

# Thoughtful faces: inferring internal states across species using facial features

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

<sup>1</sup>Ernst Strüngmann Institute for Neuroscience in cooperation with the Max Planck Society, Frankfurt am Main, 60528, Germany

<sup>2</sup>Laboratory for Clinical Neuroscience, Centre for Biomedical Technology, Universidad Politécnica de Madrid, Spain

<sup>3</sup>Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA

\*These authors contributed equally.

1 **Animal behaviour is shaped to a large degree by internal cog-** 45  
2 **nitive states, but it is unknown whether these states are similar** 46  
3 **across species. To address this question, we developed a virtual** 47  
4 **reality setup in which mice and macaques engage in the same** 48  
5 **naturalistic visual foraging task. We exploited the richness of** 49  
6 **a wide range of facial features extracted from video recordings** 50  
7 **during the task, to train a Markov-Switching Linear Regression** 51  
8 **(MSLR). By doing so, we identified, on a single-trial basis, a set** 52  
9 **of internal states that reliably predicted when the animals were** 53  
10 **going to react to the presented stimuli. Even though the model** 54  
11 **was trained purely on reaction times, it could also predict task** 55  
12 **outcome, supporting the behavioural relevance of the inferred** 56  
13 **states. The identified states were comparable between mice and** 57  
14 **monkeys. Furthermore, each state corresponded to a charac-** 58  
15 **teristic pattern of facial features, highlighting the importance of** 59  
16 **facial expressions as manifestations of internal cognitive states** 60  
17 **across species.** 61

18 **Internal states | Naturalistic behavior | Cross-species**

19 **Correspondence: [atboria@gmail.com](mailto:atboria@gmail.com)**

## 20 **Introduction**

21 In the wild, all mammals show similar behaviour: they all 65  
22 hunt or forage for food, sleep, mate, avoid predators, and ex- 66  
23 plore their environment, to name just a few. None of these 67  
24 behaviours can be simply explained as a passive reaction to 68  
25 environmental input; rather, they are crucially shaped by dy- 69  
26 namic fluctuations in internal states such as satiety, alertness, 70  
27 curiosity or attention (1, 2). So, if fundamental behaviours 71  
28 are comparable across species, how similar are the internal 72  
29 states that drive them? Is 'attention' in a monkey the same as 73  
30 'attention' in a mouse? 74

31 The common approach to investigate internal states has been 75  
32 a reductionist one: highly restrictive tasks featuring simpli- 76  
33 fied stimuli and requiring narrow behavioural repertoires (e.g. 77  
34 button presses), with little room for fluctuations over time (3– 78  
35 5). What's more, experimental paradigms diverge widely de- 79  
36 pending on the species under study. For example, attention 80  
37 studies in primates typically require the subject to fixate on 81  
38 a central fixation point while paying attention to a peripheral 82  
39 stimulus that might briefly or subtly change its appearance 83  
40 (6, 7). Attention studies in rodents, on the other hand, typ- 84  
41 ically use the 5-choice serial reaction time task (5CSRTT), 85  
42 in which the subject is required to scan a row of five aper- 86  
43 tures for the presentation of a brief light stimulus, and then 87  
44 navigate towards the light source (8, 9). Even though the 88

behaviour associated with high attention, i.e. short reac-  
tion times and accurate responses, is the same in both cases,  
clearly these tasks are too different to draw any meaningful  
cross-species comparisons.

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 re-  
flect spontaneously occurring rather than training-induced in-  
ternal 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 re-  
searcher, 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 be-  
havioural tracking is needed to generate time-resolved anal-  
yses that extract the evolution of underlying cognitive states  
from the measured behavioural parameters moment by mo-  
ment (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 cre-  
ate immersive yet highly controlled experimental settings that  
can be tailored to different species' intrinsic sensory capaci-  
ties and behavioural repertoires (14, 15). For instance, color  
spectra can easily be adapted to the visibility range of most  
species; and input devices like treadmills allow animals to  
interface with the VR in largely intuitive ways that require  
minimal training (16). As such, VR provides a powerful  
tool for studying animal behaviour in a highly controlled yet  
dynamic and ethologically sound context. Importantly, this  
maximizes adaptability across species, opening up the unique  
opportunity to record directly comparable behaviours in dif-  
ferent species.

At the same time, advances in deep-learning algorithms en-  
able us to dynamically track ongoing changes in body move-  
ment and posture from video footage (17–19). These algo-  
rithms have opened up new avenues to harvest an unprece-  
dented amount of information even from simple behavioural  
paradigms, far exceeding classical behavioural readouts such  
as button presses or saccades (20, 21). Most importantly,  
they allow for the ongoing and time-resolved tracking of be-  
havioural dynamics - a fundamental prerequisite if we aim to

89 identify the spontaneous emergence of internal cognitive and 146  
90 emotional states (22, 23). 147

91 In this study, we leverage these technological breakthroughs 148  
92 to infer and directly compare the internal states of two species 149  
93 commonly studied in neuroscience - macaques and mice. 150  
94 Specifically, we combine a highly immersive and natural- 151  
95 istic VR foraging task with a state-of-the-art deep learn- 152  
96 ing tool that allows for precise, automated tracking of be-  
97 havioural features. The features extracted in this way then 153  
98 serve as inputs to a Markov-Switching Linear Regression  
99 (MSLR) model (24), which finally captures time-varying in- 154  
100 ternal states across trials. 155

101 Importantly, such single-trial inference of internal states is 156  
102 only meaningful if the behavioural markers it relies on are 157  
103 not indirectly tracking the concrete motor outputs required 158  
104 for task performance. If the behavioural markers directly 159  
105 reflected task-related motor output (e.g. preparatory paw 160  
106 movements), then internal states inferred from this behaviour 161  
107 might be expected to trivially predict task performance. For 162  
108 instance, lack of preparatory paw movements might trivially 163  
109 predict a miss trial. To ensure that the behavioural parameters 164  
110 we chose would truly reflect internal processing, we focused 165  
111 on the animals' facial expressions. 166

112 While facial expressions have long been thought to only play 167  
113 a role in highly visual and social species like monkeys and 168  
114 humans (25–28), recent work has highlighted that also less 169  
115 social, less visual species like mice exhibit meaningful facial 170  
116 expressions (22, 29). As such, behaviourally relevant facial 171  
117 expressions seem to be much more evolutionarily preserved 172  
118 than previously expected (22, 29, 30). More specifically, they 173  
119 seem to reflect fundamental emotions like pleasure, pain, dis- 174  
120 gust and fear in a way that is not only consistent within one 175  
121 species, but also readily translatable across species (31, 32). 176  
122 This argues for an evolutionary convergent role of facial ex- 177  
123 pressions in reflecting (and potentially communicating) emo- 178  
124 tions. 179

125 Unlike these previous studies on the relation between facial 180  
126 expressions and emotions, here we for the first time analyse 181  
127 facial expressions in mice and monkeys that occur sponta- 182  
128 neously, in the absence of a pre-defined emotional context. 183  
129 Such spontaneously occurring behavioural states have so far 184  
130 mainly been tracked using single facial features to identify 185  
131 isolated cognitive states, for instance by quantifying atten- 186  
132 tion via pupil size, both in rodents (33, 34) and primates (35– 187  
133 38). Similarly, eye movements in monkeys and humans (39– 188  
134 41) and whisker movements in mice (42) have been used to 189  
135 track attention and decision-making. By focusing on entire 190  
136 facial expressions beyond individual (often species-specific) 191  
137 features, we aim to for the first time map out the spectrum 192  
138 of spontaneously occurring internal states in a way that is 1) 193  
139 agnostic, i.e. not focused on a specific cognitive process or 194  
140 facial feature, and 2) directly comparable across species. 195

141 Our approach of using facial expressions to infer internal 197  
142 states from natural behaviour constitutes a drastic move 198  
143 away from the classical approach of imposing internal states 199  
144 through restrictive behavioural paradigms (e.g. cued atten- 200  
145 tional shifts). By tying the results of this approach back 201

to known relationships between internal states and overt be-  
haviour, such as shorter reaction times during focused atten-  
tion, these data-driven, 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 inferred internal states  
across two species.

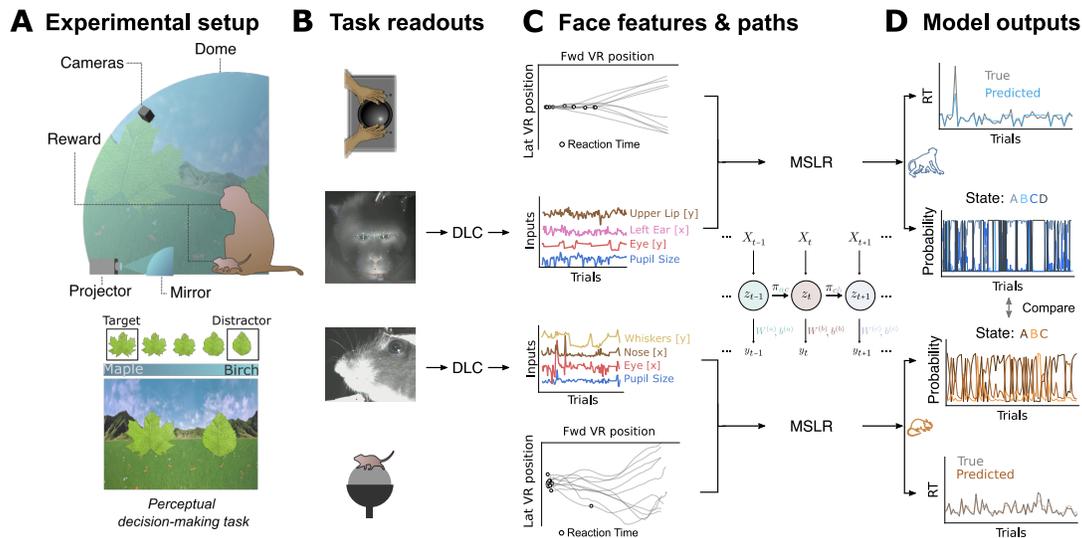
## Results

**A. Experimental set-up.** To track and compare sponta-  
neously occurring internal states of mice and macaques dur-  
ing the performance of the same naturalistic visual discrim-  
ination 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 (16). 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 sim-  
ple, 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 nat-  
ural leaf shapes integrated in a meadow landscape (Fig. 1A,  
bottom; see Methods - *Behavioral paradigm* and *Behavioral  
Training*). Their performance on this task was quantified  
first in terms of trial outcomes: hit (target stimulus reached),  
wrong (distractor stimulus reached), and miss (neither stim-  
ulus 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 deci-  
sively oriented itself towards one of the two potential targets  
(Fig. 1C; for details, see Methods - *Reaction Time*). As Fig.  
S1 shows, success rate and reaction times were largely com-  
parable 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 mon-  
key's face, as well as eye tracking output (see Methods - *Be-  
havioural 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 eye-  
brow, 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-  
dow of 250 ms before the stimuli appeared in the VR en-  
vironment. This time window was chosen to maximize the  
interpretability of the inferred hidden states: as there is no



**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.

202 task-relevant information available yet, presumably all of the 233  
 203 facial expressions that the animals make are due to internally 234  
 204 generated processes, rather than being reflective of stimulus 235  
 205 or task parameters. 236

206 **B. Model performance.** The facial features extracted in 237  
 207 this way were used as inputs to a Markov-Switching Linear 238  
 208 Regression (MSLR) model (Fig. 1C; see Methods - 239  
 209 MSLR). The MSLR manages to reflect the non-stationarity 240  
 210 and regime shifts often present in behavioural data (43- 241  
 211 46), by flexibly accommodating complex temporal dynamics 242  
 212 while keeping a relative simplicity, compared to deep 243  
 213 learning-based methods (47). Moreover, the MSLR is less 244  
 214 data-hungry than other common data-driven models (48-50). 245  
 215 The MSLR uses the 'pre-stimulus' facial features in each trial 246  
 216 to predict the animals' reaction time (RT) in the same trial by 247  
 217 assuming 'hidden' states. Each hidden state implies a different 248  
 218 linear relation between individual facial features and the 249  
 219 subsequent RT in the same trial. For instance, in one hidden 250  
 220 state, the RT might be best predicted by eyebrow movements, 251  
 221 while in another, nose sniffing might be most predictive. We 252  
 222 used cross-validation to select the number of states for each 253  
 223 species (see below). For each trial, the model then outputs the 254  
 224 predicted RT as well as the probability of each hidden state 255  
 225 (Fig. 1D). The two models (one for mice, one for monkeys) 256  
 226 were trained and tested on data from all individuals; Fig. S14 257  
 227 shows the outcomes of the models split by session and by 258  
 228 individual. 259

229 Mathematically, this model takes the form:

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

230 where  $RT_t$  is the reaction time at trial  $t$ ,  $z_t$  is the state at trial 264  
 231  $t$ ,  $W_{z_t}$  are the regression weights for state  $z_t$ ,  $x_t$  is the vector 265  
 232 of facial features at trial  $t$ , and  $\xi_{z_t}$  is a zero-mean Gaussian 266

noise with variance  $\sigma_{z_t}$ .

