

UCLA

UCLA Previously Published Works

Title

From innate to instructed: A new look at perceptual decision-making

Permalink

<https://escholarship.org/uc/item/29p6n6rz>

Authors

Oesch, Lukas T

Ryan, Michael B

Churchland, Anne K

Publication Date

2024-06-01

DOI

10.1016/j.conb.2024.102871

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

From Innate to Instructed: A New Look at Perceptual Decision-Making

Lukas T. Oesch¹, Michael B. Ryan¹, Anne K. Churchland^{1,*}

1) Department of Neurobiology, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, United States.

* Corresponding author: Churchland, Anne K. (achurchland@mednet.ucla.edu)

ORCID:

LTO: 0000-0003-2819-812X, MBR: 0000-0002-2650-0690, AKC: 0000-0002-3205-3794

Highlights

- Perceptual decision-making can be assessed with stimuli elicit either spontaneous or trained responses.
- Spontaneous behavior can be leveraged to study sensory processing and decision-making by leveraging pre-existing stimulus-response associations.
- Trained behaviors offer a way to investigate the acquisition and implementation of abstract decision rules.
- The discovery of features that lead to fast learning and generalizable task knowledge will greatly improve trained decision-making tasks.

Abstract

Understanding how sensory stimuli are processed to guide appropriate actions is a major challenge in neuroscience. To study perceptual decision-making in animals, researchers use tasks that either probe spontaneous responses to stimuli (often described as “naturalistic”) or instruct animals to associate stimuli with experimenter-defined responses. While spontaneous decisions rely on animals’ pre-existing knowledge, the versatility of instructed tasks often comes at the cost of extensive training. Here, we review emerging approaches to investigate perceptual decision-making using both spontaneous and trained behaviors, discussing their strengths and limitations. Additionally, we propose that instructed decision-making tasks could be improved by seeking to facilitate faster learning and a more generalizable understanding of task rules.

Introduction

Animals interact in complex ways with their environment when making perceptual decisions. They integrate sensory signals from diverse sources, trade off competing needs, assess potential risks and benefits arising from their choices, and accordingly decide on the best action. In some situations, animals can make use of brain circuits that have formed during development to support automatic, and often very fast, decisions. Many researchers have emphasized that studying the neural correlates of such spontaneously expressed behaviors yields insights into how brains compute decisions under more naturalistic conditions, as opposed to many highly trained laboratory tasks [1]. Indeed, the study of spontaneous behavior is regaining popularity thanks to technological developments in behavioral tracking, image segmentation, and high yield neural recordings in freely behaving animals. Some have even argued that the recent success of this growing field suggests we should “step away” from more trained paradigms, such as training an animal to use a lever or a joystick [2].

However, although properly executing innate behaviors is a vital job of the animal’s brain, it is clearly not the only job. Many state-of-the-art experiments in perceptual decision-making take advantage of animals’ ability to learn complex task structures by training them to associate abstract sensory stimuli with specific actions to receive rewards to study for example abstract sensory comparisons [3•,4], evidence accumulation [5,6] or changes in the beliefs about the frequency of stimuli in an environment [7]. Although animals can learn these abstract tasks, they do so rather slowly and only after extensive repetition. Furthermore, abstract tasks suffer from a high inter-individual variability in learning speed and success. Finally, critics have raised concerns as to how general and flexible decision-making is in these highly repetitive tasks.

Here, we highlight how new approaches in both spontaneous and trained behavior have expanded our understanding of perceptual decision-making and we discuss the strengths and weaknesses of both. We argue that trained tasks may be improved if researchers not only consider desired stimulus parameters when designing a task but also take the learning process itself into account and identify task features that might lead to faster learning and a more general understanding of the instructed decision rules. Discovering such task features could allow for more intuitive learning that might act as a bridge between current spontaneous and trained behaviors.

The promises of spontaneous behavior in perception and decision-making

Over the past few years, neuroscience has witnessed key advances through the study of behavior to understand brain function, with many neuroscientists particularly highlighting the promise of natural behaviors in understanding fundamental computations [1,8,9]. However, the term “natural behavior” is oftentimes not well-defined – after all, the laboratory is a far cry from most animals’ ethological niche. We will therefore refer to “spontaneous” behaviors as those that do not require training and contrast them with “trained” behaviors. It is important to note that both spontaneous and trained behaviors can be studied in either restrained or freely behaving animals (see for instance [10]).

