Inferring internal states across mice and monkeys using facial features

Alejandro Tlaie[1,2,✉], Muad Y. Abd El Hay[1], Berkutay Mert[1], Robert Taylor[1], Pierre-Antoine Ferracci[1], Katharine Shapcott[1], Mina Glukhova[1], Jonathan W Pillow[3], Martha N Havenith[1,*], and Marieke L Schölvinck[1,*]

[1]Ernst Strüngmann Institute for Neuroscience in cooperation with the Max Planck Society, Frankfurt am Main, 60528, Germany
[2]Laboratory for Clinical Neuroscience, Centre for Biomedical Technology, Universidad Politécnica de Madrid, Spain
[3]Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA
[*]These authors contributed equally.

Animal behaviour is shaped to a large degree by internal cognitive states, but it is unknown whether these states are similar across species. To address this question, here we develop a virtual reality setup in which male mice and macaques engage in the same naturalistic visual foraging task. We exploit the richness of a wide range of facial features extracted from video recordings during the task, to train a Markov-Switching Linear Regression (MSLR). By doing so, we identify, on a single-trial basis, a set of internal states that reliably predicts when the animals are going to react to the presented stimuli. Even though the model is trained purely on reaction times, it can also predict task outcome, supporting the behavioural relevance of the inferred states. The relationship of the identified states to task performance is comparable between mice and monkeys. Furthermore, each state corresponds to a characteristic pattern of facial features that partially overlaps between species, highlighting the importance of facial expressions as manifestations of internal cognitive states across species.

Internal states | Naturalistic behavior | Cross-species
Correspondence: ***atboria@gmail.com***

Introduction

In the wild, all mammals engage in similar foundational behaviours: they all hunt or forage for food, sleep, mate, avoid predators, and explore their environment, to name just a few. None of these behaviours can be simply explained as a passive reaction to environmental input; rather, they are crucially shaped by dynamic fluctuations in internal states such as satiety, alertness, curiosity or attention (1, 2). So, if fundamental behaviours are comparable across species, how similar are the internal states that drive them? Is 'attention' in a monkey the same as 'attention' in a mouse?

The common approach to investigate internal states has been a reductionist one: highly restrictive tasks featuring simplified stimuli and narrow behavioural repertoires (e.g. button presses), with little room for fluctuations over time (3–5). What is more, experimental paradigms diverge widely depending on the species under study. For example, attention studies in primates typically require subjects to fixate on a central fixation point while paying attention to a peripheral stimulus that might subtly change its appearance (6, 7). Attention studies in rodents, on the other hand, typically use the 5-choice serial reaction time task (5CSRTT), in which the subject is required to scan five apertures for the presentation of a brief light stimulus, and then navigate towards the light source (8, 9). Even though the behaviour associated with high attention, i.e. short reaction times and accurate responses, is the same in both cases, clearly these tasks are too different to draw any meaningful cross-species comparisons when it comes to the behavioural dynamics and neuronal mechanisms they may engage.

Breaking away from this restrictive regime towards studying internal states as they occur naturally is tricky. To tackle this challenge successfully, an ideal behavioural paradigm needs to (1) rely on innate, naturalistic behaviours to accurately reflect spontaneously occurring rather than training-induced internal states (10), (2) identify internal states in a data-driven way that is not restrained by (potentially anthropomorphising) concepts of cognitive processing imposed by the researcher, and (3) track the evolution of internal states over time to capture their intrinsically dynamic nature. For this, binary metrics of behaviour such as a button presses or nose pokes will not suffice; rather, precise, multi-parametric behavioural tracking is needed to generate time-resolved analyses that extract underlying cognitive states from the measured behavioural parameters moment by moment (11–13).

Recent technological advances have opened up new avenues to achieve these goals in a principled way. Virtual reality (VR) environments, for instance, allow researchers to create immersive yet highly controlled experimental settings that can be tailored to different species' intrinsic sensory capacities and behavioural repertoires (14, 15). Importantly, this maximizes adaptability across species, opening up the unique opportunity to record directly comparable behaviours in different species. At the same time, advances in deep-learning algorithms enable us to dynamically track ongoing changes in body movement and posture from video footage (16–18). This allows for the ongoing and time-resolved tracking of behavioural dynamics - a fundamental prerequisite if we aim to identify the spontaneous emergence of internal cognitive and emotional states (19, 20).

In this study, we leverage these technological breakthroughs to infer and directly compare the spontaneously occurring internal states of two species commonly studied in neuroscience - macaques and mice. Specifically, we combine a highly immersive and naturalistic VR foraging task solved well by both species (21–24), with a state-of-the-art deep learning tool that allows for precise, automated tracking of behavioural features. The features extracted in this way then

serve as inputs to a Markov-Switching Linear Regression (MSLR) model (25), which infers time-varying internal states across trials.

Importantly, such single-trial inference of internal states is only meaningful if the behavioural markers it relies on are not indirectly tracking the concrete motor outputs required for task performance. For instance, task-related motor output such as preparatory paw movements might trivially predict a hit trial, and a lack of such movements might predict a miss trial. To ensure that the behavioural parameters we chose would truly reflect internal processing, we focused on the animals' facial expressions.

While facial expressions have long been thought to only play a role in highly visual and social species like monkeys and humans (26–29), recent work has highlighted that also less social, less visual species like mice exhibit meaningful facial expressions (19, 30). These expressions seem to reflect fundamental emotions like pleasure, pain, disgust and fear in a way that is not only consistent within one species, but also readily translatable across species (31, 32). This argues for an evolutionary convergent role of facial expressions in reflecting (and potentially communicating) emotions.

Outside the realm of emotional processing, spontaneously occurring behavioural states have so far mainly been tracked using individual facial features to identify isolated cognitive processes. For instance, pupil size has been related to changes in attention and arousal in rodents (33–35), non-human primates (36–40), as well as humans (41, 42). Similarly, eye movements in monkeys and humans (43–45) and whisker movements in mice (46) have been used to track attention and decision-making.

In addition to such individual facial features, movement intensity across the entire face (regardless of individual features) has been used to delineate, for instance, the contribution of instructed versus uninstructed movements to task performance (47, 48), and to explain neural activity (49, 50). Orofacial motion is also a key component in explaining trial-to-trial variability in dorsal cortex activity in mice (51), especially when mice are disengaged from an instructed task (48). More detailed analysis into task performance states based on such facial features revealed that optimal visual discrimination coincides with average pupil size and moderate levels of facial movement (52).

While these findings highlight promising ways to track internal states - both in the emotional and cognitive domain - using facial expressions, it stands to reason that facial expressions contain richer and more complex information than what can be extracted using either individual facial features or overall orofacial movement. By focusing on a multi-faceted representation of entire facial expressions, we aim to for the first time map out the spectrum of spontaneously occurring internal states in a data-driven way that is directly comparable across species. Our approach of using facial expressions to infer spontaneously occurring internal states from natural behaviour constitutes a drastic departure from the classical approach, which instead imposes internal states through restrictive behavioural tasks (e.g. cued attentional shifts). By tying the results of this approach back to known relationships between internal states and overt behaviour, such as shorter reaction times during focused attention, these spontaneous, agnostically inferred internal states can be tentatively related to known cognitive processes such as attention and motivation. Importantly, this puts us in the unique position to directly compare spontaneously occurring internal states across species by minimizing confounds introduced by species-specific tasks and training procedures.

Results

A. Experimental set-up. To track and compare spontaneously occurring internal states of mice and macaques during the performance of the same naturalistic visual discrimination task, the animals were placed inside a custom-made spherical dome (Fig. 1A, top). On the inside of the dome, we projected a virtual reality (VR) environment using a custom-made toolbox called DomeVR (53). The monkeys navigated through the VR environment manually using a trackball; the mice ran on a spherical treadmill, the movements of which were translated into VR movements (for details, see Methods - *Experimental Setup*).

Two monkeys and seven mice were used in this study, comprising 18 and 29 experimental sessions (20459 and 12714 trials) respectively. The animals engaged in a simple, foraging-based two-choice perceptual decision task, in which they had to approach a target stimulus while avoiding a distractor stimulus, both of which were represented by natural leaf shapes integrated in a meadow landscape (Fig. 1A, bottom; see Methods - *Behavioral paradigm* and *Behavioral Training*). Such tasks can be performed successfully by both monkeys (23) and mice (21, 54). Their performance on this task was quantified first in terms of trial outcomes: hit (target stimulus reached), wrong (distractor stimulus reached), and miss (neither stimulus reached); as well as in reaction time (RT). For this, we identified turning points in the animals' running trajectories through the VR to define the moment when an animal decisively oriented itself towards one of the two potential targets (Fig. 1C; for details, see Methods - *Reaction Time*). As Supp. Fig. S1 shows, success rate and reaction times were largely comparable across species, although mice showed less consistent performance than monkeys, in terms of running trajectories, reaction times, and correct target choices. We hypothesize that this is due to the lack of fine motor control of the mice on the trackball.

As the animals were performing the task, we recorded their faces. For macaques, this was done by analysing video footage from one camera positioned frontally on the monkey's face, as well as eye tracking output (see Methods - *Behavioural tracking*). For mice, we analysed video footage from one camera positioned on the side of the face (Fig. 1B). From these videos, we extracted facial features such as eyebrow, nose and ear movement using DeepLabCut (Fig. 1C; see Methods - *Facial key point extraction*). For monkeys, we selected 18 features; for mice, 9 features (see Methods - *Face features* for the full list of facial features).

For each trial, facial features were averaged over a time win-

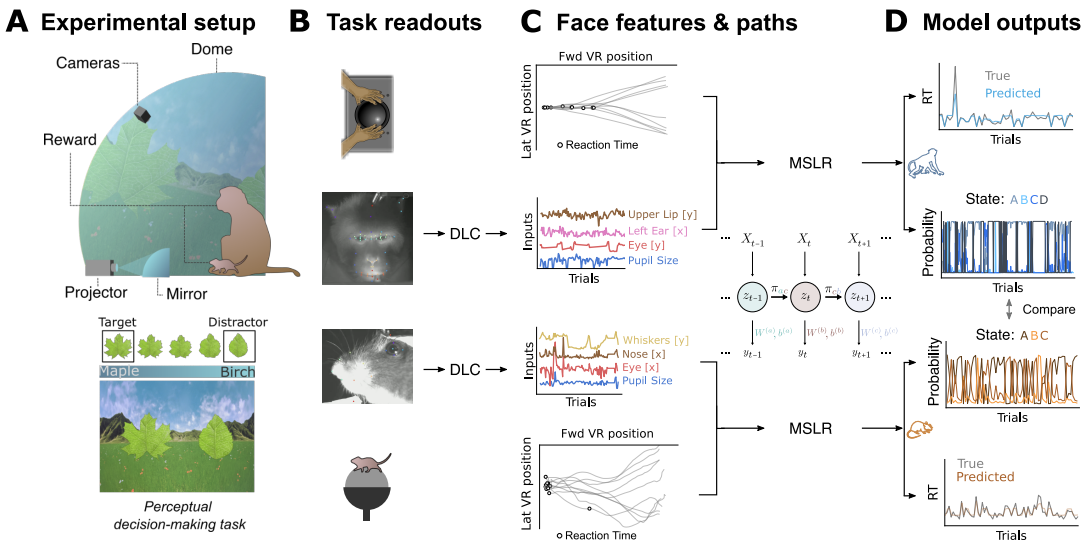

Fig. 1. Experimental setup and computational pipeline. A) Macaques and mice were seated inside a large dome on the inside of which a VR was projected via a curved mirror (top). They were rewarded for moving towards a spike-shaped leaf compared to a round-shaped leaf (bottom). **B)** As the animals were engaged in the task, behavioural data were collected: movements of the trackball (top and bottom) and videos of their faces (middle). **C)** Trackball movements were translated into paths through the virtual environment (top and bottom), from which reaction times were determined (see Methods). Individual facial features were automatically detected from the videos and tracked over time (middle). **D)** Facial features entered two separate MSLR models (one for each species), which yielded, for every trial, a predicted reaction time and internal state probabilities.

dow of 250 ms before the stimuli appeared in the VR environment. This time window was chosen to maximize the interpretability of the inferred hidden states: as there is no task-relevant information available yet, presumably all of the facial expressions that the animals make are due to internally generated processes, rather than being reflective of stimulus or task parameters.

B. Model performance. The facial features extracted in this way were used as inputs to a Markov-Switching Linear Regression (MSLR) model (Fig. 1C; see Methods - *MSLR*). We used a recently developed software package (55) to implement the MSLR. The MSLR manages to capture the non-stationarity and regime shifts often present in behavioural data (52, 56–58), by flexibly accommodating complex temporal dynamics while keeping a relative simplicity, compared to deep learning-based methods (59). Moreover, the MSLR is less data-hungry than other common data-driven models (60–62). The MSLR uses the 'pre-stimulus' facial features in each trial to predict the animals' reaction time (RT) in the same trial by assuming 'hidden' states. Each hidden state implies a different linear relation between individual facial features and the subsequent RT in the same trial. For instance, in one hidden state, the RT might be best predicted by eyebrow movements, while in another, nose sniffing might be most predictive. We used cross-validation to select the number of states for each species (see below). For each trial, the model then outputs the predicted RT as well as the probability of each hidden state (Fig. 1D). The two models (one for mice, one for monkeys) were trained and tested on data from all individuals; Fig. S20 shows the outcomes of the models split by session and by individual and Fig. ?? shows the generalisation across mice.

Mathematically, this model takes the form:

$$RT_t = W_{z_t} \cdot x_t + \xi_{z_t}, \quad (1)$$

where RT_t is the reaction time at trial t, z_t is the state at trial t, W_{z_t} are the regression weights for state z_t, x_t is the vector of facial features at trial t, and ξ_{z_t} is a zero-mean Gaussian noise with variance σ_{z_t}.

To test if this approach was appropriate for our behavioural recordings, we first checked if assuming the presence of multiple hidden states was in fact warranted by the data, or if they could also be described by one constant, uniform relationship between facial expressions and RTs over time. To this end, we determined model performance when only one internal state was permitted (Fig. 2A; see also Fig. S14). For both species, the model's predictive performance was remarkably low under these circumstances - in fact, predictions were less accurate than random guessing.

Next, we quantified model performance for different numbers of hidden states - which is the main free parameter of the MSLR. Model performance was tested by using cross-validation (see Methods - *Model tuning*). Figure 2A shows the cross-validated R^2 for different numbers of states and for optimal (solid line) and suboptimal (shaded area) constellations of hyperparameters. For both species, the cross-validated R^2 improved dramatically when allowing for more than one hidden state until reaching a plateau. Since the accuracy of RT predictions began to saturate with increased model complexity, we took the finite difference of the CV performance curve for each species and fixed the number of internal states at its maximum, in order to reach the optimal trade-off between predictive accuracy and model simplicity (Fig. 2A; see Fig. S11 for a more detailed explanation of this). This approach yielded a similar optimal number of hidden states for both species: for monkeys, the optimal number of states was 4, for mice it was 3. Tests on held-out data

showed a similar performance (Fig. 2A, insets), which, coupled with Supp. Fig. **??**, indicate that the model generalises. The CV procedure ended up selecting hyperparameters that mainly differed in the concentration parameter (α), controlling how sparse each transition matrix was. For more details, see *Methods - Model tuning*.

In both species, and across all individual animals, our models yielded remarkably accurate trial-by-trial predictions of RT (see also Fig. **??**), indicating that pre-trial facial expressions can indeed predict subsequent task performance (Fig. 2B, top row). It also suggests that the relation between facial features and task performance is dynamic rather than static over time, reflecting multiple underlying states. Finally, model performance was not impacted by removal of variables associated with pupil size, indicating that the performance did not rely trivially on known correlates of arousal (Fig S16).

Is task performance dominated by a single state at any given moment, or do several states co-exist continuously? After fitting the model parameters, we used the model to identify the animal's internal state on a trial-by-trial basis. Note that the model does not allow for the animal to be in multiple states at the same time; rather, it gives us probabilities telling how confident we can be about the state the animal is in on each trial. Specifically, we computed the posterior probability over states on each trial given all past and future observations. The probabilities of each state over time suggest that the model is highly confident about what state the animal is in on each trial (Fig. 2B, bottom row). These observations were confirmed by the highly bimodal distribution of these probabilities for both species (Fig. 2C). Crucially, in monkeys, this separation between high-certainty ($p_s \approx 1$) and low-certainty ($p_s \approx 1/n_s$) trials was particularly pronounced, while in mice, state probabilities were somewhat more mixed. Quantifying the single-trial certainty as measuring its difference with the uniform distribution –through the Kullback-Leibler divergence (KL)– corroborated these findings (Fig. 2D; Mann-Whitney U-test: $p = 1.11 \cdot 10^{-274}$). As such, the hidden states identified by our model seem to reflect largely mutually exclusive behavioural modes that animals switch in and out of. Given how consistently trials were dominated by one state, we chose to binarize hidden state outcomes by assigning each trial to its most probable hidden state.

