# Highlights

**A Semi-Automated, Parallelized, and Flexible Spatial Navigation Task for Rats**

Violet Saathoff, Charles Ian O'Leary, Anna K. Gillespie

- This behavioral apparatus enables the training of four subjects in parallel in a spatial navigation task
- Automated reward delivery reduces experimenter-induced variability and streamlines behavioral training
- Software to track subject performance and advance training parameter through a pre-set curriculum reduces experimenter error and enables robust comparison of learning metrics across cohorts.

# A Semi-Automated, Parallelized, and Flexible Spatial Navigation Task for Rats

Violet Saathoff[a], Charles Ian O'Leary[a], Anna K. Gillespie[a,b,*]

[a] *Department of Neurobiology and Biophysics, University of Washington, Seattle, USA*
[b] *Department of Lab Medicine and Pathology, University of Washington, Seattle, USA*

**Abstract**

Spatial navigation tasks for rodents have been an extremely useful tool for studying cognitive processes for decades, but are often laborious to administer, inefficient, and sensitive to experimenter-induced variability. To address these limitations, we have developed a parallelized maze system which runs four highly flexible maze environments simultaneously with automated reward delivery to reduce experimenter involvement, increase throughput, and streamline data collection. Further, our software tracks subject performance and manages the progression through a pre-set curriculum as subjects learn, reducing experimenter error and ensuring comparable learning curves across cohorts. Equipped with IR beam-break based reward ports, peristaltic pumps for liquid reward delivery, and overhead video tracking, our current maze accessories can be flexibly reconfigured to enable a vast array of behavioral tasks and support a wide range of important experimental questions. A pilot cohort of young (4-8 month) and aged (28-32 month) male rats demonstrates that this task enables the efficient training of a series of increasingly difficult visually cued and memory-dependent spatial navigation tasks which are likely to be sensitive to age-related changes.

*Keywords:* aging, memory, spatial navigation, rats, rodents, automated, high-throughput, parallelized, hippocampus

*Corresponding author. Email: annagill@uw.edu

## 1. Introduction

For decades, scientists have studied the behavior of rodents navigating maze-like environments in order to understand cognitive processes such as memory, inference, fear, sociability, and many more [1, 2, 3, 4]. Spatial navigation tasks for rodents have been particularly powerful tools due to their ethological relevance - a hungry rat will reliably seek food in a complex environment, just as it would in the wild. Spatial navigation tasks can provide a rich and sensitive behavioral readout of neural computations because they often incorporate multiple cognitive components: for instance, a subject might need to identify its location in an environment, recall past locations of food, obstacles, or threats, identify the most valuable location, plan a route towards that goal, execute the movement, and evaluate the outcome of that choice [1]) Multiple types of memory might be involved - declarative, procedural, and/or working memory, depending on the task. Spatial navigation tasks have become the gold standard for characterizing cognitive impairments in the case of neurodegenerative diseases due to their sensitivity and reliability [5, 6, 7]. Clever experimental design can allow researchers to tease apart the different cognitive components of spatial navigation challenges, and by incorporating simultaneous neural recording approaches, we can begin to describe the neural mechanisms that underlie these cognitive processes.

The collection of neural recordings from the hippocampus of behaving rodents enabled the discovery of place cells [8]. These neurons reliably fire specifically in one or a few limited locations within a given environment. Collectively, and in conjunction with other spatially tuned populations in the hippocampus and entorhinal cortex [9], these neurons provide a high-resolution signal corresponding to the location of the subject. Further, the sequences of neural firing that occur as a subject traverses through an environment provide a readout of the spatial aspect of the experience that the subject is having. This readout allows us to trace the neural representation of an initial experience as it first occurs, becomes consolidated and stored as memory, and perhaps later retrieved [9, 10]. While the spatial aspect of experience is particularly convenient to measure experimentally, hippocampal neural activity reflects far more than just spatial location. As a field we may have only scratched the surface of understanding how the hippocampus contributes to the full complexity of spatial navigation behavior, let alone other brain regions. Thus many further experiments, involving creative behavioral tasks, different neural recording modalities, cohorts of different ages, sexes,

and disease conditions, remain a high priority.

However, these types of tasks have some notable limitations which slow progress in this realm. First, spatial navigation mazes for freely moving subjects are typically quite large relative to the size of the subject, with complex maze environments often requiring entire rooms. Since space is often limited, usually only one or a few arenas are available and subjects are run serially, limiting behavioral cohort sizes and making behavioral testing slow and laborious. Second, maze "outcomes" often involve a human experimenter, for example, the delivery of a food reward or the provision of an escape opportunity in an aversive condition. The dependence on a researcher opens up opportunities for variability and error in experimenter behavior, raises the possibility of subjects viewing the experimenter as cue or stressor, and generally reduces precision and scalability of behavioral testing[11]. Third, maze apparati are typically designed for a single behavioral task, limiting the utility of each piece of equipment and the behavioral readouts collected from it. When combined with neurophysiology, the need for additional subject space, tethering, or overhead tracking can place additional constraints on maze environments. Finally, implant costs and complexity often severely limit cohort sizes for neurophysiology experiments. For these particularly labor-intensive and expensive experiments, subjects are often selected for their motivation and aptitude for the task to minimize the risk of poor performance during neural recordings. As a result, these already small cohorts may be biased towards high performers rather than reflecting the extent of natural variability in the population.

We set out to develop a spatial navigation maze apparatus that would address many of these limitations. Specifically, we set out to meet the following six criteria. First, we wanted to be able to run multiple subjects at once in order to increase cohort sizes. Second, we wanted to automate the task in order to minimize experimenter-induced variability [11]. Third, we wanted the apparatus to be flexible enough to run a variety of tasks. Not only would this provide more opportunities to use the same equipment for a range of experimental questions, but it would allow us to design a task curricula with a range of difficulty, progressing from an "easy" task variant to more challenging ones so that we could run structured, systematic behavioral testing across the full range of individual abilities [12]. Fourth, we wanted to design task options that would enable the robust study of the initial learning of a task and be amenable to repeated testing of performance (such as throughout the lifespan or with manipulations of neural activity). Fifth, our

apparatus needed to be compatible with wireless neurophysiological recordings. For hippocampal recordings that will be analyzed in relation to spatial location, this requires overhead position tracking as well as sufficiently wide corridors for an implanted rat. The training curriculum (including initial task acquisition) would ideally take less than 4 weeks, a time window similar to our periods of high quality neural recording signal. Finally, our behavioral task also needed to be highly structured, to enable the collection of many comparable behavioral trials and ensure good spatial coverage of the entire environment for subsequent analysis.

Our resultant maze apparatus is a low-cost, scalable, highly automated, and parallelized setup that can be configured to run a variety of different spatial navigation tasks. Inspired by high-throughput behavioral testing apparati in non-spatial domains [13, 14], our proof-of-principle setup can run four rat subjects simultaneously but has the potential for vertical stacking to achieve higher throughput as needed. The behavioral arena described here consists of an enclosed fan-shaped environment equipped with 7 reward ports that can light up and/or deliver liquid reward when triggered. The arenas are constructed mainly from laser-cut acrylic sheets or 3D printed parts that make construction simple, efficient, and consistent, enabling design of different environment configurations as needed. The logic of the task can be customized to test several different aspects of cognition; we here focus on a task which includes both visually cued and memory-dependent variants. We have developed a python package that automates the advancement through a series of task variants for each subject based on set performance criteria. This highly structured training program ensures that all subjects attain certain criteria before advancing, providing comparable behavioral epochs within which to study neural mechanisms of learning and performance. In addition, the automation of this process greatly reduces the opportunities for experimenter error. As much as possible, our system relies on open-source software packages and hardware designs. All custom designs, components, and software are shared here as a resource, in addition to extensive instructions and part lists, to make this system easy to scale, replicate, and implement.

## 2. Materials and Methods

### *2.1. Design Overview*

The maze apparatus consists of four fully enclosed 6-arm maze environments overlooked by a centralized control station (Fig. 1A, 1B). Each maze

environment consists of a central home area with a single reward port, and 6 radial arms each with a reward port at each arm end (Fig. 1C). Each maze is equipped with an overhead camera, a Pumps and Electronics Rack (PER), and a gutter system to handle liquid reward waste. At the control station, the experimenter can view the overhead video from each maze, control the behavior of the task, and monitor performance metrics in real time.