Two major technological advances have fueled the interest in spontaneous behavior. *First*, the miniaturization of sensors originally designed for human wearables has been used to design a broad variety of animal tracking devices. These devices include accelerometers and head-mounted cameras to sample the position of an animal’s head and both eyes [11–13], cameras to record the surrounding environment from an animal’s perspective [14••] or miniaturized

fluorescence microscopes to record neural activity in freely behaving animals [15–18]. *Second*, new machine learning tools and deep neural networks allow for unprecedented precision when tracking body parts in video data [19,20] and help researchers parse out distinct constituents of ongoing spontaneous movement [21–23].

The power of combining advanced sensors, machine learning tools, and electrophysiological recordings to study perception during spontaneous behavior was recently demonstrated. Parker and Abe recorded the neural activity in mouse primary visual cortex (V1), eye/head position and video of the mouse’s perspective [14••], **Fig. 1A**. The authors used the head and eye angles to shift the simultaneously acquired video and then reconstructed the visual input to the retina at any given moment. Using these inputs, the authors accurately measured receptive fields of V1 neurons in freely moving mice. Another important study using a similar approach demonstrated that the receptive fields of V1 neurons can be accurately recovered in freely gazing marmosets [24••]. Although not directly probing decision-making, these studies offer a way to precisely quantify visual input to an animal that might be used in future studies investigating visual decision-making in spontaneously behaving subjects.

Spontaneous behaviors also have begun to change the conventional wisdom regarding which sensory computations animals can perform. For example, the wide-set eyes and limited binocular regions of mouse V1 seem to suggest that mice largely forgo stereo vision. This assumption was tested in recent work in which mice climbed down a pole onto a transparent platform [25], **Fig. 1B**. Beneath each quadrant of the platform, a checkerboard pattern was presented at different depths. Mice with both eyes open tended to climb down to the quadrant where the checkerboard was closest to the platform. That is, mice selected the platform where they would not fall when leaving the pole. However, mice with one eye closed mostly failed to find the closest quadrant. Other examples leveraging spontaneous behaviors to study perception include continuous visual tracking of a focus of expansion in a visual flow-field in human and non-human primates [26], pursuit of prey [11,12,27] or prey-like stimuli [28], jumping over a gap [29], acoustic startle responses [30], and visual or auditory looming assays [10,31–33].




Figure 1. (A) Receptive field mapping in freely behaving mice. Schematic of the recording setup to track eye and head movements and subject-perspective visual scenery while recording from V1 neurons (left). Equipped with these recording devices, researchers can expose mice to an open field with a variety of different visual stimuli and reconstruct incoming visual information to probe V1 neural responses and receptive fields (right). Adapted from [14••]. (B) Illustration of the pole descent task assessing depth-perception in spontaneously behaving mice. Adapted from [25].

Although the examples above illustrate how spontaneous behavior can be leveraged to investigate perception and decision-making, it still suffers from a few limitations. *First*, one of the major behavioral achievements of animals is their ability to not only respond spontaneously, but also to learn new responses and associations. When focusing exclusively on spontaneous and untrained behaviors, researchers miss the chance to study a large part of an animal's behavioral capacity. *Second*, testing key hypotheses about perception and decision-making, such as changes in perceptual sensitivity for instance, requires dense sampling around the psychophysical threshold, where stimulus intensity is typically weak. Although the use of weak stimuli is technically possible for spontaneous behaviors, the tools to extract the critical parameters from these behaviors are still emerging [34]. *Third*, in many cases, a particular spontaneous behavior can only be elicited a few times for each subject before the animal stops responding due to satiation or the extinction of the behavioral response [17]. This imposes severe restrictions on the analysis of, for example, neural recordings obtained alongside behavioral data because only a few responses can be collected for each condition [35].

Trained behaviors are useful to study the discovery and behavioral implementation of defined decision rules

In contrast to spontaneous behaviors, in which one is limited to the innate (or at least extant) stimulus-response associations, behavioral training can drive entirely novel stimulus-response associations. For instance, experimenters can teach stimulus-response associations that probe specific aspects of perception and even alter those associations as needed [36–39]. Experimenters can introduce delays between the sampling of information and the reporting of a choice to disentangle the formation of evidence-driven decisions from the motor preparation and execution of the reporting action. Thanks to the freedom in designing trained tasks, it is also often possible to modify small task parameters in order to translate findings to another animal or even a human [5,40].