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). For both species, the model's predictive performance was remarkably low under these circumstances - in fact, predictions were less accurate than random guessing.

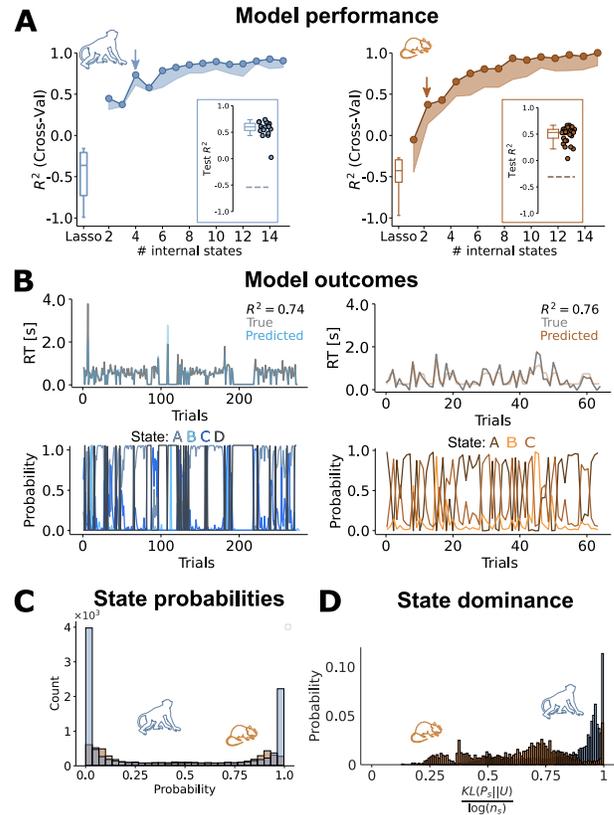
Next, we quantified model performance for different numbers of hidden states - which is the main free parameter of the MSLR. Model performance was tested by using cross-validation (see Methods - Model tuning). For both species, the cross-validated  $R^2$  improved dramatically when allowing for more 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 (Fig. S10), in order to reach the optimal trade-off between predictive accuracy and model simplicity (Fig. 2A). 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), indicating that the high predictive performance was not due to overfitting.

In both species, our models yielded remarkably accurate trial-by-trial predictions of RT, 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.

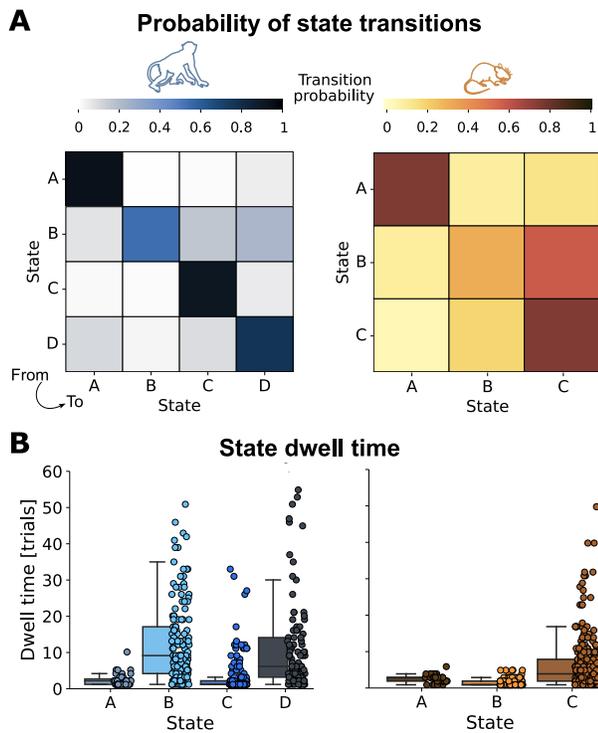
267 The coexistence of several hidden states opens up the ques-  
 268 tion whether task performance is dominated by a single state  
 269 at any given moment, or if several states co-exist continu-  
 270 ously. After fitting the model parameters, we used the model  
 271 to identify the animal’s internal state on a trial-by-trial basis.  
 272 Note that the model does not allow for the animal to be in  
 273 multiple states at the same time; rather, it gives us probabili-  
 274 ties telling how confident we can be about the state the animal  
 275 is in on each trial. Specifically, we computed the posterior  
 276 probability over states on each trial given all past and future  
 277 observations. The probabilities of each state over time sug-  
 278 gest that the model is highly confident about what state the  
 279 animal is in on each trial (Fig. 2B, bottom row). These obser-  
 280 vations were confirmed by the highly bimodal distribution  
 281 of these probabilities for both species (Fig. 2C). Crucially,  
 282 in monkeys, this separation between high-certainty ( $p_s \approx 1$ )  
 283 and low-certainty ( $p_s \approx 1/n_s$ ) trials was particularly pro-  
 284 nounced, while in mice, state probabilities were somewhat  
 285 more mixed. Quantifying the single-trial certainty as measur-  
 286 ing its difference with the uniform distribution –through the  
 287 Kullback-Leibler divergence (KL)– corroborated these find-  
 288 ings (Fig. 2D; Mann-Whitney U-test:  $p = 1.11 \cdot 10^{-274}$ ).  
 289 As such, the hidden states identified by our model seem to  
 290 reflect largely mutually exclusive behavioural modes that ani-  
 291 mals switch into and out of. Given how consistently trials  
 292 were dominated by one state, we chose to binarize hidden  
 293 state outcomes by assigning each trial to its most probable  
 294 hidden state.

295 **C. State dynamics.** To explore if the hidden states showed  
 296 attributes that could be reflective of internal cognitive states,  
 297 we first characterized their temporal dynamics. To this  
 298 end, we examined the frequency of state transitions in both  
 299 species. The state transition matrices, which show how  
 300 likely a trial of a given hidden state is followed by a trial of  
 301 any (other or same) state (Methods - *Markov-Switching Lin-*  
 302 *ear Regression*), revealed high values along the diagonal for  
 303 macaques, indicating stable states that switched rather rarely.  
 304 In mice, the diagonal of the transition matrix was slightly less  
 305 pronounced, suggesting that hidden states in mice were less  
 306 stable and more prone to transition than in macaques (Fig.  
 307 3A).

308 As a complementary analysis, we computed the dwell time  
 309 for each state. This quantity is defined as the number of con-  
 310 secutive trials that a given state is occupied for, before transi-  
 311 tioning to a different state. Supporting the previous observa-  
 312 tions, hidden states lasted generally longer in macaques than  
 313 in mice (Mann-Whitney U-test;  $n_{mac} = 4092$ ,  $n_{mice} = 2543$   
 314 trials,  $p = 0.0014$ ), suggesting that internal processing may  
 315 be more steady in macaques (Fig. 3B). This is consistent  
 316 with previous findings that behavioural dynamics may fluc-  
 317 tuate faster in mice (34, 51) than monkeys (52). Apart from  
 318 a genuinely species-driven difference, this observation may  
 319 also reflect the fact that monkeys are trained more exten-  
 320 sively and may therefore have developed more stereotyped  
 321 behavioural strategies than mice, which were trained more  
 322 briefly.



**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 5<sup>th</sup> 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 99<sup>th</sup> 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).



**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) spend more time than mice (right) in the same state, as measured by the dwell time (number of consecutive trials of each state being the most likely one). Individual dots 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; Fig. S15 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.

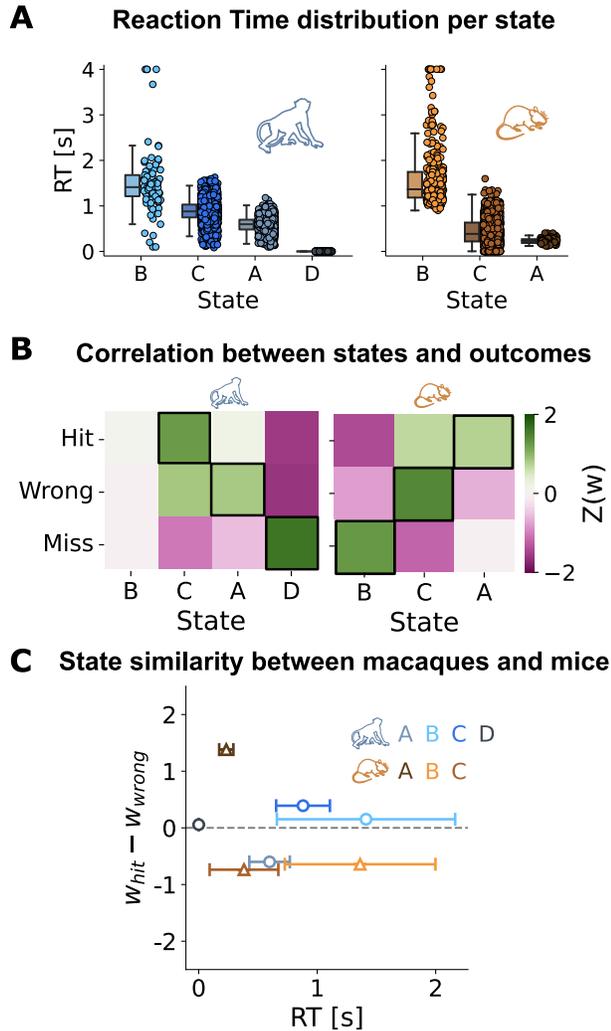
To further probe the possible link of our internal states to known cognitive processes, we related all hidden states to

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. 4B). 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.

Combining these effects of internal states on RTs and trial outcomes revealed specific combinations of speed and accuracy. We plotted mean RT per hidden state against the difference in probability of a hit versus a wrong trial, in the same state. Interestingly, the constellation of states in this space was comparable across species (Fig. 4C). Both mouse and monkey data seem to generate a 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, reflects rather fast yet often incorrect responses, potentially reflecting more impulsive decision-making (53, 54). Finally, state B for both species features particularly slow RTs, large RT variability, and mostly misses for mice and equally likely trial outcomes for monkeys, potentially signifying a state of global inattention (55, 56). 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.

**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. S8). These plots reveal highly distinct contributions of different facial features to each internal state. For example, in mice, eye movements and pupil size strongly 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 ear movements play a strong predictive role in hit and wrong, but not in miss states. One interpretation of these observations is that different senses are 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 vision-related



**Fig. 4. Internal states and task performance.** **A)** 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. **B)** 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. **C)** Conjunction of RT and excess likelihood of a hit outcome, for all states (blue circles, macaque; orange triangles, mouse).

408 facial features play a decisive role, suggesting that in other  
 409 states, behaviour may be less strongly driven by sensory sampling in the visual dimension (and more by sampling in the  
 410 olfactory dimension).  
 411

412 The overall variability of different facial features' predictive  
 413 weights across states confirms that while some facial features  
 414 had more constant predictive power across all states, others,  
 415 such as ear movement in monkeys and nose movement in  
 416 mice, are highly predictive of performance in some states,  
 417 but not in others (Fig 5B).

418 Moreover, in both species, reaction times are best predicted  
 419 by a complex constellation of facial features rather than one  
 420 feature in isolation. On the same note, different states are  
 421 consistently distinguished by more than one facial feature.  
 422 Together, these results suggest that 1) holistic analysis of  
 423 complex facial expressions is much more informative than  
 424 analysis of one isolated facial feature such as pupil size and  
 425 2) the relationship between facial features and cognitive pro-  
 426 cessing is not linear, but changes depending on the internal  
 427 state that the animal is in. For instance, in a high-performance  
 428 state, large pupil size may indeed predict trial success (as  
 429 shown e.g. by (33, 35)), whereas it may be irrelevant or anti-  
 430 correlated in a low-performance state.

431 Interestingly, the facial constellations predicting RTs in hit  
 432 and wrong states are quite similar in monkeys, but not in  
 433 mice. This resonates with the fact that 'hit' and 'wrong'  
 434 states also have more overlap in terms of trial outcomes in  
 435 monkeys than in mice (Fig 4B). This may suggest that in  
 436 monkeys, the behavioural state underlying hit and wrong tri-  
 437 als may be a more generalized engaged and high attention  
 438 state, and hit or wrong outcomes may be mainly dictated by  
 439 visual difficulty than different internal cognitive state. In con-  
 440 trast, it appears that in mice, hit and wrong trials may be  
 441 the product of more distinct underlying cognitive states (see Fig.  
 442 S13 for a summarized visualization).

