

Population-level analysis of neural activity during flexible learning of goal-directed actions

Background

A core function of the brain is to transform sensory information into useful behavior, while allowing for some degree of flexibility to adapt to an ever-changing environment. However, the precise mechanisms by which sensory representations are translated into actions, and the factors that enable this flexibility, remain poorly understood.

In the broader project, we are interested in how two different sensory stimuli that ultimately lead to the same action are represented in the dorsal striatum (dStr), a brain region implicated in the learning of goal-directed actions and in action selection, and how these representations change during acquisition and reversal. In this student project, you will work on one specific subquestion in depth and gain hands-on experience in analysing neural data at the population level.




Figure 1: **Experimental design.** (A) Task 1: shuttling during CS1 avoids the shock (goal-directed); CS2 has no effect. Task 2 reverses the CS meanings. (B) Behavioral results: mice only shuttle to the task-relevant tone. (C) dStr neural activity (GCaMP6m-expressing neurons) can be recorded with a miniscope.

The experiment and dataset

In a two-tone active avoidance paradigm (Fig. 1), mice learn to avoid a footshock by crossing to the other compartment of the cage (a “shuttle”) in response to an auditory cue. Using head-mounted miniscopes, we record hundreds of neurons per mouse in the dStr, allowing us to follow individual neurons across the whole experiment.

Mice first undergo a habituation session (day 1) where they get used to both tones (later conditioned stimuli, CS). In task 1 (days 2–4), CS1 predicts an upcoming footshock and mice learn to avoid it (CS2 has no behavioral relevance yet). In task 2 (days 5–7), CS1 loses its behavioral meaning while CS2 now predicts an upcoming footshock (contingency reversal). The last two days of the experiment (days 8–9) are extinction sessions, where both tones do not have any behavioral meaning (as during habituation).

The data is already recorded, preprocessed and ready for analysis. It consists of cross-session aligned neural data, behavioral data (DeepLabCut) as well as experimental parameters such as tone onsets/offsets and applied footshocks.

Your project

You will train a linear decoder (SVM) on avoid vs. fail trials using CS1 trials only, then test it on CS2 trials. A transferring decoder would argue for an abstract action code. Since accuracy alone does not paint the full picture, you will also analyse the weight vectors and try to identify the encoding subpopulation and investigate its properties. Further, we are also interested in identifying possible action planes, i.e. predictive activity for an upcoming shuttle, similar to what was described in Ehret et al. (2024)¹.

If time allows, the developed methods can be applied to recordings from the prelimbic cortex of mice performing the same behavioral paradigm. This will allow us to compare the two brain regions in terms of encoding of sensory cues, action plans and action.

You will learn: population-level analysis of neural data, dimensionality reduction, cross-condition decoding, and alignment of neural data across subjects.

Requirements: Python, familiarity with linear algebra and basic statistics. Prior experience with neural data and/or machine learning is a plus.

Format: Semester project or Master’s thesis, duration 3–6 months, supervised by Philipp Eugster in the [Grewe lab](#) at the [Institute of Neuroinformatics](#), UZH/ETH Zürich.

Contact

Interested? Contact Philipp Eugster, eugsteph@ethz.ch.

¹B. Ehret, R. Boehringer, E. A. Amadei et al. Population-level coding of avoidance learning in medial prefrontal cortex. *Nat Neurosci* **27**, 1805–1815 (2024). [doi:10.1038/s41593-024-01704-5](https://doi.org/10.1038/s41593-024-01704-5)