Critically, trained behaviors offer far more than the convenience of psychophysical tasks; trained behaviors also allow researchers to study learning and the discovery of abstract decision rules. Abstract rules are outside the realm of spontaneous behavior because such rules require the animal to make connections between two previously non-associated stimuli. For instance, it was recently demonstrated that mice can learn abstract decision rules based on the relative differences in visual stimuli by presenting pairs of visual gratings with different orientations [3••]. Mice then had to compare the orientation of the two gratings and select the one with the greater vertical orientation. Using this approach, the authors showed that mice can generally perceive relative angular differences of as little as 9° . Together, this work illustrates that mice can not only perform tasks requiring high visual acuity but that they are even capable of making perceptual judgments with respect to an abstract rule.

Furthermore, recent work has led to a rapid expansion of our understanding of how internal states influence perceptual decision-making. For instance, Ashwood and colleagues used a probabilistic modeling approach to determine how different states of engagement altered an animal's decision-making strategy [41]. The authors show that task engagement can rapidly fluctuate on short timescales (tens of trials) leading to markedly different task performance and weighing of sensory evidence. For example, when mice are in a state of high task engagement, their decisions are stimulus-driven and more accurate. In contrast, during states of lower task engagement, mice often employ stimulus-independent strategies, such as side biases, resulting in poor performance. This work also offers a potential explanation for previous findings regarding

lapses, or errors that occur even on the easiest stimuli. The study by Ashwood and colleagues suggests that lapses may be driven by periods of task disengagement where sensory-independent strategies dominate animal's decision-making (for an alternative account of lapses, also dependent on internal state, see [42]).

Performance on trained perceptual decision-making tasks requires animals to form new associations but open questions remain about how these associations are learned. A challenge here is the variability in learning rates across subjects. Reinert and colleagues highlight this variability by first training head-restrained mice to lick when presented with gratings with high spatial frequency but withhold licking when shown grating of low spatial frequency (e.g. “go/no-go task”), ignoring the orientation of the gratings [43••], **Fig. 2**. However, the authors note that mice showed highly variable learning rates, anywhere from ~10 sessions to almost 40 sessions to reach a defined expert performance criterion. The authors then made two clever adaptations to their task to probe the source of such variability in decision-making. *First*, they took mice trained on this spatial version of the task and instead had them classify stimuli based on their orientation (this time ignoring spatial frequency), without changing the reporting action. In this version of the task, all mice reached expert performance in less than 10 sessions. *Second*, the authors had the same subjects classify stimuli according to the same rule but respond by licking left or right to report their choice, as opposed to the previously learned go/no-go version of the task. In this variation, all animals took more than 20 sessions to perform at expert levels. This study not only highlights the variability in trained behaviors but hints that learning might occur on different timescales for acquiring stimulus-action associations given specific stimulus-action combinations and prior knowledge about the task structure.

These and other related findings about the slow and highly variable learning of trained tasks [44•] have compelled critics to question whether behavioral neuroscience should focus on these highly trained behaviors to study perception and decision-making [45]. One of the main concerns is that the extended training time might allow for plasticity mechanisms in sensory systems such that the neural activity of expert animals cannot be compared to a naïve subject. Another drawback is that perceptual tasks are often only learned through extensive repetition, possibly leading to the formation of a rigid stimulus-action association that prevents animals from generalizing their knowledge to other contexts. The findings by Reinert and colleagues that animals failed to quickly transfer the learned stimulus categories from a go/no-go to a two-alternative forced choice situation point to such rigid associations [43••]. Finally, because not only the stimulus contingencies but also the reporting actions are learned through reinforcement, it can be unclear when the perceptual rule is acquired versus behaviorally implemented [46].




Figure 2. A single set of stimuli can be associated with different decision rules. Mice are trained to distinguish visual patches with high- versus low-frequency gratings or different grating

orientations (left). It is possible to select a stimulus set such that experimenters can flip the rule from a spatial frequency rule (green separating line) to an orthogonal orientation rule (orange line) to study sensory learning (right). Adapted from [43••].

It thus becomes clear that trained tasks are often ill-suited for studying learning that is quick or that can flexibly extend to new contexts. Could researchers learn from tasks that animals learn quickly and with no explicit training to improve current psychophysical tasks? Alternatively, are there learned tasks that better balance an experimenter's need for control with an animal's innate ability to learn general principles of a task?