443 One reason why hidden states can predict trial outcomes  
 444 so accurately despite not being trained on them in any way  
 445 might be that pre-trial facial features are mostly a trivial con-  
 446 sequence of the animal's trial history. For example, facial  
 447 features might mainly reflect an animal still drinking reward  
 448 from the previous trial, which might in turn raise motivation  
 449 to perform correctly in the upcoming trial. In this case, fa-  
 450 cial features would merely be a particularly convoluted way  
 451 of quantifying the previous trial outcome, and using it to pre-  
 452 dict upcoming performance, as has been achieved previously  
 453 (57, 58). To account for this possibility, we trained an Auto-  
 454 Regressive Hidden Markov Model (ARHMM) based on RTs  
 455 (see Methods - ARHMM for details). As can be seen in Fig.  
 456 S5, the facial features model outperforms the ARHMM for  
 457 all states, for both species.

458 As an extra control, we correlated each facial feature with the  
 459 history of prominent task parameters, specifically two related  
 460 to the directly previous trial (its outcome, which might affect  
 461 motivation; and the location of its target, which might predict  
 462 side biases), and two related to the overall session history  
 463 (the cumulative amount of reward and the time that passed  
 464 since the start of the session, as proxies for satiety and fa-

465 tigue, respectively). Correlations between task variables and 521  
466 facial features were sparse in both species (Fig. S11). In 522  
467 fact, attributes of the previous trial did not relate significantly 523  
468 to facial features at all, and more sustained session attributes 524  
469 modulated facial features merely somewhat. This suggests 525  
470 that facial features may be modulated by ubiquitous internal 526  
471 processes like fatigue and satiety, which are in turn impacted 527  
472 by task parameters, but they are not a trivial reflection of task 528  
473 history. Rather, the fact that facial expressions are modulated 529  
474 by the overall task context makes them a more plausible re- 530  
475 flection of realistic fluctuations in cognitive processing. 531

## 476 Discussion 532

477 Internal cognitive states are known to substantially shape 534  
478 overall brain activity (59, 60) as well as behavioural decision 535  
479 making (1, 2), yet they are notoriously difficult to identify. 536  
480 As a result, it is even less clear to what extent they converge 537  
481 across species. To infer hidden cognitive states in mice and 538  
482 monkeys, we harnessed an MSLR model (24) trained on their 539  
483 facial features while they were engaged in an immersive VR 540  
484 foraging task. Specifically, we trained the MSLR to predict 541  
485 an animal's reaction time (RT) in a given trial based on its 542  
486 facial expressions prior to stimulus presentation. For both 543  
487 species, RTs could be predicted with high accuracy from pre- 544  
488 ceding facial features only, suggesting that facial expressions 545  
489 reflect parameters that are directly relevant to task perfor- 546  
490 mance. These parameters were only minimally shaped by 547  
491 task history, suggesting that they were not a trivial reflection 548  
492 e.g. of the previous trial outcome. 549  
493 Even more surprisingly, this approach revealed multiple dis- 550  
494 tinct hidden states, which were characterized by equally dis- 551  
495 tinct relationships between a complex constellation of facial 552  
496 features, and subsequent task performance. Moreover, in dif- 553  
497 ferent states, performance seemed to be dominated by spe- 554  
498 cific sensory modalities, e.g. eyes versus nose for hit versus 555  
499 wrong states in mice. This suggests that, depending on an an- 556  
500 imal's internal state, the relation between facial features and 557  
501 subsequent task performance can shift dramatically. 558  
502 These findings stand in marked contrast to previous research, 559  
503 which has mainly highlighted linear relationships between 560  
504 single facial features (e.g. pupil size or eye movements) and 561  
505 isolated cognitive states (e.g. attention) (33–41). Our find- 562  
506 ings imply that such analyses miss out on a large portion 563  
507 of the information available through complex facial expres- 564  
508 sions. Our findings were not dependent on the use of one spe- 565  
509 cific MSLR model, as we repeated our analyses with a differ- 566  
510 ent model and training pipeline (a GLM-HMM, training an 567  
511 individual model per animal and experimental session), with 568  
512 very similar results (61). 569  
513 Most importantly, the internal states revealed in this manner 570  
514 mapped robustly onto behavioural trial outcomes (i.e. hit, 571  
515 wrong and miss trials) - even though this information had 572  
516 been in no way part of the inputs the MSLR received. This 573  
517 suggests that the hidden states highlighted by the MSLR were 574  
518 not simply 'computational devices' increasing the model's 575  
519 predictive power. Instead, they appear to reflect genuine, dy- 576  
520 namically fluctuating cognitive states, which result in distinc- 577

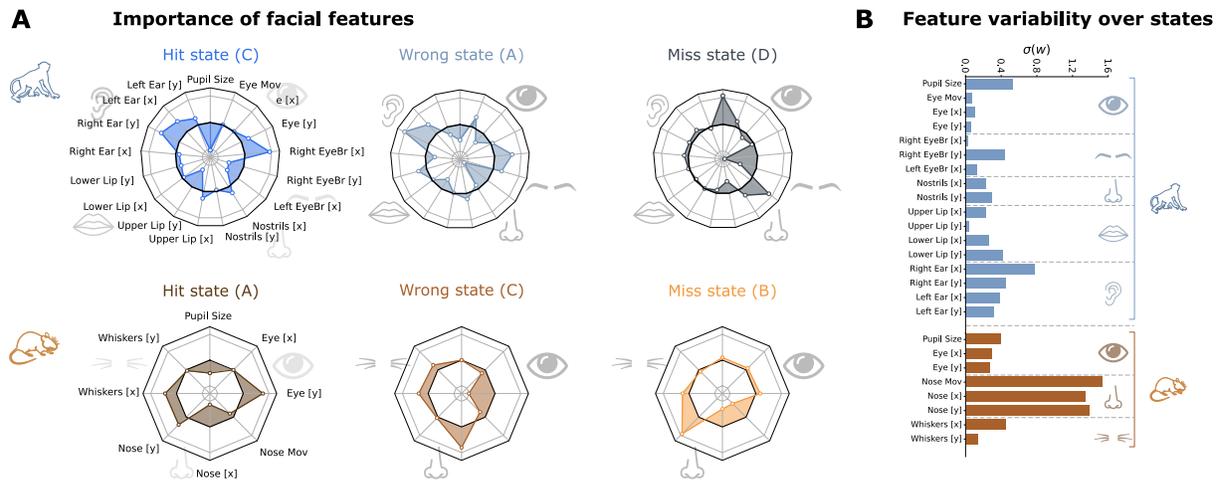
521 tive behavioural outcome profiles.

522 Interestingly, despite the fact that the optimal number of 523  
524 states was determined separately for each species and in a 524  
525 purely data-driven way, our approach converged onto a low 525  
526 and noticeably similar number of internal states for both 526  
527 species: three states for mice, four for macaques. How com- 527  
528 parable are these internal states of mice and monkeys? 528

529 We found that in terms of the dynamics by which animals 529  
530 traversed different internal states, results diverged across 530  
531 species. Specifically, mice appeared to transition more fre- 531  
532 quently between states than monkeys. A control analysis that 532  
533 matched the number of subjects, trials and facial parameters 533  
534 across species before fitting the MSLR models showed that 534  
535 this difference is not a trivial result of divergences in data 535  
536 structure (see Fig. S12). Given that mice have previously 536  
537 been shown to alternate between strategies during perceptual 537  
538 decision-making (43), this finding may point at a genuine dif- 538  
539 ference in the cognitive dynamics of mice and monkeys. Al- 539  
540 ternatively, the prolonged training time of the monkeys com- 540  
541 pared to the mice may have given them the chance to con- 541  
542 verge on more stable behavioural strategies over the course 542  
543 of training (62, 63). Recordings from mice that experienced 543  
544 a more prolonged training scheme and/or from more naive 544  
545 monkeys will give fascinating insights into the role of ex- 545  
546 pertise in fostering more stable transitions between cognitive 546  
547 states. 547

548 Beyond state dynamics, the constellation of behavioural pro- 548  
549 files covered by different states was also largely comparable 549  
550 across species (64, 65). Each hidden state predicted only a 550  
551 narrow range of reaction times; and when relating the in- 551  
552 ferred hidden states to task performance beyond the RTs 552  
553 that the model was trained to predict, we found that states 553  
554 mapped onto the behavioural outcomes (hit, wrong, miss) 554  
555 with distinct probabilities. Moreover, each hidden state cov- 555  
556 ered unique combinations of RT ranges and trial outcomes 556  
557 (hit, miss and wrong trials), despite the fact that trial out- 557  
558 comes had not been part of the MSLR in any way. Specif- 558  
559 ically, both monkeys and mice display a state where trial 559  
560 outcome is typically slow and unsuccessful (which could be 560  
561 interpreted as 'inattentive'), as well as several states where 561  
562 performance is largely fast and correct, with a preference for 562  
563 thoroughness in one state, and a preference for speed (and 563  
564 potentially impulsivity) in the other. These states potentially 564  
565 map onto various levels of task-related attention, and further 565  
566 support the notion that classical concepts of attention can in- 566  
567 deed reflect much of the internal structure of goal-directed 567  
568 behaviour, also in naturalistic settings. 568

569 The fact that different states are associated with distinct con- 569  
570 stellations of facial features points to a role of facial expres- 570  
571 sions beyond emotional expression. Facial expressions have 571  
572 so far been mostly studied in a social or emotional context, 572  
573 and mostly in social species such as monkeys (28, 66) and 573  
574 humans (25, 67). In mice, until recently facial expressions 574  
575 were thought to mainly reflect pain (29, 68, 69), until careful 575  
576 analyses using machine-learning algorithms identified their 576  
577 facial expressions as innate and accurate reflections of sev- 577  
578 eral emotional states as well (22, 31, 70). Our results suggest



**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 Fig4B). Central circle indicates a predictor weight of zero, inside this circle are negative predictor weights, outside are positive weights. Each state has its own characteristic facial expression pattern. **B)** 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.

578 that similarly to humans, facial expressions in monkeys and 616  
 579 mice also convey cognitive and motivational variables such 617  
 580 as focus or cognitive strain, even in the absence of a particu- 618  
 581 lar emotional or social context. 619

582 The fact that such performance-related states are equally ap- 620  
 583 parent in both species is particularly surprising, since one 621  
 584 can assume that the prominent differences between the two 622  
 585 species (such as the acuity and dominance of their visual 623  
 586 system) would imply that they are likely to solve tasks us- 624  
 587 ing different strategies. Yet both species are subject to task- 625  
 588 independent internal states, such as slow fluctuations in at- 626  
 589 tention (52–54, 59). We believe that the internal states we 627  
 590 tap into with the current approach are more reminiscent of 628  
 591 such fluctuations, and as such, can be important indicators 629  
 592 of underlying brain-wide activity fluctuations. Global in- 630  
 593 ternal states such as arousal, motivation, and attention typ- 631  
 594 ically manifest themselves via brain-wide dynamics, and it 632  
 595 will be an exciting endeavour to investigate how well those 633  
 596 neuronally defined internal states correspond to the ones we 634  
 597 here identified behaviourally. 635

598 The MSLR model that we used yielded single estimates of 636  
 599 the internal states per trial. This constitutes a great basis for 637  
 600 time-resolved tracking of internal states, which can be further 638  
 601 extended in future, using MSLR models with higher temporal 639  
 602 resolution. Such MSLR models will be able to align identi- 640  
 603 fied internal states with specific events within each trial, such 641  
 604 as the appearance and disappearance of stimuli, thereby al- 642  
 605 lowing for more precise characterisation of their dynamics 643  
 606 and functional roles. 644

607 Perhaps even more importantly, such a time resolved MSLR 646  
 608 would also allow us to link cognitive processes to neural ac- 647  
 609 tivity on a moment by moment basis, without the need for 648  
 610 repeatedly presenting identical trials and then doing exten- 649  
 611 sive post-hoc averaging. As the MSLR model yields a time- 650  
 612 resolved estimate of cognitive states, these time courses can 651  
 613 be directly compared to continuous neural activity. As such, 652  
 614 this approach opens up a much more naturalistic view of the 653  
 615 neuro-behavioural dynamics involved in spontaneous cogni- 654  
 655  
 656  
 657  
 658

513 tive states than traditional approaches can offer (44, 71).  
 514 These findings suggest that in an ecologically valid frame-  
 515 work that applies across species (in this case, a foraging-  
 516 based task set in a naturalistic, immersive visual environ-  
 517 ment), many features of cognitive processing are more sim-  
 518 ilar than classical paradigms might have suggested. At the  
 519 same time, presumably genuine cross-species differences,  
 520 e.g. in the transition frequency between cognitive states, also  
 521 become more apparent.

