

**INVESTIGATING AUDITORY SPATIAL BEHAVIOR USING
CLOSED-LOOP MODULATION OF AUDITORY STIMULI AND
NEURAL ACTIVITY**

A Dissertation
Presented to
The Academic Faculty

by

Sukrith Sriram Vedapuram

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
Wallace H. Coulter Department of Biomedical Engineering

Georgia Institute of Technology and Emory University
May 2025

COPYRIGHT © 2025 BY SUKRITH SRIRAM VEDAPURAM

**INVESTIGATING AUDITORY SPATIAL BEHAVIOR USING
CLOSED-LOOP MODULATION OF AUDITORY STIMULI AND
NEURAL ACTIVITY**

Approved by:

Dr. Christopher Rodgers, Advisor
Department of Neurosurgery
School of Medicine
Emory University

Dr. Garrett Stanley
Department of Biomedical Engineering
College of Engineering
Georgia Institute of Technology

Dr. Alan Emanuel
Department of Cell Biology
School of Medicine
Emory University

Date Approved: April 25, 2025

ACKNOWLEDGEMENTS

I want to express my sincere gratitude to my thesis advisor, Dr. Christopher Rodgers, for his unwavering support, insightful guidance, and consistent encouragement throughout my time in the lab. His mentorship and the positive lab environment he fostered were instrumental in helping me carry my research forward and gain confidence with new techniques that I was unfamiliar with.

I would also like to thank the members of the Rodgers lab for helping with various aspects of the project. I would like to thank Cici Liang for helping me train mice for behavioral experiments and helping me with the histology for all the mice included in this study. I want to thank my lab mates Rowan Gargiullo, Abigail McElroy, Cedric Bowe, and Lucas Williamson, for training me in various procedures and offering me constructive criticism and feedback when needed.

I am also deeply grateful to my thesis committee members Dr. Garrett Stanley and Dr. Alan Emanuel for their valuable feedback and encouragement ahead of my defense.

Finally, I would like to thank my friends and family for their constant support and encouragement throughout this journey. Their belief in me has been a steady source of strength and has motivated me to continue pursuing research.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	v
LIST OF SYMBOLS AND ABBREVIATIONS	vii
CHAPTER 1. INTRODUCTION	1
1.1 Introduction	1
1.2 Sound-Seeking Task	2
1.2.1 Measuring Performance Metrics	4
1.2.2 Observations from previous results	5
1.3 Objectives and Aims	6
CHAPTER 2. METHODOLOGY	8
2.1 Developing Behavioral Control Software	8
2.2 Training mice using Octopilot	11
2.3 Developing closed-loop system using video tracking	13
2.4 Designing closed-loop experiments	16
2.5 Analysis of behavior using pose-tracking on video data	17
2.6 Animal surgeries to induce optogenetic expression in the mouse brain	18
2.7 Testing for optogenetic responses using electrophysiological recordings	20
2.8 Verifying viral expression using histology	21
CHAPTER 3. RESULTS	22
3.1 Training mice on the newly developed behavioral control software	22
3.2 Training mice on closed-loop manipulation task	23
3.3 Video analysis of closed-loop task	27
3.4 Verifying Effectiveness of the Closed Loop System	31
3.5 Optogenetic Injections and Histology	34
3.6 Optogenetic Responses from electrophysiological recording	36
CHAPTER 4. DISCUSSION	41
4.1 Interpretation of results from closed-loop experimental data analysis	41
4.2 Interpretation of results from optogenetics	42
CHAPTER 5. CONCLUSION	43
5.1 Future Directions	43
5.2 Final Acknowledgements	44
REFERENCES	45

LIST OF FIGURES

Figure 1.1 - Adapted from Mai et al. 2024. Freely moving mice learn to seek out an ongoing sound stream.	3
Figure 1.2 - Adapted from Mai et al. 2024. Explanation of performance metrics used in sound-seeking tasks.	4
Figure 1.3 - Left – Adapted from Mai et al. 2024. Example of a trial where the mouse ducks multiple times using position data from SLEAP.	5
Figure 2.1 - Behavioral Control Software – Octopilot	8
Figure 2.2 - Basic schematic that explains how Octopilot functions.	10
Figure 2.3 - Block Diagram of the closed-loop system	13
Figure 2.4 - Example of how the video is processed in Bonsai to differentiate regions.	14
Figure 2.5 - Bonsai Workflow to monitor video of the session	15
Figure 2.6 - Schematic explaining how closed-loop experiments work.	16
Figure 2.7 - Pose-tracking using SLEAP.	17
Figure 3.1 – Behavioral Performance over time in terms of NPP for 4 mice over n = 26 sessions.	23
Figure 3.2 - Illustration of how the region that triggers sound manipulation was changed after observing no significant change in behavior	24
Figure 3.3 - Performance on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials.	26
Figure 3.4 - Performance over time on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials.	27
Figure 3.5 - Comparison of three different trial traces, for example trigger trials the mouse completes in a single session.	29
Figure 3.6 - Comparison of sounds played per second per trigger trial and non-trigger trial, averaged over all trials on all sessions for two mice.	30

Figure 3.7 - Plot of the mouse position whenever the last sound is played for a trigger trial (blue) compared to non-trigger trials (red).	31
Figure 3.8 – Comparison of performance over time on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials. Significance was calculated using a paired T-test.	33
Figure 3.9 - Comparison of average latency of the last sound played in a trial.	34
Figure 3.10 - Left – Photo of gap in viral expression in auditory cortex near injection site. Right – Viral expression in auditory cortex with 5x ChRmine dilution.	35
Figure 3.11 - Example of the optical artefact that arises in the channel traces during LED pulse	36
Figure 3.12 - Optogenetic LFP Responses from 5x ChRmine dilution recorded using dual shank probe with corresponding heatmaps.	38
Figure 3.13 - Initial auditory response with corresponding heatmap.	39
Figure 3.14 - Optogenetic LFP Responses from Channelrhodopsin injection in AC recorded using dual shank probe with corresponding heatmap.	40
Figure 3.15 - Example of spikes observed after the photovoltaic artefact that arises in the channel traces after the LED pulse	40
Figure 5.1 - Wireless Logger with LEDs for optogenetic stimulation	44

LIST OF SYMBOLS AND ABBREVIATIONS

BMI	Brain Machine Interface
GUI	Graphical User Interface
ROI	Region of Interest
LFP	Local Field Potential
AC	Auditory Cortex
IC	Inferior Colliculus
GPIO	General Purpose Input/Output
LED	Light-Emitting Diode
PFA	Paraformaldehyde
NPP	Number of Ports Poked
PBST	Phosphate Buffered Saline
DAPI	4',6-diamidino-2-phenylindole
EEG	Electroencephalogram

SUMMARY

Advances in brain-machine interface (BMI) technology, including functional electrical stimulation and real-time decoding algorithms, have enabled the stimulation of sensory feedback in individuals with sensorimotor disorders such as stroke or spinal cord injury. To develop auditory BMIs capable of providing more naturalistic stimulation for hearing, it is essential to understand how auditory pathways in the brain respond to real-time stimulation while integrating feedback from behavioral inputs. This study presents a closed-loop behavioral control system designed to modulate activity in the auditory pathways of lab mice in real time within the context of an auditory sound-seeking task in which they must navigate an octagonal maze to locate sounds. The closed-loop control software processes live video of the experiments, using the mice's position in the arena to manipulate when sounds play during the task. Additionally, the study examines the effects of optically stimulating neural activity in brain regions such as the auditory cortex using optogenetic agents in a head-fixed setup with LED light sources. After conducting several closed-loop sound-seeking experiments, we were able to induce a significant change in the behavioral performance of mice on this task, implying that they use auditory information they obtain from movement in some capacity to be able to localize sound. Upon testing the optogenetic stimulation on the brains of mice injected with agents such as ChRmine and Channelrhodopsin, we were able to trigger an electrophysiological response to LED pulses, which would allow for neural stimulation during behavioral sound-seeking experiments in the future. The outcomes of this study could contribute to the future development of auditory BMIs designed to assist individuals with conditions such as hearing loss.

CHAPTER 1. INTRODUCTION

1.1 Introduction

Sensory perception relies on a continuous feedback loop between the brain and body. For example, when we reach out to grab an object, we use visual and tactile input to guide how our hand moves. By actively interacting with our surroundings, we learn more about our environment and this information can shape future interactions. For example, if you collide with furniture while trying to navigate through a dark room, then you move away from it to prevent that unpleasant experience from happening again. Similar to how sensory processing can influence movement in this example, it can also in turn be modulated by movement. It has been proven that during locomotion mice can filter out noise generated by their own movements by modulating activity in auditory cortex and can anticipate these sounds in the future [1]. Since the brain can actively predict future sensory input using information it obtains in real-time from the environment, we can compare think of sensory processing as a closed feedback loop. When this feedback loop is disrupted due to injury, neurological disorders, or neurodegenerative diseases, sensorimotor deficits can arise, impairing an individual's ability to perceive and respond to stimuli. For example, in Parkinson's disease, proprioceptive deficits can reduce spatial awareness of limb position and cause an over-reliance on visual cues for gait [2].

Understanding how the brain encodes, and processes sensory feedback is critical for developing interventions to restore function in individuals with sensorimotor deficits. Brain-machine interfaces (BMIs) and real-time neuromodulation techniques are technologies that can be used to decode neural signals and provide artificial sensory

feedback to enhance or restore perception and motor control. For example, there are tactile BMIs that use implantable electrodes to electrically stimulate the somatosensory cortex, enabling users to feel tactile sensations such as edges, shapes, and motion [3]. There are also auditory BMIs that use real-time EEG data to classify auditory stimuli during motion [4]. Thus, delivering neural feedback in real-time during task performance will provide insights into how to develop more precise neuro-prosthetic interventions and refine rehabilitation strategies to allow for adaptive control for assistive technologies.

1.2 Sound-Seeking Task

As described in the previous work from the lab (Mai et al., 2024) [5], the sound-seeking task is an active auditory task designed by the Rodgers lab to understand how mice coordinate hearing and movement. In this task, mice must navigate a multi-chambered octagonal arena with speakers on each side, to track where a sound is coming from for any given trial in the experiment. To motivate the mice to perform the task, water in their cages was replaced with water mixed with citric acid, which gives it a sour, unpleasant taste without affecting the mouse's health in the long term. A reward port (or nosepoke) is present in each chamber of the arena underneath the speaker and rewards the mouse with fresh water whenever it detects that a mouse is attempting to drink water from the port when reward conditions are met.