Identifying task features that invite fast learning and a generalizable understanding of perceptual decision-making tasks

The long training times with high subject-to-subject variability and rigid stimulus-response associations might seem at first an inevitable consequence of psychophysical tasks. After waiting weeks for the gradual learning that is typical of some tasks, one might be tempted to wonder whether rapid insight and a general understanding of perceptual decision rules are simply absent in animals. Here, we argue instead that innovations in the design and the training pipeline of instructed tasks may improve the learning speed and generality of instructed behaviors.

A first piece of evidence that animals can quickly learn general rules about their environment comes from navigation. Rodents tend to have well-developed internal maps, presumably because they must navigate circuitous and serpentine routes to acquire food and then return to the safety of their nest, avoiding predators all the while. Laboratory mazes, while not at all ethological, tap into this ability to learn a representation of paths through an environment. For instance, recent work demonstrated that mice in mazes learn many bits of information in far less time than a typical 2-choice laboratory task [47]. Further, observations of "sudden insight" suggest that rodents sometimes gain an understanding of the principles by which the system (in this case, the maze) operates [48].

Foraging is a second example of fast learning of the structure of a particular environment. Foraging decisions seek to maximize the payoff of actions based on an understanding of the reward structure of the environment. The computations needed to inform such decisions are non-trivial: one must estimate and compare the value of multiple uncertain options. Sophisticated foraging behavior is observed in many animals [49–51]. Features of foraging can be replicated in a laboratory setting by, for instance, manipulating the size and probability of rewards associated with particular decisions. Rats, mice and monkeys are often keenly sensitive to reward manipulations, even when head-fixed [42,44•,52,53]. Animals' ability to track foraging parameters with ease stands in contrast to animal's ability to learn other parameters of trained tasks, like abstract associations between stimulus and reward. Such associations can take months to learn and are often easily disrupted [42].

Maze-learning and foraging, though interesting behaviors in and of themselves, are of limited usefulness for perceptual decision-making studies, which often rely on systematically varying carefully parameterized stimuli. Still, they do serve to demonstrate that animals are capable of fast, robust learning of the sort that would benefit sensory processing as well. What features of those tasks nurture fast learning? Might we, as a field, identify these features and then port them to other behaviors? Some pertinent examples hint that we might.

Head-restrained, visual psychophysics tasks were once thought to be beyond the capabilities of mice. One modification that improves learning of such tasks substantially is the coupling between a subject's actions and stimulus motion [44•,54], **Fig. 3A**. The modification was to provide mice with a small steering wheel coupled to movement of the stimulus. Mice are rewarded for detecting the stimulus on the right or left side of a screen and moving it to the center using the wheel. Using this task and a detailed training protocol for it, many laboratories have reported high training success rates and relatively fast task learning [44•].

Researchers might also tackle the high inter-individual variability of learning trajectories and speeds by designing tasks that can be acquired in multiple steps [55•] or through the combination of sub-tasks [56••]. Recent work shows that mice can combine knowledge that was independently acquired for two different tasks to solve a new composite task [56••]. In this study head-restrained mice separately learned to move an illuminated spout holder to center of a surface in front of them using a joystick or to simply lick when the LED on a spout was switched on. Despite rather slow learning and high inter-individual variability at the acquisition of the joystick subtask the mice were able to combine the knowledge of the two tasks to quickly learn to move the holder now equipped with the spout close to their snout and to lick for water rewards.

Why might coupling the movement of a stimulus to the turning of a wheel to report a choice lead to fast learning while, i.e., reporting a choice with a saccade or licking movement leads to slower learning? Is it possible that animals may better leverage “structural knowledge” about the connection between their actions and sensory stimuli when learning to turn the wheel than when licking to the right or left as might be suggested by the sub-task learning? Could there be a neural basis for this difference in learning? One possibility is that easily learnt tasks recruit neural activity that is within the subset of naturally occurring covariance patterns. This subset of patterns is called the intrinsic neural manifold (**Fig. 3B**). In two recent studies, macaques were trained to move a virtual cursor via a Brain Computer Interface (BCI) by producing a specific neural activity pattern in M1 [57,58•]. After the initial association, the authors changed the mapping between the neural activity and the cursor. Surprisingly, the monkeys learned certain new mappings very quickly while taking a very long time for other mappings. It turns out that the macaques can relatively easily re-associate neural activity patterns and thus learn fast if the new mapping falls within their intrinsic manifold [57]. In contrast, when the new mapping is outside the animal's intrinsic manifold, the monkey has to learn how to produce the new activity patterns leading to a slower increase in their performance [58•]. What neural covariance patterns are readily available to the animal may be largely shaped by evolutionary constraints [59].




