
Automated Animal Training and Iterative Inference of Latent Learning Policy

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Progress in understanding how individual animals learn requires high-throughput
2 standardized methods for behavioral training and ways of adapting training. During
3 the course of training with hundreds or thousands of trials, an animal may change
4 its underlying strategy abruptly, and capturing these changes requires real-time
5 inference of the animal’s latent decision-making strategy. To address this challenge,
6 we have developed an integrated platform for automated animal training, and an
7 iterative decision-inference model that is able to infer the momentary decision-
8 making policy, and predict the animal’s choice on each trial with an accuracy of
9 80%, even when the animal is performing poorly. We also combined decision
10 predictions at single-trial resolution with automated pose estimation to assess
11 movement trajectories. Analysis of these features revealed categories of movement
12 trajectories that associate with decision confidence.

13 1 Introduction

14 Learning – the change of neural representation and behavior that results from past experience and
15 the consequences of actions – is important for animals’ survival and is a central topic in biological
16 intelligence [1]. Different individuals may apply different strategies to the learning process, reflecting
17 their individual biases [2–4]. Indeed, substantial differences in sensory biases, locomotion, motivation,
18 and cognitive competence have been observed in populations of fruit flies, rodents and primates.
19 Thus, it is critical to investigate learning at the individual level.

20 Rodents, especially the mouse, have become popular experimental animals for studying associative
21 learning and decision-making, due to the widely available transgenic resources [5]. They can learn to
22 perform complex decision-making tasks that probe cognitive components such as working memory
23 and selective attention [6–8]. However, differences in learning strategies across individuals have
24 rarely been addressed, partly owing to the limitations of data gathering and methods for data analysis.

25 Studying differences among individuals requires training and collecting data from multiple animals
26 in a standardized and high-throughput fashion. The training procedures are often time-consuming,
27 requiring several days to many weeks, depending on the task [9, 10]. Human intervention in this
28 process causes additional variability, therefore an automated animal training procedure is preferential
29 for the purpose. In addition, analysis approaches tend to focus on the averaged performance over
30 many trials, while changes in behavior may deviate at any trial, and thus necessitating modeling
31 behavior at the single-trial level.

32 To address these challenges, we developed an integrated platform for automated-training of group-
33 housed mice and analysis of behavioral changes in learning a decision-making task. We designed
34 hardware that makes use of implanted radio frequency identification (RFID) chips to identify each
35 mouse, and guides the animal into a behavior training box. Synchronized video recordings and

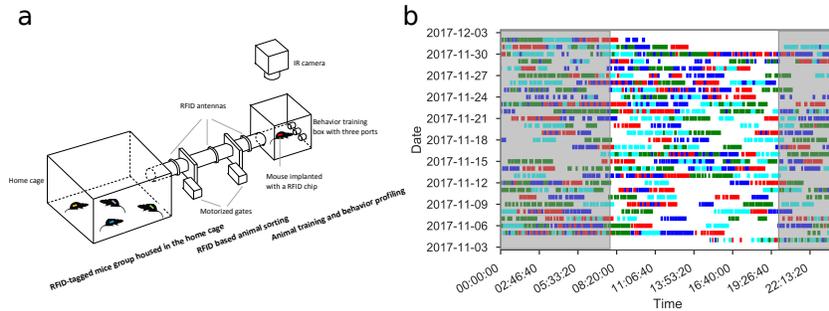


Figure 1: Automated training system and its performance. (a) Schematic of the automated training system. (b) Activity trace of a cohort of five animal in the behavior box, colorcoded by animal.

36 decision-making sequences were acquired during animal learning. To analyze the decision-making
 37 sequences, we developed an iterative generalized linear model, which predicts the animal’s choice in
 38 each trial, and is updated immediately based on the animal’s actual choice. To analyze the animal’s
 39 behavior during the task in greater detail, we extracted the animal’s pose from video data. We built an
 40 model on these extracted features to predict the animal’s decisions. Interestingly, we found pre-stimuli
 41 behavior were predictive of the correctness of the animal’s choice.

42 2 Automating Tasks and Animal Identification

43 We designed our automated training system as follows (Fig. 1a): RFID-tagged mice are grouped in
 44 a common home cage where food and bedding is supplied. The home cage connects to a behavior
 45 training box through a gated tunnel. The gates are controlled by a RFID animal sorting system. The
 46 behavior box is outfitted with three ports, each of which contains a photo-transistor to detect snout
 47 entry, a solenoid valve to deliver water reward, and a light emitting diode (LED) to present visual
 48 cues. The entire apparatus is orchestrated by a master program that coordinates the RFID sorting
 49 device, the behavioral box control system, synchronized video recording, data management and
 50 logging. We tested the automated RFID sorting and animal training system by training group-housed
 51 mice to learn a variety of decision tasks. The training period lasted 28 days, with up to five mice in
 52 the common home cage. Each animal occupied the training box for 3-4 hours per day throughout the
 53 entire training period (Fig. 1b). Each animal was trained for over 900 trials (10 sessions)