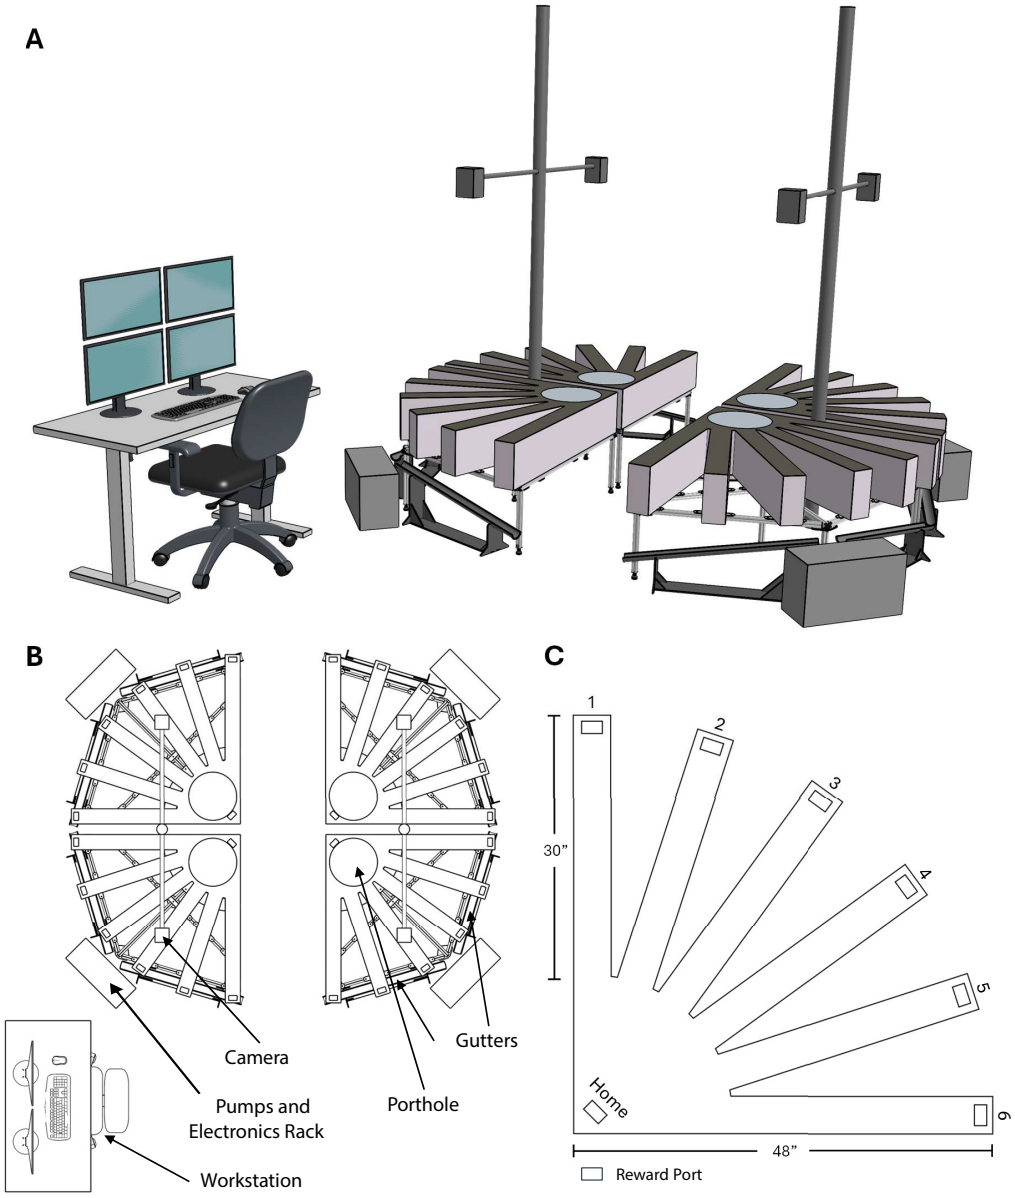


Figure 1: Room Overview:

A 3D view of the entire room. Note that colors/opacity of various components are not true to the actual design, and are instead intended to help with visual clarity.

B Top-down view of the entire room, showing the layout of all four maze environments, as well as the workstation (lower left).

C Top-down schematic view of the maze environment, with reward ports indicated by rectangles.

The maze is laid out as 6 straight arms which radiate away from a central quarter-circle home area, with reward ports at the ends of each of the arms, and an additional reward port located in the corner of the home area. Each arm is 30" long by 4" wide, and the home area has a radius of 18". The environment is fully-enclosed to prevent subjects from escaping, and has a ceiling height of 9". The ceiling over the home area is made of clear 3/16" acrylic, with a 10.5" porthole for putting subjects into the maze, and a clear 1/8" acrylic porthole cover to prevent escapes (Fig. 2A,B). The ceiling over each arm is a removable hexagonally perforated aluminum lid, affixed onto a stainless steel rod frame (Fig. 2C,D). The perforated lids allow for air to enter the environment, introduce only minimal occlusion for overhead camera tracking, and are easily removable to allow the experimenter to access maze arms when needed. All exposed perforations were hemmed to protect both the experimenter and the subject from sharp edges. To secure the lids, one lid end slides under an overhang in the ceiling of the home area, while the outer end has an end cap added. Fully enclosing the maze also ensures that subjects are unable to escape, allowing behavior to run securely without experimenter supervision.

The environment is primarily constructed out of laser cut acrylic which rests on an 8020 extruded aluminum base. Acrylic was selected because it comes in a variety of colors (both transparent and opaque), and because we can manufacture and assemble the parts in house easily and consistently. Where pieces are joined at 90° angles, box joints were used to carefully and reliably control the positions of various components. Box joints between pieces that were both 1/8" or thinner were avoided. The floor of the maze is made out of 3/16" clear acrylic which provides satisfactory rigidity. The transparent floor did not appear to cause any anxiety in rat subjects; they freely explore immediately when placed in the environment and did not show any preference for maze areas with an opaque floor. Maze walls were made from 1/8" acrylic, which provides sufficient rigidity. While most of the walls were made from clear acrylic to create a complex visual environment, the outer walls were made from white acrylic to prevent the subjects from seeing each other or the experimenter. To secure the maze to the base, the main body of the maze was epoxied to a layer of 1/16" acrylic, which was then bolted to the 8020 base (Fig. 2E), additionally, the outer walls of the maze were themselves bolted directly to the 8020 base, providing a very solid connection between the base and the maze. Acrylic pieces were bonded using a clear, slow-cure epoxy (BSI-205), or acrylic cement (Weld-On 3). For

better gluing surfaces, the walls of adjacent arms are joined with a rounded 3D-printed resin corner piece (Fig. 2F).

8020 was selected for the base as it is modular, strong, and easily attainable; easily supporting the maze while giving us options for routing tubing, cables, and drainage. When installing each maze in the behavior room, care was taken to adjust the feet on the 8020 base to distribute the weight of the maze as evenly as possible.

