrupted, the sensitivity of the face to sound intensity, suggesting a diversity of cortical and subcortical pathways converging on the musculature of the face to report brain activity. Together, our findings suggest that the M2 might be involved in generating the facial expressions associated with decision variables. By contrast, decoding decision variables from the OFC and OC was only possible after their detection in facial expressions and with significantly lower accuracy than facial movements. This suggests that the presence of decision variables in the OFC and OC activity might be driven by proprioceptive feedback. Together, widespread motor signals related to efference copies, as well as driven by proprioceptive feedback, contribute to the prevalence of movement-related activity across various brain regions[35].

The ability of video and other digital readouts to disclose information that has heretofore been considered private or internal is rapidly being exposed. Unconscious patterns of gaze and head movements detected by virtual reality headsets can be used to rapidly reveal individual identity[36,37]. Thus, understanding the limits of these technologies is important for shaping our views on the protection of privacy. Similar to facial expressions associated with emotions[9], we found that the different decision-related expressions correlated with characteristic features of the face and appeared stereotyped across mice, potentially facilitating their recognition through generalizable methods. This, together with the remarkable expressivity of facial muscles (facial expressions equaled ensembles of multiple tens of M2 neurons in their ability to reveal both the decision variable guiding the animals' choices and the alternative decision variables considered), suggests ample opportunities for noninvasive methods, like facial analysis, for inferring hidden brain computations. This highlights the need to devise effective regulatory policies to address the misuse of biometric technologies for privacy invasion[38–40].

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41593-025-02071-5.

References

1. Cowen, A. S. et al. Sixteen facial expressions occur in similar contexts worldwide. *Nature* **589**, 251–257 (2021).
2. Roberts, L. W., Chan, S. & Torous, J. New tests, new tools: mobile and connected technologies in advancing psychiatric diagnosis. *NPJ Digit. Med.* **1**, 20176 (2018).
3. Euston, D. R. & McNaughton, B. L. Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *J. Neurosci.* **26**, 13143–13155 (2006).
4. Cowen, S. L. & McNaughton, B. L. Selective delay activity in the medial prefrontal cortex of the rat: contribution of sensorimotor information and contingency. *J. Neurophysiol.* **98**, 303–316 (2007).
5. Selen, L. P. J., Shadlen, M. N. & Wolpert, D. M. Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J. Neurosci.* **32**, 2276–2286 (2012).

6. Flavell, S. W., Gogolla, N., Lovett-Barron, M. & Zelikowsky, M. The emergence and influence of internal states. *Neuron* **110**, 2545–2570 (2022).
7. Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. & Pollak, S. D. Emotional expressions reconsidered: challenges to inferring emotion from human facial movements. *Psychol. Sci. Public Interest* **20**, 1–68 (2019).
8. LeDoux, J. E. What emotions might be like in other animals. *Curr. Biol.* **31**, R824–R829 (2021).
9. Dolensek, N., Gehrlach, D. A., Klein, A. S. & Gogolla, N. Facial expressions of emotion states and their neuronal correlates in mice. *Science* **368**, 89–94 (2020).
10. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740–754 (2015).
11. McGinley, M. J., David, S. V. & McCormick, D. A. Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron* **87**, 179–192 (2015).
12. Cazettes, F., Reato, D., Morais, J. P., Renart, A. & Mainen, Z. F. Phasic activation of dorsal raphe serotonergic neurons increases pupil size. *Curr. Biol.* **31**, 192–197 (2021).
13. Bimbard, C. et al. Behavioral origin of sound-evoked activity in mouse visual cortex. *Nat. Neurosci.* **26**, 251–258 (2023).
14. Clayton, K. K. et al. Sound elicits stereotyped facial movements that provide a sensitive index of hearing abilities in mice. *Curr. Biol.* **34**, 1605–1620 (2024).
15. 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).
16. Yin, C. et al. Spontaneous movements and their relationship to neural activity fluctuate with latent engagement states. *Neuron* https://pubmed.ncbi.nlm.nih.gov/40602403/ (2025).
17. 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).
18. Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J. & Petrides, M. Neural cognitive signals during spontaneous movements in the macaque. *Nat. Neurosci.* **26**, 295–305 (2022).
19. Syeda, A. et al. Facemap: a framework for modeling neural activity based on orofacial tracking. *Nat. Neurosci.* **27**, 187–195 (2024).
20. Talluri, B. C. et al. Activity in primate visual cortex is minimally driven by spontaneous movements. *Nat. Neurosci.* **26**, 1953–1959 (2023).
21. Hasnain, M. A. et al. Separating cognitive and motor processes in the behaving mouse. *Nat. Neurosci.* **28**, 640–653 (2025).
22. Reato, D., Steinfeld, R., Tacão-Monteiro, A. & Renart, A. Response outcome gates the effect of spontaneous cortical state fluctuations on perceptual decisions. *eLife* **12**, e81774 (2023).
23. Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D. A. & Jaramillo, S. Decision-making dynamics are predicted by arousal and uninstructed movements. *Cell Rep.* **43**, 113709 (2024).
24. Cazettes, F. et al. A reservoir of foraging decision variables in the mouse brain. *Nat. Neurosci.* **26**, 840–849 (2023).
25. Vertechi, P. et al. Inference-based decisions in a hidden state foraging task: differential contributions of prefrontal cortical areas. *Neuron* **106**, 166–176 (2020).
26. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
27. Schwartz, A. B. Movement: how the brain communicates with the world. *Cell* **164**, 1122–1135 (2016).
28. Cowen, A. S. & Keltner, D. Self-report captures 27 distinct categories of emotion bridged by continuous gradients. *Proc. Natl Acad. Sci. USA* **114**, E7900–E7909 (2017).
29. Langford, D. J. et al. Coding of facial expressions of pain in the laboratory mouse. *Nat. Methods* **7**, 447–449 (2010).
30. Mangin, E. N., Chen, J., Lin, J. & Li, N. Behavioral measurements of motor readiness in mice. *Curr. Biol.* **33**, 3610–3624 (2023).
31. Lottem, E. et al. Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nat. Commun.* **9**, 1000 (2018).
32. Dotan, D., Meyniel, F. & Dehaene, S. On-line confidence monitoring during decision making. *Cognition* **171**, 112–121 (2018).
33. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
34. 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).
35. Zagha, E. et al. The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. *J. Neurosci.* **42**, 1375–1382 (2022).
36. Miller, M. R., Herrera, F., Jun, H., Landay, J. A. & Bailenson, J. N. Personal identifiability of user tracking data during observation of 360° VR video. *Sci. Rep.* **10**, 17404 (2020).
37. Asish, S. M., Kulshreshth, A. K. & Borst, C. W. User identification utilizing minimal eye-gaze features in virtual reality applications. *Virtual Worlds* **1**, 42–61 (2022).
38. Smith, M. & Miller, S. The ethical application of biometric facial recognition technology. *AI Soc.* **37**, 167–175 (2022).
39. Pereira, R. S. in *The Legal Challenges of the Fourth Industrial Revolution* (eds Moura Vicente, D., de Vasconcelos Casimiro, S. & Chen, C.) 193–209 (Springer International Publishing, 2023).
40. Faraldo Cabana, P. in *Artificial Intelligence, Social Harms and Human Rights* (eds Završnik, A. & Simončič, K.) 35–54 (Springer International Publishing, 2023).

