The RIFF: an automated environment for studying the neural basis of auditory-guided complex behavior

Maciej M. Jankowski*, Ana Polterovich*, Alex Kazakov*, Johannes Niediek*, Israel Nelken
The Edmond and Lily Safra Center for Brain Sciences and the Dept. of Neurobiology, Silberman Institute of Life Sciences, the Hebrew University of Jerusalem, Israel
*These authors had equal contributions

IN conceived and supervised the project. MMJ and AP designed and constructed the RIFF, designed the behavioral paradigms with a contribution of JN, and collected data. AP wrote the software for implementing Markov Decision Processes in the RIFF. MMJ designed and developed the chronic neural implants. AK developed the rat tracking, spike-sorting, image processing modules and the experiment visualization GUI. JN, AK and AP developed the analysis pipeline for the data from the RIFF. MMJ, AP, AK and JN jointly brought the system to life. AP and JN analyzed the behavioral and neural data presented in this paper.

Funding

This work was supported by AdERC grant GA-340063 (project RATLAND), by F.I.R.S.T. grant no. 1075/2013 from the Israel Science Foundation. JN was supported by a DFG Research Fellowship (ref. NI 2012/1-1; project number 442068558). IN holds the Milton and Brindell Gottlieb Chair in Brain Sciences.

Abstract

Behavior consists of the interaction between an organism and its environment, and is controlled by the brain. However, while brain activity varies at fast, sub-seconds time scales, behavioral measures tend to be temporally coarse, often limited just to the success or failure in a trial. The large gap between the temporal resolutions at which brain and behavior are observed likely impedes our understanding of the neural mechanisms underlying behavior.

To overcome this problem, we developed the RIFF: an interactive arena for rats that has multiple feeding areas, multiple sound sources, and high-resolution tracking of behavior, with concomitant wireless electrophysiological recordings. The RIFF can be flexibly programmed to create arbitrarily complex tasks that the rats have to solve. It records unrestrained rat behavior together with neuronal data from chronically implanted electrodes.

We present here a detailed description of the RIFF. We illustrate its power with results from two exemplary tasks. Rats learned within two days a complex task that required timed movement, and developed anticipatory behavior. Rats found solution strategies that differed between animals but were stable within each animal. We report auditory responses in and around primary auditory cortex as well as in the posterior insular cortex, but show that often the same neurons were also sensitive to non-auditory parameters such as rat location and body orientation. These parameters are crucial for state assessment and the selection of future actions.

Our findings show that the complex, unrestrained behavior of rats can be studied in a controlled environment, enabling novel insights into the cognitive capabilities and learning mechanisms of rats. This combination of electrophysiology and detailed behavioral observation opens the way to a better understanding of how the brain controls behavior.

1. Introduction

Behavior involves brain-controlled, complex interactions between an organism and its environment. Full natural behavior has many components. Examinations of single components of the interaction between organisms and their environments shape much of what we understand about brains and their functions, including sensation and perception [1]; motor control [2]; decision making [3]; and memory [4]. However, these successes have been achieved at a price – animals are often studied using behavioral tasks that are much simpler than those they face in their natural habitats.

The main observable output of behavior is movement – animals select and time their actions. Patterns of movements may be rich: Predictive movements may precede trial events when there is a temporal structure to the task; movements may differ between individuals; and movements may be used for solving cognitive tasks – for example, rats exert active control on selected motor variables in order to enable consistent perception of location when using their whiskers [5].

Evidently, even in simple tasks, animals perform much more than a single action at each trial. In more complex, naturalistic settings, the amount of movements (and decisions) required from animals is substantial. Brain activity is expected to reflect and shape the full complexity of the concomitant behavioral strategies [6] . Nevertheless, brain activity and behavior are often measured at very disparate rates. While brain activity changes at rates reaching 100/s or faster, behavior is often quantified by coarse measures at the rate in which trials are presented, often slower than 1/s. Correct/incorrect decisions, response latencies and similar outcome-driven measures do not lend access to the entire decision process occurring from trial onset to outcome. This is a substantial limitation. During a trial, neural activity throughout the brain runs its fast course, before, during and after relevant decision points. For example, movements that vary from trial to trial account for much of the inter-trial variability in wide-field calcium imaging of mouse cortex [6]. Thus, a large amount of information in the neural responses about behavior, which holds valuable insights on the links between behavior and neuronal processing, is likely wasted.

The goal of this work is the development of a behavioral environment for studying the behavior of rats in complex, auditory-guided tasks while keeping good experimental control over the resulting behavior. Exploration of animal behavior in complex settings progressed significantly in recent years. Virtual reality setups [7–15] are often used for that purpose. They are nevertheless limited in the realism that they achieve, and many experimental questions require the use of freely moving

animals. Recently, many environments for studying the behavior of freely moving animals in complex environments have been described but they often do not allow concomitant neuronal recordings [16–28]. A small number of publications describe complex environments that include wireless recording of neural activity in freely moving animals [29–31], but allow only a limited behavioral range.

We designed and constructed an arena (the 'RIFF') that allows high-resolution behavioral data acquisition in time and space with concomitant electrophysiological recordings in freely moving animals. The RIFF is controlled by a real-time loop that tracks the rat, identifies its actions and reacts to it. Because of the generic nature of this loop, it is flexible and can implement a rich set of behavioral scenarios. Neural data is collected from tens of electrodes by telemetry or with a data logger. Rats can be trained in the RIFF to perform a wide range of complex tasks.