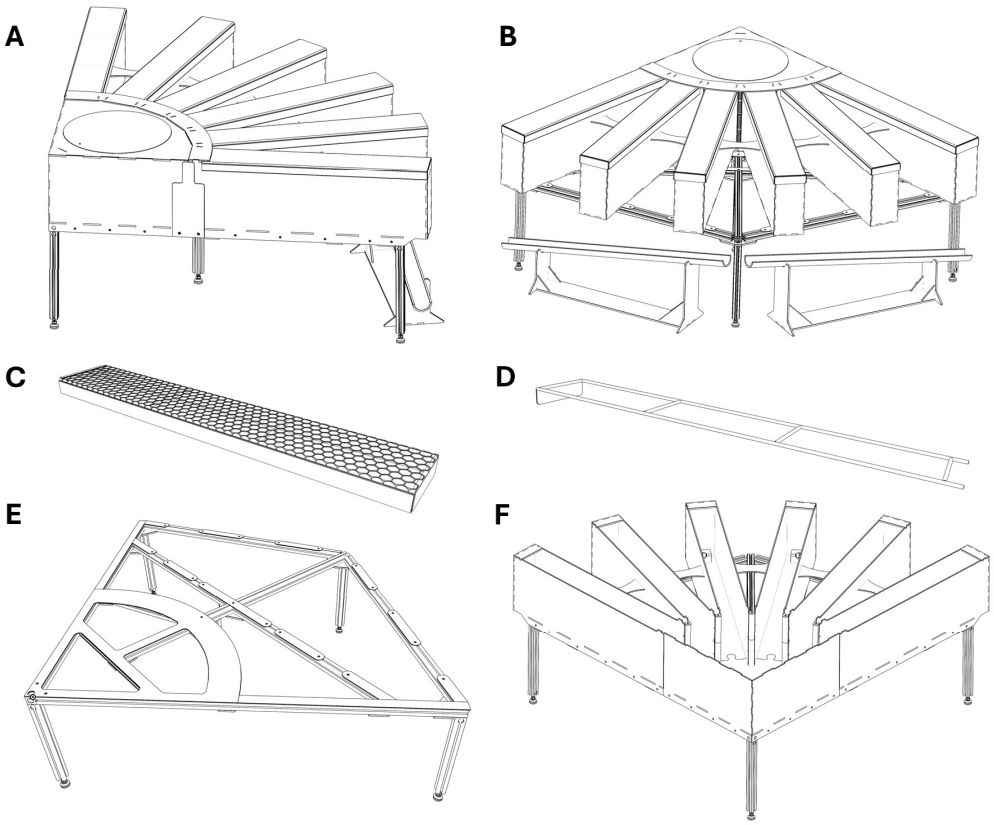


Figure 2: 6-Arm Maze:

- A 3D view of the fully-assembled maze (side).
- B 3D view of the fully-assembled maze (back), with the drainage gutters clearly visible.
- C 3D view of an arm lid. Besides this panel, lids are displayed as being solid when in reality they are made of a hexagonally-perforated aluminum affixed to a stainless steel frame.
- D 3D view of the frame for an arm lid. This frame adds rigidity and robustness to the lids, which helps to prevent the lids from becoming damaged over time due to normal use.
- E 3D view of the base (side). The base is made from 8020 extruded aluminum, and has laser-cut, 1/16" acrylic tabs bolted to it. These tabs are then glued to the floor of the maze to help secure it to the 8020 base.
- F 3D view of the maze (front) with the lids and ceiling removed. Note the 3D-printed corner pieces which join the walls between adjacent arms.

### *2.2. Reward Ports*

The most critical features of the maze are the reward ports, which dispense a liquid food reward (5% sucrose in evaporated milk). Each port con-

sists of an IR photogate (emitter and receiver; 890nm), a white flatmount LED, an 1/8" inner diameter inlet port, a 1/4" inner diameter outlet port, and an RJ-45 port for communication with the ECU (Fig. 3A-D). The reward ports' body and their electronics are based on a design by Demetris Roumis [15]. Compared to the original design, barbed tube fittings were added to make it easier to connect tubing to the reward ports, the drainage port was enlarged to help avoid overflows during cleaning, and the footprint was streamlined to remove nooks and crannies where refuse could easily accumulate. The body of the reward ports were printed on a FormLabs Form 3 printer using clear resin and a custom PCB is populated and installed inside. Printing with clear resin allows the white LED to shine through the translucent body of the reward port, allowing the port to illuminate and serve as an easily visible cue. When a rat breaks the IR beam, a signal is recorded and can trigger the delivery of liquid through the port via the associated peristaltic pump.

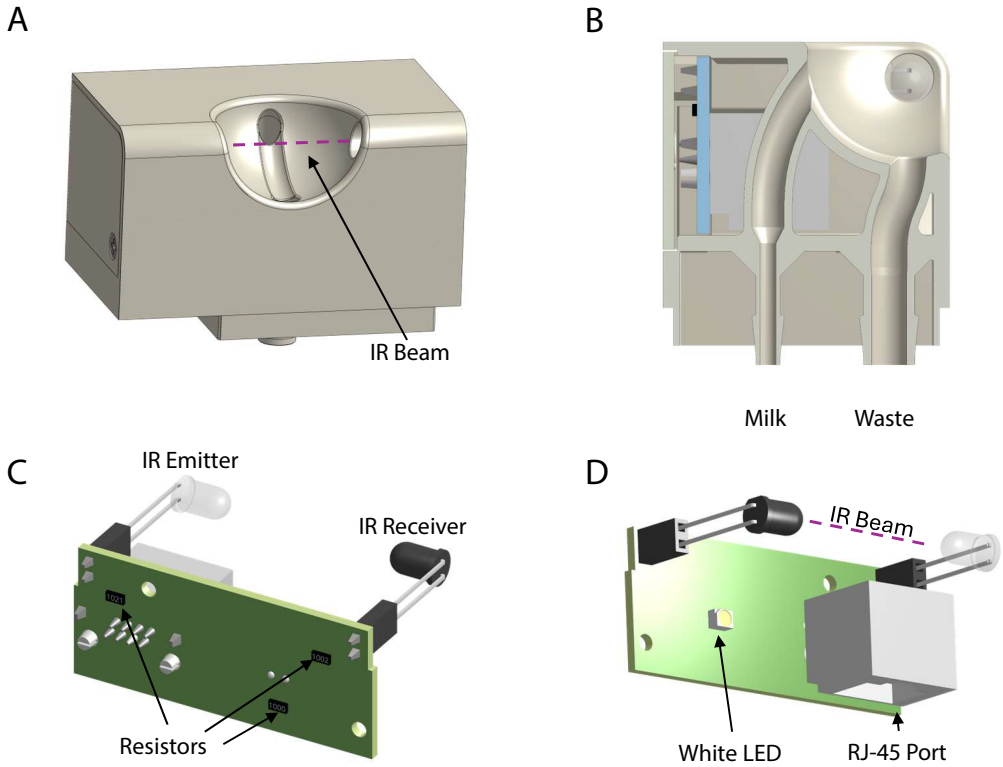


Figure 3: Reward Port:

A Reward port, front view. When the rat's nose enters the cavity, the IR beam is broken, and the port is "triggered". This behavioral event can trigger milk delivery and/or other environmental actions.

B Cutaway view, showing the flow of liquid through the reward port. Milk enters the port through an 1/8" tube, and drains out through a 1/4" tube, both connected to the well with a barb fitting.

C Reward port PCB, back view.

D Reward port PCB, 3D front view. The IR emitter is the clear LED, the IR Receiver is the black LED-shaped component. The RJ-45 port allows the port to be connected to the pump control electronics via an ethernet cable, enabling readout of the IR beam state and control over the state of the white LED.

### *2.3. Pumps and Electronics Rack (PER)*

Each maze includes several key electronic accessories, which we combine into a single organized 4U server rack and refer to as the pumps and electronics rack (PER; Fig. 4A). The PER houses a compact computer (Intel NUC), a router for connecting to the overhead cameras, a pump control/reward

port breakout board, and peristaltic pumps for liquid reward delivery. It also houses a neural data acquisition system and environmental control unit. We chose products from SpikeGadgets (an ECU and either a Main Control Unit or LoggerDock) for their price point and data collection capabilities, but other systems such as OpenEphys and/or Bpod or Bonsai environmental control strategies could be used. Each PER is connected to the central command station via the NUC, which sends a DisplayPort cable and a USB cable to a KVM switch so that a single keyboard/mouse can control any of the maze stations. The PER for each maze is located close to the maze to minimize tubing/cable lengths, and to minimize the tripping hazard that tubing and cables can cause. A custom mounting plate for the pumps and the pump control PCB was designed to mount to the server rack and secure all pumps, mitigating any vibration caused by pump activity.

We use peristaltic pumps to deliver small amounts of liquid reward to reward ports. After extensive testing of many pump models, we selected Boxer 9QX peristaltic pumps. When delivering rewards, the ECU sends a TTL pulse train signal to the pump controller, which then modulates the output of solid-state relays which gate power to each pump. Pumps are powered using a standard 12V/2A power supply, which is routed through a switch to control the pumps' direction. While the pump controller is compatible with any pump that uses a standard DC motor (i.e. not a stepper motor, or a DC brushless motor), the pump plate that integrates the pump controller and the pumps is specifically designed for the Boxer 9QX pumps.