Figure 3. (A) Mice rotate a steering wheel that causes simultaneous movement of a visual stimulus. Mice learn with experience that centering the stimulus (left) will yield a drop of water as a reward, while sending the stimulus in the opposite direction (right) will not. Adapted from [44]. (B) Conceptual schematic of neural activity patterns. Each point shows the joint firing rate (FR) of 3 neurons. These points are largely confined to the intrinsic manifold (gray plane). The black arrow indicates an “intuitive” mapping between the monkeys’ neural activity and the position of a BCI cursor. Green points: activity patterns that would move the cursor to the right; purple points: activity patterns that would move the cursor to the left. The blue arrow indicates a new mapping. This mapping requires “off manifold activity” (activity patterns not on the intrinsic manifold) These off-manifold perturbations were much slower for the monkey to learn. Adapted from [58].

Currently, our ability as experimenters to design tasks that recruit within- vs. outside-manifold activity patterns can feel a bit like grasping in the dark. It will be critical to obtain neural recordings during task learning to uncover which tasks generate entirely new activity patterns and which ones repurpose existing patterns. In-depth analysis of animal movements might likewise assist the discovery of better training regimes for instructed behaviors, since animal movement patterns might indicate when animals’ engagement waxes and wanes [60,61]. Finally, we call on researchers to share detailed reports about all the used training techniques and interventions as well as training times including summaries of animals excluded from further training. Sharing more nuanced training procedures will be transformative for instructed behaviors.

Concluding remarks

Here, we have reviewed recent innovations and novel approaches to studying perceptual decision-making using animals’ spontaneous behaviors or instructed tasks, both of which have contributed invaluable new insights into how animals use sensory information to guide their choices. We have argued that trained behaviors could greatly be improved if researchers more deeply considered an animal’s learning process and would select task features that invite fast learning of generalizable knowledge for their tasks. To achieve this goal, it is key to investigate animal learning on the behavioral and neural level and to provide detailed accounts of training procedures and task learning progress even if learning is not the focus of the study. Although for the purpose of the review we have mainly focused on visual decision-making in mice, we hypothesize that incorporating learning considerations into task design may be transformative for other decision tasks and animal species.

Conflict of interest

Nothing declared.

Acknowledgements

We thank João Couto, Xiaonan Richard Sun, Gabriel Rojas-Bowe, Eliah Djedjro, and all the members of the Churchland lab for their advice and feedback on early versions of the manuscript. This work was supported by the Swiss National Science Foundation P2BEP3 (LTO), the AP Giannini Foundation (MBR), National Science Foundation NCS-FO2219876, U19NS1237 and the Simons Collaboration on the Global Brain (AKC).

References

1. Miller CT, Gire D, Hoke K, Huk AC, Kelley D, Leopold DA, Smear MC, Theunissen F, Yartsev M, Niell CM: **Natural behavior is the language of the brain.** *Curr Biol* 2022, **32**:R482–R493.
2. Smith K: **Neuroscience goes wild.** *Nature* 2023, **618**:448–450.
3. ••Lyamzin DR, Aoki R, Abdolrahmani M, Benucci A: **Probabilistic discrimination of relative stimulus features in mice.** *Proc Natl Acad Sci U S A* 2021, **118**:1–11.

In this work the authors introduce a truly abstract visual orientation discrimination task that requires animals to compare stimuli presented on their left and the right side to a non-displayed reference orientation.