C. State dynamics. To explore if the hidden states showed attributes that could be reflective of internal cognitive states, we first characterized their temporal dynamics. To this end, we examined the frequency of state transitions in both species. The state transition matrices, which show how likely a trial of a given hidden state is followed by a trial of any (other or same) state (Methods - *Markov-Switching Linear Regression*), revealed high values along the diagonal for macaques, indicating stable states that switched rather rarely. In mice, the diagonal of the transition matrix was slightly less pronounced, suggesting that hidden states in mice were less stable and more prone to transition than in macaques (Fig. 3A). This was confirmed by the learned hyperparameters of both models: These were highly comparable across mice and

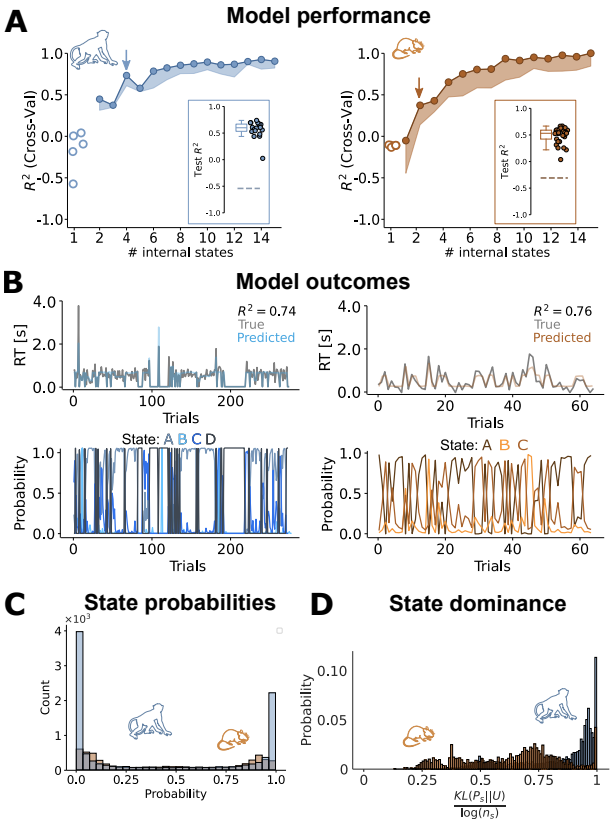

Fig. 2. Model performance and state probabilities. A) Cross-validation performance for various numbers of states, for macaques (left) and mice (right). Circles indicate the maximum CV R^2 and the shaded region extends until the 5th percentile. For both species, increasing the number of states improves model performance to a plateau at an $R^2 \approx 0.8$. *Lasso* is a regularized Linear Regression (i.e., a MSLR with 1 internal state). The arrows indicate the number of states we selected, based on the maximum difference of the CV performance curve (see Fig S10). Insets show model performance for held out data at the selected number of states; dashed horizontal lines indicate the 99th percentile of the surrogate performances (see Methods). Note that the shuffled R^2 is negative, because only uncorrelated predictors are expected to be centered at 0, and due to finite sampling effects, there is always a non-zero correlation between the shuffling and the ground-truth. Furthermore, as we are dealing with skewed distributions (see Fig. S1), the null tendency is not captured by the mean, as assumed by the default R^2. **B)** Predicted RTs (top) and state probabilities (bottom) for an example stretch of data (left, macaques; right; mice). **C)** Probabilities of all states over all trials, regardless of state identity (blue, macaques; orange, mice). The bimodal distribution suggests that states are either absent or dominant on any given trial. **D)** Kullback-Leibler divergence (KL) for monkey (blue) and mouse (orange) internal states. KL quantifies the difference between the posterior state probability under the model and the uniform distribution, normalizing by the number of states. A KL value close to 1 indicates maximally dissimilar distributions (i.e., only one present state at a time), while a value close to 0 indicates indistinguishable distributions (i.e., equally likely states).

monkeys with the exception of the concentration parameter, which governs the sparseness of transition probabilities and was markedly lower for monkeys than mice (see Methods - *Model tuning* for details and for the values of the parameters we selected).

As a complementary analysis, we computed the dwell time for each state. This quantity is defined as the number of consecutive trials that a given state is occupied for, before transitioning to a different state. As dwell times are dependent on absolute numbers of trials spent in each state (a state with very few trials overall cannot have long dwell times), we plotted the dwell times alongside the number of trials spent in each state (Fig. 3B). Supporting the previous observations, hidden states lasted generally longer in macaques than in mice (Mann-Whitney U-test; $n_{mac} = 4092$, $n_{mice} = 2543$ trials, $p = 0.0014$), suggesting that internal processing may be more steady in macaques. This is consistent with previous findings that behavioural dynamics may fluctuate faster in mice (34, 63) than monkeys (64). Apart from a genuine species-driven difference, this observation may also reflect the fact that monkeys are trained more extensively and may therefore have developed more stereotyped behavioural strategies than mice, which were trained more briefly.

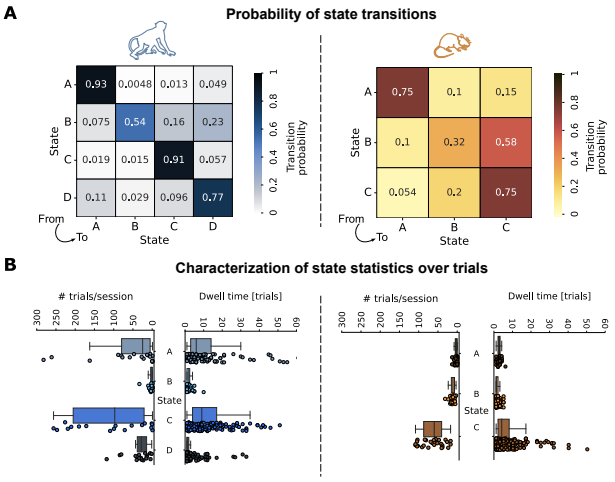

Fig. 3. State dynamics. A) State transition matrices for macaques (left) and mice (right), that show the probability, at any one trial, of transitioning from a certain state (rows) to any other state (columns). Transitions between different states (off-diagonal terms) are more frequent for mice than for macaques. **B)** Macaques (left plots) spend more time than mice (right plots) in the same state, as measured by the dwell time (number of consecutive trials of each state being the most likely one). Context for these dwell times is given by the trial counts (per session) on the left of each dwell time plot. Individual dots in the dwell times reflect sequences of consecutive trials of a particular state.

D. Hidden states as performance states. To link the identified hidden states more concretely to internal cognitive processing, we set out to investigate how each hidden state related to behavioural outcomes, starting with the RTs that the model was trained to predict. There are two potential scenarios for how the model might partition RT variability: on the one hand, it is possible that each hidden state covers the full range of RTs, but predicts them from a different constellation of facial features. Alternatively, each hidden state might 'specialize' on predicting specific ranges of RTs. For example, one hidden state might cover facial features that distinguish between fast and extremely fast RTs, while another state mainly predicts variations between slower RTs. This second scenario would make it more likely that the identified hidden states reflect genuinely distinct performance states.

To distinguish between these scenarios, we plotted the overall state-specific RT distributions, pooling trials across all sessions and animals, for each hidden state (Fig. 4A; top panel; Fig. S21 shows the same plot for individual sessions and animals). The resulting distributions support the second scenario: while one hidden state (state B in both monkeys and mice) covered a rather broad range of RTs, all other states showed a distinct profile of response speeds. This implies that the hidden states relate to distinct performance regimes (in this case in terms of response speed), making them viable candidates for defining specific internal states of cognitive task processing.

To further probe the possible link of our internal states to known cognitive processes, we related all hidden states to the three possible trial outcomes of the task (hit, wrong, and miss; see Methods - *Task performance and internal states*). Crucially, given that we trained the model to predict RTs, it never received any explicit information about trial outcome. Furthermore, RTs were only marginally related to trial outcomes (Fig. S1), so that trials with a specific RT would not be significantly more likely to result e.g. in a hit or a miss trial. Finally, as we only used information about facial features in the pre-stimulus phase of the trial to train the model, it cannot reflect stimulus features.

Even though information about trial outcomes was not part of the MSLR model, the resulting hidden states were consistently predictive of specific trial outcomes (Fig. 4A, bottom panel). For instance, in monkeys, trials that were classified as belonging to state C were most likely to result in a hit, while trials from state A often resulted in incorrect responses, even though the RT distributions of both states overlapped strongly. The same dynamic can be observed in states A and C in mice.

Adding more states to the models and subsequently predicting RTs and task outcome invariably resulted in several states strongly overlapping in their predictions. For example, using 10 states in monkeys and 8 states in mice (models with best cross-validated performance, see Fig 2A) resulted in two groups (of 2 and 6 states) in monkeys and two groups (of 2 and 3 states) in mice, where the relationship of the states to RTs and task outcome was very similar (Fig S10 and S12). We therefore conclude that the task performance space is best covered by the three states for mice and four states for monkeys.

Linking internal states to RTs and trial outcomes revealed that individual states covered unique combinations of speed and accuracy. To visualize these combinations, we plotted mean RT per hidden state against the difference in probability of a hit versus a wrong trial for each state. Interestingly, the constellation of states in this space was comparable across species (Fig. 4B). Both mouse and monkey data seem to generate one hidden state (state A in mice, state

C in monkeys) that is associated with fast RTs and largely successful trial outcomes - a performance regime that could be interpreted as globally attentive. Conversely, state C and A in mice and monkeys, respectively, map onto rather fast yet often incorrect responses, potentially reflecting more impulsive decision-making (21, 65). Finally, state B for both species features particularly slow RTs, large RT variability, and mostly incorrect trials for mice and equally likely trial outcomes for monkeys, potentially signifying a state of global inattention (66, 67). The only state that appears in monkeys but not mice (state D) features no reactions at all (i.e. no change in path direction) and only misses; a sign of complete task disengagement.

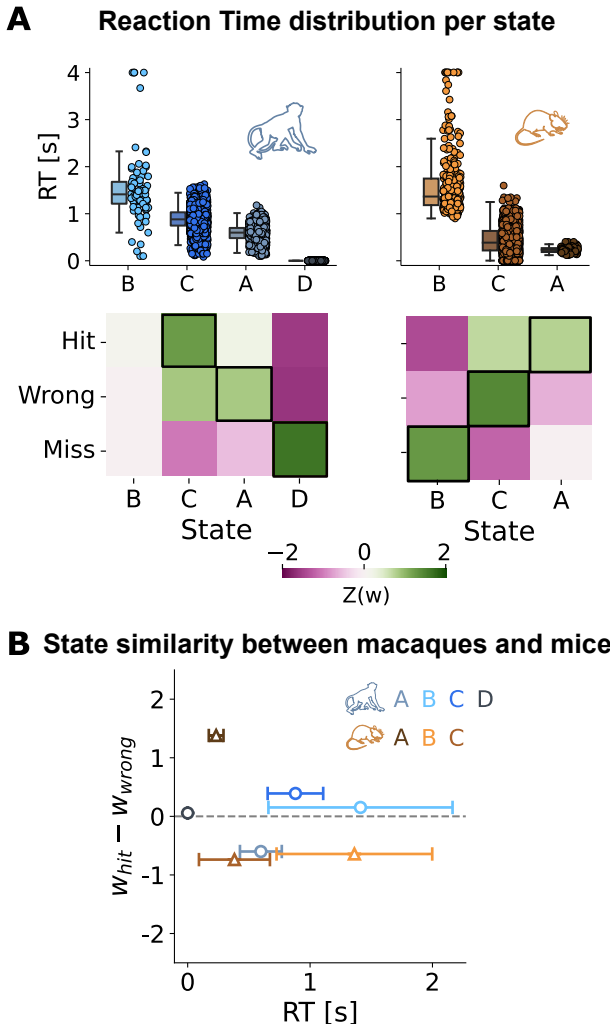

Fig. 4. Internal states and task performance. A) (Top) Splitting the RTs over internal states shows large diversity for both macaques (left) and mice (right), from fast reaction-states to extremely slow ones. Individual dots reflect trials. (Bottom) Correlations of state probabilities with the three task outcomes (hit, wrong, miss), for macaques (left) and mice (right). Black boxes indicate the states most strongly associated with a certain task outcome. **B)** Conjunction of RT and excess likelihood of a hit outcome, for all states (blue circles, macaque; orange triangles, mouse).

E. Relationship to facial features. A final clue towards the interpretation of our internal states might be given by the facial features from which they are inferred. To explore this possibility, we plotted the regression weights of all facial features for the hidden states associated with hit, wrong, and miss trials (Fig. 5A; for the facial features comprising the fourth state in the monkey, see Fig. S13). These plots reveal highly distinct contributions of different facial features to each internal state. For example, in mice, pupil size and whisker movements predict reaction speed in the 'hit' state, whereas nose movements predict reaction speed in wrong and miss states. Similarly, in monkeys, large pupil size predicts fast reactions in hit and wrong states, but slow reactions in miss states, and eye movements play a strong predictive role in hit and wrong, but not in miss states. One interpretation of these observations is that different sensory modalities may be more dominant in driving decision making (and thereby decision speed) in different states. Especially in mice, one of the hallmarks of the hit state is that it is the only state in which pupil size plays a decisive role, suggesting that in other states, behaviour may be less strongly driven by sensory sampling from the visual domain (and more by sampling e.g. from the olfactory domain).

Interestingly, the facial constellations predicting RTs in the states most closely associated with hit and wrong trials, respectively, are quite similar in monkeys, but not in mice. This may suggest that in monkeys, the behavioural state underlying correct and wrong trial outcomes may be a generally engaged and high-attention state, and hit or wrong outcomes are mainly dictated by visual difficulty than different internal cognitive state. In contrast, in mice, hit and wrong trials may be the product of more distinct underlying cognitive states, e.g. in terms of attentive capacity (see Supp. Fig. S19 for a summarized visualization).

This notion becomes even more apparent when focusing only on the facial features that both species have in common (pupil size, and eye and nose movements), as shown in Fig 5B. Surprisingly, the direct comparison of shared facial features across species shows that the contribution of individual facial features to the three performance states is highly overlapping, particularly in hit and miss states (Fig 5B). Shuffling the features and computing the mouse-monkey overlap 10.000 times, revealed that this overlap is extremely strong in the hit state, quite strong in the miss state, and indistinguishable from chance in the wrong state (at 97.9, 86.2, and 56.7 percentile of the distribution, respectively). This result further underlines the notion that unlike hit and miss trials, wrong trials may be the outcome of genuinely distinct cognitive states in mice and monkeys. More globally, the strong match of facial features especially in hit and miss states indicates that facial expressions may be evolutionarily preserved across species not only in terms of emotional (32), but also in terms of cognitive and attentional processes.

Do the internal states reflect global facial constellations, or could they be inferred equally well from individual facial features such as pupil size? To answer this question, we plotted the contribution of individual facial features across all inferred internal states. (Fig 5C shows the distribution of model weights across all facial features. Some features (e.g. the upper lip in monkeys or vertical whisker movement in mice) did not contribute significantly to predicted performance in any internal state. For more prominently predictive facial features

(e.g. right ear or pupil size in monkeys and vertical eye movements or nose movements in mice), there was a wide spread of model weights across states. This indicates that no single facial feature was predictive of performance across all states. Rather, in both species, reaction times are best predicted by a complex and variable constellation of facial features that cannot simply be reduced to individual components. This finding is further supported by the observation that different states are consistently distinguished by more than one facial feature (Fig 5A; see also Fig.S9), and that removing one of the most prominent predictive facial features - pupil size - together with all facial variables significantly correlated to it, did not impair model performance in any way (see Figs. S15 and S16).

Together, these results suggest that 1) holistic analysis of complex facial expressions is much more informative than analysis of one isolated facial feature such as pupil size and 2) the relationship between facial features and cognitive processing is not linear, but changes depending on the internal state that the animal is in. For instance, in a high-performance state, large pupil size may indeed predict trial success (as shown e.g. by (33, 36)), whereas it may be irrelevant or anti-correlated in a low-performance state (see e.g. contribution of pupil size to hit versus miss states in monkeys, shown in Fig. 5A).

Lastly, we aimed to establish whether the constellations of facial features associated with each internal state were sufficiently distinct to infer internal states in individual trials. To this end, we classified the facial read-out from each trial as part of a hit, miss or wrong state, based on its correlation to the corresponding hit, miss, and wrong facial profiles (Fig 5D). In both species, trials were classified as belonging to their rightful internal state far above chance level (Fig 5D). This suggests that in both species, there are indeed characteristic facial expressions reflecting e.g. attentive or disengaged cognitive states, that are distinctive enough to be successfully identified on a single-trial basis in a majority of trials.