To control the pumps, we designed a pump control PCB (Fig. 4B) which interfaces directly with the SpikeGadgets ECU and operates as a breakout board (all digital outputs are buffered). DIO ports 1-8 are each connected to an RJ-45 port, which we typically connect to the reward ports, but are fairly generic and can be used in a variety of ways. Ports 9-16, meanwhile, control the solid state relays used to gate power to the pumps, with filtering capacitors to protect the board from back EMF. By using standard JST-PH 2 connectors to connect to the peristaltic pumps, this board can be easily used to modulate power to a range of devices as necessary.

Critically, in order to achieve consistent reward delivery of an appropriate amount of milk at an appropriate speed, we designed a variety of pulse trains rather than simply turning the pump on for a precise duration. This modulates the duty cycle, and results in a slower, more consistent milk delivery. For an approximately $30\mu L$ "small" reward we used 3 10ms pulses separated by 20ms (33% duty cycle). For an approximately $130\mu L$ "large"

reward we used 10 10ms pulses separated by 20ms . In addition to increasing consistency, this strategy is preferable to slowing the pumps by modulating their voltage, because it retains the ability to run pumps at full speed during milk loading and cleaning. The pumps for each maze draw from a single 1L reservoir of milk. After daily behavioral training is complete, pump direction is reversed via a switch on the pump control PCB which allows us to collect the milk for overnight storage at 4°C. We then replace the reservoir with cleaning solution, return the pumps to their default direction, flush the tubing until clean, run ethanol through the tubing to sanitize, and then let dry.

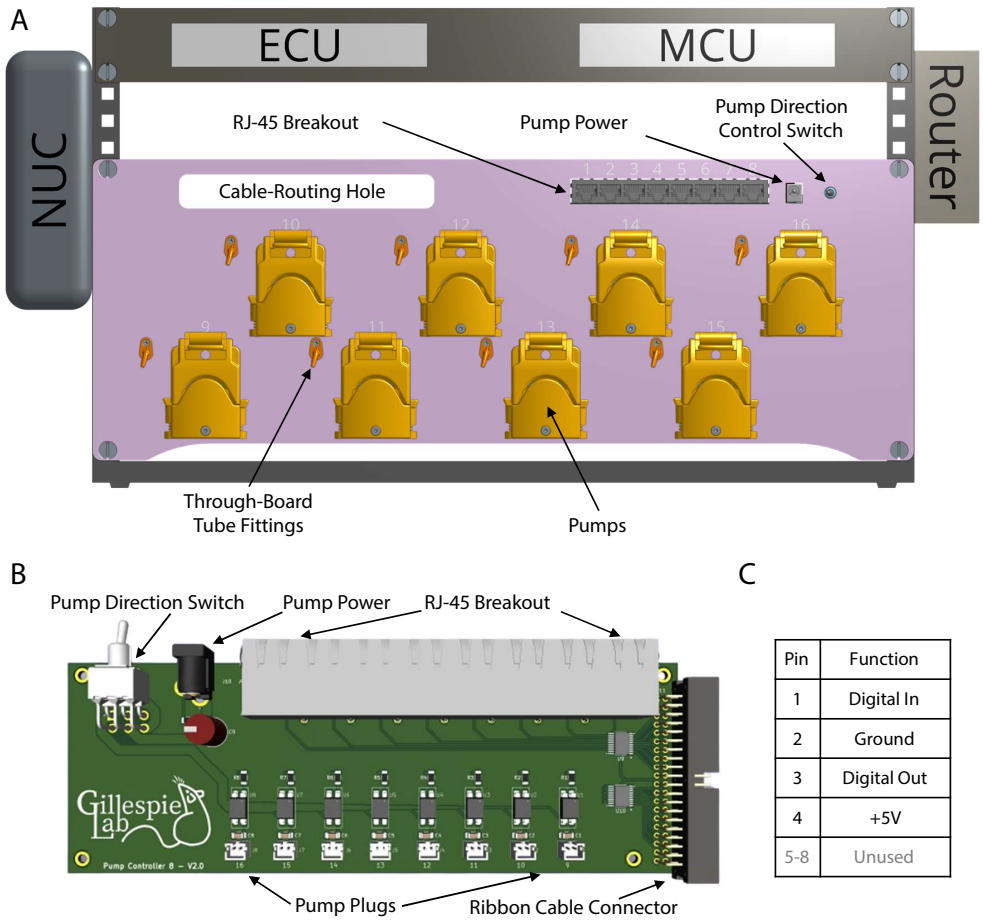


| Pin | Function |
|---|---|
| 1 | Digital In |
| 2 | Ground |
| 3 | Digital Out |
| 4 | +5V |
| 5-8 | Unused |

Figure 4: Pumps and Electronics Rack (PER):

A Front-view of the Pumps and Electronics Rack (PER); a server rack which integrates the computer (NUC), the SpikeGadgets Main Control Unit (MCU), the SpikeGadgets Environmental Control Unit (ECU), and Pump Plate (pink). The Pump Plate, made from 4mm anodized aluminum, can hold up to 8 Boxer 9QX-9007.730 peristaltic pumps, and mounts a Pump Controller PCB. The pump control electronics also include an RJ-45 breakout board which routes reward port and camera sync signals.

B Top-down view of the pump controller/RJ-45 breakout PCB. Each board connects to the ECU via a ribbon cable, and can individually power up to 8 pumps via the JST 2-PH connectors. Each board can also connect to 8 accessories (typically reward ports) via ethernet cables.

C Pinout for each RJ-45 breakout port. When a breakout port is connected to a reward port, pin 1 represents the state of the IR beam, pin 3 controls the output of the white LED, and pin 4 powers both the IR LED, and the IR receiver.

### *2.4. Drainage*

In an attempt to reduce the number of refuse containers, gutters were made so that up to 3 reward ports could share a single waste container (Fig. 2B). Additionally, some containers could be shared between 2 gutters, or between 2 home ports, reducing the total number of refuse containers for all 4 mazes to 8. Gutters were made from 2" ABS tubing, cut lengthwise, and epoxied to a simple acrylic frame. It is critical to ensure that the tubing leading from the reward ports to the gutters maintains a steep angle so that liquid does not pool and clog. Alternatively, longer tubes leading to shared refuse containers could also be used as long as the tubing angle is steep.

## 3. Maze Task and Curriculum

The spatial memory task that we ultimately wanted subjects to perform is a challenging, memory-dependent six-arm win-stay/lose-switch task, related to prior work [16]. Reliable performance is best achieved by using several simpler task variants to gradually shape behavior toward the final task [12], which we refer to as our "curriculum". The curriculum is designed to balance subjects' frustration and complacency as they progress through a series of incremental steps with clearly defined advancement criteria, and learning targets which build upon knowledge from previous steps. The first few steps are visually cued, and intended to teach subjects the trial structure and to only visit lit reward ports. Then, subjects transition to the memory-dependent task phase, during which subjects must remember which one out of a gradually increasing number of lit ports is delivers reward (Figure 5A). Breaking down the learning of the full task into simpler components is not only easier for the rats, but also opens up the possibility of studying the learning process with more structure and specificity - thus both the learning variants and the final task are scientifically valuable. Furthermore, by ensuring that all subjects meet standardized performance thresholds before advancing to the next step of the training process, this approach allows us to quantify measures of performance at each step, giving us valuable insight into the learning process. We focus on our 6-arm task variants here, but many other types of tasks could be implemented, each with their own training curriculum, within this or other maze environments.

### *3.1. Task*

For all 6-arm task variants, the subject must perform structured trials which consist of a nose poke at the home port, then a visit to a single arm of the maze, then a return to the home port. Each trial begins when the subject triggers the illuminated home port and receives a small reward. Next, some number of the outer ports illuminate, and the subject must choose one to visit; receiving a large reward if the arm is the goal arm, and nothing otherwise. After visiting a single arm the subject must return home to begin the next trial. During visually-cued task variants, any illuminated outer ports are goal arms. In contrast, during memory-dependent variants, not all illuminated ports are in goal arms, so subjects must use trial and error over multiple trials to identify the current goal location and remember it for future trials. Goal arm locations are changed in a "block" structure (after 15 goal arm visits). Within each goal block, the "search phase" is the set of trials until the goal arm has been identified, and the "repeat phase" includes all subsequent trials; the search phase includes the first rewarded trial. When the goal block is complete, new goal arm(s) are drawn, and the subject must begin a search phase to find the new goal arm(s). Any deviation from the described trial structure, such as going from one arm port to another, initiates a "lockout", where all wells are de-illuminated/deactivated for a predetermined interval (typically 15s). After the lockout is over, the home well illuminates, and the subject must return home to start the next trial. Depending on the maze parameters, subjects must either complete a set number of trials per epoch, or a set number of goal blocks per epoch. The epoch ends after all trials/goal blocks are completed, or after a set time interval has elapsed (typically 70 minutes from the epoch start). At this time, all wells are automatically de-illuminated/deactivated, and subjects are left to wait for the experimenter to remove them from the maze.