We describe here the design principles and implementation of the RIFF. We illustrate the use of the RIFF by describing data from two very different tasks. Rats that were trained inside the RIFF learned these behavioral tasks efficiently, sometimes within a few hours. For example, they learned to move in anticipation of trial events. Different rats developed individual patterns of behavior, which can be detected early in the learning process and tracked over months. We show that neural activity recorded in primary and higher-order auditory fields tracks animal location and pose in addition to the auditory signals.

2. Design principles

The RIFF was designed to study the behavior of freely-moving animals interacting with a rich environment that is nevertheless amenable to good experimental control, while simultaneously recording neural data. We selected rats as the model organism, as they can learn fairly complex tasks and can easily carry a wireless recording system. We decided to strictly avoid any form of tethering.

The use of the RIFF requires integrating a large number of interacting systems and processes (Fig. 1), which we kept as separate and modular as possible in order to allow for modifications and extensions as experience is gained and new technology becomes available. While the logic used to operate the RIFF can be very flexible, we intended it to operate mostly as a Markov Decision Process[32] (MDP). MDPs are 'state machines'. They are defined by a set of states (defined, for example, by the location of the animal, the current stimulus that is presented, and

the correct ports to poke in order to get a reward), a set of actions that an animal can take in each state, and a set of actions that can be taken by the environment (typically providing rewards and punishments to the animal). The MDP is governed by two (potentially stochastic) rules. The first is the rule by which one state follows another. In an MDP, transitions depend only on the current state and on the current action of the animal. The second rule prescribes whether and which actions the environment takes, depending only on animal action and the consequent state transition.

MDPs have a rich theory that facilitates the analysis of their behavior [32]. For example, many MDPs have 'optimal policies' - rules for state-dependent action selection - which achieve maximal reward in the long term. An optimal policy can be compared with animal behavior, supplying an absolute scale to which animal behavior can be compared. While the Markovian property, requiring that the future states depend only on the current state and not on the way by which the current state has been reached, may seem to limit the scope of tasks that such an environment can implement, in practice essentially all behavioral tasks used in neuroscience can be formalized within this framework.

A number of additional constraints influenced the design of the RIFF, in order to achieve high throughput data production. Around the experimental core, running an experiment in the RIFF requires some design choices (Fig. 1). These include the specification of the experimental paradigm and its implementation as a state machine; the preparation of the animal - initial habituation to the RIFF, as well as the preparation of the chronic implant for data collection and its implantation. Following the data collection, common offline analysis of the data includes posture extraction and spike detection and sorting. Finally, all this information, collected with different data acquisition systems, and potentially stored on different computers, has to be integrated and aligned to a single time axis for further analysis.

To implement the RIFF, we had to make specific choices for each of these processes, and our specific implementation of many of them is described in the Materials and methods section.

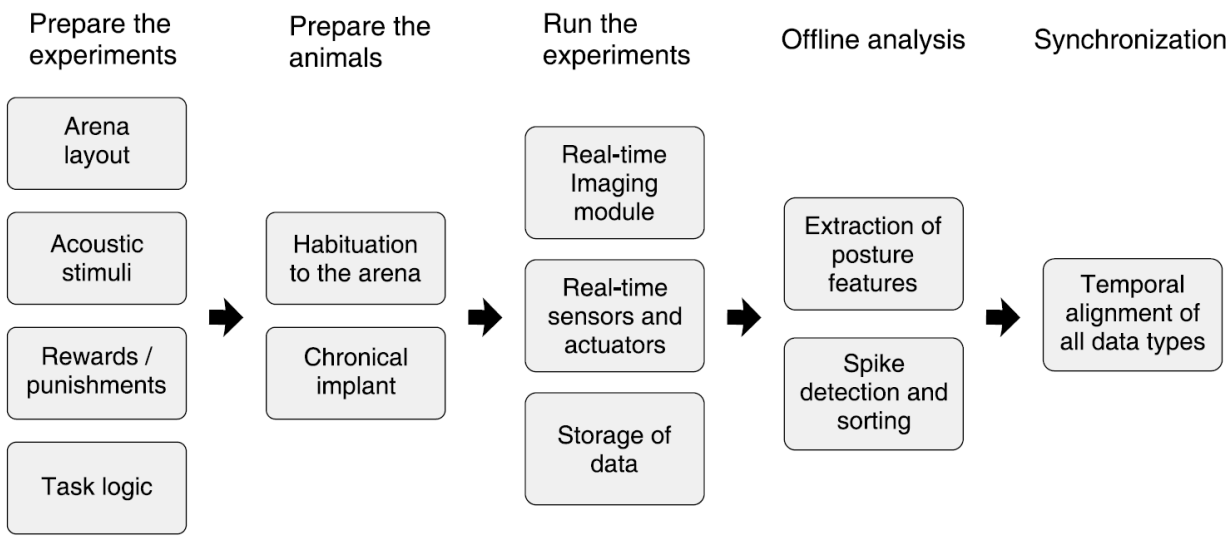

Fig. 1. Workflow for the design, preparation, execution and initial analysis of behavioral experiments as they pertain to the RIFF. The design of an auditory-guided task includes the development of task logic, the nature of rewards and punishments, as well as the selection of the auditory stimuli. The flexibility of the RIFF allows for substantial freedom in these choices. Rat preparation includes habituation to the RIFF and electrode implantation. While running the experiment, the RIFF is designed to keep a tight synchronization of all real-time events, and to ensure efficient data storage. The stored data allows for offline analysis that includes spike sorting and animal pose estimation, which are then precisely aligned with the rest of the information from the experiment.