F. Influence of trial history. One reason why hidden states can predict trial outcomes so accurately despite not being trained on them in any way might be that pre-trial facial features are mostly a trivial consequence of the animal's trial history. For example, facial features might mainly reflect an animal still drinking reward from the previous trial, which might in turn raise motivation to perform correctly in the upcoming trial. In this case, facial features would merely be a particularly convoluted way of quantifying the previous trial outcome, and using it to predict upcoming performance, as has been achieved previously (68, 69). To account for this possibility, we trained an Auto-Regressive Hidden Markov Model (ARHMM) based on RTs (see Methods - *ARHMM* for details). As can be seen in Supp. Fig. S6, the facial features model outperforms the ARHMM for all states, for both species.

As an extra control, we correlated each facial feature with the history of prominent task parameters, specifically two related to the directly previous trial (its outcome, which might affect motivation; and the location of its target, which might predict side biases), and two related to the overall session history (the cumulative amount of reward and the time that passed since the start of the session, as proxies for satiety and fatigue, respectively). Correlations between task variables and facial features were sparse in both species (Supp. Fig. S17). In fact, attributes of the previous trial did not relate significantly to facial features at all, and more sustained session attributes modulated facial features merely somewhat. This suggests that facial features may be modulated by ubiquitous internal processes like fatigue and satiety, which are in turn impacted by task parameters, but they are not a trivial reflection of task history. Rather, the fact that facial expressions are modulated by the overall task context makes them a more plausible reflection of realistic fluctuations in cognitive processing.

Discussion

Internal cognitive states are known to substantially shape overall brain activity (70, 71) as well as behavioural decision making (1, 2), yet they are notoriously difficult to identify. As a result, it is even less clear to what extent they converge across species. To infer hidden cognitive states in mice and monkeys, we harnessed an MSLR model (25) trained on their facial features while they were engaged in an immersive VR foraging task. Specifically, we trained the MSLR to predict an animal's reaction time (RT) in a given trial based on its facial expressions prior to stimulus presentation. For both species, RTs could be predicted on a trial by trial basis with high accuracy from preceding facial features only. This suggests that facial expressions reflect internal parameters that are directly relevant to task performance from moment to moment. These parameters were only minimally shaped by task history, suggesting that they were not a trivial reflection e.g. of the previous trial outcome.

Even more surprisingly, this approach revealed multiple distinct hidden states, which were characterized by equally distinct relationships between a complex constellation of facial features, and subsequent task performance. In different states, performance seemed to be dominated by specific sensory modalities, e.g. eyes versus nose for hit versus wrong states in mice. This suggests that, depending on an animal's internal state, the relation between facial features and subsequent task performance can shift dramatically, and there are no individual facial features that reliably predict behaviour across different internal states. As such, our results emphasize how data-driven characterization of internal states can highlight crucial features that might be missed when using experimenter-defined states (see also (72)).

These findings stand in marked contrast to previous research studying mostly linear relationships between single facial features (e.g. pupil size or eye movements) and isolated cognitive states (e.g. attention) pre-defined by the experimenter (33, 34, 36–39, 43–45). They also go beyond recent studies highlighting the relationship between motion of the entire face (regardless of the facial features), and task performance (Musall et al 2019; Hulsey et al 2024) or neural activity (Syeda et al 2023; Salkoff et al 2020; Talluri et al 2023). This highlights the fact that information reflected in the entire

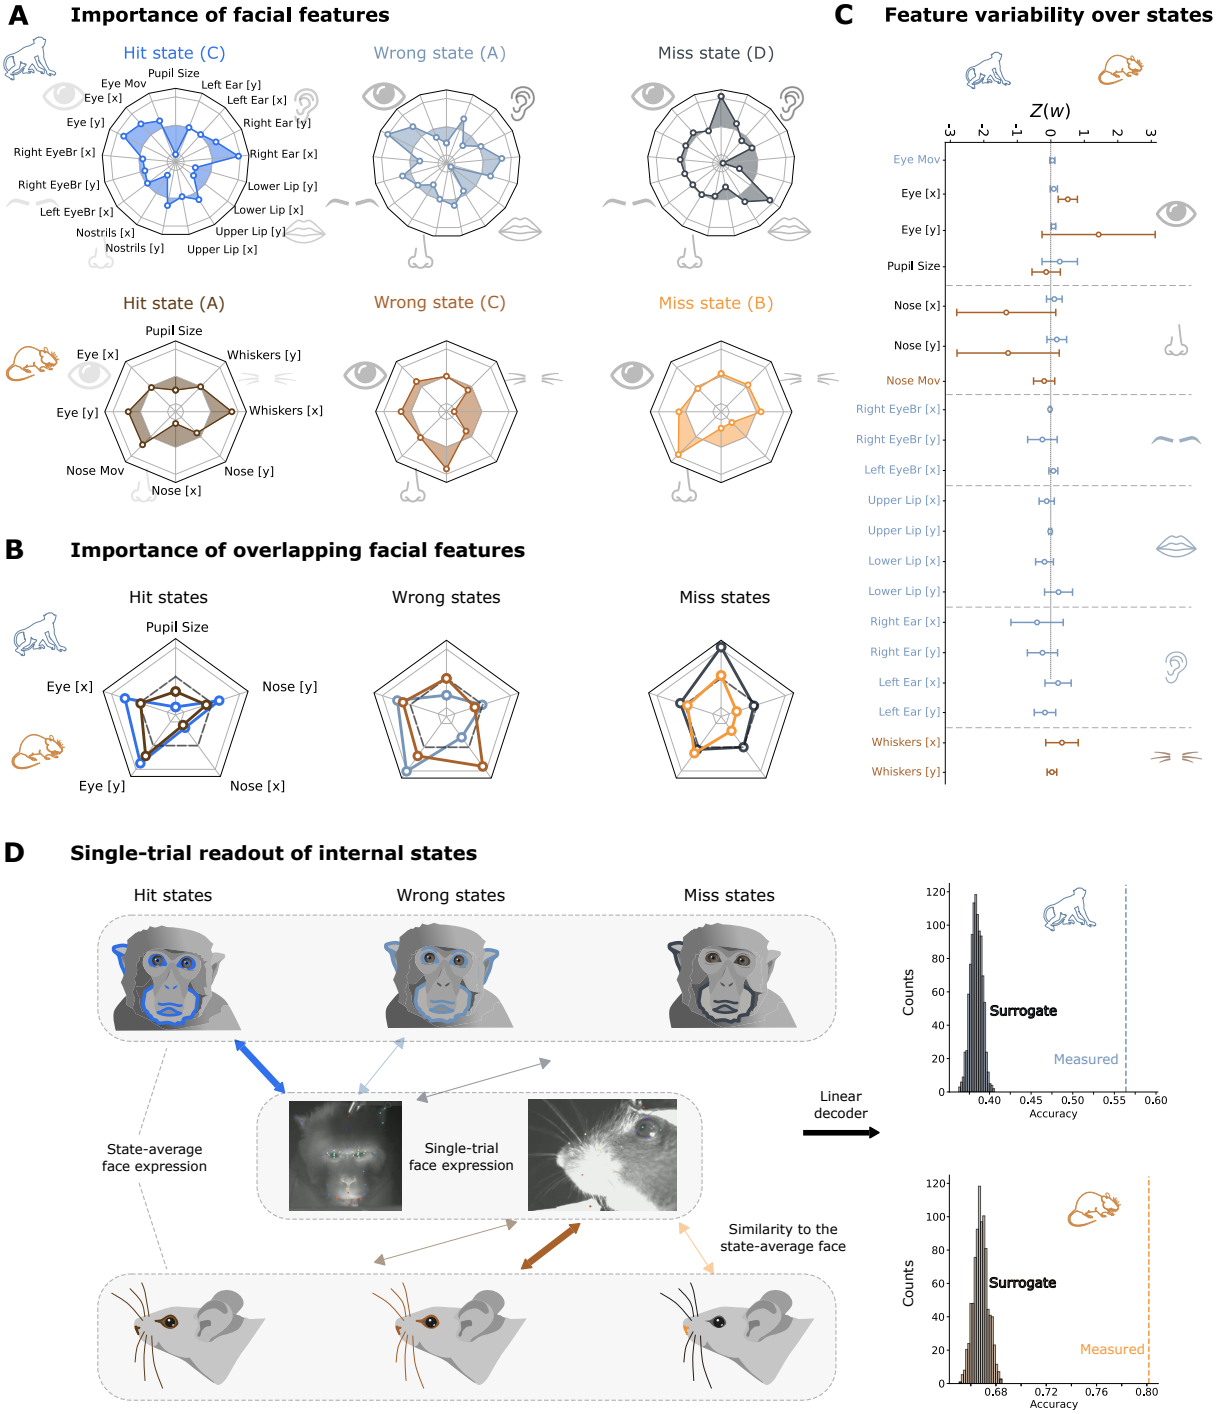

Fig. 5. Informativeness of facial features A) Predictor weights of the facial features for the macaque (top) and mouse (bottom) model in the hit, wrong and miss states (see black boxes in Fig 4A, bottom). Central circle indicates a predictor weight of zero, inside this circle are negative predictor weights, outside are positive weights. Note that a positive weight implies a negative relationship with performance and vice versa (for example, the strong negative weight on pupil size for monkeys in the hit state means that the bigger the pupil, the slower RTs it predicts). Each state has its own characteristic facial expression pattern. **B**) Plotting only the facial features that overlap between the species reveals strong similarities in their importance to the three performance states. **C)** Variability of all facial features over states. Although some features contribute more than others, clearly all features contribute to the model distinguishing between the various internal states.**D)** As seen in panel **A)**, each state has a distinctive facial expression pattern. Further supporting this, we show that we can decode states on single trial states based on similarity to each state-averaged template.

face is not only richer, but also qualitatively different from information extracted from individual facial features.

Most importantly, the internal states revealed by the MSLR model mapped distinctly onto behavioural trial outcomes (i.e. hit, wrong and miss trials) - even though this information had been in no way part of the inputs the model received. This suggests that the hidden states highlighted by our model were not simply 'computational devices' increasing its predictive power. Instead, they appear to reflect genuine, dynamically fluctuating cognitive states, which result in distinctive behavioural outcome profiles.

Our findings were not dependent on the use of one specific MSLR model, as we repeated our analyses with a GLM-HMM model (training an individual model per animal and experimental session), with very similar results (73). In two separate studies, such a GLM-HMM model was applied to behavioural choices in sensory discrimination tasks in mice. Both studies revealed distinct cognitive states that were identified as states of high or low engagement in the task (48, 52). The states were correlated with classic measures of arousal: pupil diameter, spontaneous facial movements and locomotion (52). Moreover, brain wide imaging signals recorded during a disengaged state showed higher trial-to-trial variability, which was due to task-independent (partly facial) movements(48). By taking into account comprehensive measurements of the entire face rather than mainly behavioural outcomes, our study could not only infer performance states on a single trial level (Fig 5D), but also show that disengaged states can reflect the dominance of less relevant sensory modalities (Fig 5A), and that similar behavioural states engage similar facial features in monkeys and mice (Fig 5C), suggesting strong evolutionary preservation of the underlying mechanisms.

Interestingly, despite the fact that the optimal number of states was determined separately for each species and in a purely data-driven way, our approach converged onto a low and noticeably similar number of internal states for both species: three states for mice (in agreement with (48, 52)), four for macaques. How comparable are these internal states of mice and monkeys?

We found that in terms of the dynamics by which animals traversed different internal states, results diverged across species. Specifically, mice appeared to transition more frequently between states than monkeys. A control analysis that matched the number of subjects, trials and facial parameters across species before fitting the MSLR models showed that this difference is not a trivial result of divergences in data structure (see Supp. Fig. S18). Given that mice have previously been shown to alternate between strategies during perceptual decision-making (56), this finding may point at a genuine difference in the cognitive dynamics of mice and monkeys. Alternatively, the prolonged training time of the monkeys compared to the mice may have given them the chance to converge on more stable behavioural strategies over the course of training (74, 75). Recordings from mice that experienced a more prolonged training scheme and/or from more naive monkeys will give fascinating insights into the role of expertise in fostering more stable transitions between cognitive states.

Beyond state dynamics, the constellation of behavioural profiles covered by different states was largely comparable across species (76, 77). Each hidden state predicted only a narrow range of reaction times; and when relating the inferred hidden states to task performance beyond the RTs that the model was trained to predict, we found that states mapped onto the behavioural outcomes (hit, wrong, miss) with distinct probabilities. Moreover, each hidden state covered unique combinations of RT ranges and trial outcomes (hit, miss and wrong trials), despite the fact that trial outcomes had not been part of the MSLR in any way. Specifically, both monkeys and mice display a state where trial outcome is typically slow and unsuccessful (which could be interpreted as 'inattentive'), as well as several states where performance is largely fast and correct, with a preference for thoroughness in one state, and a preference for speed (and potentially impulsivity) in the other. These states potentially map onto various levels of task-related attention, and further support the notion that classical concepts of attention can indeed reflect much of the internal structure of goal-directed behaviour, also in naturalistic settings.

Different internal states were associated with distinct constellations of facial features, as evidenced by the facts that states were highly separable based on the associated facial features, and that they could be successfully inferred through facial template matching on a single-trial level. This strong relation points to a role of facial expressions beyond emotional expression. Facial expressions have so far been mostly studied in a social or emotional context, and mostly in social species such as monkeys (29, 78) and humans (26, 79). In mice, until recently facial expressions were thought to mainly reflect pain (30, 80, 81), until careful analyses using machine-learning algorithms identified their facial expressions as innate and accurate reflections of several emotional states as well (19, 31, 82). Our results suggest that similarly to humans, facial expressions in monkeys and mice also convey cognitive and motivational variables such as focus or cognitive strain, even in the absence of a particular emotional or social context. What's more, the facial features predictive of comparable behavioural outcomes in mice and monkeys were highly overlapping (Fig 5C). This suggests that the mechanisms mapping behavioural states onto facial expressions are robust, and highly preserved across the evolutionary ladder.

The fact that such performance-related states are equally apparent in both species is intriguing, since one might have assumed that the prominent differences between the two species (such as the acuity and dominance of their visual system) would imply that they are likely to solve tasks using different strategies. And while both species have been shown to experience spontaneously shifting internal states, such as slow fluctuations in attention (21, 64, 65, 70), it has been difficult to make accurate comparisons between species based on available work. Studies of internal states have typically employed species-specific tasks that require extensive training (56, 64), which may have imposed cognitive dy-

namics that would not occur in the wild. As such, similarities and differences in the internal states identified in different species might have reflected similarities and differences in behavioural settings and training schemes just as much as innate differences in cognitive processing. Our study minimizes such confounds compared to previous research (56, 64), and thus adds one of the first rigorous estimates of which aspects of internal states truly generalize across species - and which do not. The fact that with this approach we were able to identify strong similarities in spontaneously occurring internal states across species opens the door to further inquiries into the common and evolutionary preserved principles of these states. Specifically, the spontaneous state fluctuations identified here can be important indicators of underlying brain-wide activity fluctuations that may well generalize across species to a large extent. Global internal states such as arousal, motivation, and attention typically manifest themselves via brain-wide dynamics, and it will be an exciting endeavour to investigate how well those neuronally defined internal states correspond to the ones we here identified behaviourally.

The MSLR model that we used yielded single estimates of the internal states per trial, and our subsequent analysis using 'facial templates' of internal states for single-trial predictions of behavioural outcomes was highly successful in both species (Fig 4). This constitutes a great basis for accurate, time-resolved tracking of ongoing dynamics in internal states, which can be further extended in future, using MSLR models with higher temporal resolution. Such MLSR models will be able to align identified internal states with specific events within each trial, such as the appearance and disappearance of stimuli, thereby allowing for more precise characterisation of their dynamics and functional roles.

Perhaps most importantly, such a time-resolved MSLR would also allow us to link cognitive processes to neural activity on a moment by moment basis, without the need for repeatedly presenting identical trials and then doing extensive post-hoc averaging. As the MSLR model yields a time-resolved estimate of cognitive states, these time courses can be directly compared to continuous neural activity. As such, this approach opens up a much more naturalistic view of the neuro-behavioural dynamics involved in spontaneous cognitive states than traditional approaches can offer (47, 57).

These findings suggest that in an ecologically valid framework that applies across species (in this case, a foraging-based task set in a naturalistic, immersive visual environment), many features of cognitive processing are more similar than classical paradigms might have suggested. At the same time, presumably genuine cross-species differences, e.g. in the transition frequency between cognitive states, also become more apparent.

In summary, we have shown here that in both monkeys and mice, facial features can be used to infer internal cognitive states, and to track their spontaneous dynamics over time. With this approach, we find that the basic attributes of such internal states map onto known cognitive states such as attention in both species in a translatable way, but that the dynamics by which mice and monkeys traverse these states is somewhat different. This highlights the crucial importance of using naturalistic behavioural paradigms, especially in cross-species research, in order to discern truly species-specific results from differences induced by restrictive testing methods.