### *3.2. Curriculum*

Before beginning the 6-arm curriculum, subjects are first trained to trigger reward wells in order to receive milk reward on a simple walled linear track. Subjects must first earn >50 rewards in a 45-60 minute training session before advancing to the maze; this typically takes 3-5 training sessions.

The full 6-arm curriculum consists of 10 increasingly difficult steps as outlined in Figure 5B, and requires a minimum of 18 epochs to complete (though subjects are expected to require approximately 30 epochs). Each step in the curriculum is characterized by a few parameters: the number of

goal ports, the number of lit ports, the number of rewards per goal block, the number of trials in the epoch, and the goal rate threshold. The number of goal ports determines how many ports will be chosen as goals each block. The number of lit ports is the total number of reward ports with their white LEDs turned on during the outer arm phase of the trials. These ports are selected randomly once per block alongside the goal port, and exist as all options the subject must explore when searching for the goal (goal ports are always lit). Ports which are lit but unrewarded will never give a reward during the current goal block. The number of rewards per goal block determines how often the goal changes. And lastly, for an epoch to be considered "successful", the subject must complete at least the minimum number of trials (typically 50 trials) while achieving a goal rate (#rewarded_trials / #trials) greater than or equal to the threshold goal rate. At each step in the curriculum, subjects must complete a specified number of successful epochs before moving on to the next step. The final step in the curriculum is the full task, but subjects are not considered fully trained until after they have completed step 10, at which point their performance is assumed to have plateaued, and final behavior data can be collected.

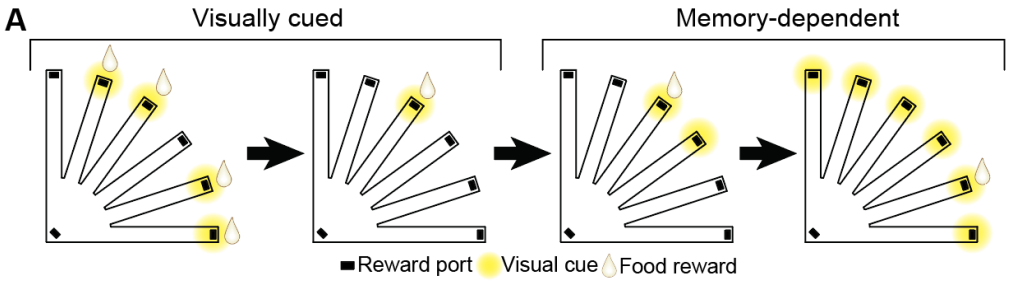


| Step | Reward Port States | Rewards Per Block | Epochs | Target Goal Rate | Learning Target |
|---|---|---|---|---|---|
| 1 | ⋆⋆⋆⋆○○ | 1 | 2 | 0.50 | acclimatize to the maze, and do trials |
| 2 | ⋆⋆⋆○○○ | 1 | 2 | 0.80 | only go to lit ports |
| 3 | ⋆⋆○○○○ | 1 | 2 | 0.80 | incremental practice |
| 4 | ⋆○○○○○ | 15 | 4 | 0.80 | learn to repeat, improve precision |
| 5 | ⋆●○○○○ | 15 | 2 | 0.66 | learn search and repeat by memory |
| 6 | ⋆●●○○○ | 15 | 1 | 0.61 | incremental practice |
| 7 | ⋆●●○○○ | 15 | 2 | 0.59 | learn to switch goals without a visual cue |
| 8 | ⋆●●●○○ | 15 | 1 | 0.57 | incremental practice |
| 9 | ⋆●●●●○ | 15 | 1 | 0.54 | incremental practice |
| 10 | ⋆●●●●● | 15 | 1 | 0.51 | incremental practice/full task |

⋆ Rewarded/Lit (Goal) ● Unrewarded/Lit ○ Unrewarded/Unlit

Figure 5: Training Curriculum

A Visual schematic of training, illustrating steps 1, 4, 5, and 10 of the training curriculum.

B A table giving an overview of the 10-step training curriculum. Outer reward ports exist in one of 3 states: Lit and Rewarded, Lit but Unrewarded, and Unlit and Unrewarded. Each well state pictogram represents how many of each wells are in each state, and does not represent which wells are in which states. In steps 5 and 6, when switching goal blocks, the previous goal well is never lit, giving the subjects a visual cue when it is time to switch goals. In steps 7 through 10, however, the previous goal well is always lit (unrewarded), forcing the subjects to learn to switch goals on their own. Goal rate thresholds are generally calculated as 75% of the rate a perfectly-trained subject would achieve if they were maximally unlucky, minus an additional 0.05 buffer for the memory-based steps.

### *3.3. Curriculum Implementation*

Setting task parameters manually, especially with increasing cohort size and parallel training sessions, is extremely prone to experimenter error, while also being tedious, stressful, and time consuming for the experimenter. To counter this, tracking each subject's progress and appropriately setting their environment's parameters for each epoch was fully automated by storing all parameter sets for the entire curriculum in a text file. In addition to being less error-prone, this strategy enforces uniform and objective environment parameter selection between training stages across subjects, removing any experimenter variability in deciding when to advance subjects to the next task variant.

To implement the curriculum, each subject has a parameter file associated with their name which contains all parameters and persistent state information required to select and run the appropriate step of the curriculum. This file is updated throughout the epoch to ensure that if errors occur, the parameter file is as up-to-date as possible. A copy of the parameter file is also added to the behavior log at the beginning/end of each epoch, which makes the parameter set accessible for downstream data processing. This also preserves a record of which changes were made to the file by the maze code, and which changes were made by the experimenter between epochs (atypical, but occasionally necessary).

### *3.4. Arm Selection*

Due to thigmotaxis, subjects tend to begin behavioral training with an innate bias towards arms 1 and 6. Additionally, as subjects explore the maze and learn the task, they tend to develop arms they prefer, and arms they are averse to. These biases interfere with task learning, confound behavioral analyses, and reduce uniform maze coverage, which is important for any subsequent analysis of neural data with respect to spatial location. To counter these biases, an arm selection algorithm was implemented which tracks the subject's arm visits, and draws from a probability distribution which favors arms that the subject has visited less in recent trials. As subjects improve at the visually cued task, the algorithm encourages them to visit all arms much more evenly. There are parameters which control how much it weights rewarded visits vs unrewarded visits, how much it weights recent data vs older data, and how aggressive it is at pushing subjects towards the less visited arms (Appendix B). The algorithm also incorporates a floor, so that even arms that are relatively unfavored by the algorithm are still chosen as the

goal occasionally, to ensure that subjects do not learn to avoid those arms entirely.

### *3.5. Data Outputs*

For post-hoc analysis, we need reliable records of all behavioral events, such as port triggers, port illumination, and reward delivery. Further, we need a system that is resilient to common experimenter errors such as typos in filenames and recording mistakes (either missing pokes, or including pokes triggered by the experimenter when removing the subject from the maze, or while cleaning). We developed a system of multiple logs, each with varying levels of detail, which are initialized at the start of each epoch and which automatically close themselves at the end of the epoch.

The first is a highly detailed log containing the maze parameters, initialization statements, and all maze events. Furthermore, it is preemptively segmented into goal blocks to make data processing as easy as possible. The second log is much more compact, containing one line for the maze parameters and a table of trials. Each trial contains the data for the home poke, the outer arm poke, goal block details, and the details for any lockout pokes. The data recorded for each poke includes the time and duration of the IR beam break, which port was triggered, and whether or not a reward was delivered. While not as detailed as the full log, this format is extremely easy to process, and requires no detective work to assemble the trial landmarks, making it very convenient for data processing.