522 In summary, we have shown here that in both monkeys and  
 523 mice, facial features can be used to infer internal cognitive  
 524 states, and to track their spontaneous dynamics over time.  
 525 With this approach, we find that the basic attributes of such  
 526 internal states map onto known cognitive states such as at-  
 527 tention in both species in a translatable way, but that the dy-  
 528 namics by which mice and monkeys traverse these states is  
 529 somewhat different. This highlights the crucial importance of  
 530 using naturalistic behavioural paradigms, especially in cross-  
 531 species research, in order to discern truly species-specific re-  
 532 sults 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. 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.

- 659 7. Georgia G Gregoriou, Stephen J Gotts, Huihui Zhou, and Robert Desimone. High- 745  
660 frequency, long-range coupling between prefrontal and visual cortex during attention. *sci-* 746  
661 *ence*, 324(5931):1207–1210, 2009. 747
- 662 8. Trevor Robbins. The 5-choice serial reaction time task: behavioural pharmacology and 748  
663 functional neurochemistry. *Psychopharmacology*, 163:362–380, 2002. 749
- 664 9. Samuel K Asinof and Tracie A Paine. The 5-choice serial reaction time task: a task of 750  
665 attention and impulse control for rodents. *JoVE (Journal of Visualized Experiments)*, (90): 751  
666 e51574, 2014. 752
- 667 10. Ann Kennedy. The what, how, and why of naturalistic behavior. *Current opinion in neurobi-* 753  
668 *ology*, 74:102549, 2022. 754
- 669 11. Alex Gomez-Marín, Joseph J Paton, Adam R Kampff, Rui M Costa, and Zachary F Mainen. 755  
670 Big behavioral data: psychology, ethology and the foundations of neuroscience. *Nature* 756  
671 *neuroscience*, 17(11):1455–1462, 2014. 757
- 672 12. Alex Gomez-Marín and Asif A Ghazanfar. The life of behavior. *Neuron*, 104(1):25–36, 2019. 758
- 673 13. Yael Niv. The primacy of behavioral research for understanding the brain. *Behavioral Neu-* 759  
674 *rosience*, 135(5):601, 2021. 760
- 675 14. Christopher D Harvey, Phillip Coen, and David W Tank. Choice-specific sequences in pari- 761  
676 etal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62–68, 2012. 762
- 677 15. Roberto A Gullí, Lyndon R Duong, Benjamin W Corrigan, Guillaume Doucet, Sylvain 763  
678 Williams, Stefano Fusi, and Julio C Martínez-Trujillo. Context-dependent representations 764  
679 of objects and space in the primate hippocampus during virtual navigation. *Nature neuro-* 765  
680 *science*, 23(1):103–112, 2020. 766
- 681 16. Katharine A Shapcott, Marvin Weigand, Iuliia Glukhova, Martha N Havenith, and Marieke L 767  
682 Schölvinck. Domevr: A setup for experimental control of an immersive dome virtual envi- 768  
683 ronment created with unreal engine 4. [Manuscript submitted for publication], 2022. doi: 769  
684 10.1101/2022.04.04.486889. 770
- 685 17. Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, 771  
686 Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estima- 772  
687 tion of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 773  
688 2018. 774
- 689 18. Alexander B Wiltschko, Tatsuya Tsukahara, Ayman Zeine, Rockwell Anyoha, Winthrop F 775  
690 Gillis, Jeffrey E Markowitz, Ralph E Peterson, Jesse Katon, Matthew J Johnson, and 776  
691 Sandeep Robert Datta. Revealing the structure of pharmacobehavioral space through mo- 777  
692 tion sequencing. *Nature neuroscience*, 23(11):1433–1443, 2020. 778
- 693 19. Kevin Luxem, Petra Mocellin, Falko Fuhrmann, Johannes Kürsch, Stephanie R Miller, 779  
694 Jorge J Palop, Stefan Remy, and Pavol Bauer. Identifying behavioral structure from deep 780  
695 variational embeddings of animal motion. *Communications Biology*, 5(1):1267, 2022. 781
- 696 20. Michael Pereira, Pierre Megeand, Mi Xue Tan, Wenwen Chang, Shuo Wang, Ali Rezaei, 782  
697 Margitta Seeck, Marco Corniola, Shahan Momjian, Fosco Bernasconi, et al. Evidence ac- 783  
698 cumulation relates to perceptual consciousness and monitoring. *Nature communications*, 784  
699 12(1):3261, 2021. 785
- 700 21. Michael H McCullough and Geoffrey J Goodhill. Unsupervised quantification of naturalistic 786  
701 animal behaviors for gaining insight into the brain. *Current Opinion in Neurobiology*, 70: 787  
702 89–100, 2021. 788
- 703 22. Nate Dolensek, Daniel A Gehrlach, Alexandra S Klein, and Nadine Gogolla. Facial expres- 789  
704 sions of emotion states and their neuronal correlates in mice. *Science*, 368(6486):89–94, 790  
705 2020. 791
- 706 23. Sandeep Robert Datta, David J Anderson, Kristin Branson, Pietro Perona, and Andrew 792  
707 Leifer. Computational neuroethology: a call to action. *Neuron*, 104(1):11–24, 2019. 793
- 708 24. Richard E Quandt. A new approach to estimating switching regressions. *Journal of the* 794  
709 *American statistical association*, 67(338):306–310, 1972. 795
- 710 25. Paul Ekman and Harriet Oster. Facial expressions of emotion. *Annual review of psychology*, 796  
711 30(1):527–554, 1979. 797
- 712 26. Ralph Adolphs, Hanna Damasio, Daniel Tranel, and Antonio R Damasio. Cortical systems 798  
713 for the recognition of emotion in facial expressions. *Journal of neuroscience*, 16(23):7678–799  
714 7687, 1996. 800
- 715 27. Robert Aubrey Hinde and Thelma E Rowell. Communication by postures and facial expres- 801  
716 sions in the rhesus monkey (*macaca mulatta*). In *Proceedings of the Zoological Society of* 802  
717 *London*, volume 138, pages 1–21. Wiley-Blackwell, 1962. 803
- 718 28. Lisa A Parr and Matthew Heintz. Facial expression recognition in rhesus monkeys, *macaca* 804  
719 *mulatta*. *Animal behaviour*, 77(6):1507–1513, 2009. 805
- 720 29. Dale J Langford, Andrea L Bailey, Mona Lisa Chanda, Sarah E Clarke, Tanya E Drummond, 806  
721 Stephanie Echols, Sarah Glick, Joelle Ingrassia, Tammy Klassen-Ross, Michael L LaCroix- 807  
722 Fralish, et al. Coding of facial expressions of pain in the laboratory mouse. *Nature methods*, 808  
723 7(6):447–449, 2010. 809
- 724 30. Christian L Ebbesen and Robert C Froemke. Body language signals for rodent social com- 810  
725 munication. *Current opinion in neurobiology*, 68:91–106, 2021. 811
- 726 31. Olivia Le Moëne and Max Larsson. A new tool for quantifying mouse facial expressions. 812  
727 *Eneuro*, 10(2), 2023. 813
- 728 32. Anna D Zych and Nadine Gogolla. Expressions of emotions across species. *Current opinion* 814  
729 *in neurobiology*, 68:57–66, 2021. 815
- 730 33. Dan Alin Ganea, Alexander Bextert, Mathias Günther, Pierre-Marie Gardères, Björn M 816  
731 Kampa, and Florent Haiss. Pupillary dilations of mice performing a vibrotactile discrimi- 817  
732 nation task reflect task engagement and response confidence. *Frontiers in Behavioral Neu-* 818  
733 *rosience*, 14:159, 2020. 819
- 734 34. Jacob Reimer, Emmanouil Froudarakis, Cathryn R Cadwell, Dimitri Yatsenko, George H 820  
735 Denfield, and Andreas S Tolias. Pupil fluctuations track fast switching of cortical states 821  
736 during quiet wakefulness. *neuron*, 84(2):355–362, 2014. 822
- 737 35. Paola Binda, Maria Pereverzeva, and Scott O Murray. Pupil size reflects the focus of feature- 823  
738 based attention. *Journal of neurophysiology*, 112(12):3046–3052, 2014. 824
- 739 36. R Becket Ebitz, John M Pearson, and Michael L Platt. Pupil size and social vigilance in 825  
740 rhesus macaques. *Frontiers in neuroscience*, 8:100, 2014. 826
- 741 37. Chin-An Wang and Douglas P Munoz. A circuit for pupil orienting responses: implications 827  
742 for cognitive modulation of pupil size. *Current opinion in neurobiology*, 33:134–140, 2015. 828
- 743 38. Siddhartha Joshi and Joshua I Gold. Pupil size as a window on neural substrates of cogni- 829  
744 tion. *Trends in cognitive sciences*, 24(6):466–480, 2020. 830
39. Tirin Moore and Mazyar Fallah. Control of eye movements and spatial attention. *Proceed-* 831  
39 *ings of the National Academy of Sciences*, 98(3):1273–1276, 2001. 832
40. Burkhardt Fischer and B Breitmeyer. Mechanisms of visual attention revealed by saccadic 833  
40 eye movements. *Neuropsychologia*, 25(1):73–83, 1987. 834
41. Miriam Spering. Eye movements as a window into decision-making. *Annual review of vision* 835  
41 *science*, 8:427–448, 2022. 836
42. Sina E Dominiak, Mostafa A Nashaat, Keisuke Sehara, Hatem Oraby, Matthew E Larkum, 837  
42 and Robert NS Sachdev. Whisking asymmetry signals motor preparation and the behavioral 838  
42 state of mice. *Journal of Neuroscience*, 39(49):9818–9830, 2019. 839
43. Zoe C Ashwood, Nicholas A Roy, Iris R Stone, International Brain Laboratory, Anne E Urai, 840  
43 Anne K Churchland, Alexandre Pouget, and Jonathan W Pillow. Mice alternate between 841  
43 discrete strategies during perceptual decision-making. *Nature Neuroscience*, 25(2):201–842  
43 212, 2022. 843
44. Luca Mazzucato. Neural mechanisms underlying the temporal organization of naturalistic 844  
44 animal behavior. *Elife*, 11:e76577, 2022. 845
45. Adam J Calhoun, Jonathan W Pillow, and Mala Murthy. Unsupervised identification of the 846  
45 internal states that shape natural behavior. *Nature neuroscience*, 22(12):2040–2049, 2019. 847
46. Daniel Hulsev, Kevin Zumwalt, Luca Mazzucato, David A McCormick, and Santiago 848  
46 Jaramillo. Decision-making dynamics are predicted by arousal and unstructured move- 849  
46 ments. *bioRxiv*, pages 2023–03, 2023. 850
47. Pedro J Gonçalves, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, 851  
47 Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, 852  
47 Tim P Vogels, et al. Training deep neural density estimators to identify mechanistic models 853  
47 of neural dynamics. *Elife*, 9:e56261, 2020. 854
48. Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and 855  
48 Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis 856  
48 of neural population activity. *Advances in neural information processing systems*, 21, 2008. 857
49. Mohammad Reza Keshkaran, Andrew R Sedler, Raed H Chowdhury, Raghav Tandon, 858  
49 Diya Basrai, Sarah L Nguyen, Hansom Sohn, Mehrdad Jazayeri, Lee E Miller, and Chethan 859  
49 Pandarinath. A large-scale neural network training framework for generalized estimation of 860  
49 single-trial population dynamics. *Nature Methods*, 19(12):1572–1577, 2022. 861
50. Sean M Perkins, John P Cunningham, Qi Wang, and Mark M Churchland. Simple decoding 862  
50 of behavior from a complicated neural manifold. *bioRxiv*, pages 2023–04, 2023. 863
51. James FA Poulet and Sylvain Crochet. The cortical states of wakefulness. *Frontiers in* 864  
51 *systems neuroscience*, 12:64, 2019. 865
52. Marlene R Cohen and John HR Maunsell. Using neuronal populations to study the mecha- 866  
52 nisms underlying spatial and feature attention. *Neuron*, 70(6):1192–1204, 2011. 867
53. Marieke N Havenith, Petra M Zijdereld, Stefan van Heukelum, Rogier B Poorthuis, Hui- 868  
53 bert D Mansvelder, and August B Smit. The virtual-environment-foraging task enables rapid 869  
53 training and single-trial metrics of attention in head-fixed mice. *Scientific Reports*, 8(1): 870  
53 17371, 2018. doi: 10.1038/s41598-018-34966-8. 871
54. Martha N Havenith, Peter M Zijdereld, Sabrina van Heukelum, Shaghayegh Abghari, Paul 872  
54 Tiesinga, and Jeffrey C Glennon. The virtual-environment-foraging task enables rapid train- 873  
54 ing and single-trial metrics of rule acquisition and reversal in head-fixed mice. *Scientific* 874  
54 *reports*, 9(1):4790, 2019. 875
55. Ayumu Yamashita, David Rothein, Aaron Kucyi, Eve M Valera, Laura Germine, Jeremy 876  
55 Wilmer, Joseph DeGutis, and Michael Esterman. Variable rather than extreme slow reaction 877  
55 times distinguish brain states during sustained attention. *Scientific reports*, 11(1):14883, 878  
55 2021. 879
56. Michael Esterman, Sarah K Noonan, Monica Rosenberg, and Joseph DeGutis. In the zone or 880  
56 zoning out? tracking behavioral and neural fluctuations during sustained attention. *Cere-* 881  
56 *bral cortex*, 23(11):2712–2723, 2013. 882
57. Matteo Carandini. Amplification of trial-to-trial response variability by neurons in visual cor- 883  
57 tex. *PLoS biology*, 2(9):e264, 2004. 884
58. Anne E Urai, Jan Willem De Gee, Konstantinos Tsetsos, and Tobias H Donner. Choice 885  
58 history biases subsequent evidence accumulation. *Elife*, 8:e46331, 2019. 886
59. Kenneth D Harris and Alexander Thiele. Cortical state and attention. *Nature reviews neu-* 887  
59 *roscience*, 12(9):509–523, 2011. 888
60. Marieke L Schölvinck, Aman B Saleem, Andrea Benucci, Kenneth D Harris, and Matteo 889  
60 Carandini. Cortical state determines global variability and correlations in visual cortex. *Jour-* 890  
60 *nal of Neuroscience*, 35(1):170–178, 2015. 891
61. Alejandro Tlaie, Katharine Shapcott, Muad Abd el Hay, Berkutay Mert, Pierre-Antoine Fer- 892  
61 racci, Robert Taylor, Iuliia Glukhova, Martha N Havenith, and Marieke Schölvinck. Thought- 893  
61 ful faces: Using facial features to infer naturalistic cognitive processing across species. 894  
61 *Cosyne Abstracts*, 2023. doi: 10.1145/1217935.1217951. 895
62. Patricia D Stokes. Learned variability. *Animal Learning & Behavior*, 23:164–176, 1995. 896
63. Patricia D Stokes, Francis Mechner, and Peter D Balsam. Effects of different acquisition 897  
63 procedures on response variability. *Animal Learning & Behavior*, 27:28–41, 1999. 898
64. Ben Mitchinson and Tony J Prescott. Whisker movements reveal spatial attention: a unified 899  
64 computational model of active sensing control in the rat. *PLoS computational biology*, 9(9): 900  
64 e1003236, 2013. 901
65. Ehud Fonio, Goren Gordon, Noy Barak, Yonatan Winetraub, Tess Baker Oram, Sebastian 902  
65 Haidarliu, Tali Kimchi, and Ehud Ahissar. Coordination of sniffing and whisking depends on 903  
65 the mode of interaction with the environment. *Israel Journal of Ecology & Evolution*, 61(2): 904  
65 95–105, 2015. 905
66. Xin-He Liu, Lu Gan, Zhi-Ting Zhang, Pan-Ke Yu, and Ji Dai. Probing the processing of facial 906  
66 expressions in monkeys via time perception and eye tracking. *Zoological Research*, 44(5): 907  
66 882, 2023. 908
67. Alan J Fridlund. *Human facial expression: An evolutionary view*. Academic press, 2014. 909
68. Alexander H Tittle, Mark J Molinaro, Jasmine F Jethwa, Susana G Sotocinal, Juan C Prieto, 910  
68 Martin A Styner, Jeffrey S Mogil, and Mark J Zyka. A deep neural network to assess 911  
68 spontaneous pain from mouse facial expressions. *Molecular pain*, 14:1744806918763658, 912  
68 2018. 913
69. Amy L Miller and Matthew C Leach. The mouse grimace scale: a clinically useful tool? *PLoS* 914  
69 *one*, 10(9):e0136000, 2015. 915
70. Erwin B Defensor, Michael J Corley, Robert J Blanchard, and D Caroline Blanchard. Facial 916  
70 917