Methods

Behavioral approach

Mice ($n = 17$, male and female, 2–9 months old) were successfully trained in a head-fixed version of a probabilistic foraging task. To ensure that the dataset consisted only of animals actively performing the task, mice that failed to learn to run and lick in the head-fixed foraging setup were excluded from the study. No statistical methods were used to predetermine sample sizes, but our sample sizes are similar to those reported in previous publications[15,26]. All mouse strains were purchased from The Jackson Laboratory and bred in-house. Mice used in the electrophysiology experiments ($n = 9$) were of the C57BL/6J background (The Jackson Laboratory, 000664). For the optogenetics experiments, B6.Cg-Tg(*Slc32a1*-COP4*H134R/EYFP)8Gfng/J mice (also known as VGAT–ChR2–YFP BAC; The Jackson Laboratory, 014548) were used. This included hemizygous VGAT–ChR2–YFP mice ($n = 5$) and wild-type littermate controls ($n = 3$). Mice were housed in individually ventilated cages under a normal 12-h light/12-h dark cycle, temperature was maintained between 19 °C and 23 °C, and humidity was maintained between 50% and 65%. All experimental procedures were approved and performed in accordance with the Champalimaud Centre for the Unknown Ethics Committee guidelines and by the Portuguese Veterinary General Board (Direção Geral de Veterinária, approval 0421/000/000/2016).

Mice were first implanted with a head bar under standard aseptic surgical procedures. After a recovery period (5 to 10 days), mice were water restricted and only received sucrose water (10%) during the task. Mice were given 1 ml of water or 1 g of hydrogel (Clear H_2O) on days when no training or recording occurred or if they did not receive enough water during the task.

During the task, mice were head fixed and placed on a linear treadmill. The behavioral apparatus was controlled by microcontrollers (Arduino Mega 2560) and scientific boards (Champalimaud Hardware platform), which recorded the time of the licks and the running speed. Running on the treadmill could activate the movement of two arms, which materialized two different foraging sites, via a coupling with digital Servo motors (Hitec HS-5625-MG). At the extremity of each arm, water flowed through a lick port by gravity through water tubing and was controlled by calibrated solenoid valves (Lee Company). To lick at an arm and receive rewards, mice had to decrease their running speed for more than 250 ms below a threshold for movement (3 cm s^{-1}). Licks were detected in real time with a camera (FLIR Chameleon-USB3, 60 fps) located on the right side of the treadmill using a thresholding method in BONSAI[41]. Each lick was rewarded in a probabilistic fashion by a small drop of water (1 µl) with a probability of 0.9 in the rewarding state. The small reward size ensured that there was no strong difference in licking rate between rewarded and unrewarded licks. Each lick could cause the state of the currently exploited site to transition from rewarding to unrewarding with a probability of 0.3. When such a transition occurred, the currently unexploited site (the distal one) transitioned from an unrewarding state to a rewarding state. To leave the currently exploited site and reach the other one, mice had to restart running above the threshold for movement for more than 150 ms and travel a fixed distance on the treadmill (around 16 cm).

The sequence of events was randomized within each session using a two-part method. First, the trial structure was probabilistic, governed by set statistics. Second, the progression and duration of behavioral bouts were contingent on the mouse's real-time responses. This design ensured that every session generated a unique, nonrepeating sequence of events.

Neuropixels recording and processing

Recordings ($n = 10$ recording sessions in 9 mice) were made using electrode arrays with 374 recording sites (Neuropixels 'Phase3A'). Before each recording session, the shank of the probe was stained with red fluorescent dye (DiI; Thermo Fisher Vybrant, V22885) to allow later track localization. Before the first recording session, mice were briefly anesthetized with isoflurane and administered a nonsteroidal analgesic (Carprofen) before drilling one small craniotomy (1 mm in diameter) over the M2 (+2.5 mm anterior and +1.5 lateral relative to bregma) from either hemisphere. The craniotomy was cleaned with a sterile solution and covered with silicone sealant (Kwik-Sil, World Precision Instruments). After several hours of recovery, mice were placed in the behavioral setup, and the probe was slowly advanced through the dura and slowly lowered to its final position (3 to 3.5 mm inside the brain). The probe was allowed to settle for at least 10 min before starting the recording to ensure better stability of the recorded units. Recordings were acquired with a SpikeGLX neural recording system (https://billkarsh.github.io/SpikeGLX/) using the external reference setting and a gain of 500 for the AP band (300-Hz high-pass filter).

After the recording session, mice were deeply anesthetized with ketamine/xylazine and perfused with 4% paraformaldehyde. The brain was extracted and fixed for 24 h in paraformaldehyde at 4 °C and then washed with 1% phosphate-buffered saline. The brain was sectioned at 50 µm, mounted on glass slides and stained with DAPI. Images were taken at ×5 magnifications for each section using a Zeiss AxioImager at two different wavelengths (one for DAPI and one for DiI). To determine the trajectory of the probe and approximate the location of the recording sites, we used SHARP-Track[42]. Characteristic physiological features were also used to refine the alignment procedure (especially the absence of spikes between frontal and olfactory cortical boundaries and the local field potential signature in deep olfactory areas).

Neural data were first automatically spike sorted with Kilosort2 (ref. 43) using MATLAB (MathWorks). To remove artifacts, traces were 'common average referenced' by subtracting the median activity across all channels at each time point. Second, the data were manually curated using an open-source neurophysiological data analysis package (Phy: https://github.com/kwikteam/phy). Clusters were discarded as noise if they exhibited nonphysiological waveforms, inconsistent spike shapes, low amplitudes, or high refractory period violations. Subsequently, neighboring units with similar waveforms were merged based on their cross-correlograms and drift patterns. Units that passed this curation process were classified as well-isolated single neurons for analysis. For all analyses, otherwise noted, we averaged for each neuron the number of spikes into bins by considering a 200-ms window centered around each lick. Because the interval between each lick was on average around 150 ms, there was little overlap between two consecutive bins, and each bin typically contained the number of spikes associated with only one single lick.

Video monitoring and processing

Facial movements were monitored at 60 fps using an infrared camera (the FLIR Chameleon-USB3 for 56 sessions and the Sony PlayStation 3 Eye Camera for 2 sessions) positioned on the side of the animal. The videos were synchronized with the electrophysiological and behavioral data for subsequent analysis. To extract movement signals from the videos, we used an open-source toolbox: FaceMap[19,26]. FaceMap was applied to all video frames of each session, including both laser ON and OFF data, for the optogenetics datasets to capture the full range of facial expressions. Specifically, we performed SVD on the motion energy (the absolute value of the difference of two consecutive frames) to compute the 500 highest dimensions of video information. For each dimension, we averaged the signal in a 200-ms window centered around each lick, similar to the processing of the neural data (see Extended Data Fig. 1a for a schematic description of this process).

To facilitate comparisons and averaging of images across different recording sessions (for example, 2D representations of decision variables on the mouse face), we performed image registration (see Extended Data Fig. 4). We first computed the average image for each video. One video's average image was selected as the reference image. A set of eight control points was then manually positioned on the average image. Corresponding points on the reference image were also

manually placed, mirroring the locations on the target image. These point pairs were used to estimate an affine transformation between the two images using MATLAB's 'fitgeotrans' function. The transformation was then applied to images derived from the videos, including, but not limited to, 2D representations of decision variables. The same transformation was also applied for each frame and session for the analysis where we concatenated videos from all sessions. To reconstruct the motion energy reflecting each decision variable, we weighed each eigenface by the corresponding model weights (see Fig. 2c for a schematic description of this process) and then performed a linear combination of all the weighted facial masks ($n = 100$).

Photoinhibition

To optically silence M2 activity by stimulating ChR2-fused-VGAT-expressing GABAergic interneurons, we used blue light from a 473-nm laser (LRS-0473-PFF-00800-03, Laserglow Technologies, or DHOM-M-473-200, UltraLasers). Light was emitted from the laser through an optical fiber patch cord (200 µm, 0.22-NA; Doric lenses) connected to a second fiber patch cord with a rotary joint (FRJ 1 × 1, Doric lenses), which in turn was connected to the chronically implanted optic fiber cannulas (M3 connector, Doric lenses). The power of the laser (5 mW) was calibrated before every session using an optical power meter kit (Digital Console with Slim Photodiode Sensor, PM100D, Thorlabs). The optical stimulation (10-ms pulses, 75 s^{-1}, 5 mW) was turned on during 30% of randomly interleaved bouts. Light delivery started when the first lick was detected and was interrupted if the animal did not lick for 500 ms (which was in 98% of bouts after the last lick of the bouts). The experimenter was blind to the genotype of the mice when performing optic fiber implantations and running the optogenetic experiments.

