A Parallelized, Automated System for Higher-Throughput Spatial Navigation Testing of Rats

Violet T. Saathoff[a], Charles I. 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 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, enabling reduced experimenter involvement, increased throughput, and streamlined 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 experimental questions. A pilot cohort of young (4-8 month) and aged (28-32 month) male rats demonstrates

*Corresponding author. Email: annagill@uw.edu

that this apparatus enables the efficient training of a series of increasingly difficult visually cued and memory-dependent spatial navigation task variants across a wide range of performance abilities.
Keywords: aging, memory, learning, neurophysiology, electrophysiology, hippocampus, spatial navigation, rats, rodents, automated, parallelized

---

Significance Statement

To increase the efficiency and reliability of spatial navigation task administration, we have designed a behavioral apparatus which enables the training of four subjects in parallel on a variety of tasks with automated reward delivery. Further, we developed software to track subject performance and advance training parameters through a pre-set curriculum to reduce experimenter error and enable robust comparison of learning metrics across cohorts.

1. Introduction

For decades, scientists have studied the behavior of rodents navigating maze-like environments in order to understand diverse cognitive processes such as memory, decision-making, inference, fear, regret, sociability, and many more [1, 2, 3, 4, 5]. 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]. Depending on the task, this might require using multiple types of memory - declarative, procedural, and/or working memory. Due to their sensitivity and reliability, spatial navigation tasks have become the gold standard for characterizing cognitive impairments in the case of neurodegenerative and other disease models [6, 7, 8]. Clever experimental design can allow researchers to tease apart the different cognitive components of spatial navigation challenges, and by incorporating simultaneous neural recording approaches, researchers have been able to begin describing the neural mechanisms that underlie these cognitive processes. Perhaps the most well known finding stemming from this approach was the discovery of place cells - spatially tuned neurons in the hippocampus[9] - that has since inspired decades of new research into the neural activity underlying the encoding of spatial experiences and their subsequent storage as memory [10, 11]. However, there is still much that we don't understand about how the brain enables spatial navigation behavior. Thus, many further experiments, involving creative task variants, different neural recording modalities, cohorts of different ages, sexes, and disease conditions, remain a high priority.

However, spatial navigation tasks have some notable limitations which can slow progress. They are typically large and often designed for a single

behavioral task, limiting the utility of each piece of equipment. Subjects are usually run serially, limiting behavioral cohort sizes and making behavioral testing slow and laborious. A human experimenter is often an essential part of the training/experimentation process (for example, delivering a food reward or providing an escape opportunity in an aversive condition), adding an experimental variable which is hard to track and which limits the scalability of behavioral testing [12, 13]. When combined with neurophysiology, the need for additional subject space, tethering, or overhead tracking can place additional constraints on maze environments. Finally, because these experiments can be so costly and time-consuming, subjects are often selected for their motivation and aptitude for the task to maximize the chance of successful data collection. As a result, these 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 five 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 [12, 13, 14, 15]. Third, we wanted the apparatus to be flexible enough to run a variety of tasks[16]. 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 [17]. Fourth, our apparatus needed to be compatible with wireless neurophysiological recordings. For hippocampal recordings to be analyzed in relation to spatial location, this required 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 [18, 19], our proof-of-principle apparatus can run four rat subjects simultaneously. The behavioral arenas described here are designed to be simple, efficient, and consistent to construct. The logic of the task can be easily customized; 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, automated training program reduces the opportunities for experimenter error and ensures that all subjects attain certain criteria before

advancing, providing comparable behavioral epochs within which to study neural mechanisms of learning and performance. 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 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.