3. The setup

The RIFF consists of the following components (Fig 2): the physical environment itself; a set of sensors that track animal behavior; a set of actuators; electrophysiological recording components; and the control logic that connects all of these components, which runs on the main control computer of the RIFF. In the following, we describe each of these components in turn.

The behavioral environment is an 18-sided polygon (160 cm diameter, Fig. 2.a) with six interaction areas (IAs), each consisting of two loudspeakers, a water port, a food port, and air-puff valves (Fig. 2.b). The main sensor for animal location in the current implementation is a ceiling-mounted camera with a dedicated computer that performs real-time image processing. The camera tracks the location of the rat with a temporal resolution of 30 ms. The location information is transmitted to the main computer on a dedicated link. Additional sensors include the nose-poke detectors at each feeding port that report their state in real time; the status of the food and water ports; and analog sensors carried by the animals such as a microphone, accelerometers, and gyroscopes[33].

The RIFF interacts with the behaving animal through the actuators (Fig. 2b). In the present implementation, the response repertoire of the RIFF includes food or water delivery, air-puff delivery, and sound presentation by any combination of the 12 speakers. The RIFF includes grids for providing foot shock as aversive stimuli in three of the six IAs, although we have not used them in the experiments described here.

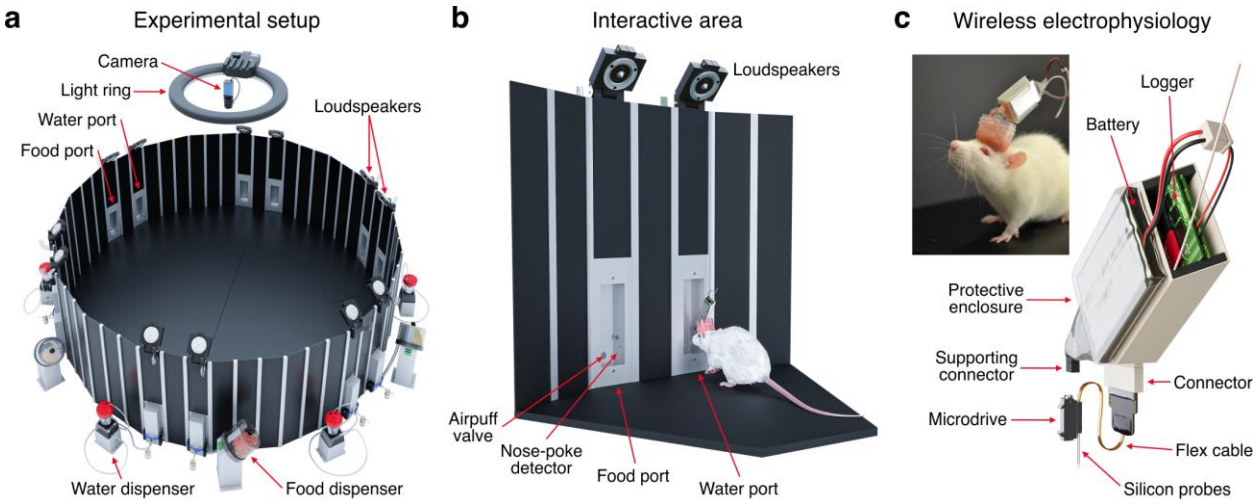

Fig. 2. An experimental arena for studying the neural basis of auditory-guided complex behavior. (a) The RIFF is a large arena (160 cm in diameter), equipped with 12 loudspeakers, 6 food dispensers and 6 fluid dispensers, 12 airpuff outlets, and a ceiling mount camera. The wall panels can be removed, the ports can be closed by blinds, and the spatial arrangement of the environment can be adjusted for different experiments. **(b)** Close-up on one interaction area. An adult female rat with a chronic implant and a neural logger is shown for scale. Each interaction area consists of two ports, one for food rewards and the other for fluids. Each port has a nose poke detector and airpuff outlet. Two speakers are mounted above each port. **(c)** An approach for chronic wireless recordings in rats. The neural logger and battery are in the protective plastic case. A multi-contact silicon probe is mounted on the microdrive. The device is inclined to the back in order to allow rats to move naturally and an undisturbed access to the ports of the interaction area. The picture shows a rat with a 32-channel moveable silicon probe implant and wireless data logger. The battery used here allows for 3 hours of recordings.

Neural responses are recorded using chronically implanted silicon probes (Fig. 2c). In conjunction with the use of state-of-the-art electrodes or silicon probes, it is possible to record neural activity in the RIFF for several months. Depending on the telemetry/logging system used, the implant is attached to a head pack that may include preamplifiers, a battery, and the storage or transmission capabilities. To illustrate the modularity in the RIFF, we used both a telemetry data acquisition system (TBSI W64, Triangle BioSystems International, Durham, NC, USA attached to an AlphaLab SNR data acquisition system, Alpha Omega, Nazareth, Israel) and a neurologger (RatLog-64, Deuteron Technologies, Jerusalem, Israel). The headstage of both systems include a microphone for recording the sound field near the head of the rat simultaneously with the

electrical recordings. The neurologger headstage also has 3-axis accelerometers and gyroscopes. The two systems can be switched seamlessly. Our current equipment supports continuous recordings in freely behaving animals for up to 24 hours using the telemetry data acquisition system, using a 500 mAh battery weighing ~10.5 g. Thus, the RIFF can be used to conduct long-term studies of the neural correlates of behavior.