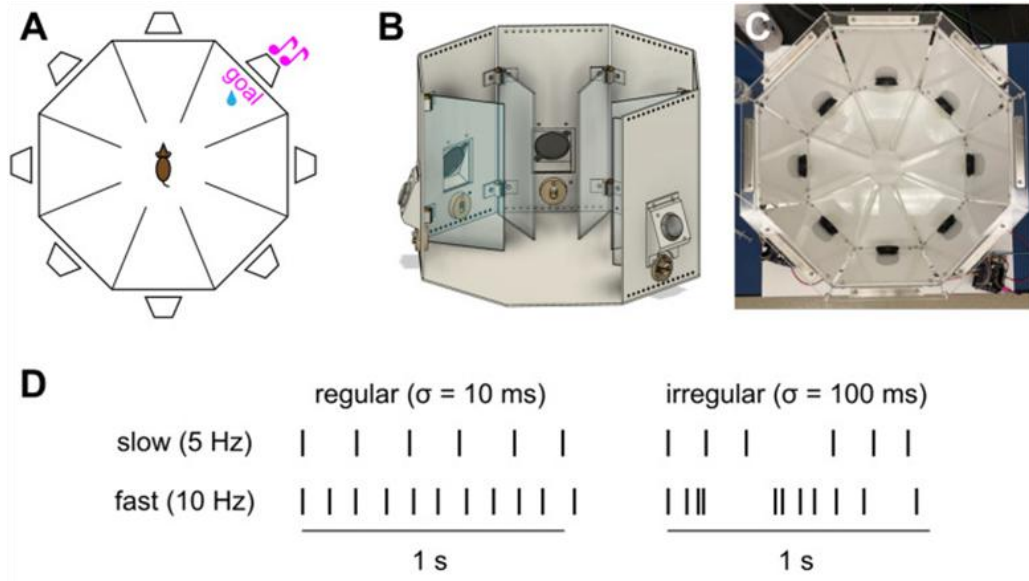


Figure 1.1 - Adapted from Mai et al. 2024. Freely moving mice learn to seek out an ongoing sound stream. (A) Multi-chambered octagonal arena for sound-seeking. A speaker (trapezoid) was mounted on each wall above a nosepoke. On each trial, one speaker was randomly selected as the goal (pink). Mice received a water reward for poking the nosepoke below the goal speaker. (B) CAD schematic of the arena with two external walls cut away to provide a view from the perspective of the mouse. Speakers are black; nosepokes off-white; and dividers transparent. (C) Top-down photograph of the arena. Each side of the octagon was 15 cm and its longest diagonal was 40 cm. Transparent dividers between chambers are nearly invisible in this photo. The entryway into each chamber was 12 cm from the speaker and 5 cm wide. (D) The auditory stimulus was an ongoing stream of intermittent noise bursts. The interval between each noise burst was independently drawn from a gamma distribution with mean μ controlling the rate and standard deviation σ controlling the irregularity.

Each session lasts approximately 25 minutes and comprises a series of trials. For any given trial, one of the eight speakers (the “goal”) plays a continuous stream of intermittent white noise bursts, which continues until the mouse pokes the correct goal port. If the port below the goal speaker is poked, the mouse receives a reward for poking the port below the goal speaker. The trial continues through incorrect pokes, and each trial lasts until the mouse pokes the correct port. After a mouse completes a trial, the next trial automatically gets initiated 1 second after the previous reward has been delivered. On each trial the goal

speaker was chosen at random. However, the same speaker was never the goal on two consecutive trials to ensure that the mice do not persist in poking the previously rewarded port.

1.2.1 Measuring Performance Metrics

E

	score on a single trial							average over session	
ports poked	1	2	3	4	5	6	7	ports poked per trial	
trial is correct	1	0	0	0	0	0	0	fraction correct	

best chance worst

Figure 1.2 - Explanation of performance metrics used in sound-seeking tasks. Adapted from Mai et al. 2024. (E) We scored behavioral performance as the fraction of correct trials and the mean number of ports poked per trial. Pokes into the previously rewarded port were never rewarded and did not affect scoring.

The performance of the mice on the sound-seeking task was quantified with primarily two metrics, after disregarding pokes in the previously rewarded port. The first metric is the fraction of correct trials, which calculates the ratio of correct trials (trials in which the first port poked by the mouse is the reward port) to incorrect trials (all other trials where a port other than the reward port is poked before the mouse correctly pokes the goal port). The second metric is the average number of pokes per trial. All duplicate pokes made to ports in the trial are ignored, and the number of unique ports visited for that trial. This metric varies from 1 to 7 for any given trial, with 1 being best performance, where the mouse only visits one port to get a reward. The worst performance that can be achieved for a given trial is 7, as the mouse visits every other port before visiting the goal port. Chance performance on the sound-seeking task would be 4 ports poked per trial, and a fraction correct of 1/7.

We mainly quantified performance as ports poked per trial because this metric distinguishes between incorrect trials with few versus many incorrect ports poked.

1.2.2 Observations from previous results

Previous work by the Rodgers lab in Mai et al. (2024) demonstrated that mice with unilateral hearing loss were eventually able to recover their performance on the sound-seeking task. On analyzing the videos of the sessions, we observed that during the learning stages of the task, the mice would visit each chamber and poke the port to see if it has a water reward. This movement was called cycling, and it was typically observed in the early stages of learning the task. It was also observed after hearing loss, when they were not able to hear sounds in the same way they were while learning. As the mice become more proficient in the task, they start to pivot away from cyclically visiting every port to check for water and instead adopt a “ducking” strategy. We observed that the mice ducked briefly into each chamber to check the sound level, before deciding whether or not to poke the goal port.



Figure 1.3 - Left – Adapted from Mai et al. 2024. Example of a trial where the mouse ducks multiple times using position data from SLEAP. Ducks are highlighted as red circles. Right – Images of the mouse doing a ducking motion before moving to the next chamber.

The ducking strategy helps them simplify their decision because they can make their decision based on the soundlevel in a specific chamber of the arena as opposed to determining the direction of a sound from the center of the arena, which has 8 potential options for rewards. They can eliminate the possibility of a chamber being the reward chamber if the volume of sound within that chamber becomes lower as they move closer to it. In this way, they use movement to cause a change in the auditory information from the stimulus, and in turn use this information to simplify their decision making, allowing themselves to complete trials more quickly.

This study was motivated by wanting to understand if the mice would still be able to do the sound-seeking task if they had no auditory information during this ‘ducking’ movement described earlier. This will give us insight into whether mice strictly rely on a change in the overall sound level to make their decisions during a specific trial to locate sound. If the mice are still able to do the task with no difference in their performance, we can hypothesize that they might be using some information from binaural cues that they receive while being in the center of the arena before they decide to move towards a specific chamber. This suggests that they could be using a difference in volume in both ears while moving their head to decide which side to move towards.

1.3 Objectives and Aims

One of the primary objectives of this study was to develop systems to allow for real-time adaptive closed-loop stimulation of neural activity for auditory sound-seeking experiments. Another parallel, but related goal for this study is to establish preliminary results and find the best practice for the stimulation of neural activity in the auditory pathway during the

sound-seeking task. This is important because it will give us insight into the type of auditory input that is required to make decisions relating to sound-location. If a general change in the sound level is necessary to determine the location, then the development of future auditory BMIs will need to focus on stimulating the brain to modulate the volume of sounds based on how the user moves their body in space. In contrast, if binaural cues play a more prominent role in localization, then an alternative strategy to stimulate activity would need to be investigated. To stimulate brain regions involved in auditory processing, this study employed optogenetics, a technique commonly used in neuroscience research to control neuronal activity using light sources such as LEDs or lasers. Thus, it provides preliminary results to eventually be able to develop a wireless closed-loop optogenetic system that uses the mouse's behavior in the arena as an input to control neural activity in the auditory pathway in real-time. By interfacing wireless optogenetic hardware with the developed closed loop system, we can manipulate neural activity in auditory brain regions while the mouse is moving freely, which could lead to novel future studies in auditory stimulation. Thus, this project has two primary aims:

1. *Test how closed-loop manipulation of the auditory stimulus affects the ability of mice to locate sounds*
2. *Influence neural activity in the auditory pathway using optogenetic stimulation*

CHAPTER 2. METHODOLOGY

2.1 Developing Behavioral Control Software

To control behavioral experiments, the lab was previously using a behavioral system called Autopilot [6]. Autopilot is a Python framework designed for performing complex, hardware-intensive behavioral experiments with groups of networked Raspberry Pi microcontrollers. While Autopilot was intended as a flexible framework for various behavioral paradigms, its adaptability was limited when modifications were required to introduce new functionalities or integrate with external software. This constraint necessitated the development of a new behavioral control system specifically tailored for sound-seeking experiments in the octagonal arena.




Figure 2.1 - Behavioral Control Software – Octopilot. Displayed are four windows of the Octopilot GUI corresponding to four different behavioral arenas.

To address these limitations, we developed Octopilot as a new behavioral software system, building upon the codebase of Autopilot. Octopilot consists of a graphical user interface (GUI) running on a desktop PC that interfaces with Raspberry Pi microcontrollers within the sound-seeking arena. The PyZMQ library was used for networking, and message queuing between different systems on the local network. The PyQt library was used to develop the GUI, and the PiGPIO library was used to control the hardware on the Raspberry Pi through its GPIO pins. The jack daemon was used to control how audio was being played from the HiFiBerry Amp2 audio card attached to all Raspberry Pis in the experimental setup.

Another key structural change was introduced with Octopilot: whereas Autopilot relied on a "parent" Raspberry Pi to determine stimulus parameters and distribute them to four "child" Raspberry Pis, Octopilot eliminates the need for a parent node. We made this decision to simplify interactions between systems and reduce the latency that comes from messages travelling over an additional system on the same wireless network. It also removes one potential point of failure in the overall behavioral system since the desktop can handle setting parameters for experiments with no performance overhead. Instead, the GUI directly communicates with the four Raspberry Pi microcontrollers present in the arena, which now function as independent agents. The GUI is responsible for setting and transmitting trial parameters to all child Raspberry Pis.