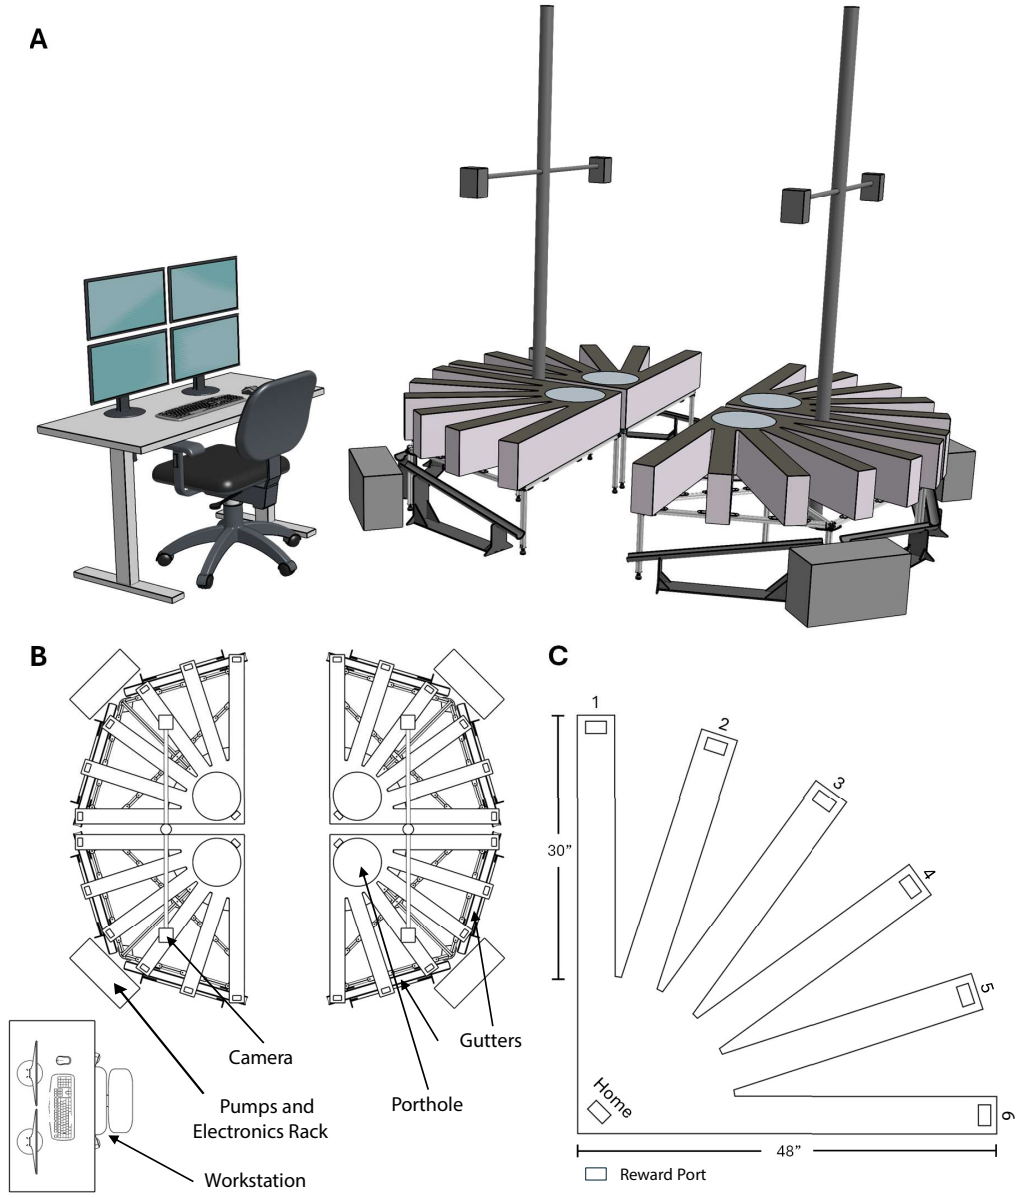

Figure 1: Room Overview:

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.

2.2. Maze Construction

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.

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.

For the maze base, we selected 8020 extruded aluminum framing 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. 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).