54 3 Online Latent Strategy Inference

55 We developed a generalized linear model (GLM) to map decision policies relevant to the animal’s
 56 decision-making to its choices through logistic regression. Because a change of policy can happen at
 57 any trial during learning, we developed the model to make online trial-by-trial choice predictions
 58 based on various strategies the animal might plausibly use. The model works in an iterative two-step
 59 process. In the prediction step, the model makes a prediction for the next decision based on the
 60 input factors. Once the outcome of the animal’s decision is observed, an error term between the
 61 model’s prediction and the observation is computed. This error, after weighting by a reward factor
 62 and a temporal discount factor, is fed back to the loss function. In the update step, the model is
 63 updated by minimizing the regularized loss function. The temporal discount factor accounts for the
 64 possibility that the most recent trials impact the current decision more than remote trials. The reward
 65 factor accounts for the fact that water rewards and timeout punishments may have effects of different
 66 magnitude on the updates of the animal’s policy (Fig. 2).

67 We illustrate the utility of this model by fitting results from a binary visual task, in which one of the
 68 two side ports lights up to indicate the location of the reward, and the optimal policy is to simply
 69 poke the port with the light. The GLM makes a prediction for the outcome of each trial based
 70 on a weighted combination of several input variables: the current visual stimulus, a constant bias
 71 term, and three terms representing the history of previous trials. These inputs from a previous trial
 72 include the port choice, whether that choice was rewarded, and a term indicating the multiplicative
 73 interaction between the choice and reward (Choice x Reward). This term supports a strategy called

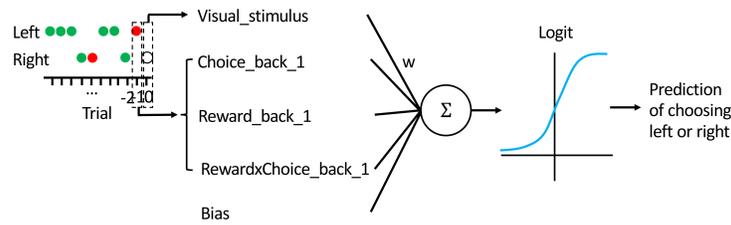


Figure 2: Illustration of the iterative inference model used to infer policies for the visual discrimination task. The model’s prediction is based on the output of a logistic function whose input is the weighted sum of a visual stimulus term, a bias term, and three history-dependent terms. The stimulus can be on the left or right and the choice can be rewarded (consistent with the stimulus, indicated by a green dot) or unrewarded (opposite to the stimulus, indicated by a red dot).

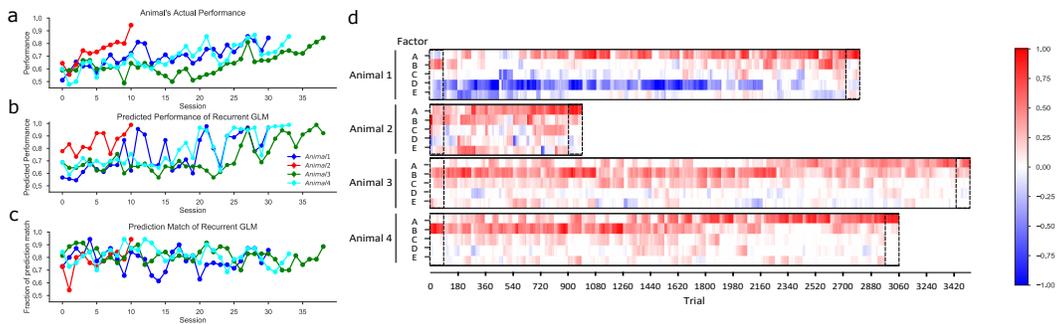


Figure 3: GLM performance. (a) The animal’s actual performance over time in the visual task. (b) Performance as predicted by the GLM. (c) Fraction of choices predicted correctly by the GLM. (d) Interpretation of policies during learning: Policy matrices recovered for the four animals show distinct individual learning processes. Dashed rectangles highlight the first and last sessions of each animal

74 win-stay-lose-switch (WSLS), which chooses the same port if it was rewarded previously and the
 75 opposite one if not. Predictions from the iterative GLM matched 80% of the animals’ actual choices
 76 (Fig. 3c), and the predicted accuracy of each animal captured the actual fluctuations of its learning
 77 curve (Figs. 3a and 3b). In addition, the iterative GLM serves to infer what policy the animal follows
 78 in making decisions, reflected by the weight of each input term. By following this weight vector
 79 across trials one obtains a policy matrix that documents how the animal’s policy changes during
 80 learning (Fig. 3d).