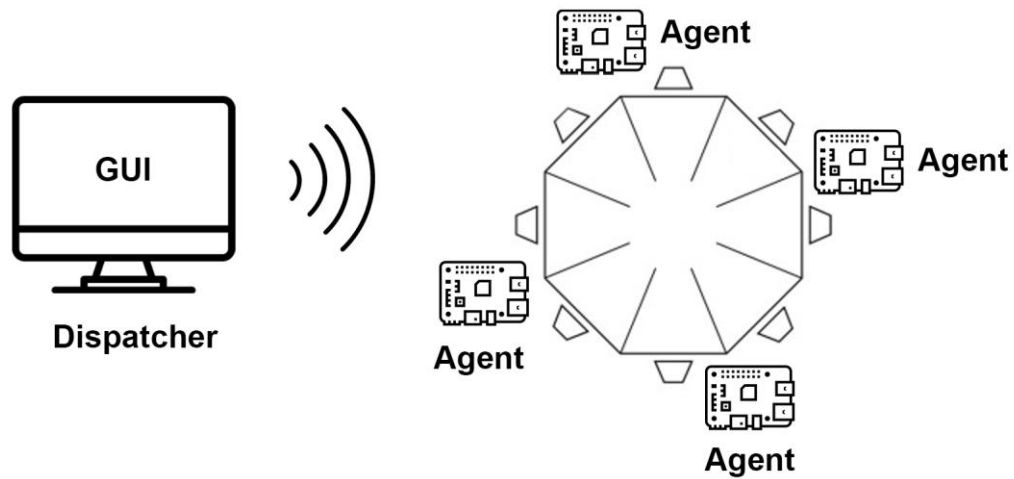


Figure 2.2 - Basic schematic that explains how Octopilot functions. There is code unique to Dispatcher and Agents, as well as shared functions they use for applications like networking. Some free-use images were adapted from NounProject.com and Flaticon.com.

Each agent Pi controls two speakers via the audio card and manages two nosepoke sensors using custom-designed printed circuit boards. The agent Raspberry Pis generate auditory stimuli for each trial and detect changes in the state of the nosepoke sensors and solenoid valves, which are responsible for delivering water rewards. Additionally, the nosepoke circuit boards are equipped with RGB LEDs to indicate trial start and end times in recorded video data. All detected hardware state changes are transmitted to the GUI, which logs experimental data in .csv format. These log files include trial parameters, individual nosepokes, sound playback details, and LED flash events. The recorded data is subsequently used for performance evaluation and synchronization with video recordings.

We developed Octopilot to enhance the robustness of the previously used behavioral control software, allowing for more precise maintenance, adjustments, and the implementation of new functionalities and tasks that were not previously feasible. It is now

employed for behavioral training of mice in all variations of the auditory sound-seeking task and provides a more adaptable system for experimental control and data collection. The development of this software serves as the foundation on which the closed-loop experiments were built upon.

2.2 Training mice using Octopilot

Before introducing the mice to the octagonal arena, their behavioral training begins with 5–7 days of handling, in which the mice are habituated to being picked up and weighed. After this stage, the water bottle in the mouse’s cage is removed and replaced with a bottle that contains citric acid water. These mice were motivated to perform the task for the water rewards because the fresh water tasted better than their home cage water. Typically, mice began at 1.5% citric acid and were increased to higher concentrations if their trial count was substantially lower than usual. We lowered the water reward size as mice became better at the task to encourage them to perform more trials.

After 1–2 days on this water regimen, mice began “pre-training” which lasted about 3–7 days (one session per day) and taught them to use the nosepekes. The first stage of pre-training was “manual baiting” in which they were introduced to the behavioral arena. In this stage, approximately 0.1 mL of water was placed inside the nosepekes using a syringe so that they would learn to associate the nosepekes with reward. After mice learned to reliably obtain their daily ration from the nosepekes, they began “automated poke training.” In this stage, no sounds were ever played, and every poke into a nosepoke triggered the release of a water reward, with the exception that the same nosepoke never delivered a reward twice in a row. Mice quickly learned to visit nosepekes serially, often

in a cycling pattern within a couple of days being put on this task. The mice are introduced to this task with their cage mates initially, before having to do the task independently in a different arena. Once they learned to reliably obtain their daily ration in this way (which typically takes 2—5 days), pre-training is concluded, and they begin training on the sound-seeking task.

There are multiple variations of the sound seeking task, which were all implemented in Octopilot, using the parameters that were used previously. The implemented tasks either had variable acoustic parameters (randomly chosen on every trial) or fixed acoustic parameters (same on every trial). In all cases, the noise bursts were narrowband white noise with 10 ms duration and 3 kHz bandwidth. The intervals between noise bursts were drawn from a gamma distribution, which allowed us to independently vary repetition rate (mean number of noise bursts per second) and irregularity (standard deviation of the interval between noise bursts). To prepare the mice for the closed-loop task, they were trained on variable acoustic parameters, in which the center frequency was selected from the range 5 to 15 kHz, sound level from 65 to 80 dB SPL, repetition rate from 3 to 10 Hz, and irregularity from 1 to 100 ms. All these acoustic parameters were constant within each trial but could vary across trials.

2.3 Developing closed-loop system using video tracking

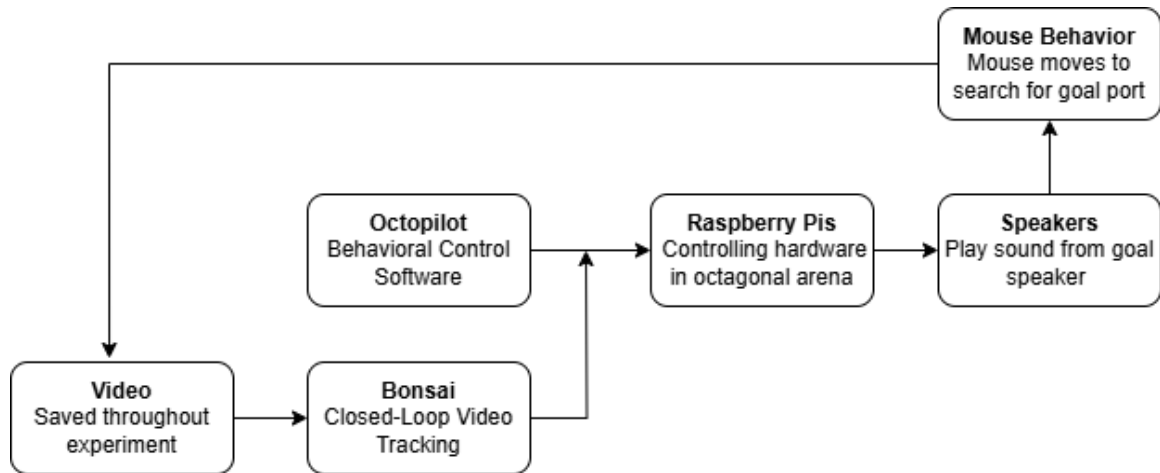


Figure 2.3 - Block Diagram of the closed-loop system.

One of the key advancements introduced during the development of Octopilot was the capability to implement new tasks that could interact with external software within the experimental system. This paved the way for the development of the closed-loop version of the sound-seeking task, which manipulates auditory stimuli in real-time using the mouse's position within the octagonal maze.

Each experimental setup includes an overhead camera that records video throughout the session. The real-time MP4 video stream from each camera is accessible via a browser-based URL, enabling experimenters to monitor the session. To implement closed-loop stimulus control, we integrated these video streams into a real-time control system developed using Bonsai, an open-source visual programming language.

2.3.1 Real-time implementation using Bonsai

Bonsai is based on the ReactiveX .NET framework, designed to write asynchronous event-based programs to interact with sequences of data points [7]. Programs in Bonsai can be written by connecting graphical elements called “Operators” to make a sequence. In this study, Bonsai was used to perform continuous image processing to determine whether the mouse was in a specific region of the octagonal arena during a session. This was achieved by applying a threshold to a grayscale version of the video after defining a region of interest within the arena. The resulting processed video stream consisted of binary pixel values, where each pixel was either black or white (corresponding to 0 or 1). The sum of these pixel values was then computed across the entire frame.




Figure 2.4 - Example of how the video is processed to differentiate regions. (A) Video of the mouse performing the sound-seeking task is converted to grayscale. (B) The region of interest is made at the center of the octagonal arena, subtracting its pixel values from the video to set threshold. (C) Overlaying region of interest on the binary version of the video with only white and black pixels. (D) Pixel values detected to be lower than threshold because mouse body is outside the center. This will trigger sound manipulation.

A Boolean operator was implemented to generate a True/False output based on whether the computed sum fell below a predefined threshold. This threshold was set at the beginning of each session using the baseline values recorded before the mouse was introduced into the arena. The resulting Boolean output was continuously streamed to the agent-side code of Octopilot, enabling real-time manipulation of auditory stimuli based on whether the mouse was within the designated region of interest. These messages were transmitted via a ZMQ Publisher socket in Bonsai and received by a ZMQ Subscriber socket in the Octopilot agent.



Figure 2.5 - Bonsai Workflow to monitor video of the session.

Although real-time pose tracking could have been integrated into Octopilot to provide a more detailed analysis of mouse movement, a simpler image processing approach was chosen to minimize additional latency. Pose tracking methods typically introduce latencies in the range of hundreds of milliseconds, whereas the Bonsai-based image processing approach was able to maintain latencies well under 50 milliseconds. A continuous stream of messages could be sent in real-time to the agent Raspberry Pis using Octopilot without any significant delay in their function. This lower latency was a critical factor in ensuring precise closed-loop control of auditory stimuli in the experiment.

2.4 Designing closed-loop experiments




Figure 2.6 - Schematic explaining how closed-loop experiments work. (A) Trigger Trials Sound will play when the mouse is in the center of the arena within the defined rectangular region. If a part of the mouse's body exits the region, Octopilot executes a command to silence the sound playing for the trial. This is to prevent any audio from being played while the mouse makes a ducking motion. **(B) Non-Trigger Trials.** These trials play sound normally irrespective of the mouse's position in the arena.

To determine whether movement strategies are necessary for sound localization, a modified version of the task was implemented in which sounds were played only when the mouse was in the center of the octagonal arena. This ensured that the mice had no auditory information while performing the ducking motion mentioned earlier. This was achieved by defining a rectangular region of interest (ROI) in Bonsai, approximately the size of the mouse's body, at the center of the arena. If the mouse moved outside this central region during a trial in which sound manipulation was applied, the auditory stimulus was immediately silenced until the mouse returned to the center of the arena.