Additionally, each subject's arm visits are plotted live at the control station, meaning that the experimenter can monitor all subjects' progress simultaneously. This both allows the experimenter to gain a deeper understanding of how subjects are performing the task without having to constantly watch the subjects and makes it easier to quickly detect any maze malfunctions.

## 4. Pilot Cohort

### *4.1. Experimental Model and Subject Details*

To evaluate the automated, parallelized maze apparatus, we used it to train a cohort of 8 male Fischer x Brown Norway F1 rats from the NIA colony (Charles River): 4 young (4-9 months) and 4 old (28-32 months). All animals were kept on a 12 hr light-dark cycle (lights on 7am-7pm) with ad libitum access to food (standard rat chow) in a temperature and humidity controlled facility. Beginning one week prior to training, rats were single-housed, and

food restricted to a maximum of 85% of free-feeding weight. All subjects were trained using the methods and environment described above. All procedures were approved by the Institutional Animal Care and Use Committee at the University of Washington, Seattle (Protocol #4566-01).

*4.2. Results*

Data was collected from 8 rats in the 6-arm mazes over 552 epochs (368 at training stages, 184 at the final task variant) and 32 days. For the first 23 training days, subjects completed 2 epochs per day, each requiring a set number of trials. The final 9 training days included 3 epochs per day, each requiring two complete goal blocks. Subjects were each assigned to a specific maze, with one young rat and one old rat in each maze. Subjects were also divided into two subgroups of 4 each, each containing two young rats and two old rats; the subgroup which ran first in the day was alternated each day.

Subjects were trained using the curriculum shown in Figure 5B, with extremely minor differences in some performance criteria and completed epochs required, as detailed in Appendix A. Subjects showed a rapid increase in reward rate during the visually cued epochs, all reaching performance levels well above chance , but also demonstrating individual variability in learning rate and ultimate performance level (Figure 6A). Subjects learned the memory-dependent version of the task much more slowly (Figure 6B), particularly steps 5 and 6, but ultimately all reached performance rates well above chance. Once subjects reached criterion at the final stage of the task (step 10), we saw no significant difference between age groups in the reward rate during the repeat phase of each goal block (Figure 6F) or in the length of the search phase (Figure 6E), indicative of search efficiency. However, the amount of training required to reach criterion for each step was different between age groups, particularly for the first two steps of the memory-dependent task variant Figure 6C). Our most drastic difference between age groups was in trial duration (Figure 6D), as aged subjects moved far slower than young subjects. To ensure that all behavioral testing could be completed daily in reasonable time, we did enforce maximum epoch lengths (typically 70 minutes) but subjects generally completed the designated number of trials or goal blocks and only rarely timed out. Two of our aged subjects failed to reach the set criterion for memory-dependent step 8 after many epochs (10 epochs for O1, and 8 epochs for O3); note that their data is not included in Figure 6C. Closer evaluation of their performance revealed that they were

performing well above chance, and nearly at criterion, suggesting not that they failed to learn but that the criterion was set unreasonably high. We automatically advanced these two subjects to step 10 after a single epoch at step 9 and their final performance subsequently was not different than the other aged subjects (Figure 6E,F, further confirming that the issue was one of parameter choice rather than a true performance difference, and has informed the improved parameters described in Figure 5B.

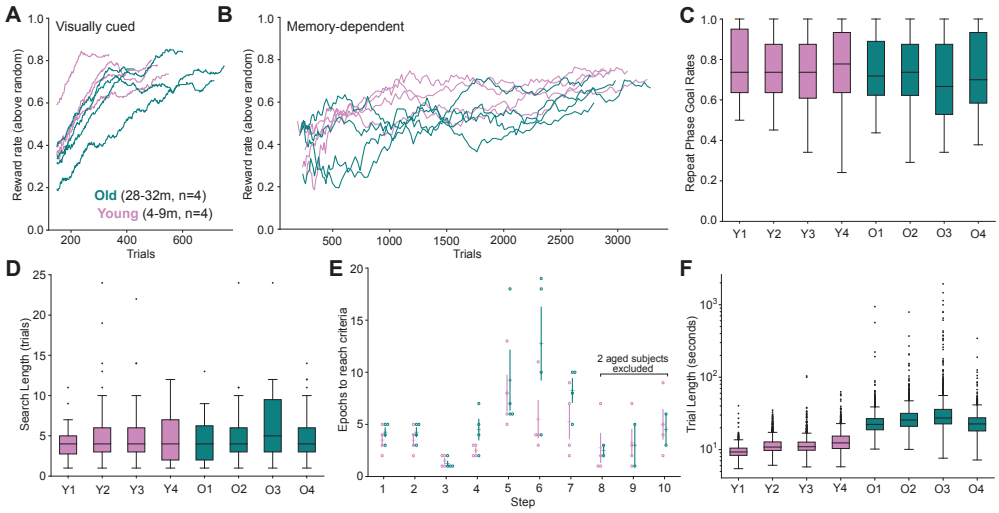


Figure 6: Pilot Cohort Results:

A Learning curves by subject for the visually cued task variants. This plot shows the reward rate for each subject minus the reward rate a random agent would achieve, linearly-normalized such that 0 is the reward rate of a random agent, and 1 is the reward rate a would achieve with a 100% goal rate. Data was smoothed with a 150 trial moving average. $n_{trials} = [Y1 : 510, Y2 : 550, Y3 : 336, Y4 : 490, O1 : 600, O2 : 424, O3 : 713, O4 : 750]$.

B Learning curves by subject for the memory dependent task variants, also above the random rate. An average reward rate was calculated for each goal block and then smoothed with a moving average of 10 goal blocks (this eliminates oscillation due to the search phase/repeat phase structure of the task). $n_{trials} = [Y1 : 3303, Y2 : 3096, Y3 : 3498, Y4 : 3172, O1 : 3206, O2 : 3537, O3 : 3044, O4 : 2990]$ (includes plateaued behavior).

C Average reward rate after finding the goal (repeat phase) during "fully trained" memory-dependent epochs, after criterion was met for step 10. $LMStatistic = 2.45, p = 0.12$, $n_{blocks} = [Y1 : 36, Y2 : 70, Y3 : 54, Y4 : 68, O1 : 32, O2 : 35, O3 : 31, O4 : 39]$ (LM statistic and p value were calculated using Powell's Method, instead of the default BFGS optimization method for the Mixed Effects Linear Regression).

D Average length of search phase per goal block during "fully trained" memory-dependent epochs, after criterion was met for step 10. $LMStatistic = 0.879, p = 0.35$, $n_{blocks} = [Y1 : 36, Y2 : 70, Y3 : 54, Y4 : 68, O1 : 32, O2 : 35, O3 : 31, O4 : 39]$.

E Average number of training epochs required to meet progression criteria for each step. Ranksum results by step (Appendix E) do not demonstrate a significant effect of age on learning rates for this task, with this curriculum.

F Average length of time for each trial (time from initial home poke to the home poke initiating the next trial). $LMStatistic = 36.8, p = 1.3 \times 10^{-9}$, $n_{trials} = [Y1 : 848, Y2 : 1711, Y3 : 1358, Y4 : 1672, O1 : 781, O2 : 960, O3 : 901, O4 : 1010]$.

## 5. Discussion

The results from the pilot cohort demonstrate that the mazes perform as expected, achieving our major goals of training a cohort of rats with minimal experimenter interference in a relatively short amount of time. In contrast to prior training regimes for a similar task [16], our apparatus allowed us to train twice as many rats in about half the time. Further, our training process now allows us to capture different learning phases in a regimented and comparable manner. The gradual advancement through increasingly complex steps allowed more rats to continue participating, rather than being dropped from the cohort due to frustration, poor motivation, or low performance. While aged subjects were certainly capable of learning and performing the final memory-dependent task variant, systematic differences in movement speed and learning rate suggest that this an appropriate task in which to study the impact of age on memory processes. Additionally, with experimenter-induced variability minimized, the individual variability remaining in both age groups further suggests that this task may be well suited to explore the neural basis of individual variability in cognitive performance.