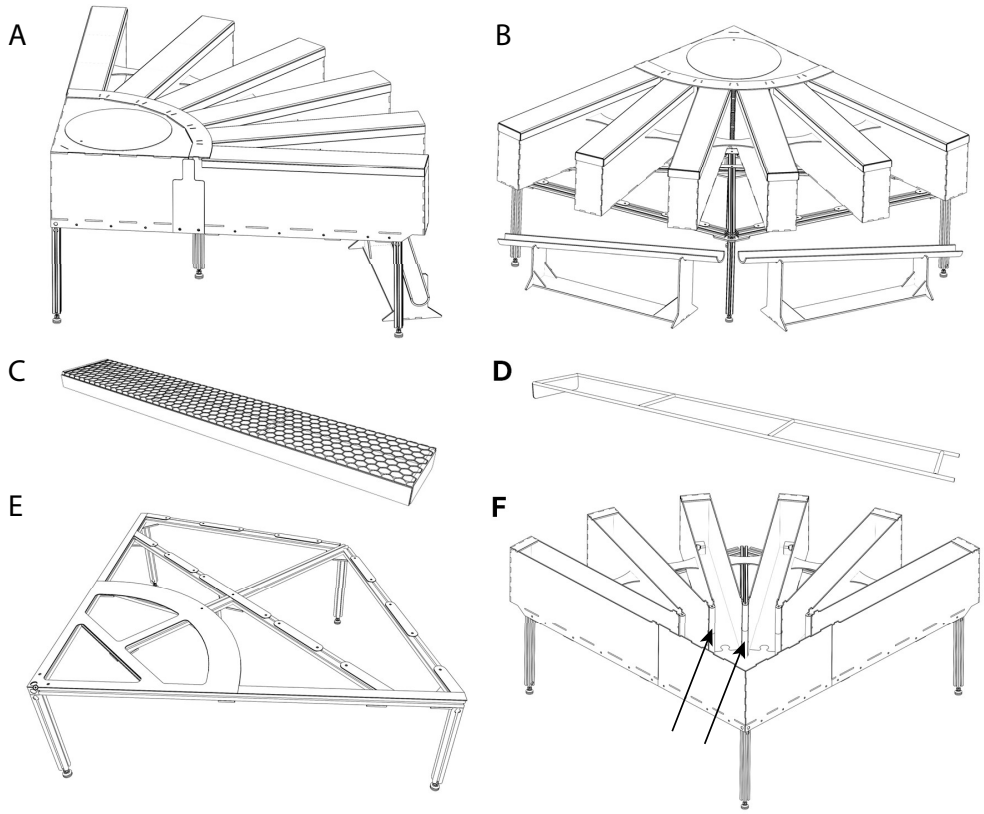

Figure 2: 6-Arm Maze:

- A 3D view of the fully-assembled maze (side) with lids and ceiling included.
- 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 (arrows) which join the walls between adjacent arms.

2.3. Reward Ports

The most critical features of the maze are the reward ports, which dispense a liquid food reward (evaporated milk with 5% sucrose added). Each port consists of an IR photogate (emitter and receiver; 890nm), a white flat-mount 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 [20]. 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 crevices 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 an associated peristaltic pump.

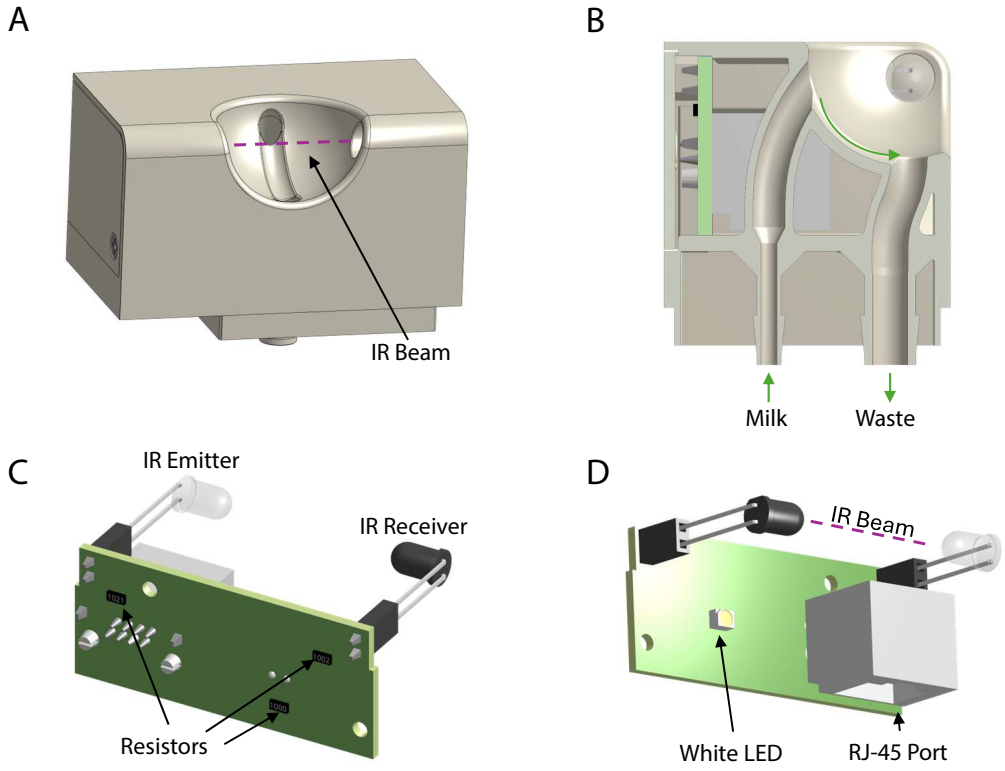

Figure 3: Reward Port:

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 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.