To make the outcomes of the closed-loop experiment interpretable, for any given experimental session, 50% of the trials would be trials in which no sound-manipulation occurs i.e., the trials proceed as normal trials in the sound-seeking task. These trials were called non-trigger trials. For the other 50% of trials, sounds would only play when the

mouse is in the center of the arena. These trials were called trigger trials. For the closed-loop experiments, Octopilot assigns a trial type to specific trials at the start of a session at random. The type of trial is logged in the data files to allow for further analysis while processing the results of the experiment.

2.5 Analysis of behavior using pose-tracking on video data



Figure 2.7 - Pose-tracking using SLEAP. Displayed is a screenshot of inference with key points corresponding to different parts of the mouse body.

During the closed-loop experiments, video data for all experimental sessions were recorded and saved. These videos were analyzed using SLEAP, an open-source deep-learning-based framework for animal pose tracking [8]. The purpose of this analysis was to examine how mice moved throughout a session and to quantify behavioral differences between trigger and non-trigger trials.

The SLEAP model used for video analysis was trained on recordings from previous sound-seeking sessions. Before conducting full-scale analysis, model inference was

verified to ensure accurate skeletal predictions in newly recorded videos. To synchronize the two data sources, Python video processing scripts were used to preprocess the video data. Trial start times recorded in the behavioral logs were aligned with the corresponding LED flashes in the video, which were used to indicate trial start and end points. This alignment produced a time offset, which followed a linear relationship between video and behavioral timestamps.

Once SLEAP was applied to all recorded videos, kinematic data for each mouse was extracted by tracking its 2D coordinates in every frame. To further refine the spatial analysis, a separate model trained on images of the octagonal arena was overlaid onto the tracked data, allowing for automatic region assignment within the arena. This enabled the distinction between the center region and the eight peripheral chambers.

2.6 Animal surgeries to induce optogenetic expression in the mouse brain

For this study, we initially used the excitatory red-shifted Channelrhodopsin ChRmine, which responds to light between wavelengths of 550–650 nm allowing for the stimulation of deeper cortical and subcortical regions [9], [10]. The optogenetic agent was delivered using the AAV8-CaMKIIa-ChRmine-mScarlet-Kv2.1-WPRE viral vector. The surgical procedure involved anesthetizing the mouse with isoflurane and administering the analgesic buprenorphine pre-surgery. Pain control was ensured by monitoring the breathing rate and performing toe-pinch reflex tests. Before making an incision, the head was shaved and sterilized using ethanol and chlorhexidine.

For this study, injections were targeted at the auditory cortex and inferior colliculus. To ensure precision, stereotactic coordinates were measured before drilling a craniotomy

at the appropriate locations. A glass micropipette was used for viral injection, which was first filled with mineral oil to prevent air bubbles before loading the virus. The virus was stored at -80°C and thawed gradually in an ice bucket during surgery.

A Nanoject injector was used to deliver the virus into the targeted brain regions. A titer of 1×10^{13} vg/mL was used for the full concentration ChRmine injections and a titer of 7×10^{12} vg/mL was used for the Channelrhodopsin injections. Across the viral injection surgeries, we varied the concentrations of the ChRmine injections and used 5x and 10x dilutions on top of the full concentration injections and used only full concentration injections for Channelrhodopsin. We also experimented with the injection volumes for both viral vectors and varied between 100-600 nL across all injections. Following the injection, the virus required three weeks for full expression. During this period, some of the injected mice underwent training on the sound-seeking task.

Aside from the mice injected with ChRmine we also injected mice in the auditory cortex and inferior colliculus with the retrograde Channelrhodopsin pAAV-hSyn-hChR2(H134R)-EYFP [11]. This optogenetic agent responds to light between 450-500 nm in wavelength. We did this to observe electrophysiological activity due to cortico-collicular projections by recording from and optically stimulating AC after injecting the viral vector in IC.

2.7 Testing for optogenetic responses using electrophysiological recordings

Before performing histological verification of viral expression, we aimed to verify the efficacy of optogenetic injections through electrophysiology by recording neural activity in mice using silicon probes in a head-fixed setup. Two types of Cambridge Neurotech silicon probes were used in this study: ASSY-325 H3 & L3 (single-shank 64-channel electrode), and ASSY-350 H12 & L13 (dual-shank 128-channel electrode). Each shank in these probes had an approximate width of approximately 80 microns per shank, with the dual-shank probe having a gap of 1000 microns (1 mm) between shanks. Each shank has a recording depth of 1300 microns (1.3 mm), with electrode sites spaced at 20-micron intervals.

To keep the mice stable in the head-fixed setup, a headplate was surgically attached to the skull after the viral injections. This required a second surgery, during which the cement over the skull was removed, exposing the previous craniotomies from the injections. Metabond, a bone cement, was applied to the skull, and once it solidified, the headplate was affixed securely. An additional craniotomy was made for a ground pin, which served as a reference electrode for neural recordings. The ground pin was inserted into another brain region and secured using UV-cured cement. For mice requiring multiple recordings, the craniotomy was sealed with KwikCast (WPI).

Once headplated, mice were anesthetized and placed in the head-fixed setup, where neural signals were recorded using the silicon probe while an LED stimulus was presented. For the ChRmine injections, we used amber and red LEDs with a power of 15 mW and for Channelrhodopsin we used blue LEDs with a power of 12.5 mW. The LED was mounted

on a manipulator, positioned above the mouse's head, and controlled via a Raspberry Pi. The Raspberry Pi ran Octopilot, enabling sound playback during a passive noise task. The electrodes were connected to an OpenEphys [12] electrophysiology acquisition board, which recorded both neural signals and analog signals from the LED, ensuring precise synchronization of optogenetic stimulation and electrophysiological responses.

2.8 Verifying viral expression using histology

To further verify the extent of the viral expression, slices of the mouse brain needed to be observed under a microscope. After recording the injected mice, these mice were euthanized using a dilution of Euthasol and perfused to preserve the brain tissue. 10mL of Paraformaldehyde (PFA) is used as a fixative agent and is directly perfused through the circulatory system so that all the tissue in the animal can be preserved in the state it was in before it was euthanized. The skull of the mouse is dissected to extract the brain, which is stored in PFA for 24 hours before being transferred to a sucrose solution to prepare it for slicing. The brain is frozen and sliced using microtome blades with the help of a cryostat machine. Once the desired number of sections have been obtained, the slides are stored in a slide box to be stained and coverslipped. The slides were washed with phosphate buffered saline (PBST) before staining. The brains used in this study were stained with DAPI (4',6-diamidino-2-phenylindole) a blue-fluorescent DNA stain that tags the nuclei of cells in the brain. The mScarlet reporter in the AAV viral vector is a red-fluorescent reporter that will express in regions of the brain with viral expression. The slides were imaged under a Keyence dark-field microscope.

CHAPTER 3. RESULTS

3.1 Training mice on the newly developed behavioral control software

From previous research, it was established that mice could successfully learn the sound-seeking task using the earlier behavioral software, Autopilot. To verify that Octopilot was also suitable for training mice on this task, one of the first experiments conducted involved running untrained mice through the training protocol using Octopilot. A cohort of four wild-type mice, two males and two females, was used for this experiment, with each mouse housed separately in individual cages and placed in separate octagonal behavioral arenas. This experiment served to confirm that Octopilot could seamlessly interact with the hardware across all available experimental setups in the lab without major issues during sessions. The primary goal was to ensure that the software functioned correctly on both the desktop client and the Raspberry Pis that controlled the experimental apparatus.

The mice in this cohort were trained on the variable auditory task, where the center frequency, amplitude, and rate of white noise bursts were altered between trials. This task was chosen over the fixed sound task to acclimate the mice to variations in sound, preparing them for future performance on the full version of the task. The mice took slightly longer than expected to reach a level of “good” performance, which is typically reflected in experimental sessions where the calculated NPP (number of ports poked) is less than 2. Further investigation revealed that this delay was due to an issue with how the speakers were playing sound during each trial. While each Raspberry Pi was connected to two speakers, only one speaker was active at any given time, meaning that only four out of the eight speakers in the octagonal arena were being used. As a result, the mice developed a

cycling strategy in a square pattern, primarily visiting chambers with the left-side speaker. In Figure 3.1, we can see that the average performance plateaued during the initial learning phase at an NPP of around 3.5 because of this issue. The issue with the speakers was corrected on the date 2024-11-10, after which we observe an improvement in learning. Once this issue was corrected, the mice rapidly learned the task and began exhibiting behavior characteristic of the later stages of learning in the sound-seeking task.

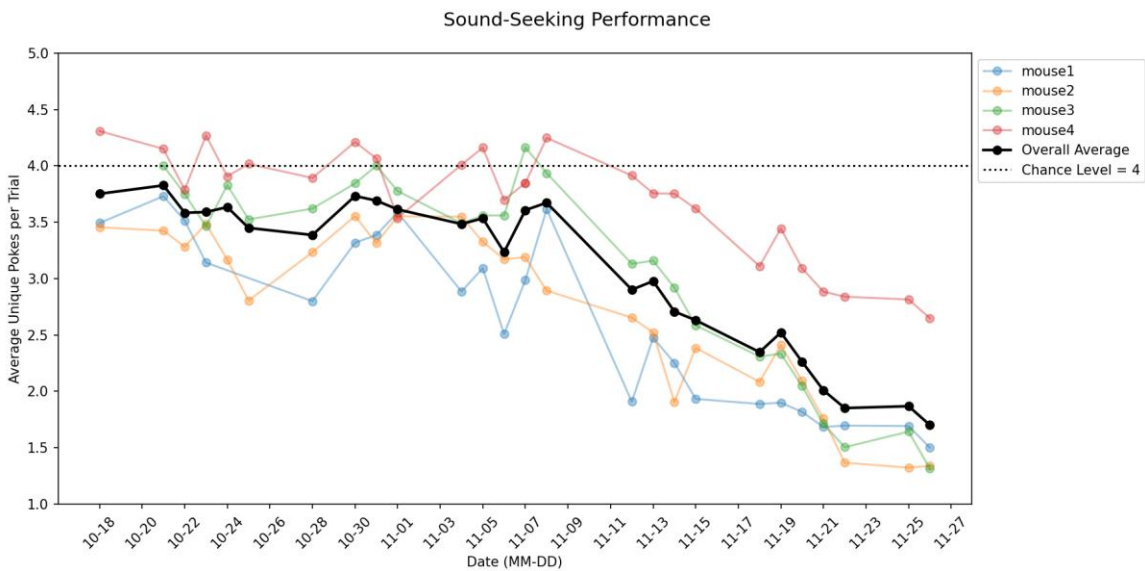


Figure 3.1 – Behavioral Performance over time in terms of NPP for four mice. Mice were trained on the variable version of the sound seeking task across n = 26 sessions, with each experimental session being conducted on a different day.