4. Zhou Y, Mohan K, Freedman DJ: **Abstract Encoding of Categorical Decisions in Medial Superior Temporal and Lateral Intraparietal Cortices.** *J Neurosci* 2022, **42**:9069–9081.
5. Raposo D, Sheppard JP, Schrater PR, Churchland AK: **Multisensory decision-making in rats and humans.** *J Neurosci* 2012, **32**:3726–3735.
6. Pinto L, Tank DW, Brody CD: **Multiple timescales of sensory-evidence accumulation across the dorsal cortex.** *Elife* 2022, **11**:1–21.
7. Pedrosa V, Menichini E, Pajot-moric Q, Vincent P, Teachen L, Zhou L, Latham P, Akrami A: **Humans, rats and mice show species-specific adaptations to sensory statistics in categorisation behaviour.** *bioRxiv* 2023,
8. Krakauer JW, Ghazanfar AA, Gomez-Marin A, Maclver MA, Poeppel D: **Neuroscience Needs Behavior: Correcting a Reductionist Bias.** *Neuron* 2017, **93**:480–490.
9. Datta SR, Anderson DJ, Branson K, Perona P, Leifer A: **Computational Neuroethology: A Call to Action.** *Neuron* 2019, **104**:11–24.
10. Fink AJ, Axel R, Schoonover CE: **A virtual burrow assay for head-fixed mice measures habituation, discrimination, exploration and avoidance without training.** *Elife* 2019, **8**:1–21.

11. Michael AM, Abe ETT, Niell CM: **Dynamics of gaze control during prey capture in freely moving mice.** *Elife* 2020, **9**:1–27.
12. Holmgren CD, Stahr P, Wallace DJ, Voit KM, Matheson EJ, Sawinski J, Bassetto G, Kerr JND: **Visual pursuit behavior in mice maintains the pursued prey on the retinal region with least optic flow.** *Elife* 2021, **10**:1–34.
13. Meyer AF, Poort J, O’Keefe J, Sahani M, Linden JF: **A Head-Mounted Camera System Integrates Detailed Behavioral Monitoring with Multichannel Electrophysiology in Freely Moving Mice.** *Neuron* 2018, **100**:46-60.e7.
14. ••Parker PRL, Abe ETT, Leonard ESP, Martins DM, Niell CM: **Joint coding of visual input and eye/head position in V1 of freely moving mice.** *Neuron* 2022, **110**:3897-3906.e5.

This paper presents an elegant approach to recovering visual information available to freely behaving mice by taking advantage of combined recordings of head/eye position, visual scenery videos and neural data.

15. Ghosh KK, Burns LD, Cocker ED, Nimmerjahn A, Ziv Y, Gamal A El, Schnitzer MJ: **Miniaturized integration of a fluorescence microscope.** *Nat Methods* 2011, **8**:871–878.
16. Cai DJ, Aharoni D, Shuman T, Shobe J, Biane J, Song W, Wei B, Veshkini M, La-Vu M, Lou J, et al.: **A shared neural ensemble links distinct contextual memories encoded close in time.** *Nature* 2016, **534**:115–118.
17. Juavinett A, Bekheet G, Churchland A: **Chronically-implanted Neuropixels probes enable high yield recordings in freely moving mice.** *Elife* 2019, **8**:1–17.
18. Luo TZ, Bondy AG, Gupta D, Elliott VA, Kopec CD, Brody CD: **An approach for long-term, multi-probe neuropixels recordings in unrestrained rats.** *Elife* 2020, **9**:1–54.
19. Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, Bethge M: **DeepLabCut: markerless pose estimation of user-defined body parts with deep learning.** *Nat Neurosci* 2018, **21**:1281–1289.
20. Pereira TD, Aldarondo DE, Willmore L, Kislin M, Wang SSH, Murthy M, Shaeviz JW: **Fast animal pose estimation using deep neural networks.** *Nat Methods* 2019, **16**:117–125.
21. Wiltschko AB, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abreira VE, Adams RP, Datta SR: **Mapping Sub-Second Structure in Mouse Behavior.** *Neuron* 2015, **88**:1121–1135.
22. Yi D, Musall S, Churchland A, Padilla-Coreano N, Saxena S: **Disentangled multi-subject and social behavioral representations through a constrained subspace variational autoencoder (CS-VAE).** *bioRxiv* 2022,
23. Biderman D, Whiteway MR, Hurwitz C, Greenspan N, Lee RS, Vishnubhotla A, Warren R, Pedraja F, Noone D, Schartner M, et al.: **Lightning Pose: improved animal pose estimation via semi-supervised learning, Bayesian ensembling, and cloud-native open-source tools.** *bioRxiv* 2023,

24. ••Yates JL, Coop SH, Sarch GH, Wu RJ, Butts DA, Rucci M, Mitchell JF: **Detailed characterization of neural selectivity in free viewing primates.** *Nat Commun* 2023, **14**.