The logic of the RIFF is implemented in a program that runs on the main control computer, collects the data from sensing hardware, and drives the actuators using a predefined logic (Fig. 3). We extensively use MDPs for that purpose [32]. Figure 3.c shows an example of an MDP that was used to formalize a sound localization/discrimination task (the L/D task, see Methods and Supp. Fig. S1b). In this task the rat learns to associate six different sounds and target locations with the six IAs from which each one of them is played. The rat self-initiates a trial by entering the central part of the arena (event No. 1 in figure 3c). Once the RIFF detects a center-crossing event, it selects a random target port and plays the corresponding sound from the two speakers above the target port (event No. 2). The same sound is repeated once every 2 seconds until the rat pokes in a port or until 10 sound presentations are made (20 s), in which case the trial is aborted. The rat is free to perform any action but will be rewarded only if it pokes in a port of the correct IA. In the example trial of FIg. 3c the rat first runs to the wrong port (event No. 3), but after hearing the target sound a second time, it makes a sharp turn and correctly nose pokes in the correct area (event No. 4), receiving a food pellet as reward (event No. 5).

Four major data types are recorded during the experiment and stored for offline analysis (Fig. 3b): time stamps, neural activity, analog sensor data, and images. Time stamps are saved for every internal state change, every behavioral event, for video frame acquisition times, and for and every sound presentation. The internal states represent the latent activity of the controlling program which transitions from state to state through its predefined logic (Fig. 3b.2). Behavioral events are produced by the interaction of the rat with the RIFF, and include location coordinates, nose pokes, food or water delivery, air puffs and foot shocks, as well as sound onsets (Fig. 3b.1). The neural activity is stored as raw voltage traces for each recording site (Fig. 3b.3). Analog sensor data from the head implant include microphone signals, accelerometer and gyroscope readings (Fig. 3b.4). Images of the rat are collected from the ceiling-mounted camera (Fig. 3b.5). The total data rate of the RIFF (with 32 electrodes and a 30/s video stream) is about 10 Gigabytes/hour.