2.4. 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[21], Arduino, BControl [22], or Bonsai[23] 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 DC 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 aluminum 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 (pinout in Fig. 5C), 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. Importantly, we have not observed electrical artifacts caused by the pumps on neural recordings made in close proximity to the PER (12").

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. All the pumps for each maze draw from a single 1L reservoir of milk, streamlining the milk loading process and reducing waste.

In addition, it greatly increases the efficiency of the cleaning process. 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 fill 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.

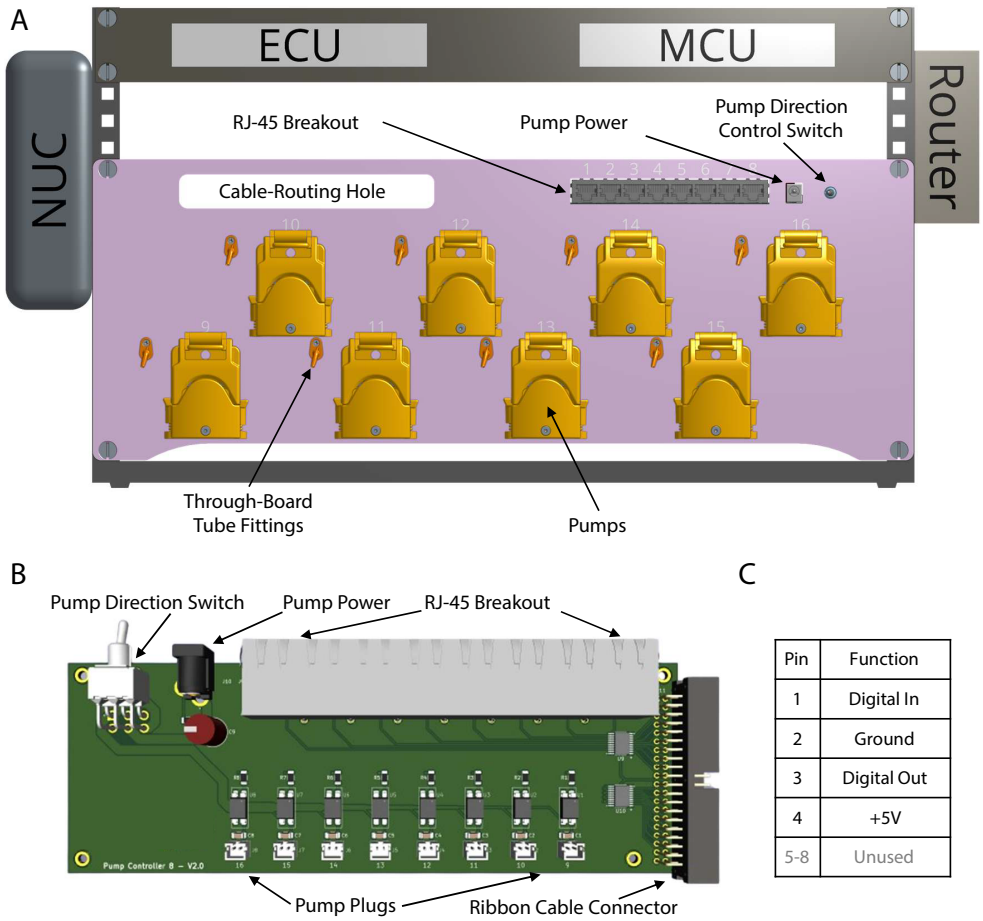

Pin	Function
1	Digital In
2	Ground
3	Digital Out
4	+5V
5-8	Unused

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. Tubing (not shown) routes liquid from a single 1L reservoir to each of the pumps, then through the pump plate via the through-board tube fittings to the reward ports.

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.

2.5. 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 in place of gutters 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 [24]. Reliable performance is best achieved by using several simpler task variants ("steps") to gradually shape behavior toward the final task [17], 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. 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 neural correlates

of the learning process with more structure and specificity[25] - thus both the training steps 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 and its training steps 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 steps, 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 (Figure 5A) . 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. The first few task steps are visually cued: rewarded outer arms are lit while unrewarded outer arms are not (Figure 5B). Later, subjects transition to memory-dependent task steps, during which both rewarded and unrewarded ports are lit and subjects must discover and remember which one(s) deliver reward. In later task steps, goal arm locations are changed in a "block" structure (after 15 goal arm visits; Figure 5C). 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, the epoch ends after a set number of trials, a set number of goal blocks, or a set amount of time (typically a maximum of 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 5D, and requires a minimum of 18 epochs to complete (though subjects typically require approximately 30 epochs). Each step in