In this important work the authors describe and validate a method to efficiently assess visual tuning properties of neurons in freely gazing marmosets using high-resolution eye tracking and neural recordings.

25. Boone HC, Samonds JM, Crouse EC, Barr C, Priebe NJ, McGee AW: **Natural binocular depth discrimination behavior in mice explained by visual cortical activity.** *Curr Biol* 2021, **31**:2191-2198.e3.
26. Knöll J, Pillow JW, Huk AC: **Lawful tracking of visual motion in humans, macaques, and marmosets in a naturalistic, continuous, and untrained behavioral context.** *Proc Natl Acad Sci U S A* 2018, **115**:E10486–E10494.
27. Hoy JL, Yavorska I, Wehr M, Niell CM: **Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice.** *Curr Biol* 2016, **26**:3046–3052.
28. Procacci NM, Allen KM, Robb GE, Ijekah R, Lynam H, Hoy JL: **Context-dependent modulation of natural approach behaviour in mice.** *Proc R Soc B Biol Sci* 2020, **287**:1–10.
29. Parker PRL, Abe ETT, Beatie NT, Leonard ESP, Martins DM, Sharp SL, Wyrick DG, Mazzucato L, Niell CM: **Distance estimation from monocular cues in an ethological visuomotor task.** *Elife* 2022, **11**:1–18.
30. Guo W, Fan S, Xiao D, Dong H, Xu G, Wan Z, Ma Y, Wang Z, Xue T, Zhou Y, et al.: **A Brainstem reticulotegmental neural ensemble drives acoustic startle reflexes.** *Nat Commun* 2021, **12**:1–12.
31. Yilmaz M, Meister M: **Rapid innate defensive responses of mice to looming visual stimuli.** *Curr Biol* 2013, **23**:2011–2015.
32. De Franceschi G, Vivattanasarn T, Saleem AB, Solomon SG: **Vision Guides Selection of Freeze or Flight Defense Strategies in Mice.** *Curr Biol* 2016, **26**:2150–2154.
33. Salay LD, Ishiko N, Huberman AD: **A midline thalamic circuit determines reactions to visual threat.** *Nature* 2018, **557**:183–189.
34. Huk A, Bonnen K, He BJ: **Beyond trial-based paradigms: Continuous behavior, ongoing neural activity, and natural stimuli.** *J Neurosci* 2018, **38**:7551–7558.
35. Urai AE, Doiron B, Leifer AM, Churchland AK: **Large-scale neural recordings call for new insights to link brain and behavior.** 2021, **25**:11–19.
36. Freedman DJ, Assad JA: **Experience-dependent representation of visual categories in parietal cortex.** *Nature* 2006, **443**:85–88.
37. Mante V, Sussillo D, Shenoy K V., Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex.** *Nature* 2013, **503**:78–84.

38. Aronov D, Nevers R, Tank DW: **Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit.** *Nature* 2017, **543**:719–722.
39. Spellman T, Svei M, Kaminsky J, Manzano-Nieves G, Liston C: **Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring.** *Cell* 2021, **184**:2750-2766.e17.
40. Brunton BW, Botvinick MM, Brody CD: **Rats and humans can optimally accumulate evidence for decision-making.** *Science (80-)* 2013, **340**:95–98.
41. Ashwood ZC, Roy NA, Stone IR, Urai AE, Churchland AK, Pouget A, Pillow JW: **Mice alternate between discrete strategies during perceptual decision-making.** *Nat Neurosci* 2022, **25**:201–212.
42. Pisupati S, Chartarifsky-Lynn L, Khanal A, Churchland AK: **Lapses in perceptual decisions reflect exploration.** *Elife* 2021, **10**:1–27.
43. •Reinert S, Hübener M, Bonhoeffer T, Goltstein PM: **Mouse prefrontal cortex represents learned rules for categorization.** *Nature* 2021, **593**:411–417.

The authors show that after highly variable initial learning rates in a visual discrimination task, expert mice quickly learn a new rule associated with the stimuli but only slowly learn a new reporting action. The researchers further report that neurons in the anterior cingulate cortex first acquire selectivity for the “go” and only later for the “no-go” category during learning.