81 4 Behavior Analysis

82 While predicting choice by inferring the underlying policy is important for training animals on a
 83 timescale of several trials, it would be more beneficial to further predict choice through pre-trial
 84 behavior. When a subject is completing a trial, momentary lapses in concentration while information
 85 related to this trial is presented can adversely affect the decision in the next trial. To explore this
 86 hypothesis, we used DeepLabCut [11] to train a pose estimator on the body landmarks, including
 87 nose, centroid, and tail-base. We fitted the pre-stimuli behavioral pose of the animal to the animal’s
 88 choice accuracy through a boosting decision tree classifier, allowing us to decode whether the animal
 89 would make a correct decision. We show the k-fold validation accuracy for several animals (Fig.
 90 4a). We find that independent of the type of choice, we can predict the correctness with above 80%
 91 accuracy. Analyzing the behavioral trajectories, we find that behaviors that suggest hesitation are
 92 captured in this model (Fig. 4b-d).

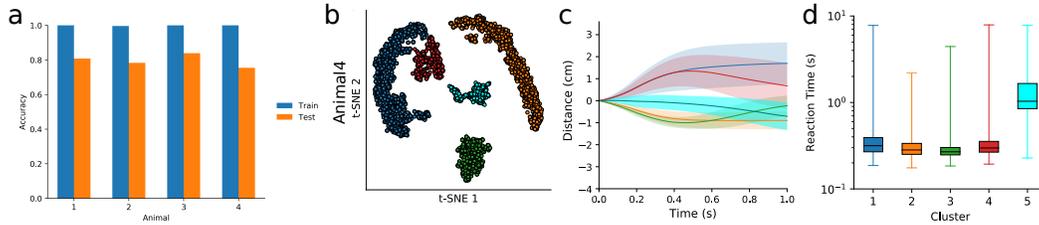


Figure 4: (a) Training and k-fold validation performance of boosting decision tree classifier trained on 6 seconds of pre-stimulus behavior pose data. (b) t-SNE of individual trial behaviors, grouped into 5 clusters. (c) Visualization of the behavior traces along the behavior-box indicate three distinct groups of movement patterns, including a cluster for hesitation. (d) Reaction plot of the 5 clusters illustrate the difference in reaction time.

93 5 Conclusion

94 We've presented a suite of hardware and machine learning tools for the automation of neuroscience
 95 experiments. Our method allows us to test multiple animals for thousands of trials without any
 96 human intervention. The policy inference model predicts an animal's choices with an accuracy of
 97 80%. And the choice prediction model using pre-stimuli behavioral trajectories informs whether the
 98 animal would make correct decisions in the future. Altogether, this method can be extended to allow
 99 adaptive training that would further reduce the training time to get to the desired behavior for neural
 100 recordings.

101 References

- 102 [1] Norman E Spear, James S Miller, and Joyce A Jagielo. Animal memory and learning. *Annual review of*
 103 *psychology*, 41(1):169–211, 1990.
- 104 [2] Kristin Branson, Alice A Robie, John Bender, Pietro Perona, and Michael H Dickinson. High-throughput
 105 ethomics in large groups of drosophila. *Nature methods*, 6(6):451, 2009.
- 106 [3] Kyle Honegger and Benjamin de Bivort. Stochasticity, individuality and behavior. *Current Biology*, 28(1):
 107 R8–R12, 2018.
- 108 [4] Louis D Matzel and Bruno Sauce. Individual differences: Case studies of rodent and primate intelligence.
 109 *Journal of Experimental Psychology: Animal Learning and Cognition*, 43(4):325, 2017.
- 110 [5] Liqun Luo, Edward M Callaway, and Karel Svoboda. Genetic dissection of neural circuits: a decade of
 111 progress. *Neuron*, 98(2):256–281, 2018.
- 112 [6] L Ian Schmitt, Ralf D Wimmer, Miho Nakajima, Michael Happ, Sima Mofakham, and Michael M Halassa.
 113 Thalamic amplification of cortical connectivity sustains attentional control. *Nature*, 545(7653):219, 2017.
- 114 [7] Ralf D Wimmer, L Ian Schmitt, Thomas J Davidson, Miho Nakajima, Karl Deisseroth, and Michael M
 115 Halassa. Thalamic control of sensory selection in divided attention. *Nature*, 526(7575):705, 2015.
- 116 [8] Athena Akrami, Charles D Kopec, Mathew E Diamond, and Carlos D Brody. Posterior parietal cortex
 117 represents sensory history and mediates its effects on behaviour. *Nature*, 554(7692):368, 2018.
- 118 [9] Matteo Carandini and Anne K Churchland. Probing perceptual decisions in rodents. *Nature neuroscience*,
 119 16(7):824, 2013.
- 120 [10] Alex Gomez-Marin, Joseph J Paton, Adam R Kampff, Rui M Costa, and Zachary F Mainen. Big behavioral
 121 data: psychology, ethology and the foundations of neuroscience. *Nature neuroscience*, 17(11):1455, 2014.
- 122 [11] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Wey-
 123 gandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts
 124 with deep learning. Technical report, Nature Publishing Group, 2018.