Data analyses

All data analyses were performed with custom-written software using MATLAB. We used GLMs to fit the switch behavior (Bernoulli distribution and a logit link function) or the decision variables (Gaussian distribution and a linear link function; these are in effect linear models) using facial movement PCs or spike counts as predictors. Model fits were performed on individual recording sessions (Figs. 2, 4 and 5) or on concatenated sessions when looking at strategies (Figs. 1 and 3). To avoid overfitting, we used nested cross-validation. We split the data into five folds. In each iteration, one fold (20% of the data) served as the test set, and the remaining four folds (80%) were used for training and hyperparameter tuning. Within this 80%, we performed an inner fivefold cross-validation to determine the optimal hyperparameters for elastic net regularization ($\alpha = 0.5$), selecting the hyperparameters that minimized deviance. After selecting the hyperparameters, we trained the Elastic Net model on the entire 80% training set to estimate the model coefficients. The trained model was then evaluated on the held out 20% test fold. This process was repeated five times, so each fold served once as the test set. The coefficients of determination (R^2) reported as a metric of the goodness of fit were calculated from the cross-validated results. A value of 0 indicates chance-level performance, whereas a value of 1 indicates perfect prediction.

The standard randomized cross-validation described above can overestimate performance when data have temporal dependencies, such as the slowly evolving latent variables (for example, trial or bout number) tested in Extended Data Figs. 2 and 3. Here, we used a fivefold sequential cross-validation procedure. The full sequence of trials (or bouts) was divided into five consecutive blocks. Each fold used a block for testing and the other four for training. This procedure prevented overfitting to slow variations across trials.

As an analytic control, we used arbitrary signals with the same power spectra as the decision variables. For each session, we generated arbitrary signals with the same second-order properties as the original time series (for example, the consecutive failures) by transforming the original data into the frequency domain (Fourier transform), randomizing the phases simultaneously across the time series and converting the data back into the time domain (inverse Fourier transform). This process ensured that the arbitrary signals shared the same power spectrum as the original data but with disrupted temporal dependencies.

To estimate the relative latencies of decision variable representation in facial movements and neural activity, we performed a time-lag analysis (Figs. 4 and 5). Our theoretical decision variable is updated at each lick. However, the actual expression of this decision variable in either neural activity or facial movements could occur before or after the lick. Therefore, to fairly compare the latencies in both modalities, for each neuron and facial movement PC, we used a 200-ms sliding window with 75% overlap to create lagged versions of the binned activity. The range of lags considered spanned from 1 s before to 1.5 s after each lick, encompassing both pre- and post-lick activity. These lagged activity series were fed into separate, cross-validated and regularized GLMs. Each GLM was trained to predict the decision variables at the time of the lick, using the lagged neural or facial activity from the corresponding window. For each decision variable, we constructed a curve depicting how decoding accuracy (R^2) changed with increasing delay, separately for neural activity and facial movements. We manually identified the peak and the trough of this curve while blinded to the experimental condition. The latency of the decoding was then defined as the time corresponding to the midpoint between these two points. For curves with an ill-defined trough, a baseline region was identified as points where the signal dropped below 5% of its peak value, and the trough was defined as the data point with the largest $x$value within this baseline region. Curves with no clear peaks, which happened when the decoding was very poor (around or below chance level), were excluded from the analysis. We then compared the median latencies between the neural and facial movement representations across sessions to determine which modality, on average, expressed the decision variables earlier.

To investigate the impact of optogenetic M2 silencing on facial expressions of decision variables, we used the regularized, cross-validated GLMs. For each behavioral session, we fit separate GLMs in laser ON and OFF conditions to predict decision variables from facial movements PCs during periods of transient M2 inactivation (laser ON) and during control periods (laser OFF).

The normality of the data distribution was not formally tested; therefore, we used nonparametric tests for all statistical analyses.

LM-HMM

To detect the precise moment when animals switch between strategies within a session, we use a previously developed model that combines a framework based on HMMs with LMs (see ref. 22). The resulting 'LM-HMM' framework takes input-driven Gaussian observations, modeling a time-varying linear dependence $\hat{F}_t = w^{(k)}\hat{R}_t + b^{(k)}$ of normalized consecutive failures $\hat{F}_t$ (observations) on normalized total rewards $\hat{R}_t$ (inputs) across bouts $t = 1,\ldots T$; ϵ_t is i.i.d. Gaussian noise with mean 0 and variance $\sigma^{(k)}$. For each session m, the normalized values $\hat{F}_t = F_t/{F_m}^{\max}$ and $\hat{R}_t = R_t/{R_m}^{\max}$ were obtained by min-maxing the raw values F_t, R_t on their within-session max $F_m^{\max}, R_m^{\max}$, allowing us to fit a single model to all sessions where both inputs and observations were bounded between 0 and 1. Here, the slope $w^{(k)}$, intercept $b^{(k)}$ and noise variance $\sigma^{(k)}$ depend on the hidden state $k = 1,\ldots K$, each state representing a different strategy. Specifically, states with $w^{(k)} = 0$ or $w^{(k)} > 0$ represent inference-based and stimulus-bound strategies, respectively. Bias $b^{(k)}$ reflects persistent (large) or impulsive (small) behavior. We fit the model to all mice using an Expectation–Maximization algorithm to optimize parameters, and we chose the best model complexity (three states) using threefold cross-validation with maximum likelihood estimation and maximum a posteriori approaches (see Cazettes et al.[24] for details).

Reporting summary
Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability
The behavioral and electrophysiological data used in this study are available on Figshare at https://figshare.com/s/924af1de619f4597f37a (ref. 44). Raw videos and electrophysiological data are too large to be shared on a public repository and are therefore available from the authors upon request.

Code availability
All analyses were performed using custom code written in MATLAB that is available upon request. The code used to process the videos is publicly available at https://github.com/MouseLand/facemap. The code used for the central GLM analyses is publicly available at https://hastie.su.domains/glmnet_matlab/. The code developed for the LM-HMM can be accessed at https://github.com/mazzulab/ssm/blob/master/notebooks/2c%20Input-driven%20linear%20model%20(LM-HMM).ipynb.

References
41. Lopes, G. et al. Bonsai: an event-based framework for processing and controlling data streams. *Front. Neuroinform.***9**, 7 (2015).
42. Shamash, P., Carandini, M., Harris, K. & Steinmetz, N. A tool for analyzing electrode tracks from slice histology. Preprint at *bioRxiv* https://doi.org/10.1101/447995 (2018).
43. Pachitariu, M., Sridhar, S., Pennington, J. & Stringer, C. Spike sorting with Kilosort4. *Nat. Methods* **21**, 914–921 (2024).
44. Cazettes, F. Facial expressions in mice and neural correlates. *Figshare* https://figshare.com/s/924af1de619f4597f37a (2025).