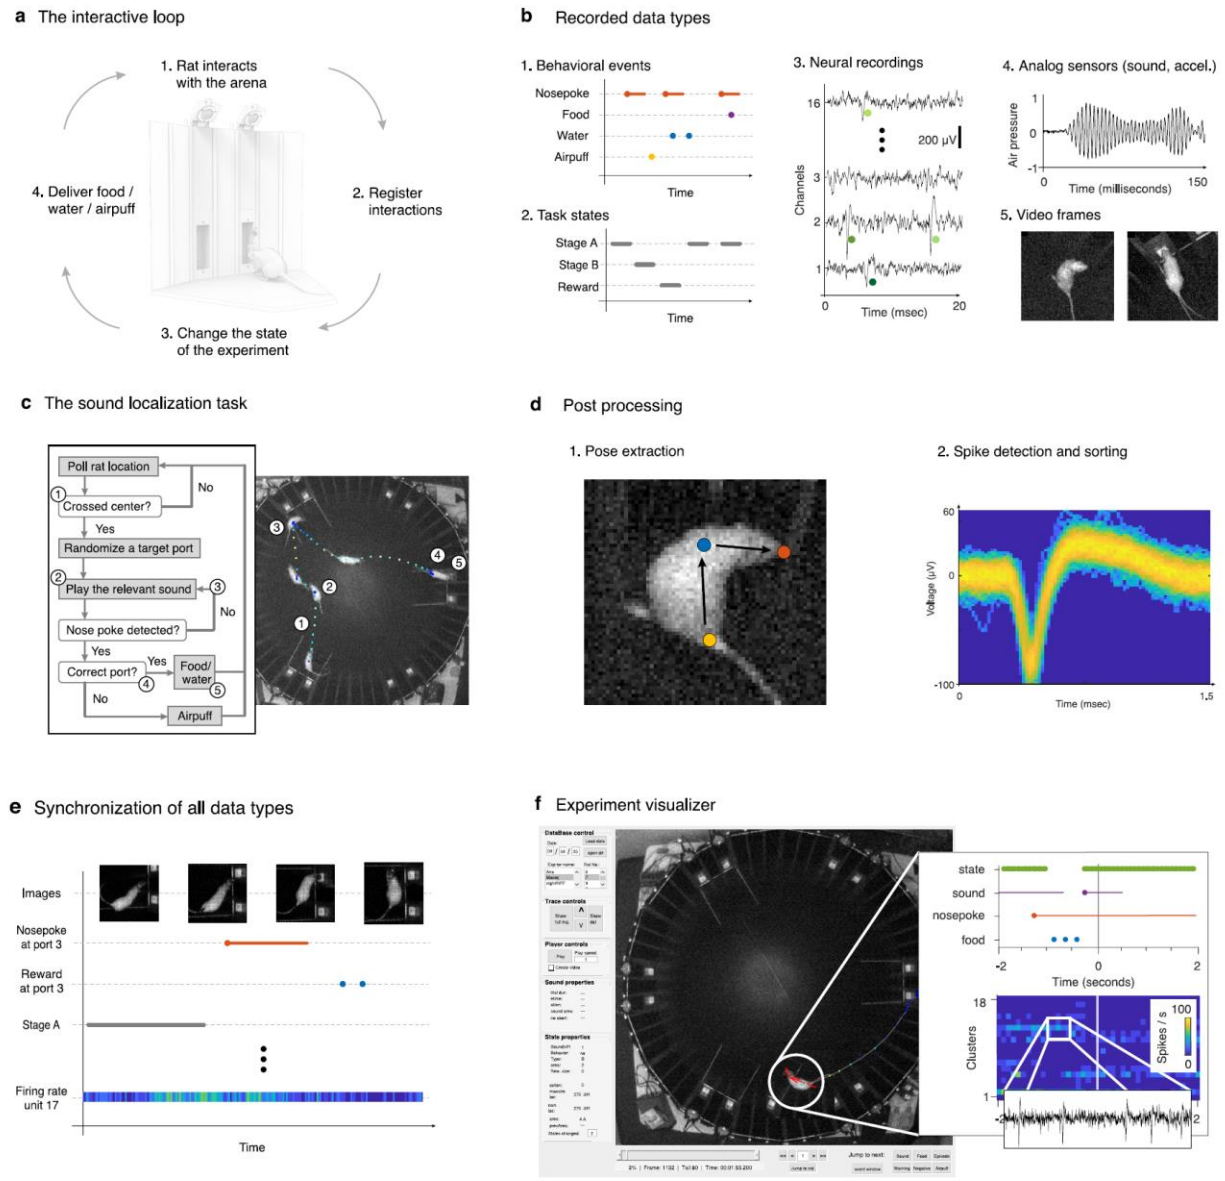

Fig. 3. Operation of the RIFF. (a) Rats interact with the RIFF through movement and nose pokes. Their location is identified online using the video camera and a dedicated computer. The experimental environment reacts to rat actions by changing its state and providing different types of feedback – sounds, rewards, or punishments. **(b)** The RIFF collects a large variety of data types: behavioral events (e.g. nose pokes, rewards, punishments); task states; neuronal activity; analog sensor signals (microphone, motion sensors); and the animal location tracked by the ceiling mount camera. **(c)** An illustration of a sequence of interactions between a rat and the RIFF during the sound localization/discrimination task. (1) The rat moved to the center of the arena in order to initialize the trial. (2) Once it crossed into the central area (a circle with a radius of 30 cm), the RIFF started sound presentation from both loudspeakers in a randomly selected interaction area. (3) Initially, in the example presented here, the rat approached a wrong port; a second sound presentation caused it to move towards the correct target port (4), and to receive the reward (5). **(d)** Data post-processing. Posture features are extracted using a custom-trained DNN. The kilosort2 program used with a custom wrapper performs largely automatic spike detection and sorting. **(e)** All data types are

synchronized on a single time axis and **(f)** can be visualized offline using custom visualizer software. The visualizer is used to browse through all synchronized data types up to the level of the raw neural signals.

4. Data integration and analysis

Here we describe the offline processing of the data collected in the RIFF. The RIFF is designed for high-throughput data acquisition, and therefore the post-processing steps are almost completely automated.

Some of these data streams require denoising, parsing and feature extraction. A convolutional neural network operates on the images of the rat, estimating the locations of the nose, neck and the base of the tail from which the head and body directions are inferred (Fig. 3d.1, Supp. Fig. S2). The neural data is denoised, and spike detection and sorting is performed by an automated custom wrapper to the Kilosort2 spike sorter (https://github.com/MouseLand/Kilosort2) (Fig. 3d.2, Supp. Fig. S3). In our implementation, the final labeling of the resulting units requires a brief manual curation, which requires a few minutes for a 3-hour experiment with 32 electrodes.

After information is extracted from each data stream, all data are synchronized to a single time axis. The precision of the synchronization varies with data type. Thus, events that are logged together with the electrophysiological data (e.g. sound stimulus onsets) are synchronized to the precision of the sampling rate of the electrophysiological data collection (32 kHz). On the other hand, the sequence of images, collected at a resolution of 30 ms, is synchronized with the electrophysiological data with a temporal precision of a few milliseconds (Fig. 3e).

The output of the analysis pipeline is a single file that combines all data types that are temporally synchronized, denoised and enhanced by additional feature extraction (e.g. rat pose determination). The post-processing steps were computationally efficient and are much faster than actual experiment time. The size of the merged file is 100MB/hour recording at a temporal resolution of 10/s, and the timing data it contains allows access to the raw data that was used in its construction.

The merged file can be replayed through an interactive graphical interface (Fig. 3f). This module loads all the data types and reproduces the state of the experiment at any specific time point, with all the peripheral information, the state of all sensors and actuators of the RIFF, and neural activity at a resolution of single spikes. This module is used for verifying the correctness of the logic, the integrity of the data, as well as for examining offline the fine details of the behavior of the animals.

5. Results

5.1 Experimental paradigms

The data presented here come from two experiments employing very different behavioral paradigms, the localization/discrimination (L/D) task and the multiple-strategy (St+) task. We describe both tasks since they are different, yet were run daily in the same setup, illustrating the flexibility of the RIFF. The behavioral data analyzed here are from the St+ task, in which the rat had to constantly shift its position in order to harvest reward. Some of the electrophysiological data presented below come from the L/D task, in which rats had to identify the IAs from which a sound was presented (Fig. 3c).

5.2 Rats rapidly learned complex behavioral paradigms

We illustrate rapid learning in the RIFF in the multiple-strategy task (St+ task; see Methods for a detailed description). Habituation of the rats to the RIFF included learning to poke for reward and habituation to the sound presentations. However, during the habituation period, no association was made between rat actions, sounds and rewards. Following the habituation, the rats were introduced to the St+ task.