- 831 expressions of mice in aggressive and fearful contexts. *Physiology & behavior*, 107(5):  
832 680–685, 2012.
- 833 71. Simon Musall, Matthew T Kaufman, Ashley L Juavinett, Steven Gluf, and Anne K Church-  
834 land. Single-trial neural dynamics are dominated by richly varied movements. *Nature neu-*  
835 *roscience*, 22(10):1677–1686, 2019.
- 836 72. Eleni Psarou, Julien Vezoli, Marieke L Schölvinck, Pierre-Antoine Ferracci, Yufeng Zhang,  
837 Iris Grothe, Rasmus Roese, and Pascal Fries. Modular, cement-free, customized headpost  
838 and connector-chamber implants for macaques. *Journal of Neuroscience Methods*, 393:  
839 109899, 2023.
- 840 73. Tim Schroeder, Robert Taylor, Muad Abd El Hay, Abdellatif Nemri, Arthur Franca, Francesco  
841 Battaglia, Paul Tiesinga, Marieke L Schoelvinck, and Martha N Havenith. The dream im-  
842 plant: A lightweight, modular and cost-effective implant system for chronic electrophysiology  
843 in head-fixed and freely behaving mice. *bioRxiv*, pages 2024–02, 2024.
- 844 74. Keiji Matsuda, Takeshi Nagami, Yasuko Sugase, Aya Takemura, and Kenji Kawano. A widely  
845 applicable real-time mono/binocular eye tracking system using a high frame-rate digital cam-  
846 era. In *Human-Computer Interaction. User Interface Design, Development and Multimodal-*  
847 *ity: 19th International Conference, HCI International 2017, Vancouver, BC, Canada, July*  
848 *9-14, 2017, Proceedings, Part I 19*, pages 593–608. Springer, 2017.
- 849 75. Shaokai Ye, Anastasiia Filippova, Jessy Lauer, Maxime Vidal, Steffen Schneider, Tian Qiu,  
850 Alexander Mathis, and Mackenzie Weygandt Mathis. Superanimal models pretrained for  
851 plug-and-play analysis of animal behavior, 2023.
- 852 76. Chang, P, Harper-Donnelly, G, Kara A, Li X, et al, and Linderman S Murphy K. Dynamax, a  
853 library for State Space Models. *GitHub*, 2022.
- 854 77. Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete  
855 data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*,  
856 39(1):1–22, 1977.
- 857 78. Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam  
858 Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems.  
859 In *Artificial Intelligence and Statistics*, pages 914–922. PMLR, 2017.
- 860 79. Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Op-  
861 tuna: A next-generation hyperparameter optimization framework. In *Proceedings of the*  
862 *25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages  
863 2623–2631, 2019.
- 864 80. Tom AB Snijders. On cross-validation for predictor evaluation in time series. In *On Model*  
865 *Uncertainty and its Statistical Implications: Proceedings of a Workshop, Held in Groningen,*  
866 *The Netherlands, September 25–26, 1986*, pages 56–69. Springer, 1988.
- 867 81. Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*,  
868 2016.

## 869 Methods

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

874 **Surgical Procedures.** All animals were fitted with custom-milled headposts for the purpose of head fixation during this  
875 experiment. The headpost design and implant procedures for the macaques have been extensively discussed in (72). Briefly, a  
876 four-legged titanium baseplate was screwed into the skull under general anesthesia. After several weeks of osseo-integration,  
877 a titanium top part was screwed onto the baseplate in a simple procedure. The headposts for the mice have been described in  
878 (73). Briefly, the animal was placed under isoflurane anaesthesia, shaved and given local analgesia on the top of the head. An  
879 incision was made and the skin on top of the cranium was removed, before the cranium was cleaned and the custom milled  
880 titanium head plate was attached using dental cement.

881 **Experimental Setup.** Experiments were carried out in a darkened room (mice) or electrically shielded booth (monkeys). The  
882 animals were in the centre of a 120-cm diameter spherical dome extending to 250 deg visual angle. The headfixed mice were  
883 positioned on a styrofoam spherical treadmill; the headfixed monkeys were seated in a monkey chair and could spin a 12-cm  
884 diameter trackball with their hands. Movements of the spherical treadmill and trackball allowed the animals to traverse a virtual  
885 reality (VR) environment projected on the inside of the dome by means of a spherical mirror. Projecting the VR environment  
886 on a dome surrounding the animals enabled both their central and peripheral view to be covered, thereby providing an immer-  
887 sive and realistic VR environment. The VR environment was created using DomeVR, our custom-made toolbox combining  
888 photorealistic graphics rendered with Unreal Engine 4, with high timing precision required for neuroscience experiments (16).

889 **Experimental Paradigm.** Mice and monkeys were required to distinguish two natural shapes at equal distance in front of them,  
890 amidst a grassy field with a blue sky above and mountains in the background (Fig 1A). The two shapes emerged out of a central  
891 shape which was either right at the starting position (for monkeys) or a short distance in front (for mice). A virtual collision  
892 with the correct shape yielded a reward ('hit'), whereas the incorrect shape yielded no reward ('wrong'), and no collision with  
893 either shape also yielded no reward ('miss') (2AFC paradigm). Rewards were drops of diluted juice for the monkeys and drops  
894 of vanilla soy milk for the mice. For the monkeys, the shapes varied smoothly between a non-rewarded, textured square and a  
895 rewarded triangle (monkey K) or between a rewarded, jagged and a non-rewarded, hour-glass shaped leaf (monkey C). On each  
896 trial, a blend between the two shapes was shown alongside the exact middle blend ('reference shape'). For the mice, the shapes  
897 and their reward contingencies were the same as for monkey C.

898 Monkey data were recorded in 7 sessions for monkey C, 11 sessions for monkey K. Each session lasted about one hour, during  
899 which the monkeys completed  $1208 \pm 186$  and  $991 \pm 492$  trials at 67 and 77 percent correct (monkeys C and K, respectively).  
900 The monkeys were both fully trained on handling the trackball to move through the VR environment, as well as the VR task.  
901 Mouse data were recorded in (12, 4, 6, 3, 2, 2) sessions for mice (001, 003, 004, 005, 012, 013), respectively. Each session  
902 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  
903  $(59, 54, 60, 77, 45, 63)$  percent correct (same mice ordering as before). Following the headpost surgery, the mice were handled  
904 for 5 days to reduce experimental anxiety due to head fixation and interaction with the experimenter, before behavioural training  
905 began. Behavioural training in the experimental setup at initial stages lasted between 3-5 sessions, before final data collection  
906 began, which lasted up to 30 sessions.

907 **Behavioural tracking.** We recorded videos of the monkeys' and mice' faces during the tasks at 60 Hz  
908 using Basler acA640-121gm infrared cameras with a modified version of PylonRecorder2 software  
909 (<https://gitlab.mpcdf.mpg.de/mpibr/scic/pylonrecorder/PylonRecorder2>). Additionally, in the monkeys, eye movements  
910 were recorded at 500 Hz using a Grasshopper3 infrared camera and the free eye tracking software *iRecHS2* (74) and  
911 synchronized with DomeVR (16).