Acknowledgements
We thank J. P. Morais for support with behavioral training. This work was funded by CNRS (F.C.), Simons Foundation (F.C.: SCGB 969875; Z.F.M.: SCGB 543011), Marie-Curie postdoctoral fellowships (F.C.: HORIZON-MSCA-2021-PF-01 101062459; D.R.: HORIZON-MSCA-2021-PF-01 101063075), Fundação para a Ciência e a Tecnologia (A.R.: LISBOA-01-0145-FEDER-032077 and PTDC/MED-NEU/4584/2021), la Caixa Foundation (A.R.: HR23-00799), the European Research Council Advanced Grant (Z.F.M.; 671251) and Champalimaud Foundation (A.R. and Z.F.M.). This work was also supported by Portuguese national funds through Fundação para a Ciência e a Tecnologia in the context of the project UIDB/04443/2020 and by the research infrastructure CONGENTO, cofinanced by Lisboa Regional Operational Programme (Lisboa2020), under the PORTUGAL 2020 Partnership Agreement, through the European Regional Development Fund and Fundação para a Ciência e a Tecnologia (Portugal) under the projects LISBOA-01-0145-FEDER-02217 and LISBOA-01-0145-FEDER-022122.

Author contributions
F.C., A.R. and Z.F.M. designed the study. F.C. and E.A. performed behavioral and optogenetics experiments. F.C. performed electrophysiological experiments and curated the data. D.R. processed the video data. R.S. collected the data used in Extended Data Fig. 3. F.C., D.R. and A.R. designed and performed the analyses. F.C., A.R and Z.F.M. wrote the paper. All authors reviewed the paper.

Competing interests
The authors declare no competing interests.

Additional information
Extended data is available for this paper at https://doi.org/10.1038/s41593-025-02071-5.

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41593-025-02071-5.

Correspondence and requests for materials should be addressed to Fanny Cazettes.

Peer review information *Nature Neuroscience* thanks Nadine Gogolla, Jeffrey Markowitz and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

https://doi.org/10.1038/s41593-025-02071-5

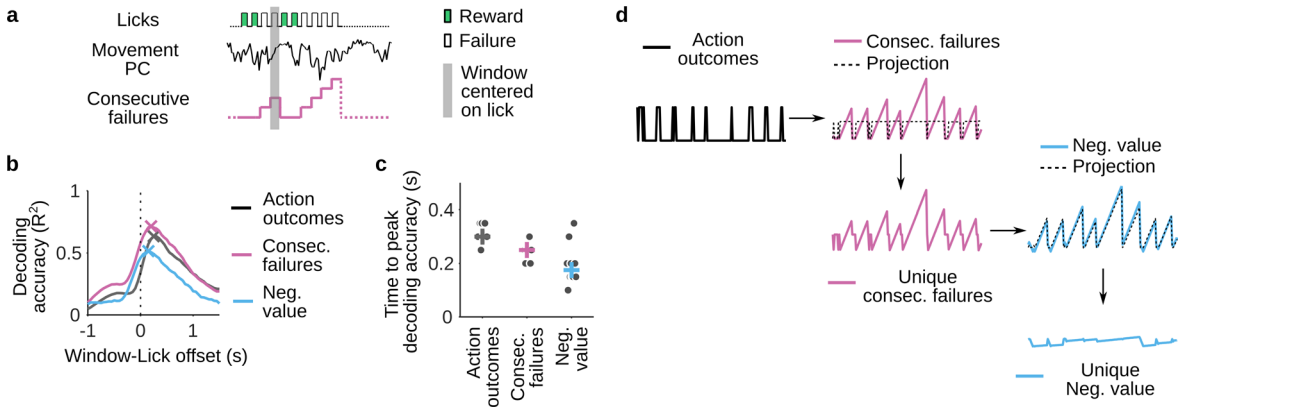

Extended Data Fig. 1 | Decoding decision variables from movement PCs. **a**, Decision variables (for example, consecutive failures) evolve after each lick outcome. Therefore, to predict these variables from movement PCs, we aligned the PCs to lick events using a 200 ms window. **b**, Decoding accuracy of various decision variables from movement PCs was estimated using a 200 ms sliding window (75% overlap) at different time lags relative to the lick event (example session shown). **c**, Peak decoding accuracy (grey dots) for each decision variable occurred between 100 ms and 350 ms after the lick. Error bars represent the median and m.a.d. across sessions (N = 10 sessions, 8 sessions from distinct animals, two from the same animal). **d**, This schematic describes the method used to partial out the linear relationship between latent variables. This approach allowed us to decompose DV1 (for example, consecutive failures) into the sum of two time series: one proportional to DV0 (for example, action outcomes) and another orthogonal (uncorrelated) to DV0, which we denote as unique DV1 (for example, unique consecutive failures). The unique consecutive failures residual (pink) is orthogonal to the action outcomes (black). Subsequently, the same procedure generated the unique negative value residual, orthogonal to both action outcomes and unique consecutive failures. The three resulting orthogonal time series (action outcomes, unique consecutive failures, and unique negative value) were then fit using movement PCs.

https://doi.org/10.1038/s41593-025-02071-5

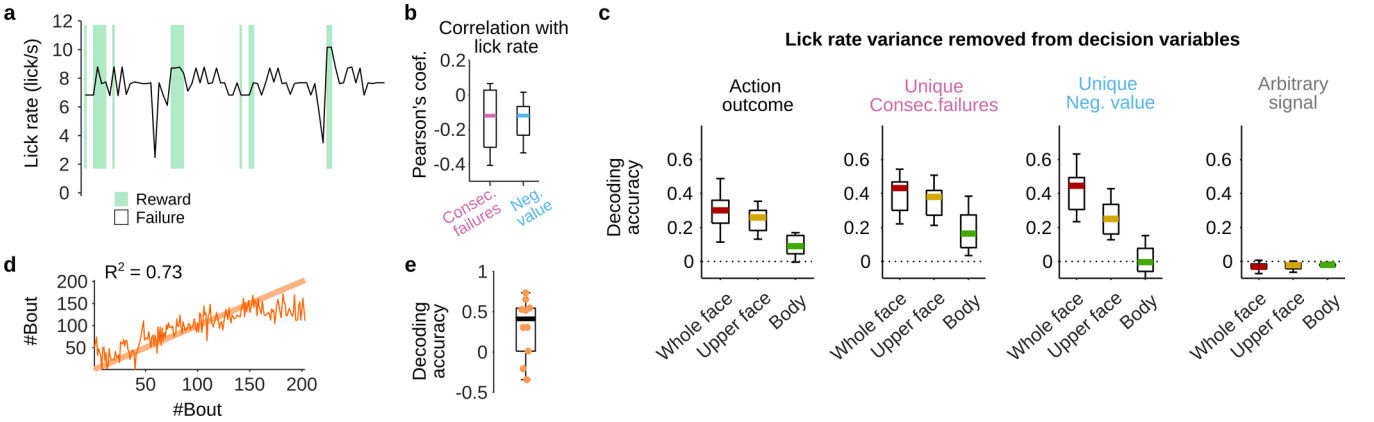

Extended Data Fig. 2 | Facial movement encodes decision variables independently of licking rate and can also reflect slow timescale processes. **a**, Changes in lick rate across 4 example bouts. **b**, Correlation (Pearson's coefficient) between lick rate and the two decision variables (median across 10 sessions with 25th and 75th percentiles; whiskers represent minimum and maximum values). A small negative correlation exists between lick rate and the decision variables, suggesting that mice tend to lick slightly faster during reward consumption and slow down towards the end of a lick bout. **c**, Decoding accuracy (cross-validated R^2) for action outcome, unique aspects of the two decision variables, and arbitrary signals in each region of interest (median across 10 sessions with 25th and 75th percentiles; whiskers represent minimum and maximum values). The movement PCs are used as predictors in multivariate regression models to predict the action outcome and the decision variables with lick rate variance removed (partialed out). The variance in action outcome, consecutive failures and negative value that is not explained by lick rate remains highly decodable. This suggests that the relationship between facial expressions and latent variables is not solely explained by licking behavior. **d**, Decoding of bout number from facial movement PCs in an example session. The actual bout number is shown as a thick light orange line, and the decoded projection is shown as a thin dark orange line. **e**, Decoding accuracy (cross-validated R^2) across all sessions (N = 10, orange dots). Boxplots show the median (center line), 25th and 75th percentiles (box edges), and minimum and maximum values (whiskers).