A behavioral episode of the St+ task started with an attention sound presented from all 12 loudspeakers simultaneously, following which the rat had 2.5 seconds to move to a position in the arena that determined the location of the next reward. The simplest such strategy consisted of moving from one IA to one of its neighbors, but other strategies could be adopted as well. Following the attention event, and depending on the location of the rat, one of the IAs was selected as the next target and its location was communicated to the rat by a target sound from the loudspeakers of that IA. The target sound was specific to each IA, but its only role was to indicate the availability of reward at the IA. The rat had to poke in one of the two ports of the target IA within 20 s of the presentation of the target sound. A poke or a timeout resulted in a feedback sound (correct feedback sound in case of a correct poke location, mistake feedback sound if the poke was in an incorrect port or in case of a timeout) followed by an inter-trial interval (3 s). After the inter-trial interval, the next attention event occurred. Despite the non-trivial structure of the task, the rats learned it very rapidly. We analyze here in detail the first two days in which rats were

exposed to the task, illustrating the ease with which rats can learn complex tasks in the RIFF. In later stages of training (not shown here) rats learned more demanding behavioral routines, including for example avoiding punishments in go/no-go schedules that punctuated the normal flow of the task.

Each rat was trained in the RIFF for 70 minutes per day. During the first two days of training, rats quickly learned to move from one IA to the next. Figure 4a shows one-minute segments of a rat's trajectory on Day 1 (4.5 minutes into the session) and Day 2 (20 minutes into the session). Erratic loops and hesitations that have been clearly present on Day 1 disappeared on Day 2. Figure 4b shows the angular position and angular speed from the same trajectory segments, along with a schematic illustration of the trial structure. The development of a stereotypic running pattern is evident especially when comparing the time courses of the angular speed in Day 1 and Day 2.

In order to analyze changes in behavior both within and across sessions, the following analyses use a subdivision of each session into two halves (the four resulting observation periods from Day 1 and 2 are called 1:1, 1:2, 2:1, and 2:2, respectively). The percentage of rewarded trials out of all trials - a simple measure of learning - improved from 32% rewarded trials in observation period 1:1 to 66% rewarded trials in observation period 2:2 (combined data from five rats; Fig. 4c). We modeled the percentage of rewarded trials as a linear function of the observation period using a linear mixed-effects model (LME) with random intercepts for rats. The model had a significant main effect of observation period with a slope corresponding to an average improvement of 13% additional rewarded trials from one observation period to the next ($t_{18} = 5.95$; $p = 1.2 * 10^{-5}$).

As rats learned to position themselves appropriately for collecting rewards, they also adapted to the timing of the task. Fig. 4d shows the time points at which rats approached the IA where a reward was upcoming, relative to the time of the corresponding attention sound (the relevant area was defined as the sector with a central angle of 30° symmetrically centered around the food/water ports of the rewarded IA). On Day 1, rats tended to arrive at the reward sectors shortly after the attention sounds. The median arrival time for successful trials on Day 1 was 0.4 seconds *after* the attention sound (interquartile range 3.7 sec). On Day 2, rats tended to arrive before the attention sound, with a median arrival time of 0.6 seconds *before* the attention sound (interquartile range 2.2 sec). Thus, the median arrival time was earlier by one second on Day 2 compared to Day 1. This suggests that rats learned the time structure of the experiment and correctly predicted the upcoming reward time. We modeled the logarithm of arrival times as an LME with the observation period as a fixed factor and with random intercepts for rats (times were shifted to start

5.5 sec before the attention sound to turn all arrival times into positive numbers; the logarithmic transformation was selected in order to reduce the skewness of the distribution of arrival times). There was a significant effect of observation period (estimate -0.069; $t_{1829} = -8.86$; $p = 1.8 * 10^{-18}$), so that the arrival time shifted to earlier times by a factor of exp(-0.069) = 0.93 from one observation period to the next.

The velocity trajectories in Fig. 4a suggest that the movement patterns of the rats became quite stereotypical very fast. To further elucidate this observation, we calculated the maximal angular speed in both the clockwise and counterclockwise directions during each trial (analysis time window, 0.5 to 9 seconds after the feedback sound of the previous trial). Plots of these two maximal speed values against each other revealed that all rats learned three clearly distinct behavioral motives: "not moving", "clockwise running", and "counterclockwise running" (See Supp. Fig. S4 for cluster analysis). Each experimental session can thus be described as a sequence of these three motives, timed by the rat to allow it to harvest rewards from the RIFF.

To understand how rat behavior changed over time, we analyzed the distribution of the three motives along the four observation periods. An important contribution to the change in rat behavior consisted of the decrease in the number of trials in which the rat did not move. We modeled the proportion of "not moving" trials as a linear function of the observation period using an LME with random intercepts for rats. Observation period had a significant main effect (estimate, -0.14; $t_{18} = -5.89$; $p = 1.4 * 10^{-5}$). Thus, on average, the proportion of "not moving" trials decreased by 14% from one observation period to the next: rats learned to move more often as training progressed.

Importantly, rats showed a clear, individual preference for running in one direction or the other. Thus, we found that the proportion of "clockwise running" out of all running trials (clockwise and counterclockwise running) depended on the rat (χ^2-test for independence; $\chi^2_4 > 701$ for the five rats together; $p < 10^{-20}$). Figure 4e illustrates these findings for two rats. Over the course of the four observation periods, the fraction of "not moving" trials generally decreased for both rats (shown in gray in Fig. 4e). Furthermore, rat 4 had a consistent preference for running in the counterclockwise direction (blue), while rat 5 strongly and consistently preferred running in the clockwise direction (red). The distribution of "clockwise"/"counterclockwise" trials significantly depended on rat identity in each observation period (for these two rats: $\chi^2_1 > 77.6$, $p < 10^{-20}$ in each observation period).