ACKNOWLEDGEMENTS

We thank Carlos Wert Carvajal for useful input on earlier versions of the manuscript. We thank Victor Geadah, Iris Stone and Matthew Creamer for their valuable comments on model training, and two anonymous reviewers for helpful comments on the manuscript. A.T. acknowledges support from the Margarita Salas Fellowship (Spanish Ministry of Economy) and from the the Add-On Fellowship for Interdisciplinary Life Sciences (Joachim Herz Stiftung).

Bibliography

1. Steven W Flavell, Nadine Gogolla, Matthew Lovett-Barron, and Moriel Zelikowsky. The emergence and influence of internal states. *Neuron*, 2022.
2. David Deutsch, Diego Pacheco, Lucas Encarnacion-Rivera, Talmo Pereira, Ramie Fathy, Jan Clemens, Cyrille Girardin, Adam Calhoun, Elise Ireland, Austin Burke, et al. The neural basis for a persistent internal state in drosophila females. *Elife*, 9:e59502, 2020.
3. John W Krakauer, Asif A Ghazanfar, Alex Gomez-Marin, Malcolm A MacIver, and David Poeppel. Neuroscience needs behavior: correcting a reductionist bias. *Neuron*, 93(3): 480–490, 2017.
4. Matteo Grasso, Larissa Albantakis, Jonathan P Lang, and Giulio Tononi. Causal reductionism and causal structures. *Nature neuroscience*, 24(10):1348–1355, 2021.
5. Gaël Varoquaux and Russell A Poldrack. Predictive models avoid excessive reductionism in cognitive neuroimaging. *Current opinion in neurobiology*, 55:1–6, 2019.
6. Conrado A Bosman, Jan-Mathijs Schoffelen, Nicolas Brunet, Robert Oostenveld, Andre M Bastos, Thilo Womelsdorf, Birthe Rubehn, Thomas Stieglitz, Peter De Weerd, and Pascal Fries. Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron*, 75(5):875–888, 2012.
7. Georgia G Gregoriou, Stephen J Gotts, Huihui Zhou, and Robert Desimone. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *science*, 324(5931):1207–1210, 2009.
8. Trevor Robbins. The 5-choice serial reaction time task: behavioural pharmacology and functional neurochemistry. *Psychopharmacology*, 163:362–380, 2002.
9. Samuel K Asinof and Tracie A Paine. The 5-choice serial reaction time task: a task of attention and impulse control for rodents. *JoVE (Journal of Visualized Experiments)*, 90: e51574, 2014.
10. Ann Kennedy. The what, how, and why of naturalistic behavior. *Current opinion in neurobiology*, 74:102549, 2022.
11. Alex Gomez-Marin, Joseph J Paton, Adam R Kampff, Rui M Costa, and Zachary F Mainen. Big behavioral data: psychology, ethology and the foundations of neuroscience. *Nature neuroscience*, 17(11):1455–1462, 2014.
12. Alex Gomez-Marin and Asif A Ghazanfar. The life of behavior. *Neuron*, 104(1):25–36, 2019.
13. Yael Niv. The primacy of behavioral research for understanding the brain. *Behavioral Neuroscience*, 135(5):601, 2021.
14. Christopher D Harvey, Philip Coen, and David W Tank. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62–68, 2012.
15. Roberto A Gulli, Lyndon R Duong, Benjamin W Corrigan, Guillaume Doucet, Sylvain Williams, Stefano Fusi, and Julio C Martinez-Trujillo. Context-dependent representations of objects and space in the primate hippocampus during virtual navigation. *Nature neuroscience*, 23(1):103–112, 2020.
16. Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 2018.
17. Alexander B Wiltschko, Tatsuya Tsukahara, Ayman Zeine, Rockwell Anyoha, Winthrop F Gillis, Jeffrey E Markowitz, Ralph E Peterson, Jesse Katon, Matthew J Johnson, and Sandeep Robert Datta. Revealing the structure of pharmacobehavioral space through motion sequencing. *Nature neuroscience*, 23(11):1433–1443, 2020.
18. Kevin Luxem, Petra Mocellin, Falko Fuhrmann, Johannes Kürsch, Stephanie R Miller, Jorge J Palop, Stefan Remy, and Pavol Bauer. Identifying behavioral structure from deep variational embeddings of animal motion. *Communications Biology*, 5(1):1267, 2022.
19. Nate Dolensek, Daniel A Gehrlach, Alexandra S Klein, and Nadine Gogolla. Facial expressions of emotion states and their neuronal correlates in mice. *Science*, 368(6486):89–94, 2020.
20. Sandeep Robert Datta, David J Anderson, Kristin Branson, Pietro Perona, and Andrew Leifer. Computational neuroethology: a call to action. *Neuron*, 104(1):11–24, 2019.
21. Martha N Havenith, Petra M Zijderveld, Stefan van Heukelum, Rogier B Poorthuis, Huibert D Mansvelder, and August B Smit. The virtual-environment-foraging task enables rapid training and single-trial metrics of attention in head-fixed mice. *Scientific Reports*, 8(1):17371, 2018. doi: 10.1038/s41598-018-34966-8.
22. LT Hunt, ND Daw, P Kaanders, MA MacIver, U Mugan, Emmanuel Procyk, AD Redish, E Russo, Jacqueline Scholl, K Stachenfeld, et al. Formalizing planning and information search in naturalistic decision-making. *Nature neuroscience*, 24(8):1051–1064, 2021.
23. Panos Alefantis, Kaushik Lakshminarasimhan, Eric Avila, Jean-Paul Noel, Xaq Pitkow, and Dora E Angelaki. Sensory evidence accumulation using optic flow in a naturalistic navigation task. *Journal of Neuroscience*, 42(27):5451–5462, 2022.
24. Ashley L Juavinett, Jeffrey C Erlich, and Anne K Churchland. Decision-making behaviors:

weighing ethology, complexity, and sensorimotor compatibility. *Current opinion in neurobiology*, 49:42–50, 2018.

25. Richard E Quandt. A new approach to estimating switching regressions. *Journal of the American statistical association*, 67(338):306–310, 1972.
26. Paul Ekman and Harriet Oster. Facial expressions of emotion. *Annual review of psychology*, 30(1):527–554, 1979.
27. Ralph Adolphs, Hanna Damasio, Daniel Tranel, and Antonio R Damasio. Cortical systems for the recognition of emotion in facial expressions. *Journal of neuroscience*, 16(23):7678–7687, 1996.
28. Robert Aubrey Hinde and Thelma E Rowell. Communication by postures and facial expressions in the rhesus monkey (macaca mulatta). In *Proceedings of the Zoological Society of London*, volume 138, pages 1–21. Wiley-Blackwell, 1962.
29. Lisa A Parr and Matthew Heintz. Facial expression recognition in rhesus monkeys, macaca mulatta. *Animal behaviour*, 77(6):1507–1513, 2009.
30. Dale J Langford, Andrea L Bailey, Mona Lisa Chanda, Sarah E Clarke, Tanya E Drummond, Stephanie Echols, Sarah Glick, Joelle Ingrao, Tammy Klassen-Ross, Michael L LaCroix-Fralish, et al. Coding of facial expressions of pain in the laboratory mouse. *Nature methods*, 7(6):447–449, 2010.
31. Olivia Le Moëne and Max Larsson. A new tool for quantifying mouse facial expressions. *Eneuro*, 10(2), 2023.
32. Anna D Zych and Nadine Gogolla. Expressions of emotions across species. *Current opinion in neurobiology*, 68:57–66, 2021.
33. Dan Alin Ganea, Alexander Bexter, Mathias Günther, Pierre-Marie Gardères, Björn M Kampa, and Florent Haiss. Pupillary dilations of mice performing a vibrotactile discrimination task reflect task engagement and response confidence. *Frontiers in Behavioral Neuroscience*, 14:159, 2020.
34. Jacob Reimer, Emmanouil Froudarakis, Cathryn R Cadwell, Dimitri Yatsenko, George H Denfield, and Andreas S Tolias. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *neuron*, 84(2):355–362, 2014.
35. Matthew J McGinley, Martin Vinck, Jacob Reimer, Renata Batista-Brito, Edward Zagha, Cathryn R Cadwell, Andreas S Tolias, Jessica A Cardin, and David A McCormick. Waking state: rapid variations modulate neural and behavioral responses. *Neuron*, 87(6):1143–1161, 2015.
36. Paola Binda, Maria Pereverzeva, and Scott O Murray. Pupil size reflects the focus of feature-based attention. *Journal of neurophysiology*, 112(12):3046–3052, 2014.
37. R Becket Ebitz, John M Pearson, and Michael L Platt. Pupil size and social vigilance in rhesus macaques. *Frontiers in neuroscience*, 8:100, 2014.
38. Chin-An Wang and Douglas P Munoz. A circuit for pupil orienting responses: implications for cognitive modulation of pupil size. *Current opinion in neurobiology*, 33:134–140, 2015.
39. Siddhartha Joshi and Joshua I Gold. Pupil size as a window on neural substrates of cognition. *Trends in cognitive sciences*, 24(6):466–480, 2020.
40. R Becket Ebitz and Michael L Platt. Neuronal activity in primate dorsal anterior cingulate cortex signals task conflict and predicts adjustments in pupil-linked arousal. *Neuron*, 85(3): 628–640, 2015.
41. Jan Willem De Gee, Tomas Knapen, and Tobias H Donner. Decision-related pupil dilation reflects upcoming choice and individual bias. *Proceedings of the National Academy of Sciences*, 111(5):E618–E625, 2014.
42. Anne E Urai, Anke Braun, and Tobias H Donner. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nature communications*, 8(1):14637, 2017.
43. Tirin Moore and Mazyar Fallah. Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences*, 98(3):1273–1276, 2001.
44. Burkhart Fischer and B Breitmeyer. Mechanisms of visual attention revealed by saccadic eye movements. *Neuropsychologia*, 25(1):73–83, 1987.
45. Miriam Spering. Eye movements as a window into decision-making. *Annual review of vision science*, 8:427–448, 2022.
46. Sina E Dominiak, Mostafa A Nashaat, Keisuke Sehara, Hatem Oraby, Matthew E Larkum, and Robert NS Sachdev. Whisking asymmetry signals motor preparation and the behavioral state of mice. *Journal of Neuroscience*, 39(49):9818–9830, 2019.
47. Simon Musall, Matthew T Kaufman, Ashley L Juavinett, Steven Gluf, and Anne K Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience*, 22(10):1677–1686, 2019.
48. Chaoqun Yin, Maxwell D Melin, Gabriel Rojas-Bowe, Xiaonan Richard Sun, Steven Gluf, João Couto, Alex Kostiuk, Simon Musall, and Anne K Churchland. Engaged decision-makers align spontaneous movements to stereotyped task demands. *bioRxiv*, pages 2023–06, 2023.
49. Atika Syeda, Lin Zhong, Renee Tung, Will Long, Marius Pachitariu, and Carsen Stringer. Facemap: a framework for modeling neural activity based on orofacial tracking. *Nature Neuroscience*, 27(1):187–195, 2024.
50. Bharath Chandra Talluri, Incheol Kang, Adam Lazere, Katrina R Quinn, Nicholas Kaliss, Jacob L Yates, Daniel A Butts, and Hendrikje Nienborg. Activity in primate visual cortex is minimally driven by spontaneous movements. *Nature Neuroscience*, 26(11):1953–1959, 2023.
51. David B Salkoff, Edward Zagha, Erin McCarthy, and David A McCormick. Movement and performance explain widespread cortical activity in a visual detection task. *Cerebral Cortex*, 30(1):421–437, 2020.
52. Daniel Hulsey, Kevin Zumwalt, Luca Mazzucato, David A McCormick, and Santiago Jaramillo. Decision-making dynamics are predicted by arousal and uninstructed movements. *bioRxiv*, pages 2023–03, 2023.
53. Katharine A Shapcott, Marvin Weigand, Iuliia Glukhova, Martha N Havenith, and Marieke L Schölvinck. Domevr: A setup for experimental control of an immersive dome virtual environment created with unreal engine 4. *[Manuscript submitted for publication]*, 2022. doi: 10.1101/2022.04.04.486889.
54. Wen-Kai You and Shreesh P Mysore. Dynamics of visual perceptual decision-making in freely behaving mice. *Eneuro*, 9(2), 2022.
55. Chang. P, Harper-Donnelly. G, Kara A, Li X, et al, and Linderman S Murphy K. Dynamax, a library for State Space Models. *GitHub*, 2022.
56. Zoe C Ashwood, Nicholas A Roy, Iris R Stone, International Brain Laboratory, Anne E Urai, Anne K Churchland, Alexandre Pouget, and Jonathan W Pillow. Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience*, 25(2):201–212, 2022.
57. Luca Mazzucato. Neural mechanisms underlying the temporal organization of naturalistic animal behavior. *Elife*, 11:e76577, 2022.
58. Adam J Calhoun, Jonathan W Pillow, and Mala Murthy. Unsupervised identification of the internal states that shape natural behavior. *Nature neuroscience*, 22(12):2040–2049, 2019.
59. Pedro J Gonçalves, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, Tim P Vogels, et al. Training deep neural density estimators to identify mechanistic models of neural dynamics. *Elife*, 9:e56261, 2020.
60. Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in neural information processing systems*, 21, 2008.
61. Mohammad Reza Keshtkaran, Andrew R Sedler, Raeed H Chowdhury, Raghav Tandon, Diya Basrai, Sarah L Nguyen, Hansem Sohn, Mehrdad Jazayeri, Lee E Miller, and Chethan Pandarinath. A large-scale neural network training framework for generalized estimation of single-trial population dynamics. *Nature Methods*, 19(12):1572–1577, 2022.
62. Sean M Perkins, John P Cunningham, Qi Wang, and Mark M Churchland. Simple decoding of behavior from a complicated neural manifold. *bioRxiv*, pages 2023–04, 2023.
63. James FA Poulet and Sylvain Crochet. The cortical states of wakefulness. *Frontiers in systems neuroscience*, 12:64, 2019.
64. Marlene R Cohen and John HR Maunsell. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*, 70(6):1192–1204, 2011.
65. Martha N Havenith, Peter M Zijderveld, Sabrina van Heukelum, Shaghayegh Abghari, Paul Tiesinga, and Jeffrey C Glennon. The virtual-environment-foraging task enables rapid training and single-trial metrics of rule acquisition and reversal in head-fixed mice. *Scientific reports*, 9(1):4790, 2019.
66. Ayumu Yamashita, David Rothlein, Aaron Kucyi, Eve M Valera, Laura Germine, Jeremy Wilmer, Joseph DeGutis, and Michael Esterman. Variable rather than extreme slow reaction times distinguish brain states during sustained attention. *Scientific reports*, 11(1):14883, 2021.
67. Michael Esterman, Sarah K Noonan, Monica Rosenberg, and Joseph DeGutis. In the zone or zoning out? tracking behavioral and neural fluctuations during sustained attention. *Cerebral cortex*, 23(11):2712–2723, 2013.
68. Matteo Carandini. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS biology*, 2(9):e264, 2004.
69. Anne E Urai, Jan Willem De Gee, Konstantinos Tsetsos, and Tobias H Donner. Choice history biases subsequent evidence accumulation. *Elife*, 8:e46331, 2019.
70. Kenneth D Harris and Alexander Thiele. Cortical state and attention. *Nature reviews neuroscience*, 12(9):509–523, 2011.
71. Marieke L Schölvinck, Aman B Saleem, Andrea Benucci, Kenneth D Harris, and Matteo Carandini. Cortical state determines global variability and correlations in visual cortex. *Journal of Neuroscience*, 35(1):170–178, 2015.
72. Annika Stefanie Reinhold, Juan Ignacio Sanguinetti-Scheck, Konstantin Hartmann, and Michael Brecht. Behavioral and neural correlates of hide-and-seek in rats. *Science*, 365 (6458):1180–1183, 2019.
73. Alejandro Tlaie, Katharine Shapcott, Muad Abd el Hay, Berkutay Mert, Pierre-Antoine Ferracci, Robert Taylor, Iuliia Glukhova, Martha N Havenith, and Marieke Schölvinck. Thoughtful faces: Using facial features to infer naturalistic cognitive processing across species. Cosyne Abstracts, 2023. doi: 10.1145/1217935.1217951.
74. Patricia D Stokes. Learned variability. *Animal Learning & Behavior*, 23:164–176, 1995.
75. Patricia D Stokes, Francis Mechner, and Peter D Balsam. Effects of different acquisition procedures on response variability. *Animal Learning & Behavior*, 27:28–41, 1999.
76. Ben Mitchinson and Tony J Prescott. Whisker movements reveal spatial attention: a unified computational model of active sensing control in the rat. *PLoS computational biology*, 9(9): e1003236, 2013.
77. Ehud Fonio, Goren Gordon, Noy Barak, Yonatan Winetraub, Tess Baker Oram, Sebastian Haidarliu, Tali Kimchi, and Ehud Ahissar. Coordination of sniffing and whisking depends on the mode of interaction with the environment. *Israel Journal of Ecology & Evolution*, 61(2): 95–105, 2015.
78. Xin-He Liu, Lu Gan, Zhi-Ting Zhang, Pan-Ke Yu, and Ji Dai. Probing the processing of facial expressions in monkeys via time perception and eye tracking. *Zoological Research*, 44(5): 882, 2023.
79. Alan J Fridlund. *Human facial expression: An evolutionary view*. Academic press, 2014.
80. Alexander H Tuttle, Mark J Molinaro, Jasmine F Jethwa, Susana G Sotocinal, Juan C Prieto, Martin A Styner, Jeffrey S Mogil, and Mark J Zylka. A deep neural network to assess spontaneous pain from mouse facial expressions. *Molecular pain*, 14:1744806918763658, 2018.
81. Amy L Miller and Matthew C Leach. The mouse grimace scale: a clinically useful tool? *PloS one*, 10(9):e0136000, 2015.
82. Erwin B Defensor, Michael J Corley, Robert J Blanchard, and D Caroline Blanchard. Facial expressions of mice in aggressive and fearful contexts. *Physiology & behavior*, 107(5): 680–685, 2012.
83. Eleni Psarou, Julien Vezoli, Marieke L Schölvinck, Pierre-Antoine Ferracci, Yufeng Zhang, Iris Grothe, Rasmus Roese, and Pascal Fries. Modular, cement-free, customized headpost and connector-chamber implants for macaques. *Journal of Neuroscience Methods*, 393: 109899, 2023.
84. Tim Schroeder, Robert Taylor, Muad Abd El Hay, Abdellatif Nemri, Arthur Franca, Francesco Battaglia, Paul Tiesinga, Marieke L Schoelvinck, and Martha N Havenith. The dream implant: A lightweight, modular and cost-effective implant system for chronic electrophysiology in head-fixed and freely behaving mice. *bioRxiv*, pages 2024–02, 2024.
85. Keiji Matsuda, Takeshi Nagami, Yasuko Sugase, Aya Takemura, and Kenji Kawano. A widely applicable real-time mono/binocular eye tracking system using a high frame-rate digital camera. In *Human-Computer Interaction. User Interface Design, Development and Multimodal-*