https://doi.org/10.1038/s41593-025-02071-5

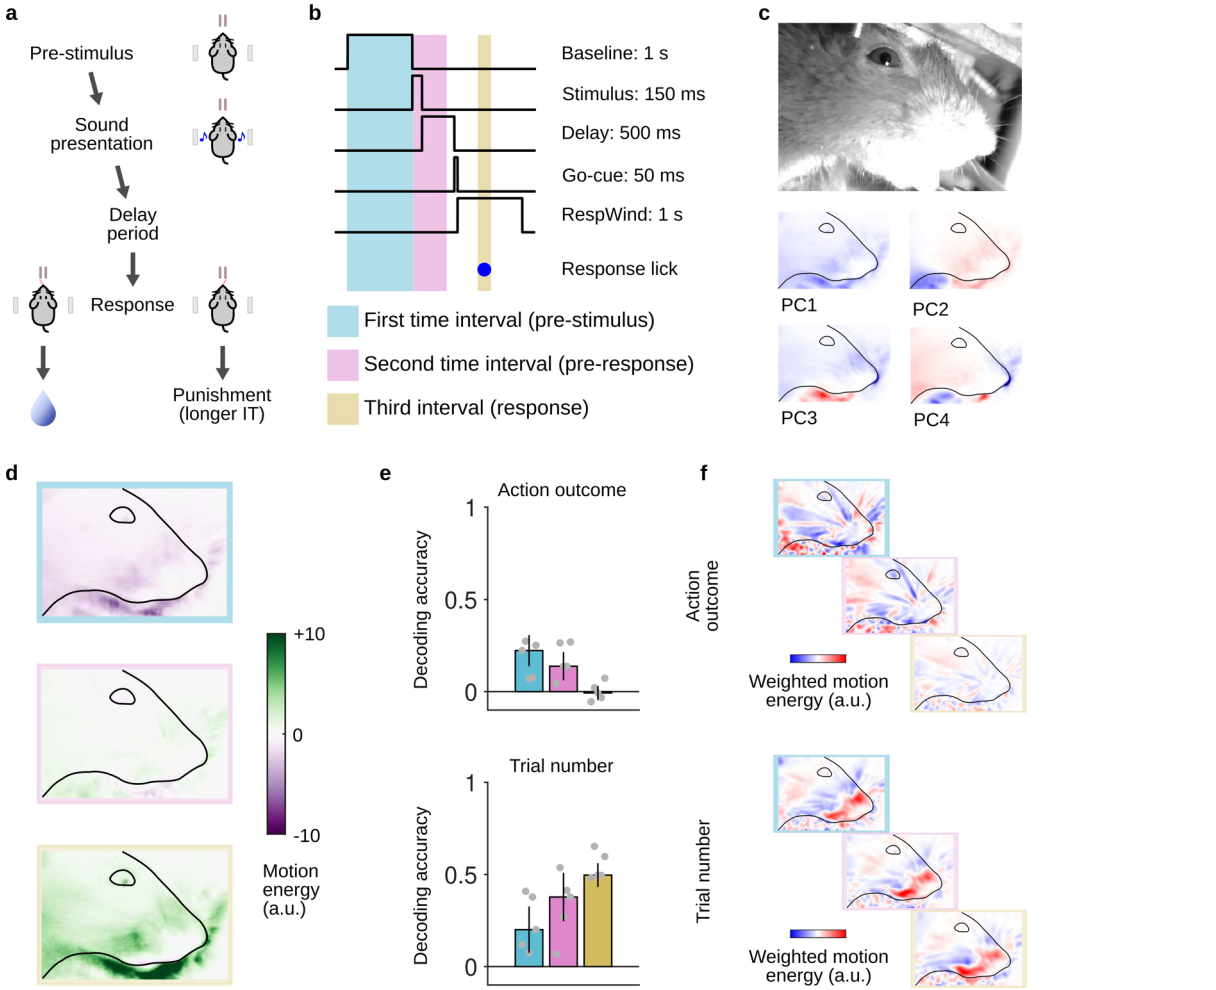

Extended Data Fig. 3 | Facial expression of different task variables in an auditory two-alternative forced choice (2AFC) task. a, Mice (N = 5 for a total of 20 behavioral sessions) were presented with single tones (150 ms) of varying frequencies (low: 9.9, 12, and 13 kHz; high: 15, 16.3, and 20 kHz) and, after a delay period (500 ms after stimulus offset), reported their perceived frequency (high or low) by licking one of two water spouts (left or right) to receive a water reward if the response was correct. Simultaneously videos were recorded at 60 fps. **b**, Task schematics and time intervals analyzed. Video analysis focused on three periods, color-coded in the schematic: "pre-stimulus" (before tone onset; 1 s), "pre-response" (including stimulus presentation and part of the delay period; 500 ms), and "response" (around the time of the lick response; 200 ms). **c**, Representative frame from the co-registered video (generated by combining the 20 sessions), along with the first four eigenfaces. Error bars represent the median and m.a.d. across mice (N = 5). **d**, Average motion energy across task periods. Motion energy (relative to the video average) is shown for each analysis period. Low motion is observed during the pre-stimulus period (top), increasing during stimulus presentation (middle), and peaking during the response (bottom), reflecting movement associated with licking. **e**, Decoding accuracy of task variables from facial movement PCs. Decoding accuracy from facial movement PCs for previous trial outcome (rewarded/unrewarded) and trial number. Each dot represents the average decoding accuracy across sessions for a single mouse; error bars indicate the median and m.a.d. across mice. The task variables, especially the slow latent variable 'Trial number', can be decoded with relatively high accuracy from facial movement PCs. **f**, Weighted masks (that is, facial representation of the decoded task variables in panel e) for the three different time intervals. The expression of the task variables on the face is highly consistent across the different time intervals, suggesting that this representation is independent of the animal's overall movement.