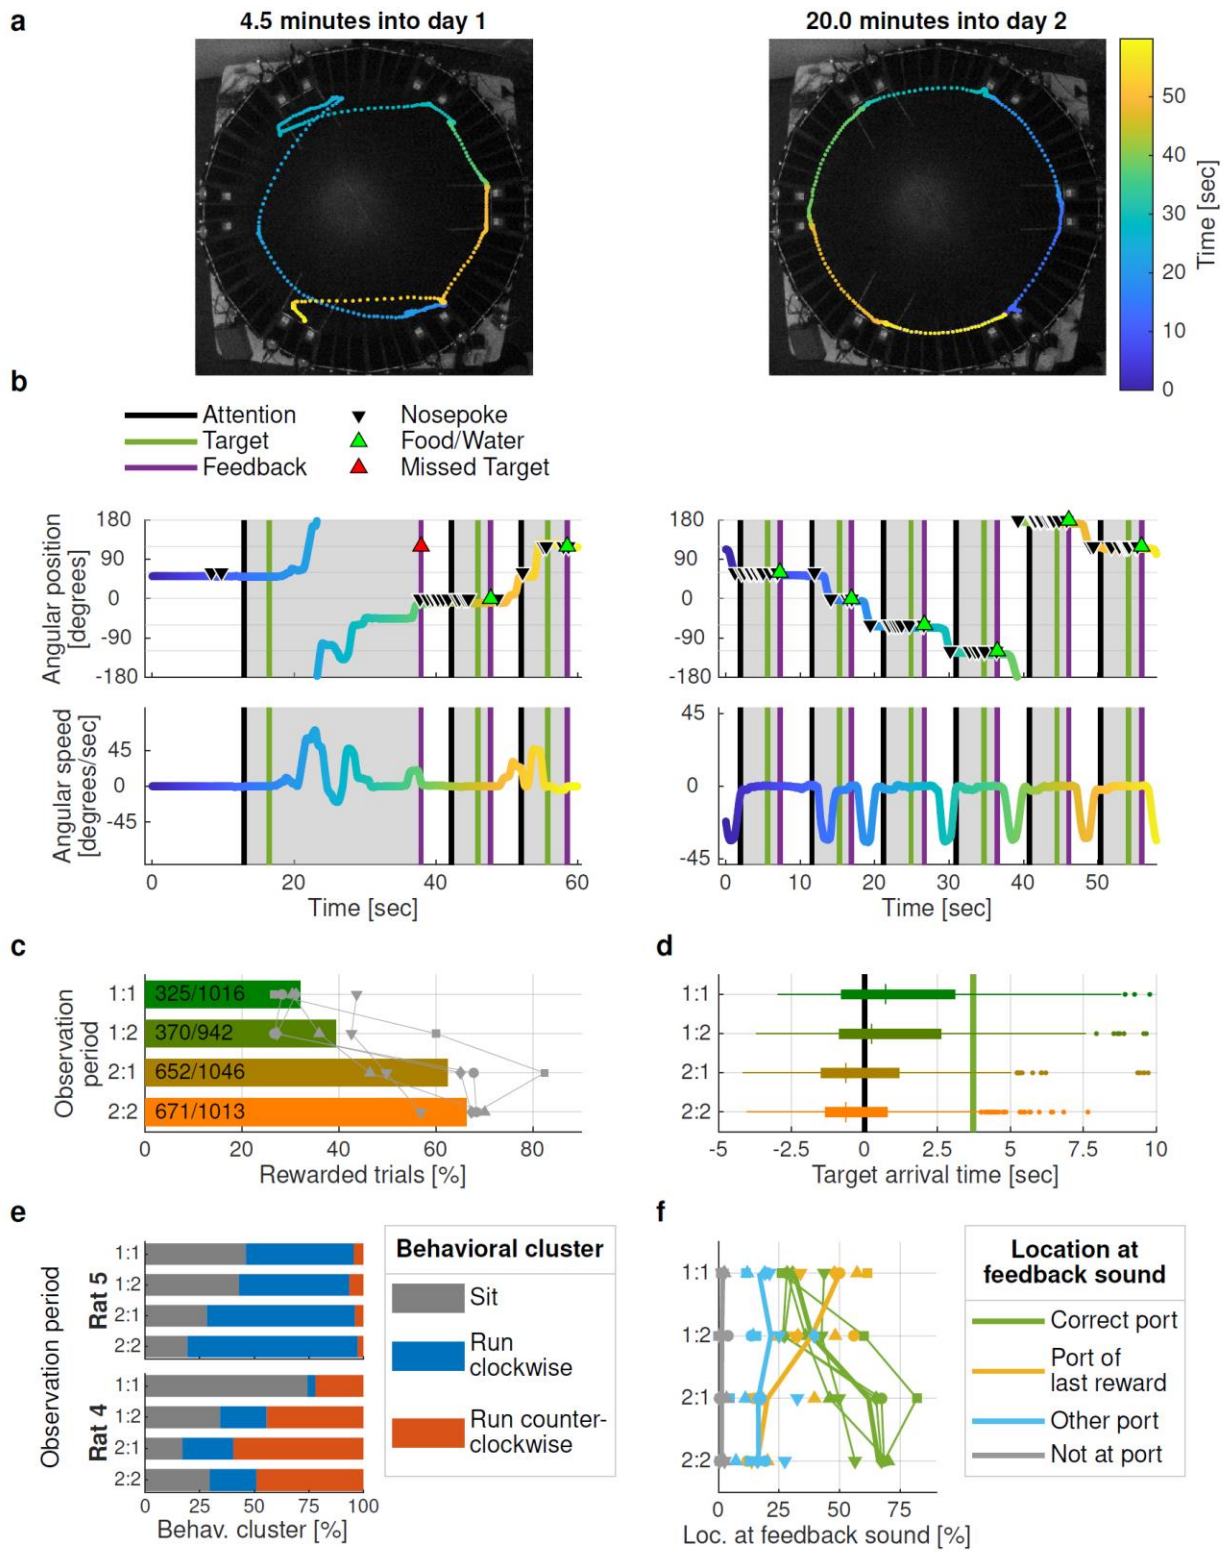

Fig. 4. Rapid learning in the RIFF. **(a)** One-minute trajectory segments taken from 70-minute training sessions on day 1 and day 2 of learning, respectively. **(b)** Angular position and speed for the same trajectories as in (a). Trial structure is indicated by blocks of Attention-Target-Feedback events. Nosepokes, food/water rewards, and missed targets are indicated by triangles. By day 2, rats have developed a

stereotypic running pattern and successfully collected rewards at a fast pace. **(c)** Success rate improved over time. Observation periods 1:1 and 1:2 refer to the first and second half of the training session on day 1, observation periods 2:1 and 2:2 refer to the first and second half of the training session on day 2. A successful trial is a trial in which the rat nose-poked at the correct location before timeout. Counts of successful trials out of all trials, summed over five rats, are displayed along with bars indicating average success rates; grey symbols indicate success rates of each individual rat. **(d)** Rats learned the temporal structure of the task. Boxplots show the distribution of arrival times relative to the attention sound (zero time, indicated by a black line; the presentation of the target sound is indicated by the green line), for each of the four observation periods, calculated from all trials of five rats. The median arrival time was 0.4 seconds after the attention sound on day 1 and 0.6 seconds before the attention sound on day 2. **(e)** Different rats learned different strategies. Each trial was classified as either sitting, clockwise running, or counterclockwise running. The proportions of each behavioral class for two example rats are displayed. Rat 4 preferred counterclockwise running over clockwise running in each observation period, while Rat 5 almost completely avoided counterclockwise running in each observation period. **(f)** Learning to move from one interaction area to its (clockwise or counterclockwise) neighbor was the most important contribution to improvement in task performance. Each trial was classified according to the location of the rat at the end of the trial as either "correct port", "port of last reward", "other port", or "not at port". Average proportions of each location across five rats are displayed as thick lines, individual rats are displayed as symbols and thin lines. The proportion of "correct" locations increased from one observation period to the next, and the proportion of "port of last reward" locations decreased over time, while the proportions of "other port" and "not at port" did not substantially change.