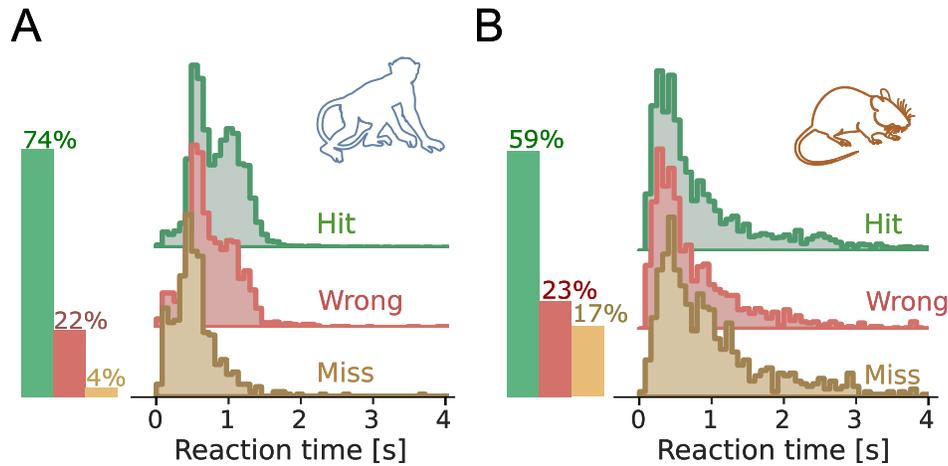
912 **Facial key point extraction.** To extract facial key points from the videos, we used markerless pose estimation on them, as  
913 implemented in DeepLabCut (17, 75). For mice, features were extracted from videos of the left side of the face using our own  
914 model to identify key points such as the coordinates of the eye, whisker pad and nose. For mouse pupillometry, we used the  
915 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  
916 the "mouse pupil vclose" Animal Zoo model (provided by Jim McBurney-Lin at the University of California Riverside, USA)  
917 included with DeepLabCut. The output of the pupil model was 8 points covering the circumference of the mouse pupil, that  
918 were then used to calculate pupil and eye summary statistics.

919 For the macaque facial key points, we used the pre-trained "primate face" model from the DeepLabCut Animal Zoo (provided  
920 by Claire Witham at the Centre for Macaques, MRC Harwell, UK) and extended it with additional points on the lips to capture

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

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

928 **Reaction Time.** In our VR setting, where animals move towards one of two stimuli rather than pressing a button or lever,  
 929 or making an eye movement, we define the reaction time (RT) as the time point of the initial substantial movement directed  
 930 towards either stimulus. While determining this time point, it is crucial to distinguish between stimulus-related movements  
 931 and minor positional adjustments. We specifically focus on the first deviation in lateral movement, while excluding forward  
 932 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.

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

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

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

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

941 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  
 942 points (we have a univariate function in this case):

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

943 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  
 944 they have a minimum prominence ( $\lambda$ ). Prominence is a measure of the significance of a peak by comparing the peak to its  
 945 surroundings:

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

946 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  
 947 by which  $r(w) \equiv -R^2$ .

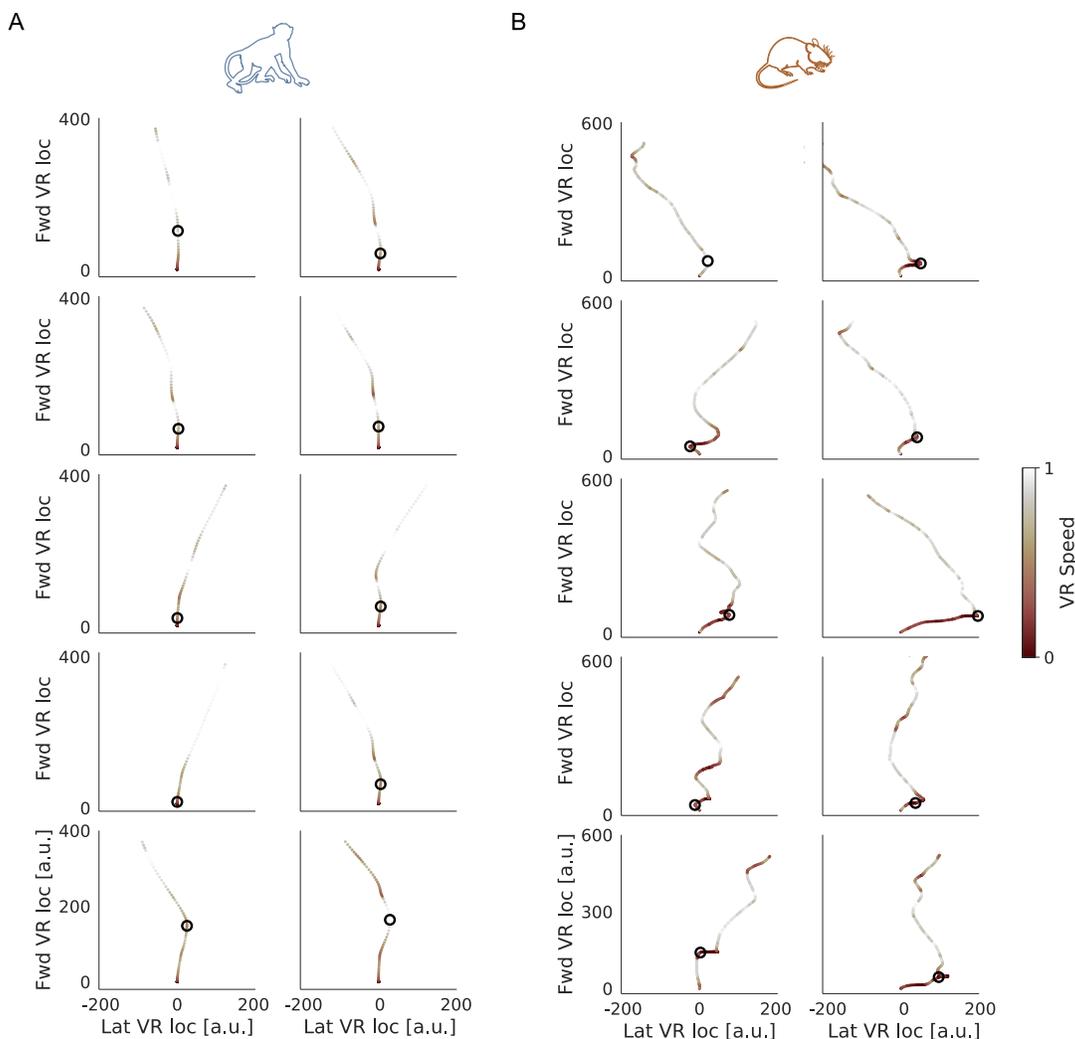
948 For each peak in  $r(w)$ , we calculate the prominence and discard the ones that are below a given threshold ( $\lambda_0$ ). The particular  
 949 value for this threshold was not critical for the overall performance of the algorithm. For the sake of stability, we use multiple  
 950 window sizes (100, 150, 200 and 250 ms) and combine the results in the following way. For each window  $k$ , we have an array  
 951 of candidate points ( $x_{cand}^k$ ). Then, we create a vector of weights ( $w_k \in \mathbb{R}^n$ ) that have a value equal to a Gaussian distribution  
 952 centered around each candidate point of each window. Mathematically:

$$w_k(x) = \begin{cases} \mathcal{N}(x - x_{cand}^k, \sigma) & \text{if } x \in \mathcal{B}_{cand}^k \\ 0 & \text{otherwise} \end{cases}$$

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

$$RT = \operatorname{argmax}_x \left[ \left( \sum_k w_k(x) \right) / x \right] \quad (5)$$

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

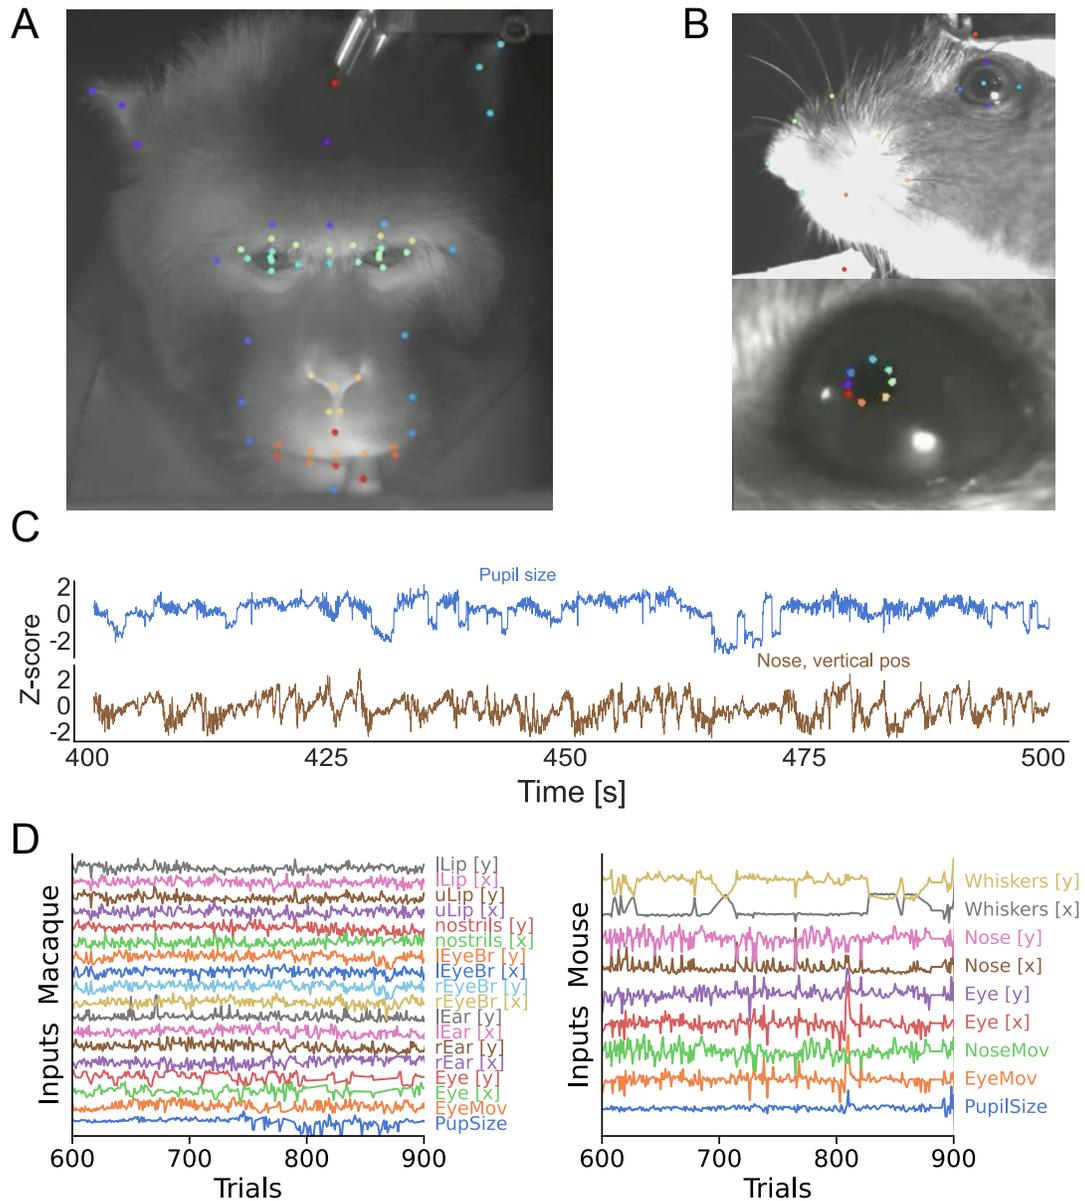


**Fig. S2. Example VR paths** Paths are colored according to the normalized running speed. **A)** Example paths for macaques, with the detected RT as circles. **B)** Same, but for mice.

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

958 First, we chose several points of interest on the animals' faces, which are then automatically identified and tracked over time  
 959 using DeepLabCut (17):

- 960 • Macaque: both ears, eyebrows, nostrils and lips (see Fig. S3 A).
- 961 • 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.

962 Once the data streams were aligned, we computed the median location  $(x, y)$  of each facial point over the 250 ms window before  
 963 the stimuli appeared on the dome. This time window was chosen to make sure that all of the facial expressions of the animals  
 964 are due to internally generated processing, rather than stimulus processing. Different window sizes (particularly: 200, 300 and  
 965 500 ms) did not yield any qualitative difference. In addition to the median location, we also computed the total velocity of each  
 966 facial point.

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

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

971 **Markov-Switching Linear Regression.** Markov-Switching Linear Regression (MSLR) models, which we ran using *Dyna-*  
 972 *max* (76), are a powerful tool for modeling time series data that exhibit regime-switching behaviour, where the underlying

973 dynamics of the system change over time.

974 The MSLR model is defined by a set of linear regressions, each associated with a particular state of a discrete Markov chain.  
975 The state of the Markov chain determines which sets of weights and biases predicts the evolution of the observed data at each  
976 time step. The transitions between states are governed by the transition probabilities of the Markov chain, which are learned  
977 from the data.