3.2 Training mice on closed-loop manipulation task

After verifying that mice could be trained using Octopilot, the next step was to test them on the closed-loop version of the task. In the first iteration of the closed-loop task, the center of the arena was defined as a smaller octagon that covered most of the areas outside the reward chambers. Sound would be silenced if the mouse was observed inside the chambers, meaning that the only place it could hear sound was in the center of the arena.

After running several sessions under these conditions, task performance was compared between “trigger” and “non-trigger” trials by calculating NPP separately for each trial type over multiple sessions. Octopilot chose the type of each subsequent trial at random and did not alternate between trial types. However, the initial comparisons revealed no consistent differences in task performance between the two trial types across several days.

This finding motivated a change in the task setup. Instead of defining the entire center of the arena as the region with sound playing, the area where the mouse had to be to hear sound was reduced to a much smaller rectangular zone in the center. It is possible that this change might have led to inconsistency while analyzing behavior, making trials with reward ports opposite to the corners of the region easier to complete than ports opposite to the edges. However, the intent behind this change was to constrain the region to be just big enough to contain the body of the mouse. Seeing the head of the mouse poke out of the center made it easier to set the threshold for each session. Having a larger octagonal area with sound playing made the task easier for the mouse to complete. Adjustment to the region in the center was necessary because we suspected that sound manipulation was occurring only after mice had already made their decision about which port to poke for a reward, rather than influencing their decision-making process.




Figure 3.2 - Illustration of how the region that triggers sound manipulation was changed after observing no significant change in behavior.

After implementing this change, a more consistent difference in performance emerged between trials where sound was manipulated versus trials without manipulation. Initially, we hypothesized that because the mice could no longer rely on a movement strategy near the chambers to make their decision, their performance on manipulated trials would drop to chance levels. However, the results contradicted this hypothesis. Even after making the task more challenging, the overall performance on "trigger" trials, though lower than on "non-trigger" trials, remained comparable to the average performance observed in late-stage training. This suggested that mice were able to adopt a strategy to localize sounds above chance level, despite the sound only playing when they remained in the center—a region they typically avoid. Further analysis of the behavior logs from Octopilot revealed that, on average, mice heard fewer sounds and took longer to complete trials in which sound was being manipulated.

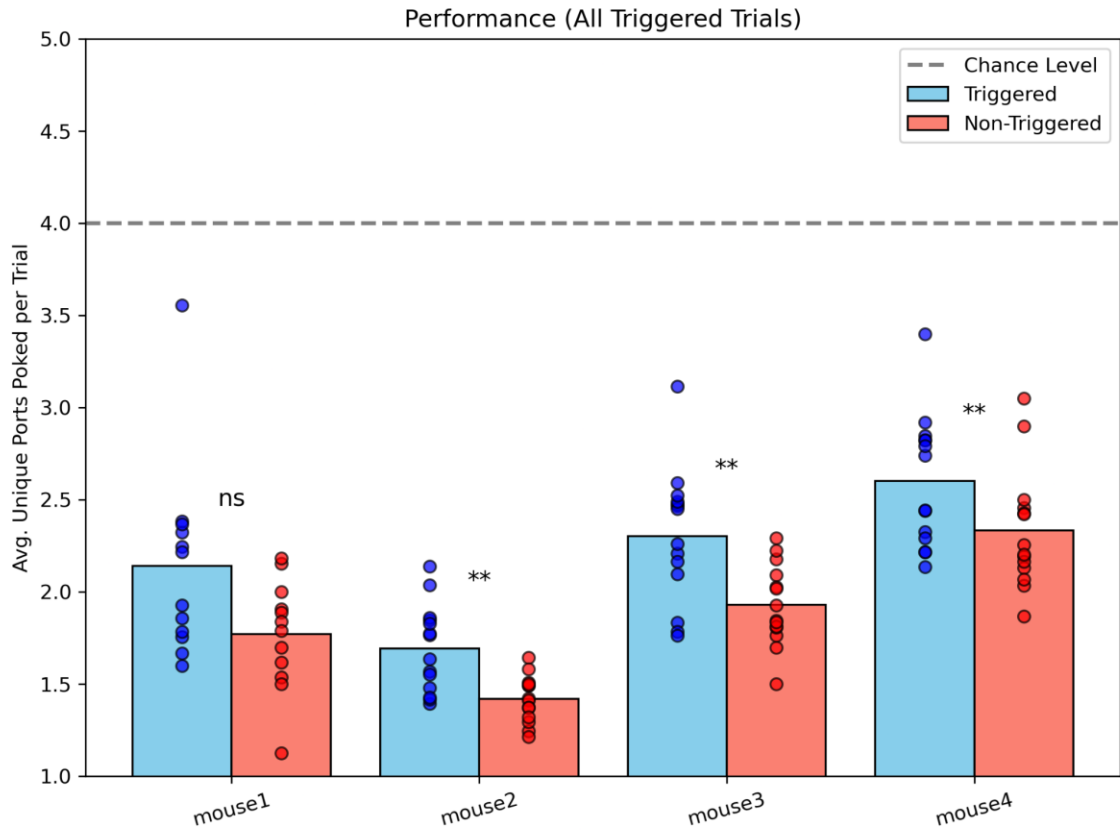


Figure 3.3 - Performance on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials. The bar graphs represent the Average Unique Ports Poked for a trial, averaged over a session. The scatter plots represent the values for individual sessions for the mouse.

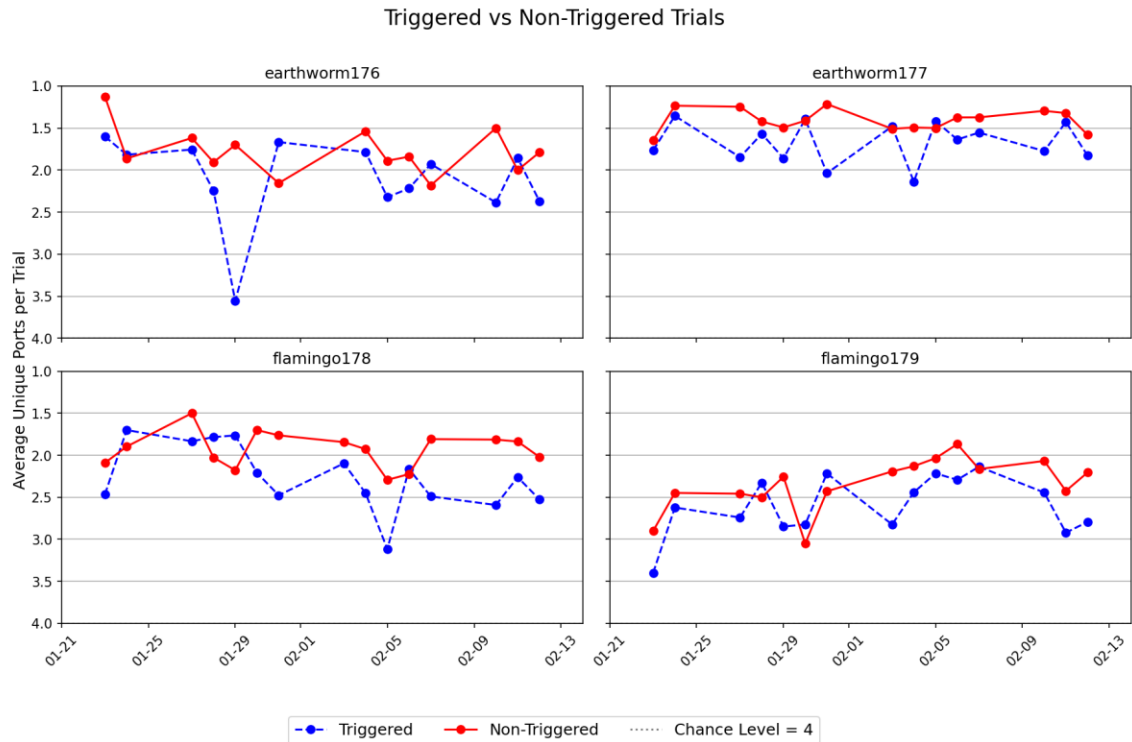


Figure 3.4 - Performance over time on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials. Behavioral difference between the two types of trials is consistent over time.

3.3 Video analysis of closed-loop task

After initial testing and troubleshooting of the closed-loop system, a new cohort of mice trained on the variable version of the sound-seeking task using Octopilot began the closed-loop task. In total, about 14 sessions were completed across $n = 4$ mice, with videos being recorded for all subjects. These experiments aimed to analyze and quantify changes in behavior using SLEAP pose estimation.[8] The goal was also to support the preliminary results from the initial testing of the first cohort of mice on this task.

Similar to the previous cohort, the new cohort initially faced some difficulty adjusting to the variability in when sound played, but they were eventually able to perform

above chance on trials where sound was manipulated. We conducted video analysis to better understand the strategies the mice used to outperform random guessing.

Initially, we started analyzing video by observing the path the mouse took for individual trials by tracking the position of a labelled key point corresponding to the location of the snout on the mouse's body across all the frames of video. We then compared these traces to see if there was an apparent difference in how the mouse moves in trigger trials compared to non-trigger trials. Upon further investigation, we observed that the results significantly varied between the trigger trials. To gain some insight into paths taken by the mouse, we decided to investigate the times when the sound was playing for specific trials. We took the logged timestamps at which independent sounds played in a specific trial and translated it into a frame index within the range of frames for that trial. This allowed us to overlay sounds over the path traces at the approximate frames that it would have played. When visualizing these trials, we could see that for most trigger trials, it looked like the mouse was moving as intended. It would cross the center to trigger the sound and visit the reward port after potentially poking adjacent chambers. However, upon further inspection it became apparent that the mouse was able to complete some trigger trials without visiting the center at all. Which suggested that something in the system was not working as intended and warranted further investigation.