Two aspects of the behavioral analysis were particularly illuminating. First, the dramatic difference in movement speed, and thus trial length, between the age groups was striking. This underscored the importance of having each epoch be limited not by time, but by trial or goal block criteria. This ensured that the aged and young subjects were getting more similar amounts of experience with each epoch and will be essential for any studies of aged subjects. It was also encouraging that, given sufficient time, aged subjects were motivated and able to engage with the task similarly to the young subjects. Second, the implementation of such a gradual curriculum was extremely effective, as indicated by the lack of any stark performance dropoff caused by progression to a subsequent task step and the continued engagement of every subject throughout the full curriculum. Further, the inclusion of learning both a visually cued and memory-dependent task separates the learning of distinct task rules in a way that will be useful for investigating the neural correlates of this learning in future.

### *5.1. Limitations*

While very infrequent, hardware errors and experimenter errors did still occasionally occur. Hardware errors tended to be reward port failure due to beam break sensors failing or leaking tubing causing a failure to deliver

reward. These issues were difficult to catch by inspection of the apparatus, but became readily apparent during behavior as subjects would rapidly learn to avoid ports that never delivered reward and these biases were clearly visible on the live performance plots. Watching for these issues required attentiveness from the experimenter, but once caught were promptly fixed, and epochs were re-done as necessary (dropping any data which was adversely affected). Experimenter error tended to be minor, and extra epochs were performed to replace impacted data.

A second limitation that our pilot cohort revealed was the importance of setting progression criteria that were achievable by all subjects, yet also require sufficient learning of the current task variant. These parameters were initially set too high, hindering subjects' progression through the curriculum, but have now been revised to improved values. Setting appropriate criteria is key for efficient training without sacrificing the mastery of each step. Careful consideration should be given when choosing progression critera and pilot training cohorts may be necessary to inform performance expectations if using the maze setup for other tasks.

### *5.2. Future Expansion*

With the current apparatus as is, cohorts of female subjects, models of neurodegenerative disease, and longitudinal behavioral assessment of cohorts across the lifespan will all be fascinating comparators to our existing pilot behavioral results. The apparatus is also fully compatible with wireless neural recordings and overhead position tracking, enabling the assessment of neural activity patterns that may underlie the learning and memory processes that we set out to study. In addition, the flexible nature of the task logic can allow for many different tasks to be implemented in the existing maze environment, such as sequence learning tasks, alternation tasks [17, 18], or even pure foraging tasks.

However, while the current apparatus substantially extends our ability to train and collect data on larger cohorts of animals, it is still far from true high-throughput behavior. The biggest limitation to higher throughput is the requirement for overhead camera based tracking; were that unnecessary (either by switching to an alternative method of position tracking, or by eliminating position tracking all together), then it would be feasible to stack 3-4 mazes vertically. Some adaptation of current drainage and tubing management strategies would also need to be updated, but all the task control and logging software would seamlessly extend to this use case. This

improvement could allow comparable throughput to other highly successful behavioral training approaches [13, 14], and would be the first to extend this high-throughput approach into the spatial navigation realm.

## 6. Author Credit Statement

A.K.G. conceived the study, acquired funding, and supervised the project. V.T.S. and C.I.O. designed the experimental apparatus. V.T.S., C.I.O., and A.K.G. built the experimental apparatus. V.T.S. wrote behavioral and analysis software and collected data. A.K.G. and V.T.S. designed the curriculum. V.T.S. analyzed the data with input from A.K.G. V.T.S. and A.K.G. prepared data visualizations and wrote the manuscript with input from all authors.

## 7. Acknowledgements

The authors are grateful to all Gillespie Lab members for their help constructing the maze apparatus and their insightful feedback on the manuscript. They are particularly grateful to Zhana Prince for her assistance during the final stages of data collection. This work was supported by the NIH Brain Initiative/NIA (grant R00AG068425 to A.K.G.), the Simons Collaboration on Plasticity and the Aging Brain (grant 1169066 to A.K.G.), and the Esther A. & Joseph Klingenstein Fund (Fellowship Award in Neuroscience to A.K.G.).

## Appendix A. Curriculum for Pilot Cohort

### *Appendix A.1. Curriculum Table and Discussion*

The pilot cohort was not trained using the curriculum as presented in Figure 5B, instead using a similar prototype curriculum (Table A.1). This version of the curriculum had two main flaws. First, the subjects did not generally spend enough time at step 4 to maintain a strong preference for lit wells during early memory training. Second, the thresholds for success during memory were too high, causing subjects to plateau below the threshold, adding several, likely unnecessary, days to training. While more pilot cohorts are required to validate/perfect the curriculum, we hope that the version in Figure 5B will train rats significantly faster than the pilot cohort.

| Step | Reward Port States | Rewards Per Block | Epochs | Target Goal Rate | Learning Target |
|---|---|---|---|---|---|
| 1 | ⋆⋆⋆⋆○○ | 1 | 2 | 0.50 | acclimatize to the maze, and do trials |
| 2 | ⋆⋆⋆○○○ | 1 | 2 | 0.80 | only go to lit ports |
| 3 | ⋆⋆○○○○ | 1 | 2 | 0.80 | incremental practice |
| 4 | ⋆○○○○○ | 15 | 1 | 0.80 | learn to repeat, improve precision |
| 5 | ⋆●○○○○ | 15 | 2 | 0.70 | learn search and repeat by memory |
| 6 | ⋆●●○○○ | 15 | 1 | 0.66 | incremental practice |
| 7 | ⋆●●○○○ | 15 | 2 | 0.63 | learn to switch goals without a visual cue |
| 8 | ⋆●●●○○ | 15 | 1 | 0.60 | incremental practice |
| 9 | ⋆●●●●○ | 15 | 1 | 0.58 | incremental practice |
| 10 | ⋆●●●●● | 15 | 1 | 0.55 | incremental practice/full task |

⋆ Rewarded/Lit (Goal) ● Unrewarded/Lit ○ Unrewarded/Unlit

Table A.1: Training Curriculum. A table giving an overview of the 10-step training curriculum used for the pilot cohort.

*Appendix A.2. Threshold Calculation*

Goal rate thresholds for the pilot cohort were calculated as 80% what a perfect rat would achieve if they were maximally unlucky (i.e. they had to search every lit well in order to find the goal). Then, an additional 5% was subtracted to give the rats a buffer.

| Step | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|
| Optimal Rate | 0.94 | 0.88 | 0.88 | 0.83 | 0.79 | 0.75 |
| 80% Rate | 0.75 | 0.71 | 0.71 | 0.67 | 0.63 | 0.60 |
| Threshold | 0.70 | 0.66 | 0.63 | 0.60 | 0.58 | 0.55 |

Table A.2: Goal Rate Threshold Calculations. A table detailing the basic goal rate threshold calculations for the pilot cohort's training curriculum. Steps 7 and 8 were adjusted to account for the expected perseveration.

## Appendix B. Arm Selection Algorithm

For most of training, the following algorithm was used to track subject's arm visits, and select which arms to use as the goal arms. The premise is to maintain a value for each arm which quantifies how much the subject has visited that arm recently compared to other arms; these value are updated every time the subject visits as well. Then, when it is time to select a goal, the array of values is passed though an algorithm to generate a probability distribution where the most visited arms are the least likely to be selected,

and the least visited arms are the most likely to be selected. Goals are then drawn from this distribution as needed.

There are four parameters which control the algorithm:

- $\alpha$ : weights between rewarded visits and unrewarded visits. $\alpha \in (0, 1)$; should be close to 0.5
- $\beta$ : a "temperature" parameter which controls how aggressively the algorithm will push subjects to explore all arms evenly. $\beta < 0$; shouldn't be too large (i.e. $\beta > -2$ is recommended
- $\delta$ : a floor on how unlikely any particular arm can become - this helps prevent subjects from learning to never visit an arm that used to be their favorite, creating a new arm they're averse to. $\delta \in [0, 1/6)$; should be small (i.e. $0 < \delta < 0.03$).
- $\gamma$ : a decay parameter which causes more recent data to be more important than older data. This decay is applied approximately once per epoch. $\gamma \in (0, 1)$; should be bigger than 0.5

Two of these parameters have derived quantities, which are useful in practice:

- $\delta' = \frac{\delta}{1-6\delta}$
- $\gamma' = \gamma^{\frac{1}{expected_trials}}$

*Appendix B.1. Weighted Visits Update Function*

*well* is the index of the well the subject just poked, and *reward* is either a 0 or 1 based on whether the subject received a reward. All vector operations are element-wise, and the $\overrightarrow{weighted_visits}$ array is updated every time a subject visits an outer arm.

$$\overrightarrow{weighted_visits}[well] \longrightarrow \overrightarrow{weighted_visits}[well] + [\alpha \times reward + (1-\alpha)(1-reward)]\gamma$$
$$\overrightarrow{weighted_visits} \longrightarrow \overrightarrow{weighted_visits} \times \gamma'$$

*Appendix B.2. Probability Distribution Function*

Arms are chosen at random based on a probability distribution from the $\overrightarrow{weighted_visits}$ array. The probability distribution is re-computed using the formula below every time new goals need to be selected (i.e. at the start of every goal block).

$$
\begin{aligned}
softmax(\vec{X}, \beta) =& \frac{e^{\beta \vec{X}}}{\sum e^{\beta \vec{X}}} \\
\vec{Q} =& softmax(\overrightarrow{weighted_visits}, \beta) \\
\vec{P} =& \frac{\vec{Q} + \delta'}{1 + 6\delta'}
\end{aligned}
$$

*Appendix B.3. Goal Selection Notes*

- When only 1 goal is being selected (steps 4+ of the curriculum), care is taken to ensure that no arm is ever selected as the goal arm twice in a row.