https://doi.org/10.1038/s41593-025-02071-5

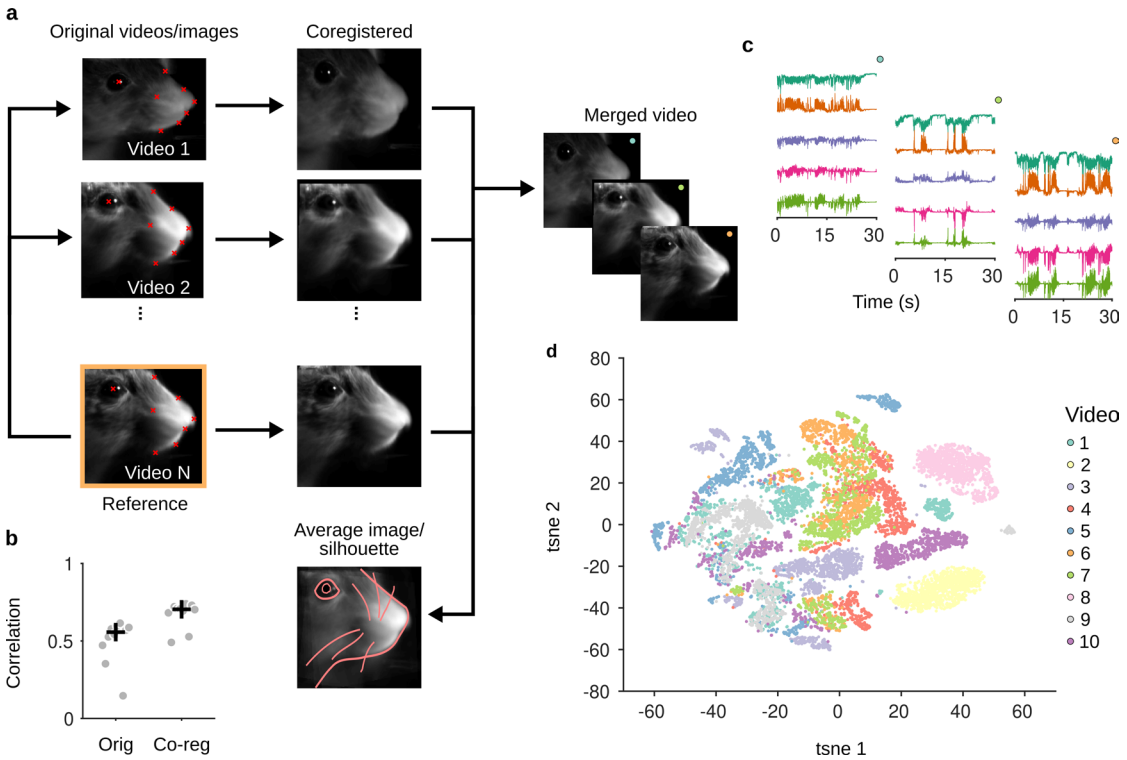

Extended Data Fig. 4 | Image co-registration across videos. a, Eight facial landmarks (red) were manually identified on average frames from each video. An affine transformation determined by using these landmarks and MATLAB's fitgeotrans function, co-registered frames from each video to a reference frame (orange, bottom). This transformation facilitated comparisons and averaging of weighted masks (Fig. 2), video concatenation (Fig. 3), and definition of an average facial silhouette (Fig. 2). **b**, Improvement in pairwise 2D cross-correlation between average video frames before and after co-registration (N = 10 videos; error bars represent the median and m.a.d. across mice). **c**, Example traces of the first five principal components (PCs) derived from motion energy analysis of three co-registered videos. Singular value decomposition (SVD) was applied to the merged video data. **d**, t-SNE visualization of the first 100 PCs of the merged video. Points are color-coded by video identity, revealing substantial overlap between videos.