44. •The International Brain Laboratory, Aguillon-Rodriguez V, Angelaki D, Bayer H, Bonacchi N, Carandini M, Cazettes F, Chapuis G, Churchland AK, Dan Y, et al.: **Standardized and reproducible measurement of decision-making in mice.** *Elife* 2021, **10**:1–28.

In a collaborative effort and through sharing open-source code and detailed methodological details the authors reproduced mouse visual decision-making behavior in seven distinct institutions.

45. Meister M: **Learning, fast and slow.** *Curr Opin Neurobiol* 2022, **75**:102555.
46. Kuchibhotla K V., Hindmarsh Sten T, Papadoyannis ES, Elnozahy S, Fogelson KA, Kumar R, Boubenec Y, Holland PC, Ostojic S, Froemke RC: **Dissociating task acquisition from expression during learning reveals latent knowledge.** *Nat Commun* 2019, **10**:2151.
47. Rosenberg M, Zhang T, Perona P, Meister M: **Mice in a labyrinth: Rapid learning, sudden insight, and efficient exploration.** *Elife* 2021, **10**:1–30.
48. Krausz TA, Comrie AE, Kahn AE, Frank LM, Daw ND, Berke JD: **Dual credit assignment processes underlie dopamine signals in a complex spatial environment.** *Neuron* 2023, **111**:1–14.
49. Stephens DW, Krebs JR: *Foraging Theory.* Princeton University Press; 2019.
50. Kennedy A: **The what, how, and why of naturalistic behavior.** *Curr Opin Neurobiol* 2022, **74**:102549.

51. Mobbs D, Trimmer PC, Blumstein DT, Dayan P: **Foraging for foundations in decision neuroscience: Insights from ethology.** *Nat Rev Neurosci* 2018, **19**:419–427.
52. Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN: **Elapsed decision time affects the weighting of prior probability in a perceptual decision task.** *J Neurosci* 2011, **31**:6339–6352.
53. Lottem E, Banerjee D, Verтеchi P, Sarra D, Lohuis MO, Mainen ZF: **Activation of serotonin neurons promotes active persistence in a probabilistic foraging task.** *Nat Commun* 2018, **9**:1–12.
54. Burgess CP, Lak A, Steinmetz NA, Zátka-Haas P, Bai Reddy C, Jacobs EAK, Linden JF, Paton JJ, Ranson A, Schröder S, et al.: **High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice.** *Cell Rep* 2017, **20**:2513–2524.
55. •Noel JP, Caziot B, Bruni S, Fitzgerald NE, Avila E, Angelaki DE: **Supporting generalization in non-human primate behavior by tapping into structural knowledge: Examples from sensorimotor mappings, inference, and decision-making.** *Prog Neurobiol* 2021, **201**:101996.

This work provides a critical look at prevailing trained tasks and introduces a set of tasks for continuous decision-making in non-human primates.

56. ••Makino H: **Arithmetic value representation for hierarchical behavior composition.** *Nat Neurosci* 2023, **26**:140–149.

This study shows that mice and artificial reinforcement learning agents can make use of knowledge from previously acquired sub-tasks to quickly learn a new composite task. The author also provides an account on how neural representations of sub-task may be combined during the learning of the full task.

57. Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu BM: **Learning by neural reassociation.** *Nat Neurosci* 2018, **21**:607–616.
58. •Oby ER, Golub MD, Hennig JA, Degenhart AD, Tyler-Kabara EC, Yu BM, Chase SM, Batista AP: **New neural activity patterns emerge with long-term learning.** *Proc Natl Acad Sci U S A* 2019, **116**:15210–15215.

This study finds that, in contrast to intuitive mappings between neural activity patterns and cursor motion in a brain computer interface, learning non-intuitive mappings requires changing the neural correlation structure and takes multiple days.

59. Safaie M, Chang JC, Park J, Miller LE, Dudman JT, Perich MG, Gallego JA: **Preserved neural population dynamics across animals performing similar behaviour.** *bioRxiv* 2022,
60. Hulsey D, Zumwalt K, Mazzucato L, McCormick DA, Jaramillo S: **Decision-making dynamics are predicted by arousal and uninstructed movements.** *bioRxiv* 2023,
61. Yin C, Melin MD, Rojas-Bowe G, Sun XR, Gluf S, Couto J, Kostiuk A, Musall S, Churchland AK: **Engaged decision-makers align spontaneous movements to stereotyped task**

demands. *bioRxiv* 2023,