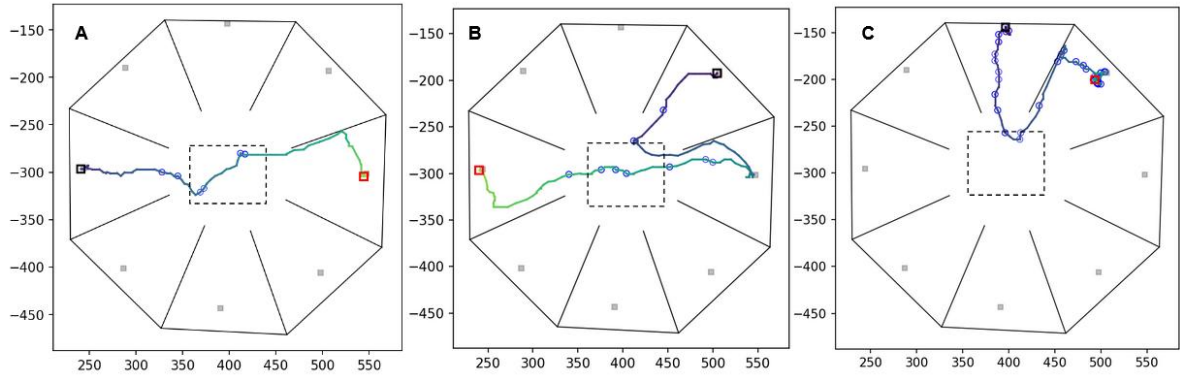


Figure 3.5 - Comparison of three different trial traces, for example trigger trials the mouse completes in a single session. Mouse location when sounds were played were overlaid onto trial paths. (A) Trial where sound manipulation is working as intended. (B) Trial where it is unclear if manipulation is working. (C) Trial where it looks like sound manipulation is not working.

Another analysis that was made possible because of pose estimation was to be able to compare the amount of time the mouse spent in different regions of the arena for both types of trials. Using the values assigned to different regions by SLEAP, we could compare time spent in the center and compare it to the time spent in individual chambers during a trial. Similarly, we can also compare how many sounds were played in different parts of the arena. Using this information, we can calculate the rate of sounds that played in different regions over several sessions and compare the values between trigger and non-trigger trials. On doing so we can see that there is a significant reduction in sound rate for trigger trials in both the center and the chambers. Upon further quantification it was revealed that the rate in these regions for trigger trials was approximately 70% lower compared to non-trigger trials. This suggested that the sound manipulation was working, albeit not as effectively as initially intended. Furthermore, multiple sounds were observed in the previous reward chamber for trials, which would be the region that the mouse starts out in for any given trial before moving to another region. If the sounds were being silenced

in the way that it was designed to be, then no sounds should have been observed in this region.




Figure 3.6 - Comparison of sounds played per second per trigger trial and non-trigger trial, averaged over all trials on all sessions for two mice. Solid color lines represent the averages across all sessions whereas the gray traces represent the average rate of individual sessions.

To further verify the method in which sounds were played, we plotted the last sound for all trigger and non-trigger trials over all sessions. If the system were working as it was designed to be, it would show more sounds near the center for trigger trials, and all sounds at the reward port for non-trigger trials because the sound is supposed to be silent outside the center on trigger trials. However, upon plotting these sounds, we noticed that there were a lot of sounds played in the chambers on trigger trials. We hypothesized that this could be because of a latency issue, which caused a delay before the sounds were silenced after the mouse exited the center of the arena. Non-trigger trials were mostly at reward ports as expected, which suggested that there was an issue with some of the trigger trials.




Figure 3.7 - Plot of the mouse position whenever the last sound is played for a trigger trial (blue) compared to non-trigger trials (red).

3.4 Verifying Effectiveness of the Closed Loop System

From the observed performance metrics, it was clear that mice were able to perform the task well above chance even with the closed-loop manipulation of the auditory stimulus. However, after performing video analysis on all the sessions, it was clear that there were some problems with how trigger trials were being implemented during the experiments. This motivated us to investigate log files to verify if the system was working as intended. When developing the closed-loop version of the task we implemented separate logfiles that logged whenever there was a change in the auditory stimulus. There was a continuous stream of messages being sent to Octopilot from Bonsai based on the mouse's position in the arena, however, it was programmed such that the sound manipulation would only be triggered if there was a change in state between the regions. Upon investigating the log files for individual sessions in certain trials, we noticed that not all the trials assigned as trigger trials were present in the log file that included all the timestamps that a manipulation

in sound occurred. After investigating the detailed log of individual sessions, we concluded that not all trigger trials were triggering sound manipulation. It was also observed that this error would occur most often when two consecutive trials were assigned as trigger trials. The stimulus would not be manipulated until the mouse visited the center of the arena where sound is supposed to play on trigger trials. The silencing effect in the chambers would not be triggered until the mouse visited the center of the arena at least once during that trial. As a consequence of this, the state of the sound would continue to be in the state that it was in at the end of the last trial. This meant that if there was no sound playing at the end of the previous trial, there would be no sound playing throughout the current trial. If there was sound playing at the end of the previous trial, then sound would continue to play, until the mouse visits the center and then enters a chamber.

We tried to quantify how often this phenomenon occurred over all the experimental sessions we conducted by calculating the percentage of total trigger trials that were not logged to manipulate sound and discovered that there was no manipulation of sound for 37% of all trigger trials. This meant that we had to reanalyze the behavioral performance of the mice after excluding trials without any sound manipulation. After doing so, we observed that there is a more significant behavioral difference when comparing trigger trials with non-trigger trials over all sessions.

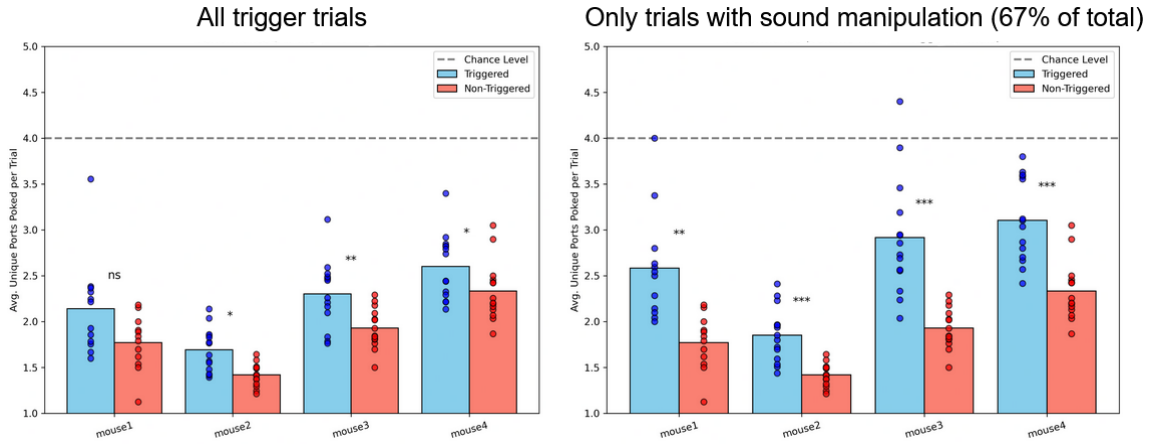


Figure 3.8 – Comparison of performance over time on the closed-loop task in terms of NPP for 4 mice, separated by trigger trials and non-trigger trials. Trigger trials without any logged sound manipulation were excluded from the performance calculation of the plot on the right, which shows a more significant behavioral effect. Significance was calculated using a paired T-test.

Another aspect of the closed-loop system that might have been lacking after observing the previous results was the latency with which the sound manipulation was functioning. To quantify this, we calculated the time difference between the last timestamp at which the mouse crossed over from the center to the reward chamber for a given trial and the timestamp of the last sound that was played for that trial. Doing so would give us an idea of how much time there was between when the mouse entered the chamber and when the sound started being silenced. On doing this analysis it was observed that there were approximately 250 milliseconds of latency on average between the chamber entry and sound manipulation, which was significantly higher than expected. Further analysis suggests that this latency could be arising from how sound is being generated in Octopilot. The time it took for Bonsai to send messages regarding state change was typically in the order of ~10 ms, which suggests that the delay is not being caused due to the network. In the current implementation of the sound manipulation task, Octopilot completely empties

all the sounds it intends to generate whenever the mouse leaves the center and generates sound according to the set parameters from scratch, which could be causing a delay whenever there is a change between states. As a future direction, a possible method to circumvent this issue would be to let the sound play continuously and introduce a scaling factor that greatly reduces the amplitude of sound only when the mouse exits the center of the arena.

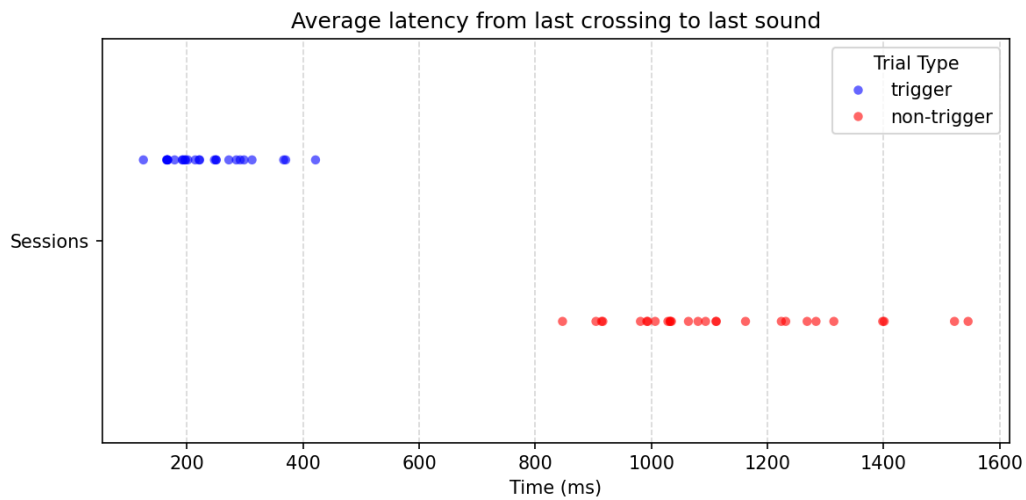


Figure 3.9 - Comparison of average latency of the last sound played in a trial. Scatter points represent the average latency values for individual sessions.