ity: 19th International Conference, HCI International 2017, Vancouver, BC, Canada, July 9-14, 2017, Proceedings, Part I 19, pages 593–608. Springer, 2017.

86. Shaokai Ye, Anastasiia Filippova, Jessy Lauer, Maxime Vidal, Steffen Schneider, Tian Qiu, Alexander Mathis, and Mackenzie Weygandt Mathis. Superanimal models pretrained for plug-and-play analysis of animal behavior, 2023.
87. Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977.
88. Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial Intelligence and Statistics*, pages 914–922. PMLR, 2017.
89. Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
90. Tom AB Snijders. On cross-validation for predictor evaluation in time series. In *On Model Uncertainty and its Statistical Implications: Proceedings of a Workshop, Held in Groningen, The Netherlands, September 25–26, 1986*, pages 56–69. Springer, 1988.
91. Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016.
92. András Vargha and Harold D Delaney. A critique and improvement of the cl common language effect size statistics of mcgraw and wong. *Journal of Educational and Behavioral Statistics*, 25(2):101–132, 2000.
93. Emily B Fox, Erik B Sudderth, Michael I Jordan, and Alan S Willsky. The sticky hdp-hmm: Bayesian nonparametric hidden markov models with persistent states. *Arxiv preprint*, 2, 2007.

Methods

Animals. This study includes data from two male macaques (*Macaca mulatta*) and six male Black6 mice (*Mus musculus*). All procedures were approved by the regional authorities (*Regierungspräsidium Darmstadt*) under the authorization number F149/2000 and were performed in accordance with the German law for the protection of animals and the "European Union's Directive 2010/63/EU".

Surgical Procedures. All animals were fitted with custom-milled headposts for the purpose of head fixation during this experiment. The headpost design and implant procedures for the macaques have been extensively discussed in (83). Briefly, a four-legged titanium baseplate was screwed into the skull under general anesthesia. After several weeks of osseo-integration, a titanium top part was screwed onto the baseplate in a simple procedure. The headposts for the mice have been described in (84). Briefly, the animal was placed under isoflurane anaesthesia, shaved and given local analgesia on the top of the head. An incision was made and the skin on top of the cranium was removed, before the cranium was cleaned and the custom milled titanium head plate was attached using dental cement.

Experimental Setup. Experiments were carried out in a darkened room (mice) or electrically shielded booth (monkeys). The animals were in the centre of a 120-cm diameter spherical dome extending to 250 deg visual angle. The headfixed mice were positioned on a styrofoam spherical treadmill; the headfixed monkeys were seated in a monkey chair and could spin a 12-cm diameter trackball with their hands. Movements of the spherical treadmill and trackball allowed the animals to traverse a virtual reality (VR) environment projected on the inside of the dome by means of a spherical mirror. Projecting the VR environment on a dome surrounding the animals enabled both their central and peripheral view to be covered, thereby providing an immersive and realistic VR environment. The VR environment was created using DomeVR, our custom-made toolbox combining photorealistic graphics rendered with Unreal Engine 4, with high timing precision required for neuroscience experiments (53).

Experimental Paradigm. Mice and monkeys were required to distinguish two natural shapes at equal distance in front of them, amidst a grassy field with a blue sky above and mountains in the background (Fig 1A). The two shapes emerged out of a central shape which was either right at the starting position (for monkeys) or a short distance in front (for mice). A virtual collision with the correct shape yielded a reward ('hit'), whereas the incorrect shape yielded no reward ('wrong'), and no collision with either shape also yielded no reward ('miss') (2AFC paradigm). Rewards were drops of diluted juice for the monkeys and drops of vanilla soy milk for the mice. For the monkeys, the shapes varied smoothly between a non-rewarded, textured square and a rewarded triangle (monkey K) or between a rewarded, jagged and a non-rewarded, hour-glass shaped leaf (monkey C). On each trial, a blend between the two shapes was shown alongside the exact middle blend ('reference shape'). For the mice, the shapes and their reward contingencies were the same as for monkey C.
Monkey data were recorded in 7 sessions for monkey C, 11 sessions for monkey K. Each session lasted about one hour, during which the monkeys completed 1208 ± 186 and 991 ± 492 trials at 67 and 77 percent correct (monkeys C and K, respectively). The monkeys were both fully trained on handling the trackball to move through the VR environment, as well as the VR task. Mouse data were recorded in $(12, 4, 6, 3, 2, 2)$ sessions for mice (*001, 003, 004, 005, 012, 013*), respectively. Each session lasted about one hour, during which the mice completed $(280 \pm 103, 514 \pm 70, 573 \pm 112, 246 \pm 87, 462 \pm 8, 394 \pm 87)$ trials at $(59, 54, 60, 77, 45, 63)$ percent correct (same mice ordering as before). Following the headpost surgery, the mice were handled for 5 days to reduce experimental anxiety due to head fixation and interaction with the experimenter, before behavioural training began. Behavioural training in the experimental setup at initial stages lasted between 3-5 sessions, before final data collection began, which lasted up to 30 sessions.

Behavioural tracking. We recorded videos of the monkeys' and mice' faces during the tasks at $60\,\mathrm{Hz}$ using Basler acA640-121gm infrared cameras with a modified version of PylonRecorder2 software (https://gitlab.mpcdf.mpg.de/mpibr/scic/pylonrecorder/PylonRecorder2). Additionally, in the monkeys, eye movements were recorded at $500\,\mathrm{Hz}$ using a Grasshopper3 infrared camera and the free eye tracking software *iRecHS2* (85) and synchronized with DomeVR (53).

Facial key point extraction. To extract facial key points from the videos, we used markerless pose estimation on them, as implemented in DeepLabCut (16, 86). For mice, features were extracted from videos of the left side of the face using our own model to identify key points such as the coordinates of the eye, whisker pad and nose. For mouse pupillometry, we used the eye coordinates from the face model to crop the video to include the entire left eye and ran it through a refined model based on the "mouse pupil vclose" Animal Zoo model (provided by Jim McBurney-Lin at the University of California Riverside, USA) included with DeepLabCut. The output of the pupil model was 8 points covering the circumference of the mouse pupil, that were then used to calculate pupil and eye summary statistics.
For the macaque facial key points, we used the pre-trained "primate face" model from the DeepLabCut Animal Zoo (provided by Claire Witham at the Centre for Macaques, MRC Harwell, UK) and extended it with additional points on the lips to capture

more precise mouth movement than in the original model. All models were further trained and refined to achieve a detection error of less than 2 pixels per tracked key point in all conditions. The macaque raw pupil size recorded by the eye-tracker was Z-scored over time within the training data set.

To synchronise the video timing with events in the virtual reality environment, we used $32\,\mathrm{ms}$ long infrared flashes emitted from an LED mounted near the camera lens. These flashes were then extracted from the face videos to be used as timestamps for synchronisation with DomeVR. Five consecutive flashes indicated the start of a behavioural session; a single flash indicated the start of a trial.

Reaction Time. In our VR setting, where animals move towards one of two stimuli rather than pressing a button or lever, or making an eye movement, we define the reaction time (RT) as the time point of the initial substantial movement directed towards either stimulus. While determining this time point, it is crucial to distinguish between stimulus-related movements and minor positional adjustments. We specifically focus on the first deviation in lateral movement, while excluding forward movement due to its susceptibility to random movements and its task irrelevance.

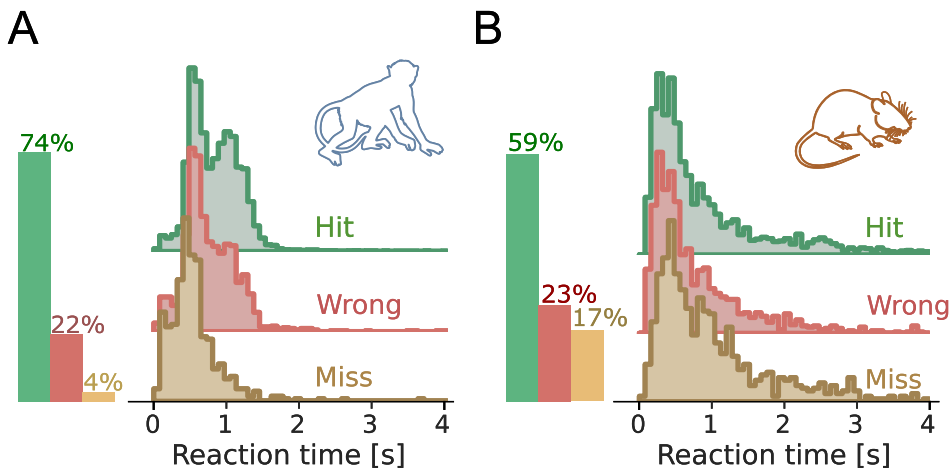

Fig. S1. Reaction Times. Distribution of Reaction Times for macaques (A) and mice (B), split by behavioral outcome; data are pooled over sessions ($n = 18$ and $n = 28$ for macaques and mice, respectively). The three distributions largely overlap.

To calculate the RT, we use a sliding window linear regression approach, incorporating a time decay mechanism. This approach enables us to detect non-linearity by examining the coefficient of determination (R^2) for each window. A low R^2 value indicates that the data deviate from linearity, and such a deviation can be interpreted as a deviation in lateral movement.

First we compute a linear regression on the time series of lateral VR movement for adjacent sliding windows i and j of a given size (n_w). Then, R_i^2 (i.e., R^2 for window i) is calculated as:

$$R_i^2 = 1 - \frac{\sum_{j=1}^{n_w} \left(l_j - \hat{l}_j^i\right)^2}{\sum_{j=1}^{n_w} \left(l_j - \bar{l}\right)^2} \tag{2}$$

where l_j is the jth element of the lateral movement observed in the second window, $\hat{l}_j^i$ is the corresponding predicted lateral movement value (based on window i) and $\bar{l}$ is the mean lateral movement within the second window. As a result, we get an array of R^2 values: $\mathbf{R}^2 = [R_1^2, \ldots, R_n^2]$.

Subsequently, we reverse the sign of the $-\mathbf{R}^2$ array and detect its local maxima. For this, we resort to the definition of extreme points (we have a univariate function in this case):

$$L = \underset{w}{\operatorname{argmax}} \left[\frac{d^2 r(w)}{dw^2} \right] \tag{3}$$

where we have simplified the notation, using $-\mathbf{R}^2 \equiv r(w)$. Once we have found the local maxima (L), we further require that they have a minimum prominence (λ). Prominence is a measure of the significance of a peak by comparing the peak to its surroundings:

$$\lambda_i = r(w_0) - \max\left[r(b_{l,i}), r(b_{r,i})\right], \quad \begin{aligned} b_{l,i} &= \underset{j\in[0,L_0]}{\operatorname{argmin}} \left[r(w_j)\right], \\ b_{r,i} &= \underset{j\in[L_0,n-1]}{\operatorname{argmin}} \left[r(w_j)\right] \end{aligned} \tag{4}$$

where $r(w_0)$ is $-R^2$ at L_0 and b_l and b_r are the arrays of left and right bases of the peaks; we are making use of the notation by which $r(w) \equiv -\mathbf{R}^2$.

For each peak in $r(w)$, we calculate the prominence and discard the ones that are below a given threshold (λ_0). The particular value for this threshold was not critical for the overall performance of the algorithm. For the sake of stability, we use multiple window sizes (100, 150, 200 and 250 ms) and combine the results in the following way. For each window k, we have an array of candidate points ($\mathbf{x}^k_{cand}$). Then, we create a vector of weights ($\mathbf{w}_k \in \mathbb{R}^n$) that have a value equal to a Gaussian distribution centered around each candidate point of each window. Mathematically:

$$\mathbf{w}_k(x) = \begin{cases} \mathcal{N}(x - x^k_{cand}, \sigma) & \text{if} \quad x \in \mathcal{B}^k_{cand} \\ 0 & \text{otherwise} \end{cases}$$

where $\mathcal{B}^k_{cand}$ denotes the vicinity of each point in x_{cand} for window k. Finally, the RT is given by:

$$RT = \underset{x}{\text{argmax}} \left[\left(\sum_k \mathbf{w}_k(x) \right) / x \right] \tag{5}$$

Fig. S1 shows the distribution of RTs split by trial outcome over sessions, for both species; Fig S2 shows example paths and detected RTs for both species.

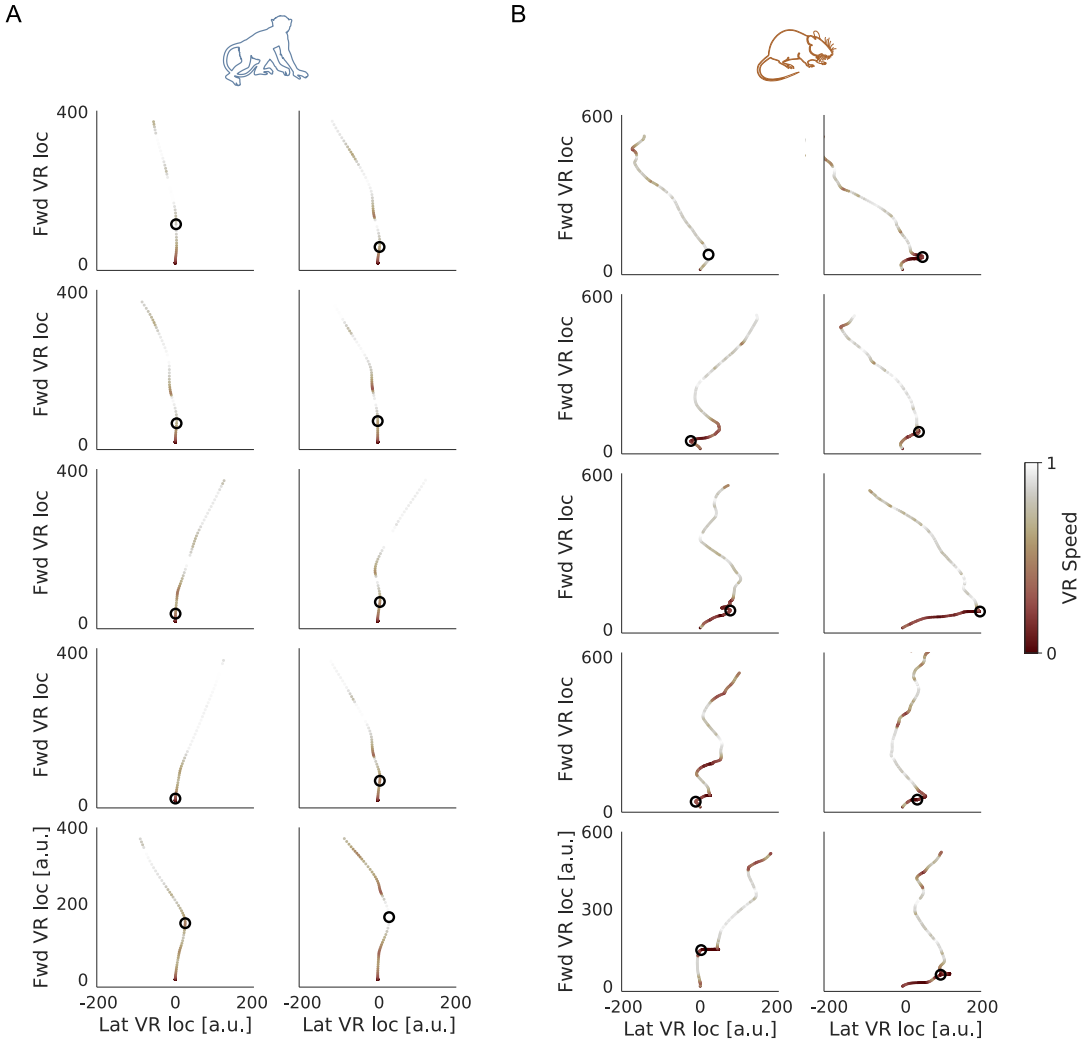

Fig. S2. Example VR paths Paths are colored according to the normalized running speed. **A)** Example paths for macaques, with the detected RT as circles. **B)** Same, but for mice.