the curriculum is characterized by a set of 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 (#rewarded trials / #trials over the entire epoch). 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. 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 greater than or equal to the threshold goal rate. Goal rate thresholds for memory-dependent steps are generally calculated as 75% of the rate a perfectly-trained subject would achieve if they were maximally unlucky, minus an additional 5% buffer for the memory-based steps. At each step in the curriculum, subjects must complete a specified number of successful epochs before moving on to the next step. After the completion of step 10, subjects are considered fully trained, although additional epochs of behavior may be collected to fully characterize final performance levels.

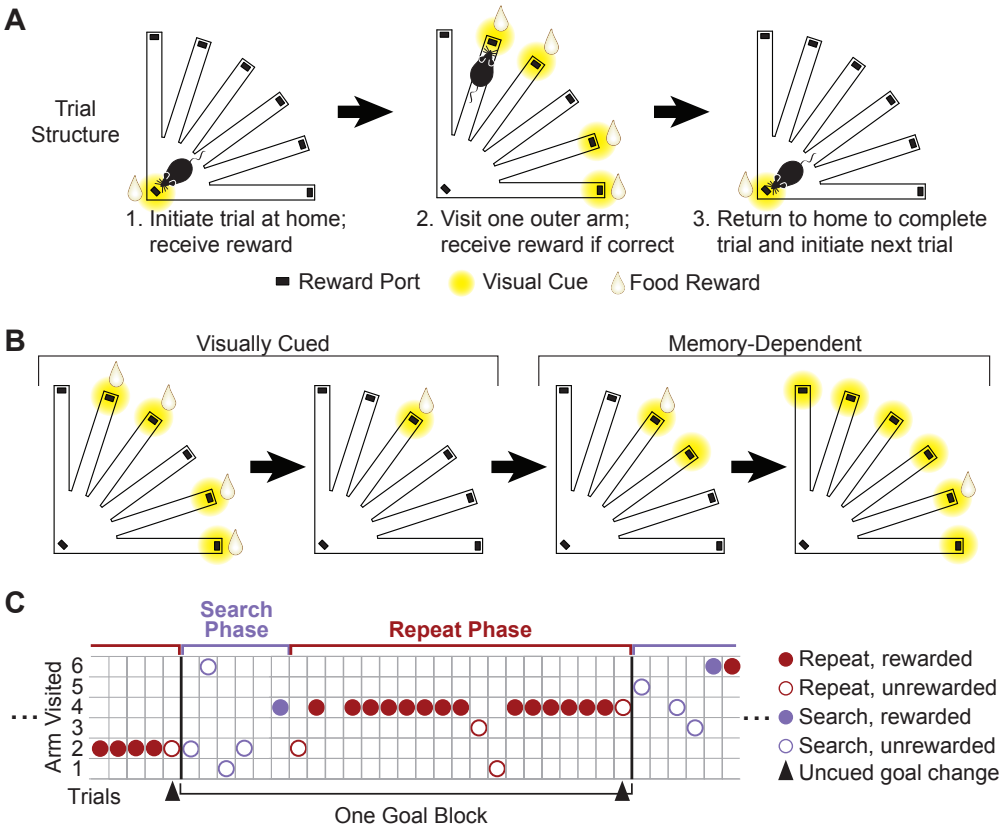

D

Step	Well state	Rewards per block	Epochs at criteria	Reward rate threshold	Learning target
1	●●●●○○	1	2	0.50	acclimate to maze; learn trial structure
2	●●●○○○	1	2	0.80	only visit 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 to repeat by memory
6	●◐◐○○○	15	1	0.61*	incremental practice
7	●◐◐○○○	15	2	0.59*	switch goals without visual cue
8	●◐◐◐○○	15	1	0.57*	incremental practice
9	●◐◐◐◐○	15	1	0.54*	incremental practice
10	●◐◐◐◐◐	15	1	0.51*	incremental practice/final task

● Lit + Rewarded ◐ Lit + Unrewarded ○ Unlit + Unrewarded

Figure 5: Training curriculum:

A Visual schematic of a single example trial at step 1, including the initiating poke at the home port, choice of one outer arm, and the return to home to complete the trial.

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

C Example trial series for step 10 behavior, illustrating the classification of search and repeat trials.