3.5 Optogenetic Injections and Histology

The initial optogenetic injections used full concentration ChRmine in the auditory cortex and inferior colliculus. For the first cohort of injected mice, 3 mice were injected into the auditory cortex and 2 were injected in inferior colliculus. These initial attempts revealed problems with viral expression, including gaps (shown in Figure 3.8) at the injection sites and poor localization of the virus.




Figure 3.10 - Left – Photo of gap in viral expression in auditory cortex near injection site. Right – Viral expression in auditory cortex with 5x ChRmine dilution compared to location of the auditory cortex in the Allen Brain Atlas.

To improve the results, we experimented with different viral concentrations and injection volumes. For concentrations, we opted to attempt injections with more diluted concentrations of ChRmine after suspecting that neurons in the injected site were dying due to excitotoxicity. The concentration of the original titer for ChRmine injections was 1×10^{13} vg/mL. Injections with 5x and 10x dilutions of this concentration were attempted with subsequent mice. We initially started with injections at 600nL but observed better localization in the images with volumes of 300 nL or lower. Injections with volumes of 100nL and 200nL were also attempted to see if the precision of the localization could be improved, preventing spread to other regions. Upon testing concentrations, we found that the 5x dilution of ChRmine showed better expression in the auditory cortex. The smaller injection volumes also produced improved results. We additionally tested retro-Ch2 injections in both the inferior colliculus and auditory cortex, which successfully showed viral expression in both brain regions.

3.6 Optogenetic Responses from electrophysiological recording

The electrophysiological recordings initially struggled to capture clear optogenetic responses. Early recordings, performed more than a month after viral injection, did not show distinct neuronal activation when using red light stimulation. In early stages, the light source was an LED powered by Raspberry Pi GPIO pins. The LED would be pulsed at an interval to observe a spiking response following the stimulus. When we attempted to record from mice with the initial full-concentration ChRmine injections, clear spiking responses were difficult to observe. A clear optical artefact could be observed for the duration of the LED pulse across all recordings. This was likely caused due to a photovoltaic effect that was induced in the electrode by the light stimulus.




Figure 3.11 - Example of the optical artefact that arises in the channel traces during LED pulse.

While we were not able to observe photo-evoked spikes in early recordings, we were able to capture a brief local field potential (LFP) response after the light stimulus. The early LFP response was mild in mice with a full concentration ChRmine injection, likely due to poor viral expression over the recording site. We were able to observe a more prominent LFP response when recording from mice with more dilute concentrations.

In early stages, the LED was powered by Raspberry Pi GPIO pins, which provided insufficient stimulation. We then switched to high-power LEDs with a dedicated power source, which finally revealed more consistent neuronal spiking patterns. We tested different LED pulse durations (1ms, 3ms, and 10ms) and found that the spiking patterns changed depending on the pulse length, with longer pulses having a more sustained response. Experiments with LEDs with wavelengths of 625 nm and 592 nm produced optically evoked spiking responses.

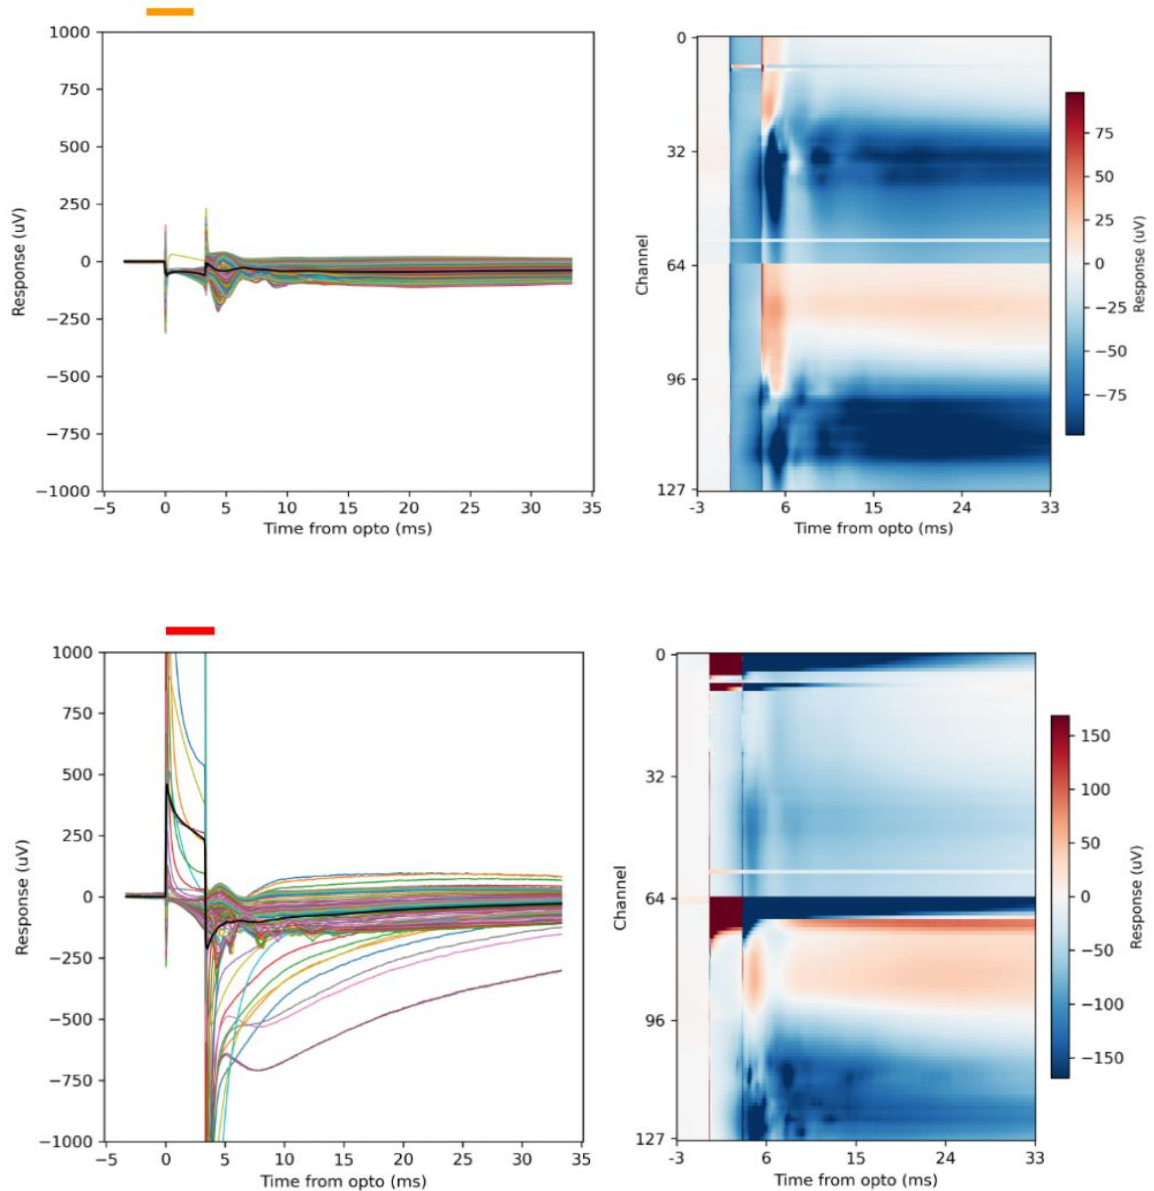


Figure 3.12 - Optogenetic LFP Responses from 5x ChRmine dilution recorded using dual shank probe with corresponding heatmaps. Top – recording using amber light source. Bottom – recording using red light source.

We also tested to see if there was any auditory response from the mouse after sound stimulus. To do this, we used a version of Octopilot on a Raspberry Pi with a passive audio task designed to play white noise bursts at continuous intervals on two speakers placed next to the mouse's ears. On doing so, we could observe an auditory response in the

recordings even with mice that had no prominent LFP or spiking patterns, which helped us verify if we were recording from the auditory cortex. A subsequent test using Channelrhodopsin with blue LED pulses showed a successful neuronal response in the auditory cortex, but did not show any prominent spiking activity in mice injected in the inferior colliculus.

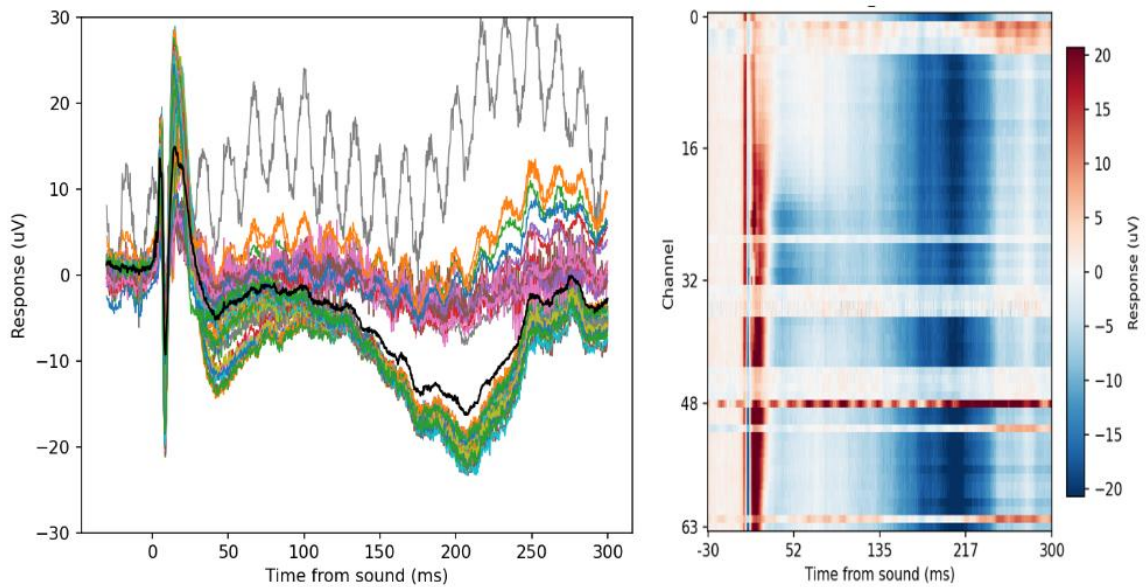


Figure 3.13 - Initial auditory response with corresponding heatmap. Average electrophysiological response is measured from the time a sound is played after being aligned with log files from Octopilot.