Facial features. The extraction of the predictors for the MSLR model involves a multi-step process to go from continuous recording time (60 Hz for video data and 500 Hz for the macaque eye-tracker) to trial-based predictions.

First, we chose several points of interest on the animals' faces, which are then automatically identified and tracked over time using DeepLabCut (16):

- Macaque: both ears, eyebrows, nostrils and lips (see Fig. S3 A).
- Mouse: nose tip, left ear, left eye and median whiskers location (since we have a side view of the face, see Fig. S3 B).

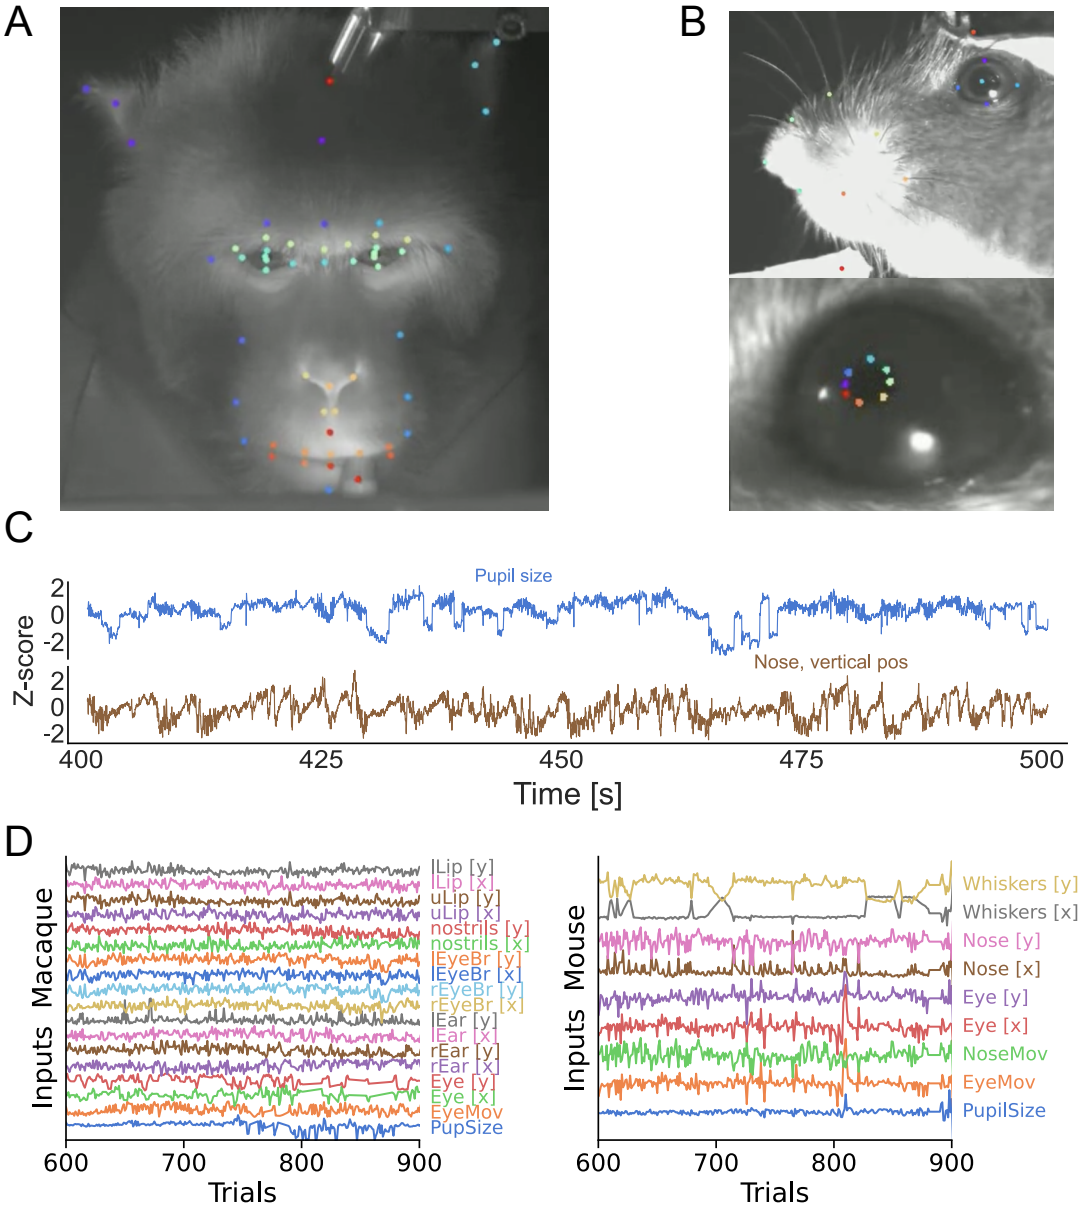

Fig. S3. Face features. A) Example frame of the macaque face camera. We have marked the key points that we used as the raw data for our pipeline. For this animal, we track a total of 73 key points. Some of them will be aggregated into centroids of interest, to minimize the influence of noise. **B)** Same as **A)**, but for the mouse. In this case, we also have a separate model for tracking pupil changes. **C)** Two example traces of a common feature for both species, over time. **D)** As described in the Methods, we use trial summaries for each of the face features of interest. Here, we show all of them, after having preprocessed them, for an arbitrary selection of 300 trials.

Once the data streams were aligned, we computed the median location (x, y) of each facial point over the 250 ms window before the stimuli appeared on the dome. This time window was chosen to make sure that all of the facial expressions of the animals are due to internally generated processing, rather than stimulus processing. Different window sizes (particularly: $200, 300$ and 500 ms) did not yield any qualitative difference. In addition to the median location, we also computed the total velocity of each facial point.

For both species, we further computed the median pupil size over the same time window. Pupil size is a well-known indicator of arousal and cognitive load, and thus provides valuable information about the internal state of the animal.

This resulted in a set of data points for each trial, corresponding to the median vertical and horizontal location, and total velocity of each of the facial features. These data points serve as the predictors for the MSLR model.

Markov-Switching Linear Regression. Markov-Switching Linear Regression (MSLR) models, which we ran using *Dynamax* (55), are a powerful tool for modeling time series data that exhibit regime-switching behaviour, where the underlying

dynamics of the system change over time.
The MSLR model is defined by a set of linear regressions, each associated with a particular state of a discrete Markov chain. The state of the Markov chain determines which sets of weights and biases predicts the evolution of the observed data at each time step. The transitions between states are governed by the transition probabilities of the Markov chain, which are learned from the data.
Formally, an MSLR model can be described as follows. If S is the total number of latent (discrete) states of a Markov process, at each time step t, a given state $z_t\big(\in\{0,1,...,S\}\big)$ will follow a Markovian evolution such that:

$$P(z_{t+1}=j|z_t=i)=\pi_{ij} \tag{6}$$

As these are stochastic matrices, $\pi_{ij}\in[0,1]$.
Let the $M-$dimensional input time series at time t be denoted by $x_t\big(\in\mathbb{R}^M\big)$. Let the $N-$dimensional output time series at time t be denoted by $y_t\big(\in\mathbb{R}^N\big)$. Then, in the case of a MSLR, the discrete latent variable at time t (z_t), will dictate which emission weights ($W\in\mathbb{R}^{N\times M}$) and emission biases ($b_s\in\mathbb{R}^N$) we will use to predict the outputs (emissions) based on the inputs (predictors). Moreover, an emission covariance matrix ($\Sigma_s\in\mathbb{R}^{N\times N}_{\succeq 0}$) will also have to be learnt. Explicitly, at time t, the emission distribution in this model is given by:

$$p(y_t\mid z_t,x_t,\theta)=\mathcal{N}(y_t\mid W_{z_t}x_t+b_{z_t},\Sigma_{z_t}) \tag{7}$$

Therefore, the problem of fitting this model amounts to finding the set of emission parameters denoted by:

$$\theta=\Big\{\big(W_s,b_s,\pi_s,\Sigma_s\big)\Big\}_{s=1}^{S} \tag{8}$$

In other words, the aim is to find the weights (W_s) and biases (b_s) for the linear regressions and the transition π_s and covariance Σ_s matrices for the Markov process.
In our case, the discrete latent variable (z_t) represents the internal state of the animal at trial t, which is inferred from the facial features (x_t) extracted using DeepLabCut (16) and the observation (y_t) that represents the RT of the animal. We trained the MSLR model using the Expectation-Maximization (EM) algorithm (87), which iteratively computes the probability over latent states given the data and updates the model parameters to maximize the likelihood of the observed data. For further details, we refer the reader to (88). We iterated the EM algorithm for 50 times, for all models. We initialized the model parameters using a normal distribution for weights and biases and we used the identity matrix as the initial covariance matrix for the emissions. We assumed a Dirichlet prior for the transition matrix. We repeated this process 10 times to increase confidence that we got the optimum value for each combination of parameters.

Training and inference. We used an $80:20$ ratio for train-test splitting and performed hyperparameter optimization by cross-validating the training set only (see *Model tuning* for details on CV and model selection). For each species, we concatenated the training sets of all sessions, with forced transitions in between the sessions (setting predictors and emissions to 0 for 50 consecutive trials), so that state probabilities are reset. Then, after optimizing each model, we performed inference on each held out test set (separately per session). We decided to take this approach for various reasons:

- Model generalization: as the model learns from potentially different faces, it is likely that it can pick up on common information between them.
- Model interpretability: given that we do not update the model parameters at the inference step, all internal states have the same meaning over subjects and, thus, are directly comparable.
- Better convergence: increasing the number of training samples (i.e. concatenating sessions as opposed to training a different model per session) allows the model to have more data to learn from.

All of the results in the main text, unless otherwise stated, are for held out data.

Model tuning. For the model we described in *Markov-Switching Linear Regression*, there are several parameters that can be tuned to explain the data better. We assumed a Dirichlet prior distribution, as it is the conjugate prior distribution of the categorical distribution. The concentration parameter (α) controls the relative density or sparseness of the resulting transition matrix. We decided to explore the influence of changing the maximum number of internal states (S) and to add sticky transitions to the Markov process (a self-bias term in the transition matrix π, making states taking longer to transition to a different one); the parameter that controls this is the stickiness (β), and to vary the transition matrix sparsity (concentration).
In order to balance model performance with scientific insights, we took a hybrid approach. We increased the number of internal states in a greedy way, to show that the error saturates and that there are diminishing returns when increasing model complexity. On the other hand, for a given number of states, we optimized the two free parameters of the Markov process (α and β).

For the sake of efficiency, we used Optuna (89), a flexible framework to implement Bayesian optimization. In Table 1 we report the relevant quantities for this process.
To select the best combination of parameters, we performed 5-fold Time-Blocked Cross-Validation (90).

Parameter	**Range**	**Fitted Values**
Concentration (α)	$[0, 100]$	$\alpha_{mac} = 19.3, \alpha_{mou} = 34.3$
Stickiness (β)	$[0, 100]$	$\beta_{mac} = 45.7, \beta_{mou} = 42.0$
Procedure	**Setting**	**Notes**
Sampler	CMA-ES (91)	–
Objective Function	R^2	Maximized
Number of Searches	100	–

Table 1. Parameter values for the Bayesian parameter optimization procedure. The fitted values are shown for macaque (mac) and mouse (mou) scenarios, independently explored for each number of internal states in the HMM.

We ended up selecting these values, after cross-validation: for macaques: $\alpha = 19.315$, $\beta = 45.705$; for mice: $\alpha = 34.312$, $\beta = 42.029$. These values explain the overall tendency we find in the main paper: macaques have a more block-like behavior (sparser and more diagonally-dominant transition matrix), whereas mice transition more and in between more states (denser and less diagonally-dominant transition matrix).

Synthetic data and ground truth states. In order to validate the retrieval of states when we do not have access to ground truth ones, we generated a time series of ground truth emissions and states based on the given inputs (by using the same input data as in the main text). To this end, we trained an MSLR model with a known given number of states and sampled some emissions and states sequence from it. We aimed to recover the appropriate number of states with the correct temporal sequence, and to correctly predict the emissions.
In Figure S4A, we show that the model's R^2 saturates at the ground-truth number of states (dashed vertical line). Nevertheless, we took the best performing number of states ($n_s = 10$) and cluster (through a hierarchical clustering algorithm, details below) these states based on their predicted emissions, and compare their similarities against a surrogate distribution (constructed via 500 shufflings of one of the input predictions, thus destroying any temporal structure). For those within or below a significance range (chosen to be the $90^{th} - 99^{th}$ percentile of the surrogate distribution), we group them together. At the end of this process, we end up selecting 5 states (thus recovering the ground-truth n_s), see Figure S4B.
For the selected number of states, we fit a MSLR using the EM algorithm for 50 steps (same as in the main paper) and plot the state sequence (true vs inferred), see Figure S4C. To quantify the observation that the inferred states closely match the ground-truth ones, we computed their correlation (one-hot encoding the true and predicted state sequences), showing a very high value for each value across the diagonal (and only for that one), see Figure S4D. There is a very close match ($\rho(s_{true}, s_{pred}) > 0.7$) between the true and inferred states, given that the 99th percentile of the surrogate correlation distribution was 0.12. We mask all values below this surrogate threshold.
As a technical note: In order to quantify the pairwise distribution similarity, we have used the Jensen-Shannon Distance, so we have turned each state-specific RT emissions distribution into a histogram (using the Freedman-Diaconis rule to automatically select the optimal number of bins).