- When subjects are searching for a single goal arm among several lit options, the goal selection algorithm presented above was found to increase frustration more than it helped minimize subjects' search preferences. So, after step 6 of the curriculum, goals are selected at random from a distribution which is forced to be flat. To do this, the next 6 goals are selected at a time such that each arm will be chosen once over the next 6 goal blocks. This information is stored in the parameter file to ensure that this arm selection persists between epochs. This is necessary because the number of goals drawn for each subject is relatively small, so random fluctuations in goal selection can introduce biases into the subjects' behavior.

## Appendix C. StateScript Interface

In developing both this environment and other environments in the past, one common source of errors has been StateScript. Communication between python and StateScript in particular has caused numerous bugs. To address this, a StateScriptInterface was developed in python which defines wrappers for all of the StateScript built-in functions, as well as a class for managing DIO ports, and a static class for tracking input/output states as reported by

the ECU. Additionally, StateScriptInterface provides a __callback__ function which supports an "⟨timestamp⟩ ⟨opcode⟩ [⟨value⟩]" to easily send callback data to the appropriate handler.

## Appendix D. Repositories

- Maze Code:
  - Main Scripts: GitHub
  - Modules: GitHub
- Maze Design:
  - Maze CAD: OnShape
  - Laser Cutter Drawings: GitHub
  - Reward Well CAD: OnShape
  - Pump Plate CAD: OnShape
  - Pump Controller PCB: GitHub
  - Bill of Materials: Google Drive

## Appendix E. Figure 6.C Ranksum Results

| step | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| $n_{young}$ | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 |
| $n_{old}$ | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 2 | 2 | 2 |
| ranksum statistic | -0.866 | -0.866 | 0.577 | -1.44 | -0.289 | -1.44 | -1.44 | -0.694 | 0.231 | 0.00 |
| p value | 0.39 | 0.38 | 0.56 | 0.15 | 0.77 | 0.15 | 0.15 | 0.49 | 0.82 | 1.0 |

## References


[1] K. Wijnen, L. Genzel, J. van der Meij, Rodent maze studies: from following simple rules to complex map learning, Brain Structure & Function 229 (4) (2024) 823–841. doi:10.1007/s00429-024-02771-x.

[2] R. G. M. Morris, Spatial localization does not require the presence of local cues, Learning and Motivation 12 (2) (1981) 239–260. doi:10.1016/0023-9690(81)90020-5.
URL https://www.sciencedirect.com/science/article/pii/0023969081900205

[3] D. S. Olton, R. J. Samuelson, Remembrance of places passed: Spatial memory in rats, Journal of Experimental Psychology: Animal Behavior Processes 2 (2) (1976) 97–116, place: US Publisher: American Psychological Association. doi:10.1037/0097-7403.2.2.97.

[4] E. C. Tolman, Cognitive maps in rats and men, Psychological Review 55 (4) (1948) 189–208, place: US Publisher: American Psychological Association. doi:10.1037/h0061626.

[5] R. D'Hooge, P. P. De Deyn, Applications of the Morris water maze in the study of learning and memory, Brain Research Reviews 36 (1) (2001) 60–90. doi:10.1016/S0165-0173(01)00067-4.
URL https://www.sciencedirect.com/science/article/pii/S0165017301000674

[6] C. V. Vorhees, M. T. Williams, Value of water mazes for assessing spatial and egocentric learning and memory in rodent basic research and regulatory studies, Neurotoxicology and Teratology 45 (2014) 75–90. doi:10.1016/j.ntt.2014.07.003.

[7] R. S. Astur, L. B. Taylor, A. N. Mamelak, L. Philpott, R. J. Sutherland, Humans with hippocampus damage display severe spatial memory impairments in a virtual Morris water task, Behavioural Brain Research 132 (1) (2002) 77–84. doi:10.1016/s0166-4328(01)00399-0.

[8] J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat, Brain Research 34 (1) (1971) 171–175. doi:10.1016/0006-8993(71)90358-1.
URL https://linkinghub.elsevier.com/retrieve/pii/0006899371903581

[9] M.-B. Moser, D. C. Rowland, E. I. Moser, Place Cells, Grid Cells, and Memory, Cold Spring Harbor Perspectives in Biology 7 (2) (2015) a021808. doi:10.1101/cshperspect.a021808.
URL http://cshperspectives.cshlp.org/lookup/doi/10.1101/cshperspect.a021808

[10] G. Findlay, G. Tononi, C. Cirelli, The evolving view of replay and its functions in wake and sleep, SLEEP Advances 1 (1) (2020) zpab002. doi:10.1093/sleepadvances/zpab002.

URL https://academic.oup.com/sleepadvances/article/doi/10.1093/sleepadvances/zpab002/6124581

[11] E. Y. Pioli, B. N. Gaskill, G. Gilmour, M. D. Tricklebank, S. L. Dix, D. Bannerman, J. P. Garner, An automated maze task for assessing hippocampus-sensitive memory in mice, Behavioural Brain Research 261 (2014) 249–257. doi:10.1016/j.bbr.2013.12.009.
URL https://www.sciencedirect.com/science/article/pii/S0166432813007523

[12] E. Holleman, J. Maka, T. Schröder, F. Battaglia, An incremental training method with automated, extendable maze for training spatial behavioral tasks in rodents, Scientific Reports 9 (1) (2019) 12589. doi:10.1038/s41598-019-48965-w.
URL https://doi.org/10.1038/s41598-019-48965-w

[13] Y. Hao, A. M. Thomas, N. Li, Fully autonomous mouse behavioral and optogenetic experiments in home-cage, eLife 10 (2021) e66112, publisher: eLife Sciences Publications, Ltd. doi:10.7554/eLife.66112.
URL https://doi.org/10.7554/eLife.66112

[14] R. Poddar, R. Kawai, B. P. Ölveczky, A Fully Automated High-Throughput Training System for Rodents, PLOS ONE 8 (12) (2013) e83171, publisher: Public Library of Science. doi:10.1371/journal.pone.0083171.
URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0083171

[15] D. Roumis, Rhythmic Action Synchronizes Memory Replay During Reinforcement Learning, Ph.D. thesis, UCSF (2020).
URL https://escholarship.org/uc/item/10s205q3

[16] A. K. Gillespie, D. A. Astudillo Maya, E. L. Denovellis, D. F. Liu, D. B. Kastner, M. E. Coulter, D. K. Roumis, U. T. Eden, L. M. Frank, Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice, Neuron 109 (19) (2021) 3149–3163.e6. doi:10.1016/j.neuron.2021.07.029.

[17] L. M. Frank, E. N. Brown, M. Wilson, Trajectory Encoding in the Hippocampus and Entorhinal Cortex, Neuron 27 (1) (2000) 169–178,

publisher: Elsevier. doi:10.1016/S0896-6273(00)00018-0.
URL https://www.cell.com/neuron/abstract/S0896-6273(00)00018-0

[18] D. B. Kastner, A. K. Gillespie, P. Dayan, L. M. Frank, Memory Alone Does Not Account for the Way Rats Learn a Simple Spatial Alternation Task, The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 40 (38) (2020) 7311–7317. doi:10.1523/JNEUROSCI.0972-20.2020.