D A table giving an overview of the 10-step training curriculum. Outer reward ports exist in one of 3 states: Lit + Rewarded, Lit + Unrewarded, and Unlit + Unrewarded. 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. * indicates suggested goal rate thresholds; those used for the pilot cohort behavior were marginally higher: 0.70, 0.66, 0.63, 0.60, 0.58, and 0.55, respectively. ** indicates the suggested number of epochs at step 4; only 1 was required for the pilot cohort.

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, each subject's environment parameters and learning progress are stored in a text file in order to fully automate the subject's progress through the curriculum. This text file is associated with the subject's name, and is automatically read/updated by the same python script which controls the maze environment. 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 step.

To safeguard against any errors which may occur during the epoch, the parameter file is updated as trials are completed, ensuring that it is as up-to-date as possible at all times. Furthermore, a copy of the parameter file is also added to the behavior log at the beginning and 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, since those arms share a continuous wall with the home area. As subjects explore the maze and learn the task, they tend to develop idiosyncratic biases towards and against specific arms, which confounds behavioral analyses. 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 less frequently visited arms, encouraging subjects to visit all arms as evenly as possible.

While simple in its aims, there were a few important considerations that we needed the algorithm to account for. First, if the algorithm ever pushes the probability of selecting a particular arm to effectively 0, subjects can and will over-correct, learning to completely avoid that arm, which is hard to recover from (in general, we found it easier to train subjects away from their favorite arms rather than towards arms they were averse to). To counter this, a lower bound was selected for the arm selection probability (typically $0.01 - 0.02$).

Another issue we encountered was the algorithm being somewhat unresponsive, and failing to select arms which used to be a subject's favorite long after the subject had learned to visit that arm less. While this could be countered by making the algorithm more aggressive, that comes at the cost of the algorithm becoming overly sensitive to the randomness inherent to goal selection and subjects' behavior. So, to help the algorithm be

more responsive, we added a parameter which weights recent arm visits more heavily than older arm visits; allowing subjects' past behavior to quickly become completely irrelevant, while also only responding to new behaviors after a certain level of consistency is demonstrated. Two additional parameters were included which control whether the rewarded visits or unrewarded visits are more influential, and how aggressively the algorithm attempts to push subjects towards even exploration.

With these four parameters, our arm selection algorithm is flexible and responsive, while avoiding some some of the pitfalls which can seriously hamper learning.

3.5. Data Outputs

For post-hoc analysis, we need reliable records of all behavioral events, such as port triggers, port illumination state, and reward delivery times. Further, we need a system that is resilient to common experimenter errors such as typos in filenames and recording mistakes. 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 param-

eters and a table of important trial events. Each trial contains the data for the home visit, the outer arm visit, goal block details, and the details for any lockout port visits. The data recorded for each port visit 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 easy to quickly detect any maze malfunctions.

4. Results

4.1. Experimental Model and Subject Details

To test the automated, parallelized maze apparatus, we trained 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 singly housed and food restricted to 85% of their free-feeding weight. All procedures were approved by the Institutional Animal Care and Use Committee at the authors university.

4.2. Behavioral Training

Data were 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 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. Each subject was assigned to a specific maze, with one young rat and one old rat assigned to each maze.

Subjects were trained using the curriculum shown in Figure 5D, with minor differences indicated by asterisks and detailed in the figure caption. The suggested values shown in the table are expected to avoid two pitfalls encountered in the pilot cohort. First, the subjects did not generally spend enough time at step 4 to build a sufficiently strong preference for lit wells to persist into early memory steps. Second, the thresholds for success during memory steps were slightly too high, causing subjects to plateau below the threshold, adding several, likely unnecessary, days to training. Different criteria may be required for different curriculum versions, ages, or strains of rats. Nevertheless, subjects showed a rapid increase in reward rate during the visually cued epochs, all reaching performance levels well above chance, but also demonstrating substantial 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, indicative of search efficiency (Figure 6E). The amount of training required to reach criterion for each step was suggestive of slower learning in the aged subjects, particularly for the first two steps of the memory-dependent task variant (Figure 6C), although this did not reach statistical significance due to our small sample size. The most drastic difference between age groups was in trial duration (Figure 6F), as aged subjects moved far slower than young subjects. Despite their slow speed, aged 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 are 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). This further confirms that the issue was one of parameter choice rather than a true performance difference, and has informed the improved parameters described in

Figure 5B.