https://doi.org/10.1038/s41593-025-02071-5

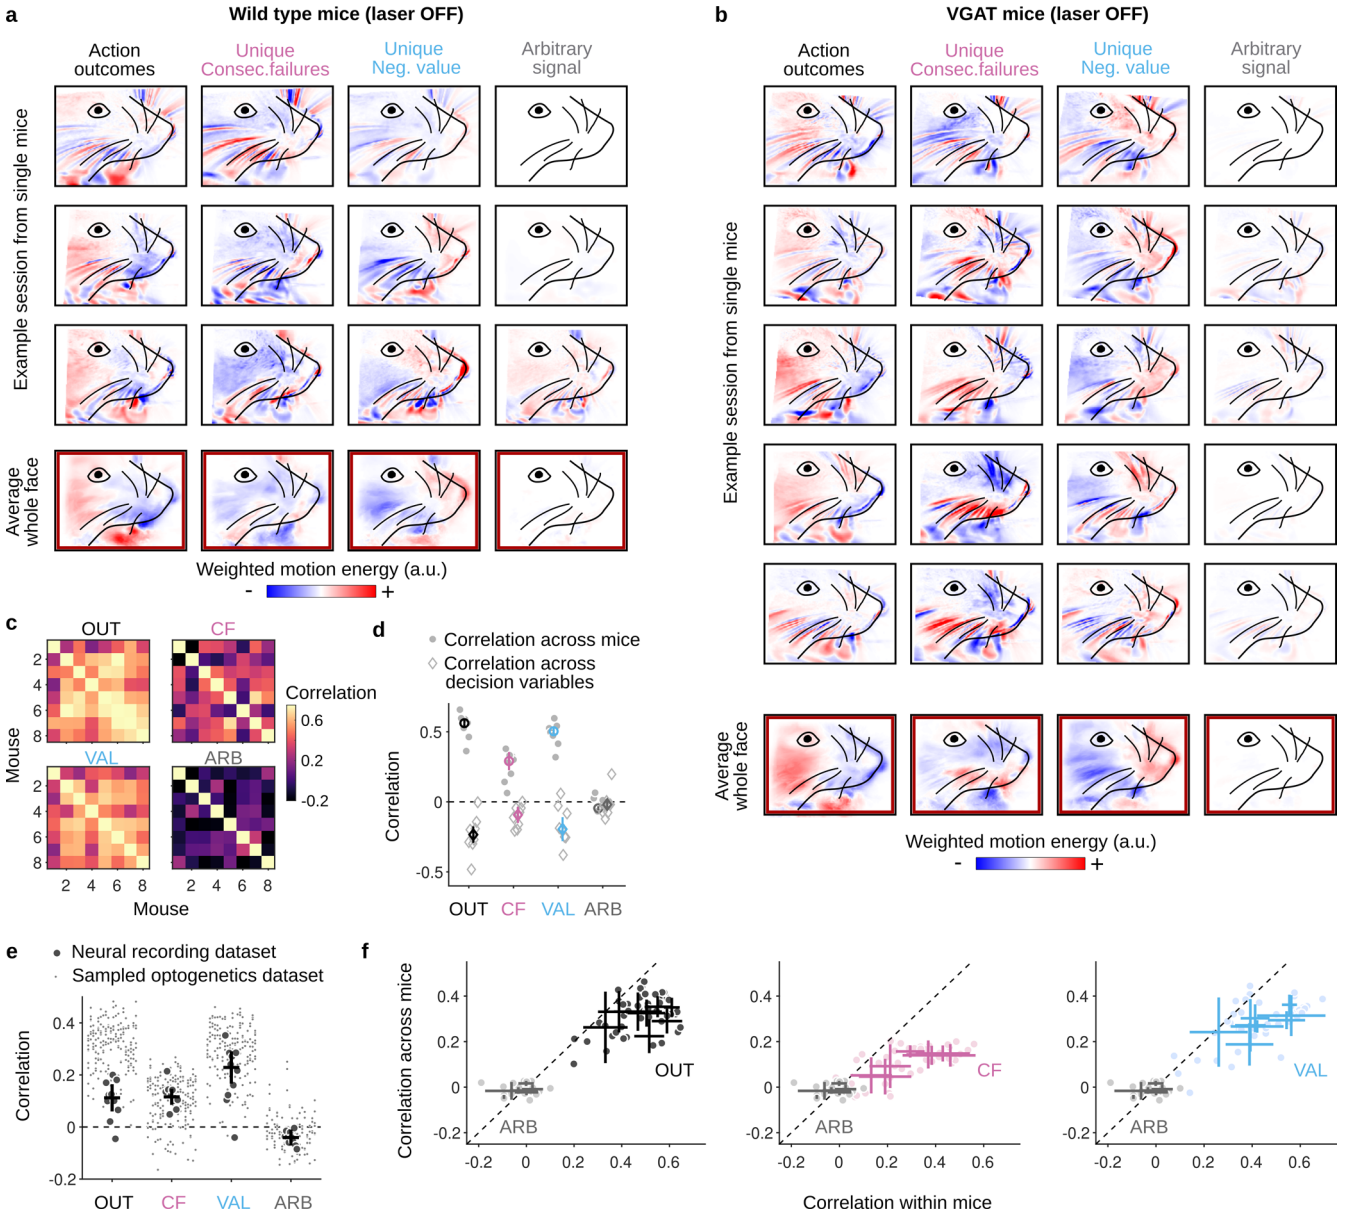

Extended Data Fig. 5 | Stereotyped facial expressions of decision variables in wild-type and VGAT mice. a, Weighted masks for a single example session (top) of each wild-type mouse (N = 3) and averaged across all sessions (N = 24, bottom) from all mice during the laser OFF condition. **b**, Same as in panel (a) but for a single example session (top) of each VGAT (N = 5) and averaged across all sessions (N = 24, bottom) from all mice during the laser OFF condition. **c**, Inter-animal similarity of facial expression of decision variable (Action outcome: OUT; Unique consecutive failure: CF; Unique negative value: VAL; Arbitrary signal: ARB). Colors represent the normalized 2D cross-correlation at zero lag between the mean weighted masks of two mice (for each mouse the mean weighted mask was the average across sessions, N = 8 mice, 6 ± 2.7 sessions per mouse). **d**, Mean weighted mask similarity for each mouse and decision variable. Each gray dot represents the average pairwise correlation of the mean weighted mask of a mouse for a given decision variable with the mean weighted masks of all the other mice for the same decision variable. Each gray diamond represents the average pairwise correlation of the mean weighted mask of a mouse for a given decision variable with the mean weighted masks of the same mouse for all the other decision variables. Color errorbars represent median and m.a.d across mice (N = 10, 8 sessions from distinct animals, two from the same animal). **e**, Distribution of average pairwise 2D correlations at zero lag for electrophysiology (data in Figs. 2–4) and downsampled optogenetics (data in Fig. 5) datasets. Pairwise correlations of facial expression of decision variables were calculated for all sessions in the electrophysiology dataset (black dots; median and median absolute deviation (m.a.d.) indicated). For comparison, 10 sessions were randomly sampled from the optogenetics dataset multiple times, and average correlations were calculated for each sample (gray dots). The sampling procedure was repeated twenty times, each represented as a row of gray dots. Note that correlations are overall smaller than the ones estimated using averages across sessions (rather than single ones) as in (d). **f**, Within- and across-mice correlation of decision variable facial expression. Each point represents the average correlation of a single behavioral session with all other sessions, either from the same mouse (within) or different mice (across). Errorbars represent mean and s.d. across sessions per mouse (N = 8 mice, 6 ± 3 sessions per mouse).

https://doi.org/10.1038/s41593-025-02071-5

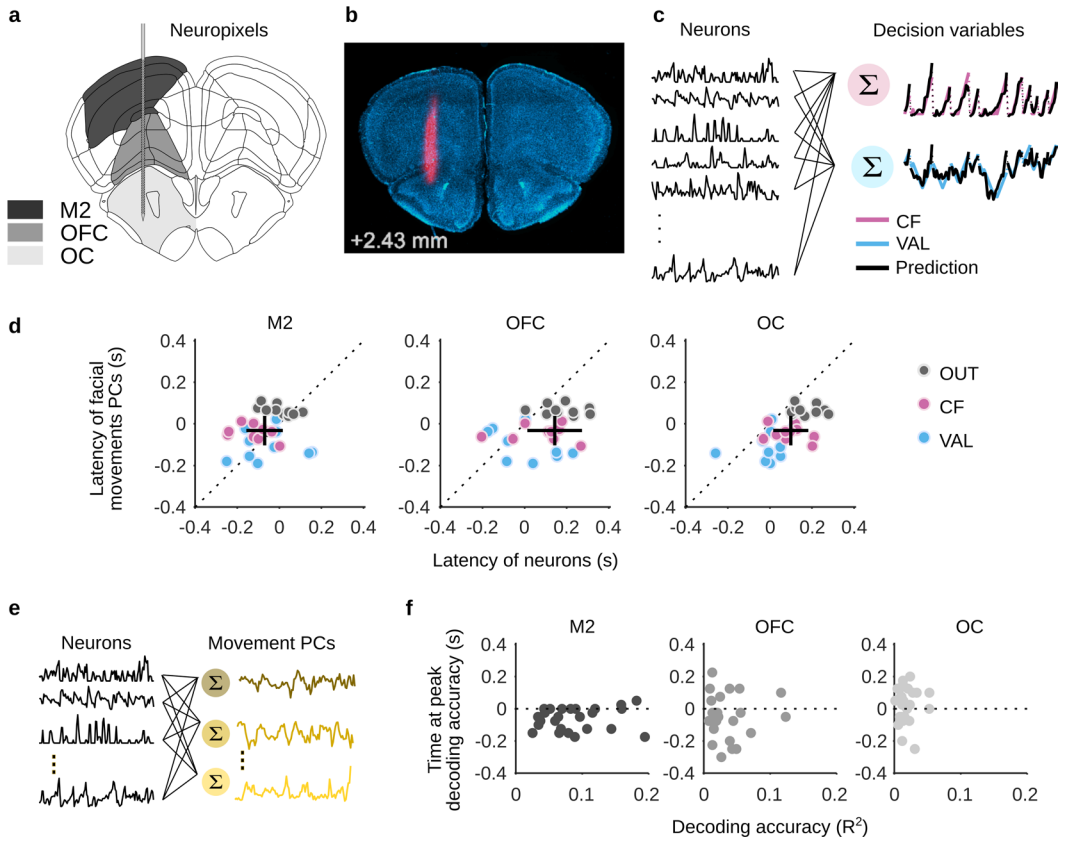

Extended Data Fig. 6 | Decoding multiple decision variables and facial movement from neural activity. a, We recorded with Neuropixels probes in multiple regions of the frontal cortex. Schematic target location of the neuropixels probe insertion. Vertical insertions were performed within a 1 mm diameter craniotomy centered around +2.5 mm anterior and +1.5 mm lateral from Bregma. **b**, An example of histology with the electrode track. We painted the probe with a red, fluorescent dye to recover the probe's location post-hoc. **c**, To decode the instantaneous value of multiple decision variables (pink & blue traces, right), we used regression models taking as predictors the activity of simultaneously recorded neurons in each brain region (black traces, left, example activity from M2). The model predictions (the weighted sums of neural activity, black trace right) overlap with the decision variables. **d**, Neural vs. facial movement PC decoding latencies for different decision variables (N = 10 sessions). Each session contributes three points (gray, pink, blue), one per decision variable. Points above the identity line indicate later facial movement representation relative to neural representation. Black cross: median ± m.a.d. across all points. **e**, Predicting facial movements from neural activity in M2, OFC, and OC. GLMs were trained to predict facial movement PCs using a 50 ms non-overlapping sliding window of lagged neural activity from M2, OFC, and OC. Facial movement PCs were derived from concatenated videos to enable cross-session comparisons. **f**, Relationship between decoding accuracy and the time of peak decoding accuracy for facial movements. Peak times (median of cross-validated R^2 across sessions, N = 10) are shown for the 25 PCs of facial movement with the largest variance, as a function of decoding accuracy from neurons in M2, OFC, and OC. Each dot represents one facial movement PC. Negative values, particularly in M2 (M2 = −0.05 ± 0.06, p = 0.002; OFC = −0.05 ± 0.10, p = 0.194; OC = 0.05 ± 0.09, p = 0.0194; median ± m.a.d, N = 10 sessions, Wilcoxon signed rank test, Holm-Bonferroni corrected), indicate that neural activity preceding the facial movement is most predictive.