Face separability. We computed a "face separability" measure, which quantifies how dissimilar the different distributions of facial features are when split over states.
Mathematically, this measure is defined as:

$$\Omega = \max(A, 1 - A) \tag{9}$$

where A is Vargha-Delaney's A-statistic (also known as measure of stochastic superiority), an effect size derived from the Mann-Whitney U-test (92). We decided to use Ω because it is especially interpretable. As it is related to the U-statistic, it can be thought of as the probability of a randomly selected point from one distribution being higher than another randomly selected point from the other distribution. This measure is bounded between 0.5 and 1. If there is no overlap, $\Omega = 1$. In this extreme case, one distribution would have complete stochastic dominance over the other. If X and Y are completely overlapping, $\Omega = 0.5$. Intermediate cases reflect the degree of dominance one distribution has over the other. The more its value deviates from 0.5, the less overlapping the distributions are.
We have shaded regions according to the common interpretation (92):

- Small separability (gray): $0.56 < \Omega < 0.64$.
- Medium separability (orange): $0.64 < \Omega < 0.71$

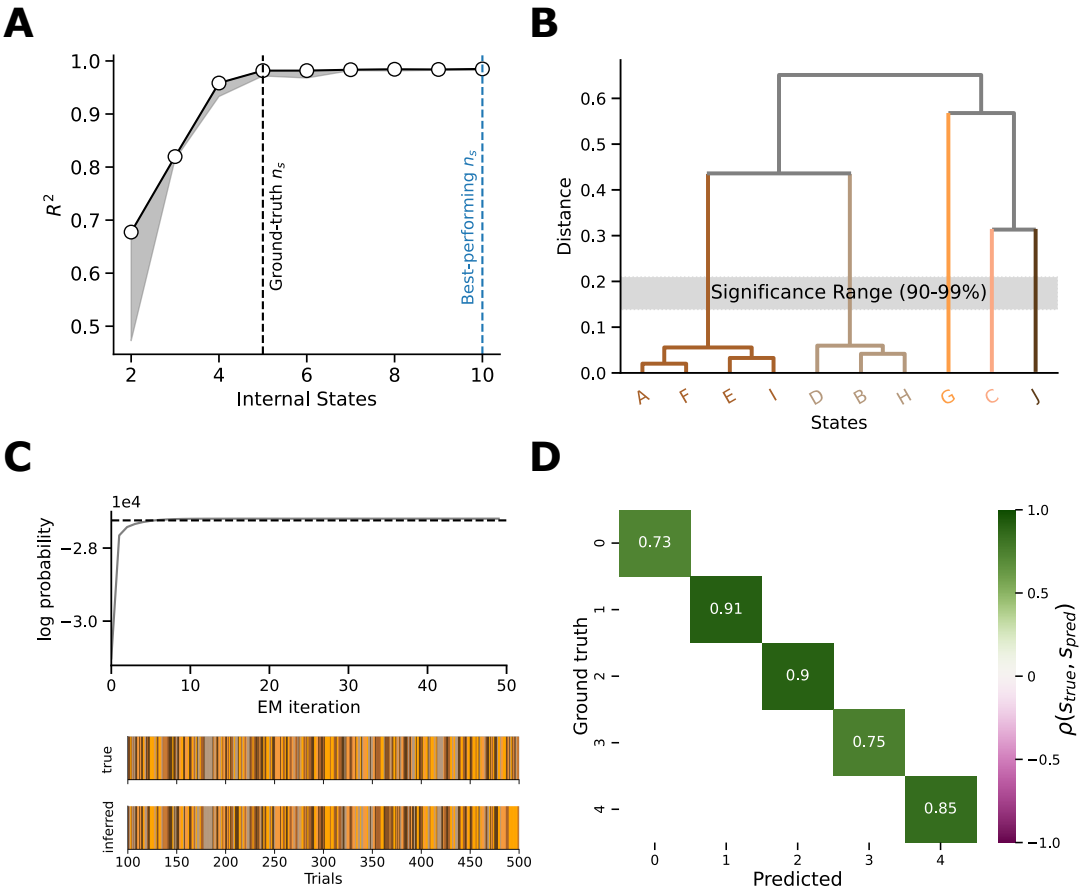

Fig. S4. Synthetic states and emissions. A) Performance when varying the number of states. We are able to recover the number of states (vertical dashed line) that generated the ground truth emissions. **B)** For the selected number of states, log-probability of the fitted parameters; it approaches to the ground truth value (horizontal dashed line). **C)** Some example trials for the true and predicted states. State transitions are correctly captured. **D)** Temporal correlation between the one-hot encoded state arrays. There is avery close match in between the predicted and the true states arrays.

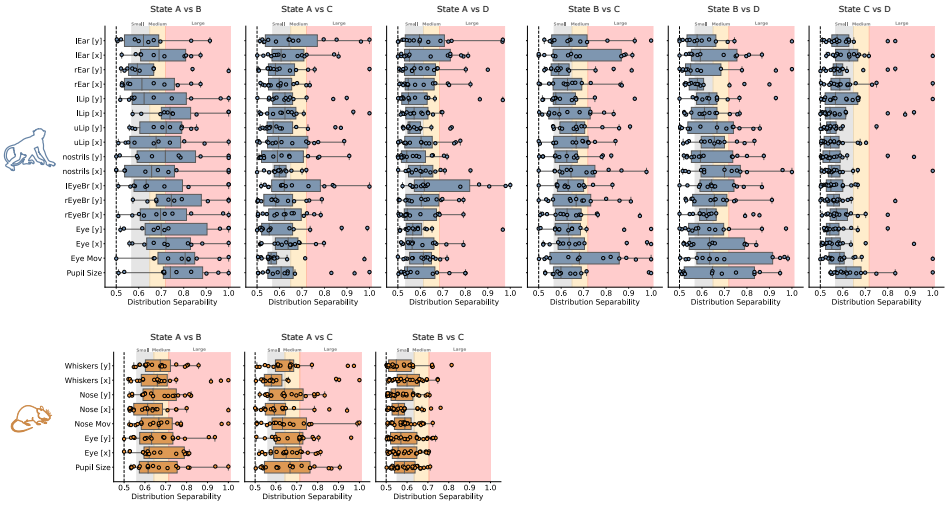

Fig. S5. Face features separability across states. We quantify how dissimilar the different distributions of facial features are when split over states. Mathematically, this measure is defined as the Vargha-Delaney's A-statistic (also known as measure of stochastic superiority). This measure is bounded between 0.5 (total distribution overlap) and 1 (no overlap). We have shaded regions according to common interpretations (92): Small separability (gray): $0.56 < \Omega < 0.64$; Medium separability (orange): $0.64 < \Omega < 0.71$; High separability (red): $\Omega \geq 0.71$. Each dot is an experimental session. No single feature appears useful on its own to disambiguate between states.

- High separability (red): $\Omega \geq 0.71$

For example, for the first column in the macaque row (Fig. S5, State A vs B), taking pupil size (bottom row) as an example, this metric quantifies how differently distributed the pupil size in state A vs state B is over sessions. Its median value is around 0.72, meaning high separability. This implies that this feature is very different across these two states, whereas the vertical left ear position (lEar [y] in the plot, top distribution) is much less useful to disambiguate between these two states. Comparing the separability (across features) between the various state combinations, two main results emerge:

- There are some states whose faces are similar to each other (Macaque: State C vs A; Mouse: State B vs C). This is in line with what we are showing in Figure 5A in the main text.
- No single feature is useful to disambiguate between states. This points back to one of our conclusions in the main paper: when accounting for internal states, the field is currently discarding potentially useful information by only relying on a single feature (most often, pupil size).

ARHMM. As we wanted to ensure that facial features were indeed informative of reaction times (RTs) beyond what is to be expected by the RT autocorrelation structure, we implemented an Auto-Regressive Hidden Markov Model (ARHMM). In this case, we used the same pipeline as we detailed in the previous sections, but substituted the facial features at the current trial t for the RT of the previous trial $(t-1)$. As it can be seen from Fig. S6, the facial features model outperforms the ARHMM for all states, for both species.

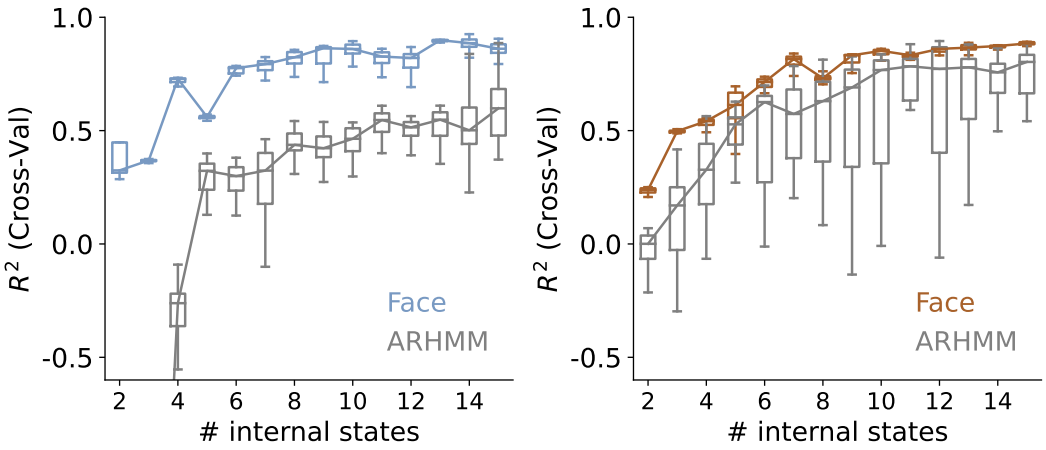

Fig. S6. Comparison of the MSLR face model and the reaction time Auto-Regressive HMM. Both face feature models outperform their Autoregressive counterparts, for any number of internal states that we swept over. Nevertheless, it can be seen that the performance gap is smaller in mice than in macaques. This is consistent with the finding that mice are more history dependent than macaques (See Fig. 5 D).

Task performance and internal states. We were interested in investigating whether the inferred internal states were correlated with task performance, even though the model had not been trained on such information. We therefore used the predicted single-trial state probabilities to decode choice, using a simple Logistic Regression model, with a $L2$ penalty term. After verifying that the model does indeed classify outcome beyond chance level (Fig. S7), we took the weight of each state as a proxy for how related it was to each outcome.

Data Availability. The data sets analysed during the current study are publicly available at https://github.com/atlaie/thoughtful-faces/.

Code Availability. All the relevant scripts are publicly available at https://github.com/atlaie/thoughtful-faces/.

Supplementary figures

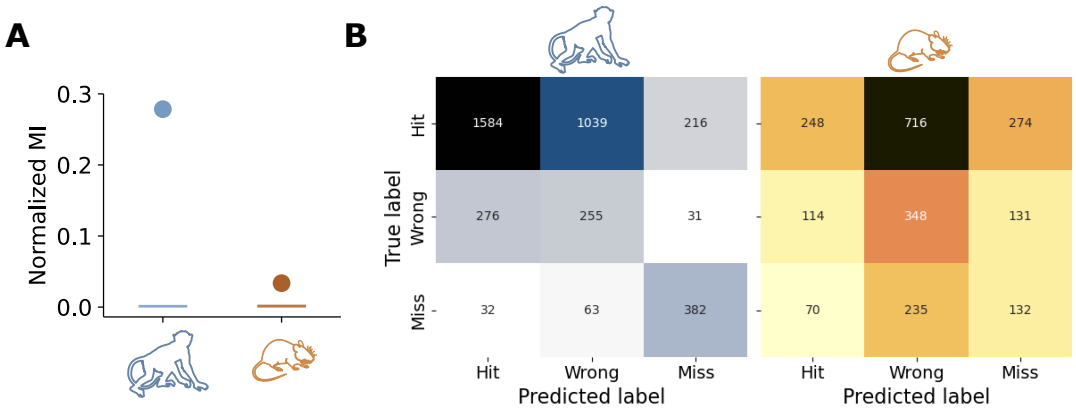

Fig. S7. Inferred state probabilities decode outcome beyond chance. We used a normalized version of Mutual Information (panel A) that already takes chance level into account and sets that as 0. The filled dots are the actual values and the shaded regions represent the surrogate distributions (built by randomly shuffling the outcomes over trials). As expected by construction of the metric, the random prediction is precisely at 0, whereas the measured value for both species is non-zero. When identifying which states were confused for another state by the model (panel B), we see that wrong and hit trials are misclassified more regularly than both of these trial types with miss trials. As before, for mice, wrong trials appear to be confused roughly equally for hit and miss trials, while in monkeys they are almost exclusively confused for hit trials, again suggesting that in monkeys hit and wrong states are more similar to each other than in mice.

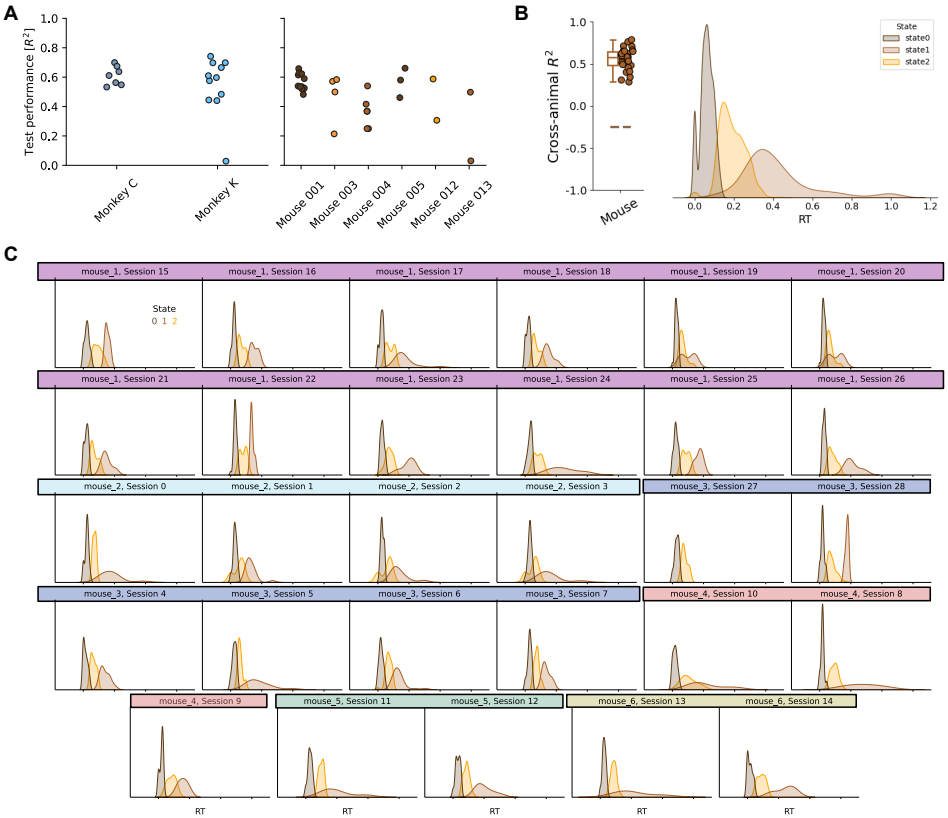

Fig. S8. Model generalization across animals. A) We show that the model performs uniformly well across animals, for both species (left, macaques; right, mice). There is no outlier animal for which the model does not perform well or just a group of them for which it does. **B)** If we follow a Leave-One-Out Cross-Validation (i.e., training a model with $n-1$ mice and testing on another one), we see that the model generalizes and there is a robust mapping of states to Reaction Times. **C)** We also show individual sessions, grouped by *validation mouse*.

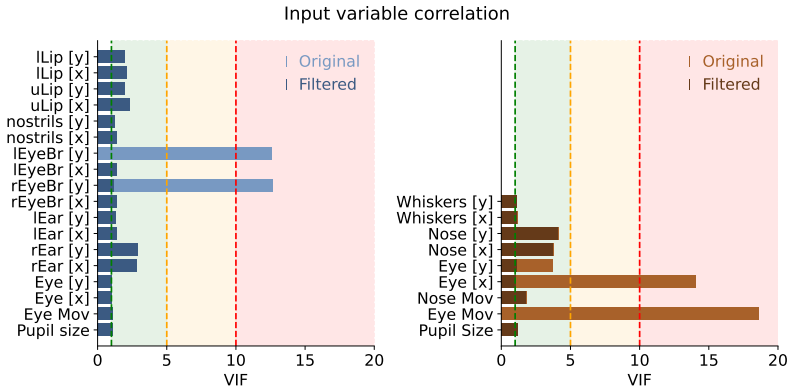

Fig. S9. Input variable correlation. We show that, out of all of the original variables (in lighter colors), we end up discarding one per animal (*Left Eyebrow [y], macaques; Eye movement, mice*), given that they were highly multi-colinear with some of the other predictors, as measured by the Variance Inflation Factor (VIF). After discarding them and recomputing the VIF, we did not find any alarming colinearity.

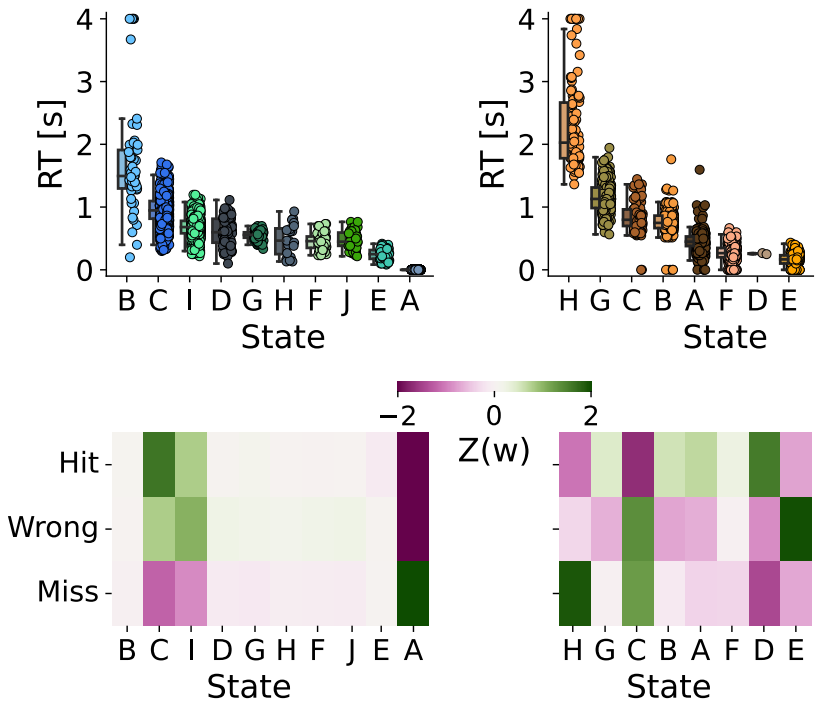

Fig. S10. Fragmented states. We show that, if we select the number of states (n_s) based on CV performance only (Fig. 2 in the main text), we would use $n_s = 10$ and $n_s = 8$ for macaque and mice, respectively. When looking at the emissions and the correlation with outcomes for each of these states, we find what others have already reported in the context of HMMs (93): if we allow states to be too fragmented, they will become redundant. In our case, we see that emissions and correlation with outcomes are highly redundant across states. Hence, both because of these reasons because of the state clustering results (Fig. S11 A), we decided to reduce the number of states for the rest of the paper.