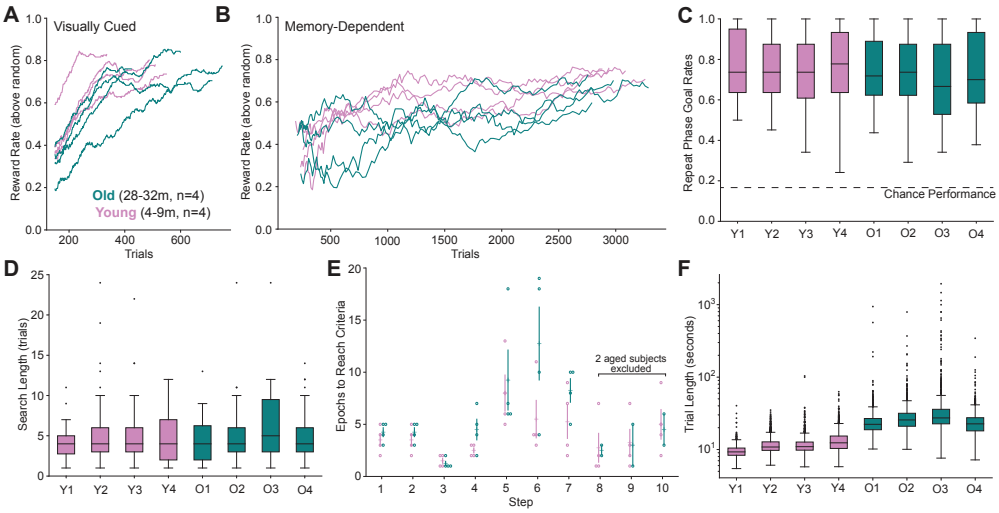

Figure 6: Pilot Cohort Results:

A Learning curves by subject for the visually cued task steps. 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 steps, 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. Linear mixed effects models were used to evaluate the significance of age effects (Python statsmodels mixedlm function). $LMStatistic = 2.45, p = 0.12$, $n_{blocks} = [Y1 : 36, Y2 : 70, Y3 : 54, Y4 : 68, O1 : 32, O2 : 35, O3 : 31, O4 : 39]$ (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 B) 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 on a complex curriculum. In contrast to prior training regimes for a similar task [24], 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, 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. While unsurprising [26, 27], this underscored the importance of having each epoch be limited by trial or goal block criteria, rather than by time (as is more common). This ensured that the aged and young subjects received similar amounts of experience with each

epoch and will be essential for any studies of aged subjects. It was also encouraging that, when 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, allowing for future investigation into the neural correlates of each rule separately.

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 tubing leaks causing failure in reward delivery. 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 criteria and pilot training cohorts may be necessary to inform performance expectations when developing new tasks or curricula.

5.2. Future Outlook

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. If position information were unnecessary, or could be collected via alternate means (via inertial sensors[28], for instance), then it would be feasible to stack 3-4 mazes vertically. Some updates to drainage and cleaning strategies would be needed in this case, but all task control and logging software would seamlessly extend to the stacked use case. Such stacking of the mazes could allow comparable throughput to other highly successful behavioral training approaches [18, 19], and would be the first to our knowledge to extend this high-throughput approach into the spatial navigation realm.

Even at its current capacity, the maze apparatus opens up exciting possibilities for efficiently testing larger cohorts of female subjects, models of neurodegenerative disease, and performing longitudinal behavioral assessment of cohorts across the lifespan — all of which will be fascinating comparators to the expansion of our existing pilot study of age-related behavioral differences. Since the apparatus is also fully compatible with wireless neural recordings and overhead position tracking, future studies will include the assessment of neural activity patterns that underlie the learning and memory processes that we set out to study. The ability to test larger cohorts of animals with minimal preselection will allow us to capture and characterize performance across a much wider range of abilities, facilitating our understanding of the neural basis of individual performance differences. Another key challenge in studying individual behavioral differences is the need to identify biologically-driven, rather than experimentally-induced, variability. The automation and parallelization of the current maze apparatus suggests that the variability we see in pilot cohort performance may reflect, and allow the study of, true individual differences.

We hope that the low cost and open source nature of many aspects of our task apparatus will make it a useful resource for others striving to scale up and increase the automation of their behavioral testing approaches. In the described maze environment, the flexible nature of the task logic can allow for many different tasks to be implemented, such as sequence learning tasks, alternation tasks [29, 30], or even pure foraging tasks. The use of alternative

maze environments equipped with our, or other, maze features [16] opens endless possibilities for structured, efficient behavioral testing. Regardless of behavioral paradigm, a shift towards parallelized training offers improved control of other influences on behavior, such as circadian rhythms, in addition to increased efficiency. Finally, the adoption of a structured curriculum with trial-based advancement criteria, rather than time-based advancement criteria, will be beneficial for accurately and rigorously capturing behavioral differences of across groups of subjects no matter the species or task.