978 Formally, an MSLR model can be described as follows. If  $S$  is the total number of latent (discrete) states of a Markov process,  
979 at each time step  $t$ , a given state  $z_t (\in \{0, 1, \dots, S\})$  will follow a Markovian evolution such that:

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

980 As these are stochastic matrices,  $\pi_{ij} \in [0, 1]$ .

981 Let the  $M$ -dimensional input time series at time  $t$  be denoted by  $x_t (\in \mathbb{R}^M)$ . Let the  $N$ -dimensional output time series at  
982 time  $t$  be denoted by  $y_t (\in \mathbb{R}^N)$ . Then, in the case of a MSLR, the discrete latent variable at time  $t$  ( $z_t$ ), will dictate which  
983 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  
984 inputs (predictors). Moreover, an emission covariance matrix ( $\Sigma_s \in \mathbb{R}_{\geq 0}^{N \times N}$ ) will also have to be learnt. Explicitly, at time  $t$ ,  
985 the emission distribution in this model is given by:

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

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

$$\theta = \left\{ (W_s, b_s, \pi_s, \Sigma_s) \right\}_{s=1}^S \quad (8)$$

987 In other words, the aim is to find the weights ( $W_s$ ) and biases ( $b_s$ ) for the linear regressions and the transition  $\pi_s$  and covariance  
988  $\Sigma_s$  matrices for the Markov process.

989 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  
990 features ( $x_t$ ) extracted using DeepLabCut (17) and the observation ( $y_t$ ) that represents the RT of the animal. We trained the  
991 MSLR model using the Expectation-Maximization (EM) algorithm (77), which iteratively computes the probability over latent  
992 states given the data and updates the model parameters to maximize the likelihood of the observed data. For further details, we  
993 refer the reader to (78). We iterated the EM algorithm for 50 times, for all models. We initialized the model parameters using  
994 a normal distribution for weights and biases and we used the identity matrix as the initial covariance matrix for the emissions.  
995 We assumed a Dirichlet prior for the transition matrix. We repeated this process 10 times to increase confidence that we got the  
996 optimum value for each combination of parameters.

997 **Training and inference.** We used an 80 : 20 ratio for train-test splitting and performed hyperparameter optimization by cross-  
998 validating the training set only (see *Model tuning* for details on CV and model selection). For each species, we concatenated  
999 the training sets of all sessions, with forced transitions in between the sessions (setting predictors and emissions to 0 for 50  
1000 consecutive trials), so that state probabilities are reset. Then, after optimizing each model, we performed inference on each held  
1001 out test set (separately per session). We decided to take this approach for various reasons:

- 1002 • Model generalization: as the model learns from potentially different faces, it is likely that it can pick up on common  
1003 information between them.
- 1004 • Model interpretability: given that we do not update the model parameters at the inference step, all internal states have the  
1005 same meaning over subjects and, thus, are directly comparable.
- 1006 • Better convergence: increasing the number of training samples (i.e. concatenating sessions as opposed to training a  
1007 different model per session) allows the model to have more data to learn from.

1008 All of the results in the main text, unless otherwise stated, are for held out data.

1009 **Model tuning.** For the model we described in *Markov-Switching Linear Regression*, there are several parameters that can be  
1010 tuned to explain the data better. In our case, we decided to explore the influence of changing the maximum number of internal  
1011 states ( $S$ ), to add sticky transitions to the Markov process (a self-bias term in the transition matrix  $\pi$ , making states taking  
1012 longer to transition to a different one), and to vary the transition matrix sparsity (concentration).

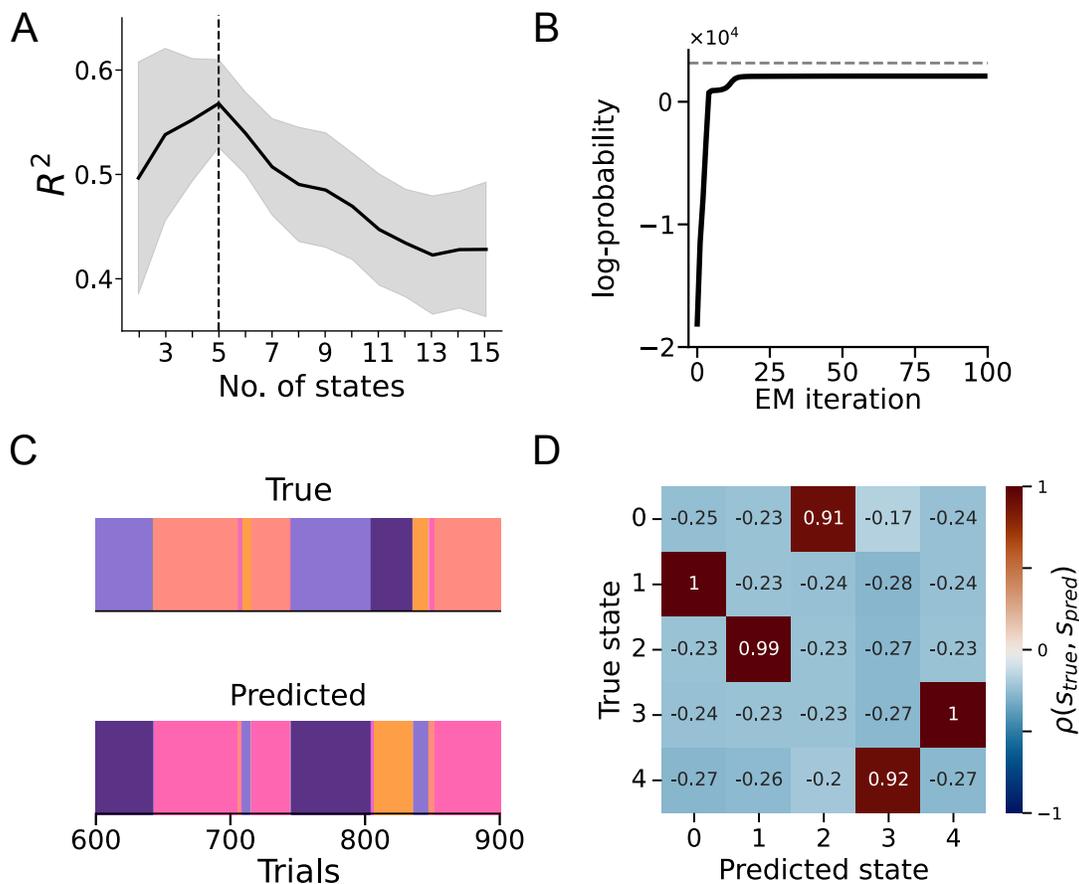
1013 In order to balance model performance with scientific insights, we took a hybrid approach. We increased the number of internal  
1014 states in a greedy way, to show that the error saturates and that there are diminishing returns when increasing model complexity.  
1015 On the other hand, for a given number of states, we optimized two free parameters of the Markov process: state stickiness and  
1016 state concentration. For the sake of efficiency, we used Optuna (79), a flexible framework to implement Bayesian optimization.  
1017 In Table 1 we report the relevant quantities for this process.

1018 To select the best combination of parameters, we performed 5-fold Time-Blocked Cross-Validation (80).

Parameter	Value
Concentration	[0, 100]
Stickiness	[0, 100]
Sampler	CMA-ES (81)
Objective function	$R^2$
Number of Searches	100

**Table 1.** Parameter values for the Bayesian parameter optimization procedure. These are independently explored for each number of internal states of the HMM.

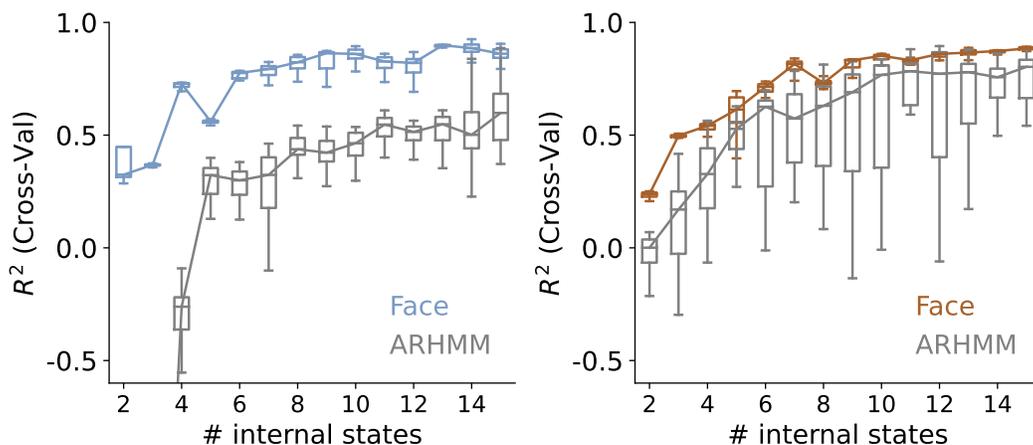
1019 **Synthetic data and ground truth states.** In order to validate the retrieval of states when we do not have access to ground  
 1020 truth ones, we generated a time series of ground truth emissions and states based on the given inputs (by using the same  
 1021 input data as in the main text). To this end, we trained an MSLR model with a known given number of states and sampled  
 1022 some emissions and states sequence from it. We aimed to recover the appropriate number of states with the correct temporal  
 1023 sequence, and to correctly predict the emissions. Figure S4A illustrates the input data (composed of session-concatenated  
 1024 mouse facial features, as described in the main text). In Figure S4C, we show that, once we have selected the appropriate  
 1025 number of states, the model's log-probability does peak at the ground-truth one (dashed vertical line). In Figure S4D, we show  
 1026 a comparison between the true and the inferred states, for some example trials. Although the temporal coincidence of the state  
 1027 transitions is very high, due to the stochastic nature of the model, some state labels might be permuted (i.e., state 1 in our  
 1028 model might correspond to state 0 in the ground truth states). Therefore, in order to quantify state similarities and to account  
 1029 for state-swapping, we one-hot encoded the true and predicted states sequences and correlated all pairs with each other (Figure  
 1030 S4D). There is an almost perfect match ( $\rho(s_{true}, s_{pred}) > 0.9$ ) between the true and inferred states (the 99<sup>th</sup> percentile of the  
 1031 surrogate correlation distribution was 0.12).



**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 converges to the ground truth value (horizontal dashed line). **C)** Some example trials for the true and predicted states. State transitions are correctly captured, but state labels might be permuted. **D)** Temporal correlation between the one-hot encoded state arrays. There is an almost perfect match in between the predicted and the true states arrays, up to a label permutation.

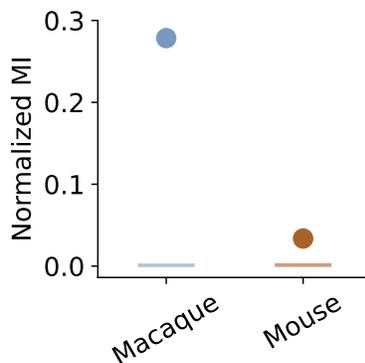
1032 **ARHMM.** As we wanted to ensure that facial features were indeed informative of reaction times (RTs) beyond what is to be  
 1033 expected by the RT autocorrelation structure, we implemented an Auto-Regressive Hidden Markov Model (ARHMM). In this

1034 case, we used the same pipeline as we detailed in the previous sections, but substituted the facial features at the current trial  $t$   
1035 for the RT of the previous trial ( $t - 1$ ). As it can be seen from Fig. S5, the facial features model outperforms the ARHMM for  
1036 all states, for both species.



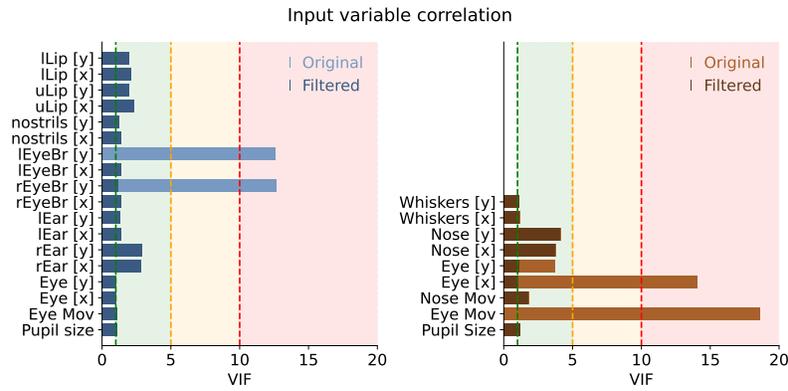
**Fig. S5. Comparison of the MSLR face model and the reaction time Auto-Regressive HMM.** Both face feature models outperform their Autoregressive counterparts, for any number of internal states that we swept over. Nevertheless, it can be seen that the performance gap is smaller in mice than in macaques. This is consistent with the finding that mice are more history dependent than macaques (See Fig. 5 D).