https://doi.org/10.1038/s41593-025-02071-5

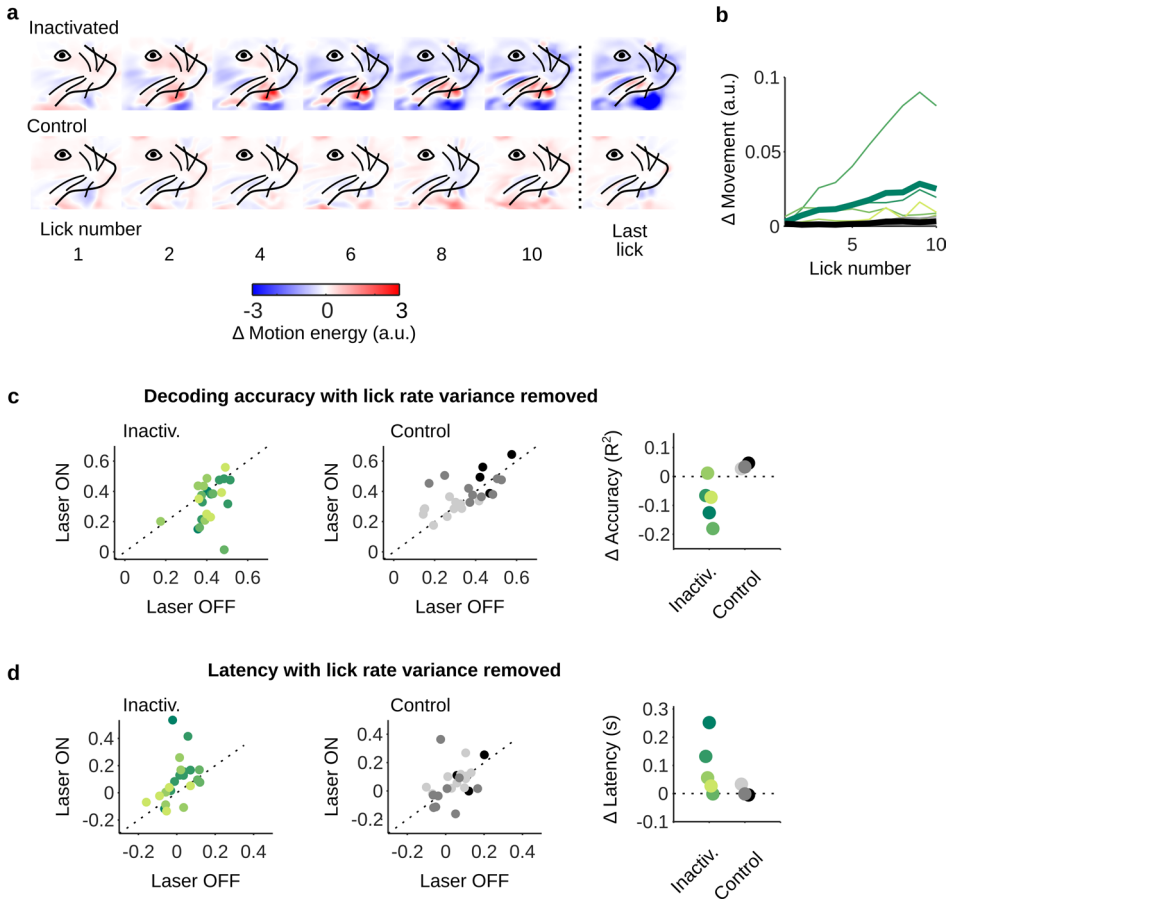

Extended Data Fig. 7 | Effect of M2 inactivation on movement and facial expressions of decision variables with lick rate variance removed. a, Laser-induced changes in facial movement patterns. 2D masks show the difference in average facial motion (calculated from movement PCs) between laser ON and OFF conditions, across mice and at different lick numbers (decision time points). **b**, Variability of laser-induced facial motion changes. The variance of the difference in motion energy between laser ON and OFF conditions (normalized by the variance in laser OFF) is shown as a function of lick number for inactivated (green, N = 5 mice) and control (black, N = 3 mice) groups. Thick lines represent group means. **c**, Decoding accuracy after removing the effect of lick rate (top & middle): Comparison of laser ON vs. laser OFF conditions. Dots below the unity line indicate that representations of decision variables derived from facial movement PCs were decoded less accurately during laser ON than laser OFF. Difference in decoding accuracy (bottom): Laser ON minus laser OFF (mean across sessions for each mouse in the inactivated and control groups). Individual mice are indicated by color. **d**, Same as in (c) but for decoding latency. Dots above the unity line indicate that representations of decision variables derived from facial movement PCs were decoded later during laser ON than laser OFF. Partial silencing of M2 reduced the accuracy and increased the latency with which facial movement PCs predicted decision variables, even after controlling for lick rate. This suggests that the latent variable represented in facial movements is not simply a consequence of the relationship between M2 activity and lick rate.

nature portfolio

Corresponding author(s): Fanny Cazettes

Last updated by author(s): Jul 16, 2025

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒ A description of all covariates tested
☐	☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☐	☒ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	We used the software package Bonsai (https://bonsai-rx.org/) for video collections and Arduino (Mega 2560), scientific boards (Champalimaud Hardware platform) and Matlab (2018a) for behavioral data collection. We used spikeGLX (https://billkarsh.github.io/SpikeGLX/, phase3A release) for electrophysiological data collection.
Data analysis	We used Kilosort2 (https://github.com/MouseLand/Kilosort) and Phy (https://github.com/cortex-lab/phy) for spike sorting analysis, FaceMap (https://github.com/MouseLand/facemap) for video analysis and custom Matlab (2019b) and Python codes for all the other analyses. For the GLM analysis we used a publicly available code (https://hastie.su.domains/glmnet_matlab/.) The code developed for the LM-HMM and the Expectation-Maximization algorithm can be accessed at: https://github.com/mazzulab/ssm/blob/master/notebooks/2c%20Input-driven%20linear%20model%20(LM-HMM).ipynb.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

The datasets generated during the current study has been deposited on Figshare: https://figshare.com/s/924af1de619f4597f37a

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	n/a
Reporting on race, ethnicity, or other socially relevant groupings	n/a
Population characteristics	n/a
Recruitment	n/a
Ethics oversight	n/a

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	A first group of animals was trained for the electrophysiology (N = 9). A second group of animals was the trained for optogenetics (N = 8). Sample sizes were selected based on our previous knowledge of mouse-to-mouse variability in the foraging behavior (Cazettes et al. 2021, 2023). No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those reported in previous publications (Musall et al., 2019 Nature Neuroscience, Stringer et al. 2019 Science).
Data exclusions	Animals that did not learn to perform the behavioral task were excluded from the study. For the electrophysiological datasets, animals were included only if all the Neuropixels probe was located in the target regions by post-hoc histological analysis. Single units (neurons) spike clustering quality were assessed manually using Phy. Units were only included if firing rate did not drift over the recording session, and spikes did not violate absolute refractory period (see Method).
Replication	All relevant behavioral effects were present in the majority of animals trained. Effects reported in neural data were consistent across animals. We replicated the results on the facial expressions of decision variable found in the first cohort of mice recorded with Neuropixels (N = 9 mice) on a second cohort (N = 8) used for optogenetics.
Randomization	The task statistics were similar across sessions but probabilistic. Thus, the sequences of events during behavior bouts were naturally randomized by the probabilities and the length of behavior bouts also depended on individual behavioral responses. Therefore, each session had a unique sequence of behavior bouts.
Blinding	The experimenter was blind to the genotype of the mice when performing optic fiber implantations and running the optogenetic experiments.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☐	☒ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other research organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research, and Sex and Gender in Research

Laboratory animals	Male and female C57BL/6J and VGAT mice (2-9 months old) were used in this study. All mouse strains were purchased from The Jackson Laboratory and bred in-house. Mice used in the electrophysiology experiments (N = 9) were of the C57BL/6J background (Jax stock number: 000664). For the optogenetics experiments, B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J mice (also known as VGAT-ChR2-YFP BAC, Jax stock number: 014548) were used. This included hemizygous VGAT-ChR2-YFP mice (N = 5) and wild-type littermate controls (N = 3). Mice were housed in individually ventilated cages under a normal 12 h light/12 h dark cycle, temperature was maintained between 19 °C and 23 °C and humidity between 50% and 65%.
Wild animals	The study did not involve wild animals.
Reporting on sex	Both male and female mice were used in this study.
Field-collected samples	The study did not involve samples collected from the field.
Ethics oversight	All experimental procedures were approved and performed in accordance with the Champalimaud Centre for the Unknown Ethics Committee guidelines and by the Portuguese Veterinary General Board (Direção Geral de Veterinária, approval 0421/000/000/2016).

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Plants

Seed stocks	n/a
Novel plant genotypes	n/a
Authentication	n/a