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. Resources

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

Appendix B. Extended Data to 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] A. P. Steiner, A. D. Redish, Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task, Nature Neuroscience 17 (7) (2014) 995–1002. doi:10.1038/nn.3740.

[6] 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

[7] 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.

[8] 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.

[9] 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

[10] 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

[11] 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

[12] 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

[13] E. J. Chesler, S. G. Wilson, W. R. Lariviere, S. L. Rodriguez-Zas, J. S. Mogil, Influences of laboratory environment on behavior, Nature Neuroscience 5 (11) (2002) 1101–1102, publisher: Nature Publishing Group. doi:10.1038/nn1102-1101.
URL https://www.nature.com/articles/nn1102-1101

[14] M. Qiao, T. Zhang, C. Segalin, S. Sam, P. Perona, M. Meister, Mouse

Academy: high-throughput automated training and trial-by-trial behavioral analysis during learning, pages: 467878 Section: New Results (Feb. 2019). doi:10.1101/467878.
URL https://www.biorxiv.org/content/10.1101/467878v2

[15] S. Zou, C. T. Li, High-Throughput Automatic Training System for Spatial Working Memory in Free-Moving Mice, Neuroscience Bulletin 35 (3) (2019) 389–400. doi:10.1007/s12264-019-00370-z.
URL https://pmc.ncbi.nlm.nih.gov/articles/PMC6527628/

[16] B. S. Porter, J. M. Olson, C. A. Leppla, Duvelle, J. H. Bladon, M. A. A. v. d. Meer, S. P. Jadhav, Adapt-A-Maze: An Open-Source Adaptable and Automated Rodent Behavior Maze System, eNeuro 12 (7), publisher: Society for Neuroscience Section: Research Article: Methods/New Tools (Jul. 2025). doi:10.1523/ENEURO.0138-25.2025.
URL https://www.eneuro.org/content/12/7/ENEURO.0138-25.2025

[17] 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

[18] 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

[19] 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

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

[21] J. H. Siegle, A. C. López, Y. A. Patel, K. Abramov, S. Ohayon, J. Voigts, Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology, Journal of Neural Engineering 14 (4) (2017) 045003, publisher: IOP Publishing. doi:10.1088/1741-2552/aa5eea.
URL https://doi.org/10.1088/1741-2552/aa5eea

[22] Bcontrol.
URL https://brodylabwiki.princeton.edu/bcontrol/index.php/Main_Page

[23] G. Lopes, P. Monteiro, New Open-Source Tools: Using Bonsai for Behavioral Tracking and Closed-Loop Experiments, Frontiers

in Behavioral Neuroscience 15, publisher: Frontiers (Mar. 2021). doi:10.3389/fnbeh.2021.647640.
URL https://www.frontiersin.org/journals/behavioral-neuroscience/articles/10.3389/fnbeh.2021.647640/full

[24] 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.

[25] M. Shujah, S. Joshi, M. Nakamura, T. R. Barkat, Curriculum Learning: sequential acquisition of task complexity enhances neuronal discriminability, pages: 2024.09.10.612027 Section: New Results (Sep. 2024). doi:10.1101/2024.09.10.612027.
URL https://www.biorxiv.org/content/10.1101/2024.09.10.612027v1

[26] J. E. Birren, L. M. Fisher, Aging and speed of behavior: possible consequences for psychological functioning, Annual Review of Psychology 46 (1995) 329–353. doi:10.1146/annurev.ps.46.020195.001553.

[27] J. A. McQuail, M. M. Nicolle, Spatial reference memory in normal aging Fischer 344 × Brown Norway F1 hybrid rats, Neurobiology of Aging

36 (1) (2015) 323–333, place: Netherlands Publisher: Elsevier Science. doi:10.1016/j.neurobiolaging.2014.06.030.

[28] R. Fayat, M. Sarraudy, C. Léna, D. Popa, P. Latouche, G. P. Dugué, DISSeCT: An unsupervised framework for high-resolution mapping of rodent behavior using inertial sensors, PLOS Biology 23 (10) (2025) e3003431, publisher: Public Library of Science. doi:10.1371/journal.pbio.3003431.
URL https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3003431

[29] 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

[30] 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.