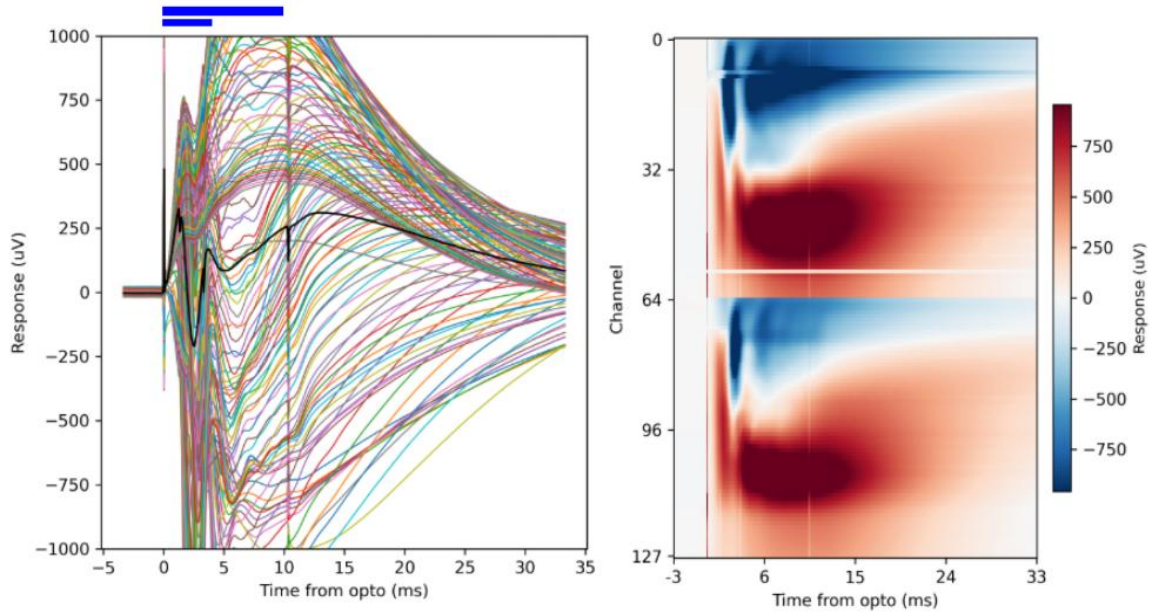


Figure 3.14 - Optogenetic LFP Responses from Channelrhodopsin injection in AC recorded using dual shank probe with corresponding heatmap. This example session used a mix of both 3ms and 10ms LED pulses.




Figure 3.15 - Example of potential spikes observed after the photovoltaic artefact that arises in the channel traces after the LED pulse

CHAPTER 4. DISCUSSION

Through this study we were able to develop a closed-loop system that used video to influence the behavior of mice performing an auditory sound-seeking task by manipulating auditory stimulus. We were also able to trigger a neural response in the auditory cortex using optogenetic stimulation in a head fixed setup.

4.1 Interpretation of results from closed-loop experimental data analysis

From the results of the closed-loop task we identified that there was a difference in behavioral performance on the sound-seeking task. The difference in performance became even more apparent after identifying that the system was not manipulating sound for all trigger trials and excluding trials where the sound manipulation failed. This suggests that mice are reliant on this movement strategy to be able to localize sound, however they are not completely dependent on it as they can perform well above chance. Since the mice could hear sound within the chamber, this means that they could still be using the movement strategy and change in sound-level to make their decisions. The results also suggest that the mice might change their decision on whether to poke a port or not depending on if they notice the sound being silenced after they enter the chamber.

However, due to the nature of the oversights while implementing the closed-loop task, it is possible that their behavioral performance could be impaired further if the closed-loop system was able to manipulate sound with a lower latency as originally intended. This latency is potentially caused due to how the functions in Octopilot were structured. In the current implementation, the queue of sounds to be played in a trial is completely emptied

whenever the mouse leaves the center of the arena. It also regenerates a new queue of sounds when the mouse enters the center, which causes a delay in the manipulation. This can be overcome in the future by implementing a scaling factor that reduces the amplitude of the sounds to be played to a null value whenever there is a change of state in the region. Furthermore, the region the mouse is in should be monitored between trials. The integration of Bonsai with Octopilot needs to be modified such that Octopilot can track the previous chamber and ensure that no sound is generated if the mouse is present in a chamber at trial start.

4.2 Interpretation of results from optogenetics

After modifying the dilutions and concentrations of the viral injections, we were able to observe a more uniform expression of virus over the auditory cortex and inferior colliculus from histology. Furthermore, in the electrophysiological recordings, we observed a more distinct LFP and spiking response from neurons in brains injected with more dilute concentrations of ChRmine virus. The failure of the initial ChRmine injections could have been caused by the death of neurons in the injection site. The initial optogenetic recordings were completed multiple weeks after viral expression had completed in those mouse brains. This was not the case in later injections with more dilute concentrations as better results were generally observed after recording soon after the expression date. As mentioned above, the power of the LEDs used for the initial recordings were also not sufficient for optical stimulation, so it could be possible that there could have been a response from the initial recordings if the LED pulses had more power.

CHAPTER 5. CONCLUSION

5.1 Future Directions

Future directions for this study include improving the closed-loop system to increase precision in manipulation of stimulus. The ambiguity in the results of the experiment will need to be eliminated to ensure that mice only get stimulated in the specified regions. Most of the mice that were tested were injected into auditory cortex, so we can try injecting more mice into the inferior colliculus to broaden the scope of the findings listed. Additionally, further recordings should be made from more ChR2 mice to contrast with results from ChRmine. Once some technicalities with the techniques that we attempted implement in the study are addressed, a long-term goal is to be able to perform closed-loop stimulation of neurons during the sound-seeking task. To do this, Octopilot needs to be interfaced with wireless optogenetic hardware such that it can interact with the real-time closed-loop system. The optogenetic hardware will need to be implanted on the skull of the mouse surgically, before testing it on the task. When all the requisite steps are complete, testing implanted mice during behavioral sessions will provide valuable insights into the real-time effects of optogenetic stimulation on sound localization and movement.




Figure 5.1 - Wireless Logger with LEDs for optogenetic stimulation.

5.2 Final Acknowledgements

All animal experiments were conducted under the guidance and approval of the Emory University Institutional Animal Care and Use Committee under protocol number 202100154. Research reported in this document was supported by the National Institute of Neurological Disorders and Stroke of the National Institutes of Health under Award Number R34NS137017. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The research was also supported by the Whitehall Foundation. We would like to acknowledge Children's Healthcare of Atlanta, Rebecca Wood MA, BS, LVT, RLATg, and Emory University's Animal Physiology Core, RRID:SCR_025835, for performing some of the surgeries. pAAV-CaMKIIa-ChRmine-mScarlet-Kv2.1-WPRE was a gift from Karl Deisseroth (Addgene viral prep #130991-AAV8; <http://n2t.net/addgene:130991> ; RRID:Addgene_130991). pAAV-hSyn-hChR2(H134R)-EYFP was a gift from Karl Deisseroth (Addgene viral prep #26973-AAVrg; <http://n2t.net/addgene:26973> ; RRID:Addgene_26973).

REFERENCES

- [1] D. M. Schneider, J. Sundararajan, and R. Mooney, “A cortical filter that learns to suppress the acoustic consequences of movement,” *Nature*, vol. 561, no. 7723, pp. 391–395, Sep. 2018, doi: 10.1038/s41586-018-0520-5.
- [2] G. Abbruzzese and A. Berardelli, “Sensorimotor integration in movement disorders,” *Mov. Disord.*, vol. 18, no. 3, pp. 231–240, Mar. 2003, doi: 10.1002/mds.10327.
- [3] S. N. Flesher *et al.*, “A brain-computer interface that evokes tactile sensations improves robotic arm control,” *Science*, vol. 372, no. 6544, pp. 831–836, May 2021, doi: 10.1126/science.abd0380.
- [4] D. E. Callan, G. Durantin, and C. Terzibas, “Classification of single-trial auditory events using dry-wireless EEG during real and motion simulated flight,” *Front. Syst. Neurosci.*, vol. 9, Feb. 2015, doi: 10.3389/fnsys.2015.00011.
- [5] J. Mai *et al.*, “Sound-seeking before and after hearing loss in mice,” *Sci. Rep.*, vol. 14, no. 1, p. 19181, Aug. 2024, doi: 10.1038/s41598-024-67577-7.
- [6] J. L. Saunders and M. WEHRmailto, “Autopilot: Automating behavioral experiments with lots of Raspberry Pis”.
- [7] G. Lopes *et al.*, “Bonsai: an event-based framework for processing and controlling data streams,” *Front. Neuroinformatics*, vol. 9, Apr. 2015, doi: 10.3389/fninf.2015.00007.

- [8] T. D. Pereira *et al.*, “SLEAP: A deep learning system for multi-animal pose tracking,” *Nat. Methods*, vol. 19, no. 4, pp. 486–495, Apr. 2022, doi: 10.1038/s41592-022-01426-1.
- [9] K. E. Kishi *et al.*, “Structural basis for channel conduction in the pump-like Channelrhodopsin ChRmine,” *Cell*, vol. 185, no. 4, pp. 672-689.e23, Feb. 2022, doi: 10.1016/j.cell.2022.01.007.
- [10] J. H. Marshel *et al.*, “Cortical layer-specific critical dynamics triggering perception,” *Science*, vol. 365, no. 6453, p. eaaw5202, Aug. 2019, doi: 10.1126/science.aaw5202.
- [11] K. Y. Chan *et al.*, “Engineered AAVs for efficient noninvasive gene delivery to the central and peripheral nervous systems,” *Nat. Neurosci.*, vol. 20, no. 8, pp. 1172–1179, Aug. 2017, doi: 10.1038/nn.4593.
- [12] J. H. Siegle, A. C. López, Y. A. Patel, K. Abramov, S. Ohayon, and J. Voigts, “Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology,” *J. Neural Eng.*, vol. 14, no. 4, p. 045003, Aug. 2017, doi: 10.1088/1741-2552/aa5eea.