1037 **Task performance and internal states.** We were interested in investigating whether the inferred internal states were corre-  
1038 lated with task performance, even though the model had not been trained on such information. We therefore used the predicted  
1039 single-trial state probabilities to decode choice, using a simple Logistic Regression model, with a  $L_2$  penalty term. After veri-  
1040 fying that the model does indeed classify outcome beyond chance level (Fig. S6), we took the weight of each state as a proxy  
1041 for how related it was to each outcome.

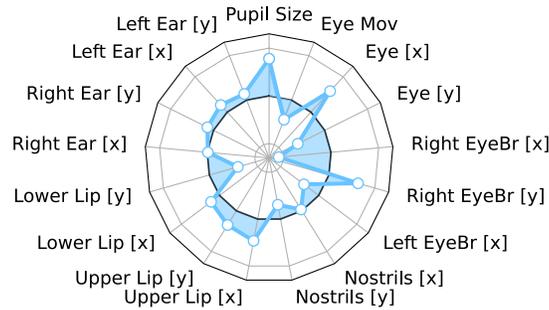


**Fig. S6. Inferred state probabilities decode outcome beyond chance.** We used a normalized version of Mutual Information that already takes chance level into account and sets that as 0.

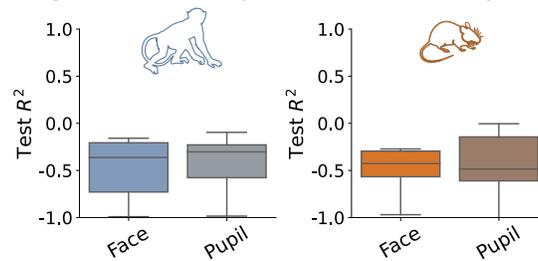
## 1042 Supplementary figures



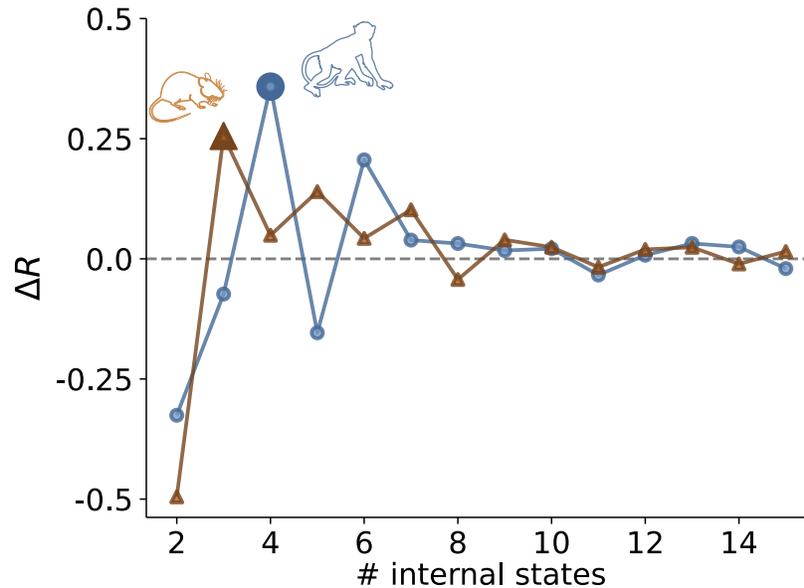
**Fig. S7. Input variable correlation.** We show that, out of all of the original variables (in lighter colors), we end up discarding one per animal (*Left Eyebrow [y]*, macaques; *Eye movement*, mice), given that they were highly multi-colinear with some of the other predictors, as measured by the Variance Inflation Factor (VIF). After discarding them and recomputing the VIF, we did not find any alarming colinearity.



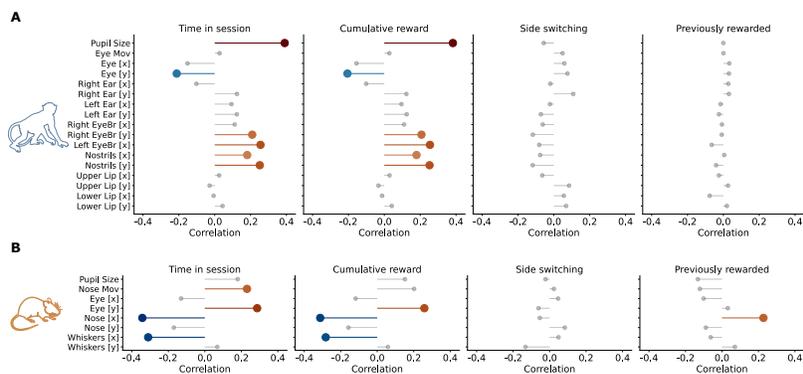
**Fig. S8. Face features importance for State B in macaques.**



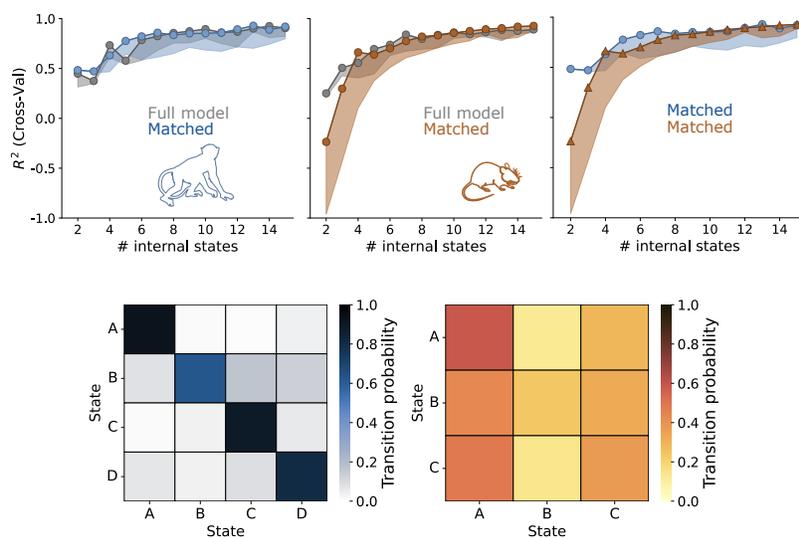
**Fig. S9. No-switching model performances.** If we assume there is a simple linear relationship (i.e. we do not allow for any switching) between the face (or the pupil) and the Reaction Time, the test performance is back at chance level.



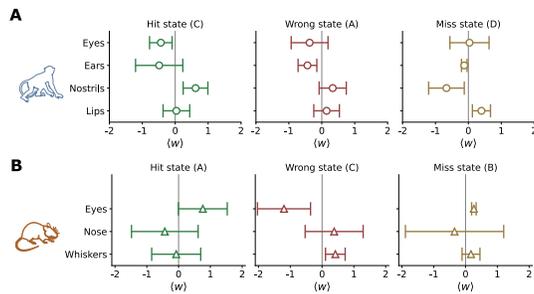
**Fig. S10. Finite difference performance curves.** We selected the number of states that was given by the maximum relative gain in performance, computed as the finite difference in Cross-Validated  $R^2$ .



**Fig. S11. 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. S12. Matched models.** Our main findings still hold if we match the number of predictors, trials and animals we use in each species.



**Fig. S13. Simplified face features weights.**

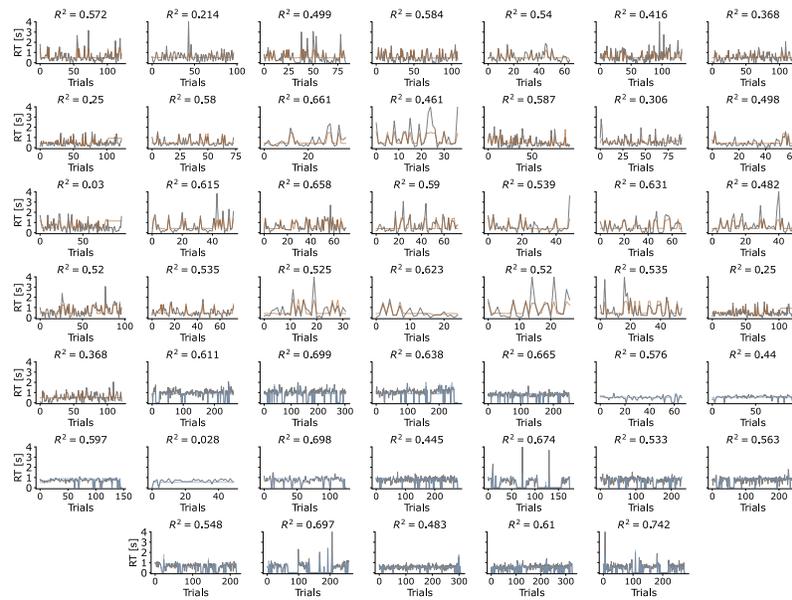


Fig. S14. Predicted and true RTs for all sessions, for the held out test set. For both species, true RTs are shown in gray; for mice, predicted RTs are shown in orange; for macaques, predicted RTs are shown in blue. Each subplot is titled with that session's model performance ( $R^2$ ).

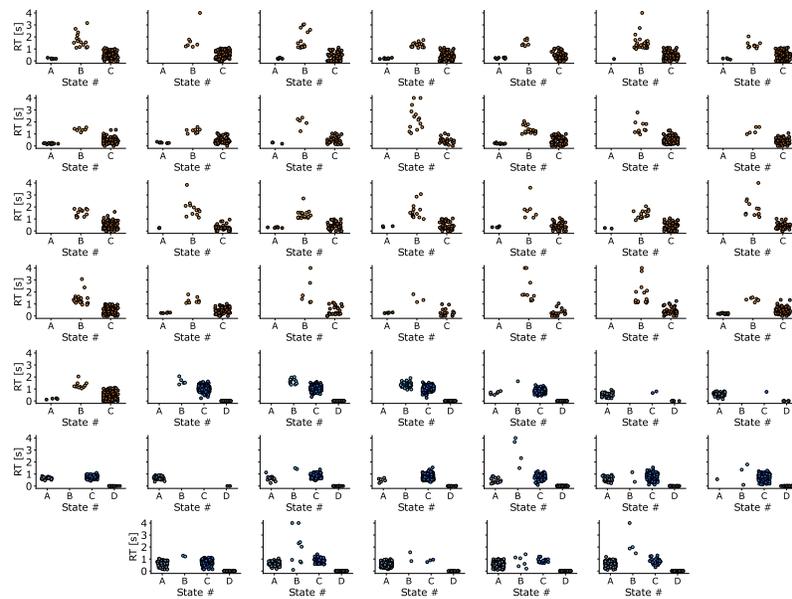
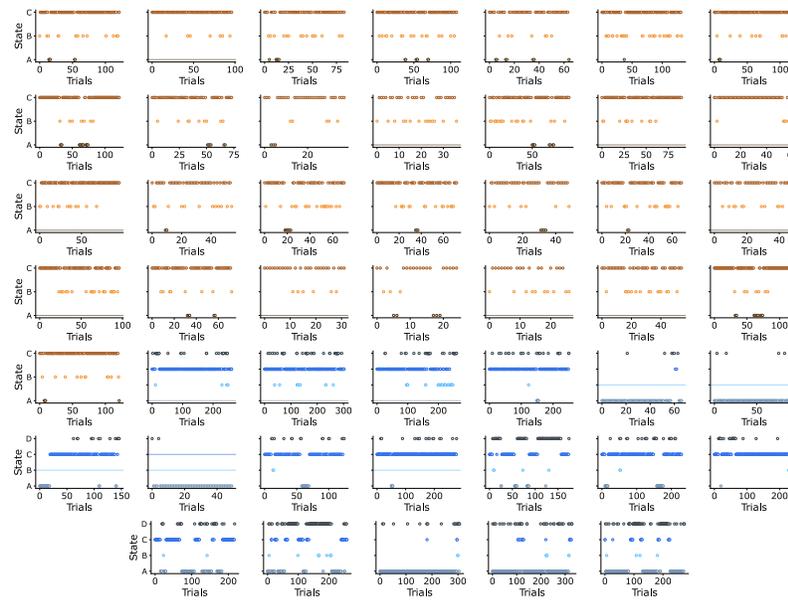


Fig. S15. RTs over states for all sessions, for the held out test set. For mice, distributions are shown in orange; for macaques, distributions are shown in green-blue.



**Fig. S16. Most likely states** for all sessions, for the held out test set. For mice, states are shown in orange; for macaques, states are shown in green-blue.