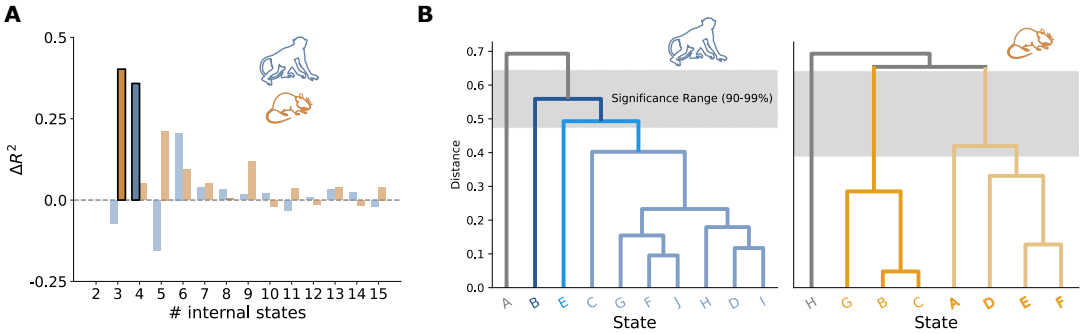

Fig. S11. State selection pipeline. We selected the number of states following the same pipeline as in Fig. S4: we first select the n_s corresponding to the highest CV R^2 ($n_s = 10, 8$ for macaques and mice, respectively); then, we inspect whether the corresponding emissions (Fig. S10) can be clustered together. This yields a potential $n_s \in \{2, 3, 4\}$ for macaques and $n_s \in \{2, 3, 4\}$ for mice (panel A). In order to disambiguate between these, we compute the finite difference in CV R^2 (panel B). We end up selecting $n_s = 4, 3$ for macaques and mice, respectively

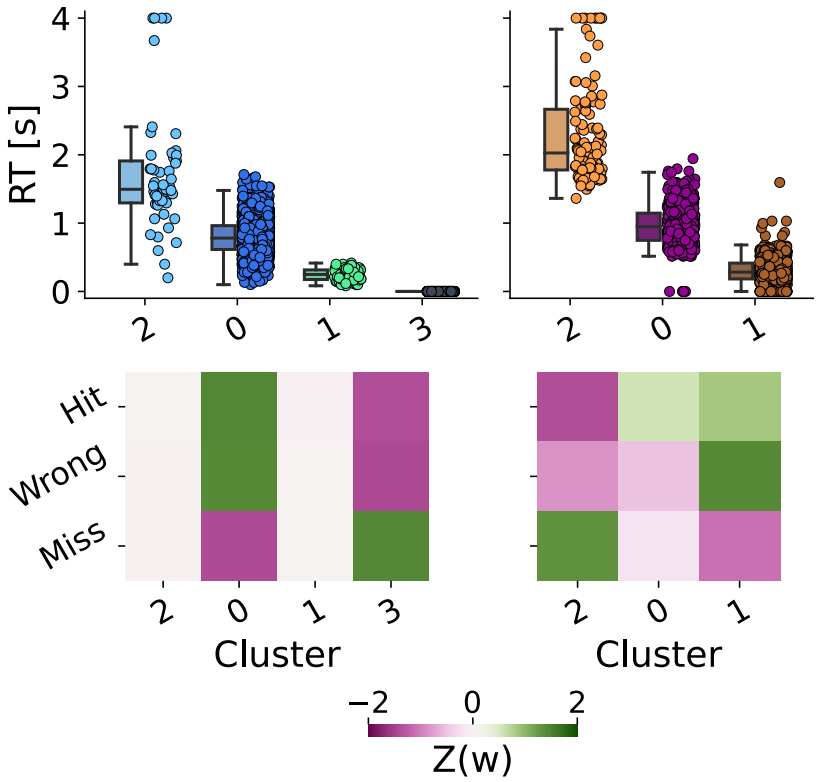

Fig. S12. Clustered states. If we cluster the emission distributions corresponding to the fragmented states we found in Fig. S10, we see that the most salient features demonstrated in the main figures are reproduced here: States cluster onto specific ranges of reaction times from slow to fast, and they spontaneously map onto trial outcomes, with hit and wrong trials being more interchangeable than miss trials with either hit or wrong trials. This is remarkable because these states and the ones from the main text do not need to be inferred from the same local minima in the loss landscape.

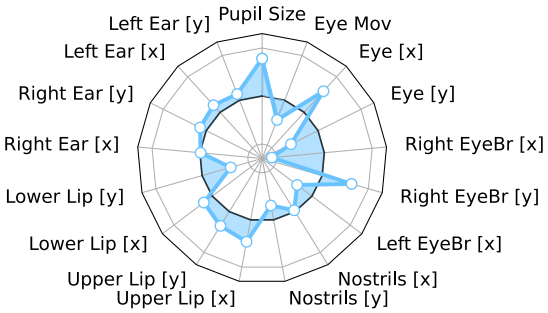

Fig. S13. Face features importance for State B in macaques. In comparison to the other three internal states in macaques, this state is associated with yet a different constellation of facial features.

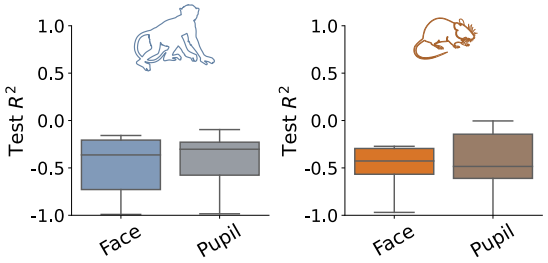

Fig. S14. No-switching model performances. If we assume there is a simple linear relationship (i.e. we do not allow for any switching) between the face (or the pupil) and the Reaction Time, the test performance is back at chance level.

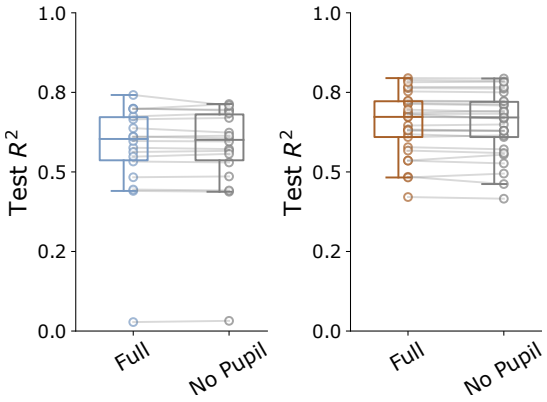

Fig. S15. Unaltered model performance without pupil size. If we remove the pupil size as a predictor for Reaction Time, model performance remains unchanged. This speaks against the usual perspective by which pupil size is paramount in inferring internal states.

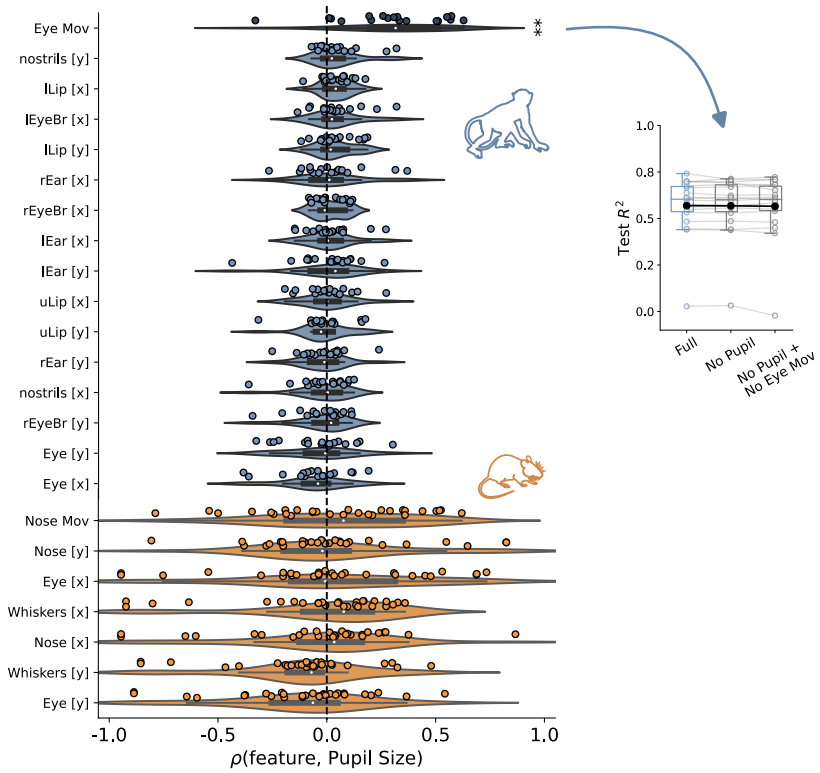

Fig. S16. Unaltered model performance without pupil size and correlated features. If, on top of removing the pupil size feature, we also remove those predictors that are significantly correlated with it, the model performance remains unchanged. This, together with Fig. S15 speaks against the usual perspective by which pupil size is paramount in inferring internal states.

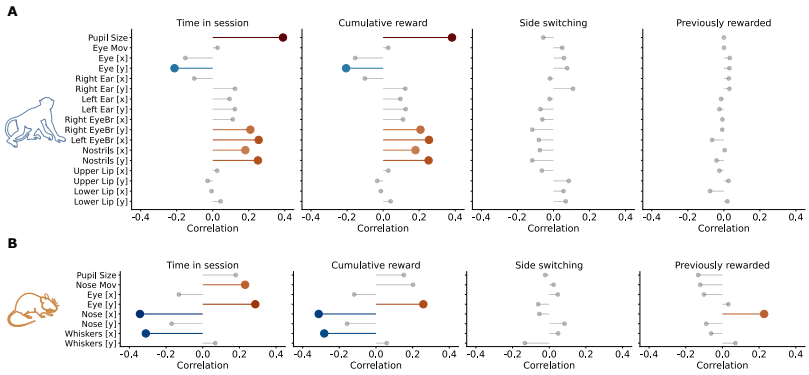

Fig. S17. Facial features and task history correlations, for macaques (top) and mice (bottom): time spent in session, cumulative reward in the session, the correct stimulus switching from being the left to the right one compared to the previous trial and vice versa, and whether the previous trial had been rewarded or not. Those highlighted in colour are significantly correlated.

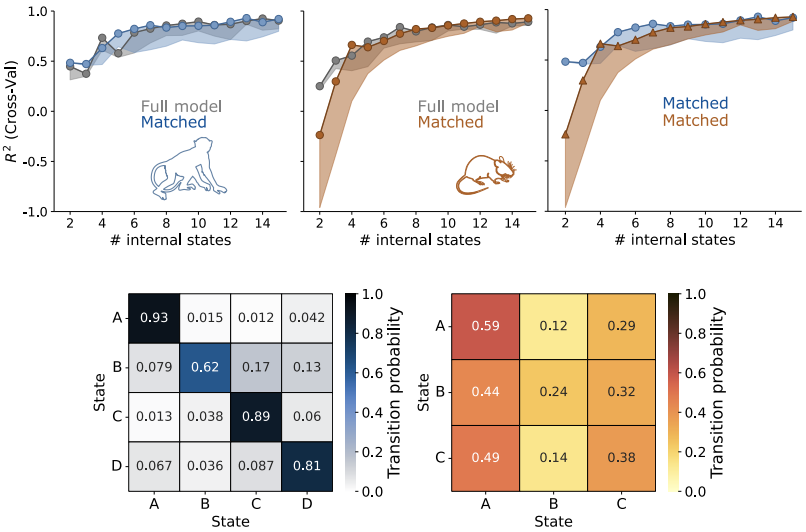

Fig. S18. Matched models. Our main findings still hold if we match the number of predictors, trials and animals we use in each species.

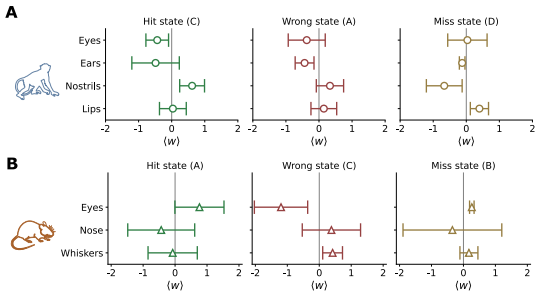

Fig. S19. Simplified face features weights. Collapsing facial features into a few general groups such as eyes, ears, and nose, again shows varying contributions of these feature groups for the different states. We show the signed $L1$ norm of the grouped feature weights.

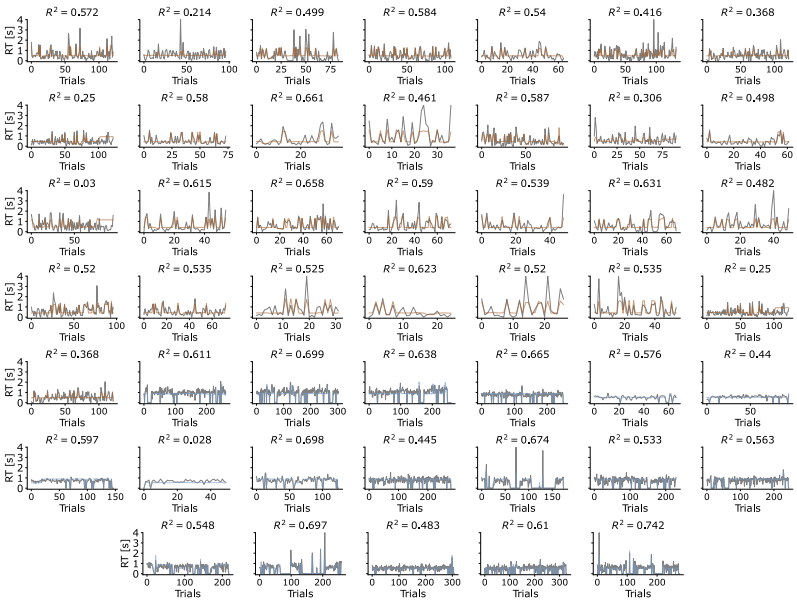

Fig. S20. Predicted and true RTs for all sessions, for the held out test set. For both species, true RTs are shown in gray; for mice, predicted RTs are shown in orange; for macaques, predicted RTs are shown in blue. Each subplot is titled with that session's model performance (R^2).

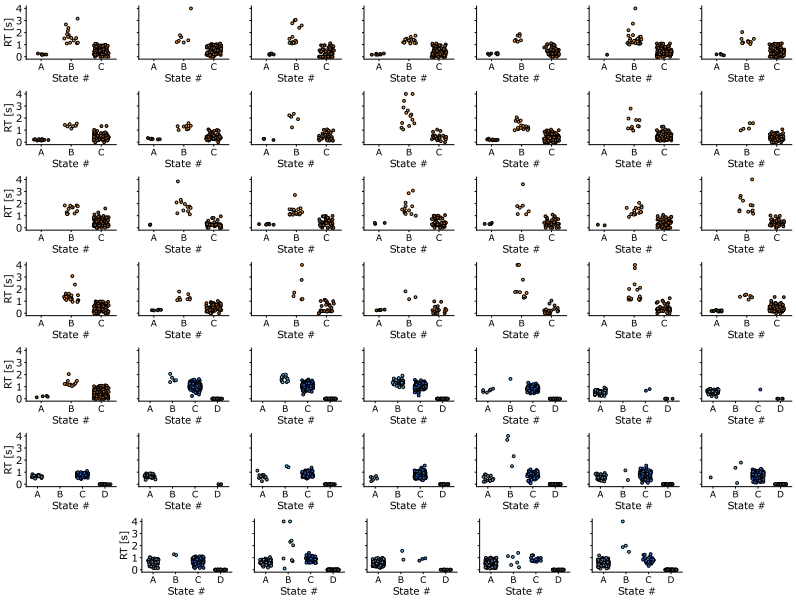

Fig. S21. RTs over states for all sessions, for the held out test set. For mice, distributions are shown in orange; for macaques, distributions are shown in green-blue.