Lastly, to better understand the nature of failed trials, we classified the rat location at the time of the feedback sound into four types: “correct port” (rewarded trials), “port of last reward”, “another port”, and “not at any port”. As expected, as learning progressed all rats reached the correct port more often (green line in Fig. 4f). The most common error consisted of rats remaining in a reward location after receiving a reward, thereby missing the next reward opportunity (orange line in Fig. 4f). This type of error decreased substantially from Day 1 to Day 2 of training. Statistical analysis confirmed that the distribution of rat location types depended on the observation period (χ^2_{45} = 617, $p < 10^{-20}$ for all rats; $\chi^2_9 > 38.4$, $p < 1.5 * 10^{-5}$ in each individual rat). To quantify the interaction between the location types and observation periods, we modeled the probability of the location types as an LME in which the probability for each location type was a linear function of the observation period with a slope that depended on the location type, but with random intercepts for each location type in each rat. This model revealed a significant interaction between observation period and location type ($F_{3,72} = 44.1$; $p = 2.8 * 10^{-16}$), confirming that over time, the rats changed their prefered location types. Indeed, the slope of the "correct port" location type was 0.13 ($t_{72} = 8.36$; $p = 3.2 * 10^{-12}$), confirming the increase in correct trials over the four observation periods. In contrast, the slope for "port of last reward" location type was -0.12 showing a decrease in this type of error. The difference between these two slopes was highly significant ($t_{72} = -11.5$; $p = 6.1 * 10^{-18}$).

5.3 Neural correlates of behaviors

The rat insular cortex possesses an auditory field which is anatomically separate from other auditory cortical fields [34, 35]. This field will be termed Ins below. Rats in the St+ task were implanted with electrodes in Ins. Rats in the L/D task were implanted with electrodes at locations that made them traverse the primary auditory field at least in part of their trajectory. We will term their implantation location AC. Both tasks are described in detail in the Methods section.

Units in both the AC and Ins showed responses to sounds. Figures 5a and 5b show responses of units in AC of two rats. The stimuli were word-like stimuli - 200 ms word excerpts, with their spectral envelopes shifted and stretched to fit a 1-40 kHz range (see Methods). These stimuli were some of the target sounds for receiving reward in the L/D task (mean response of N = 96 and N = 221 sound presentations in Figs 5a and 5b respectively, paired t-test, $p = 6.49 * 10^{-7}$ and $p = 0.021$ respectively). Both units had well-timed onset responses. Figures 5c and 5d show units recorded in Ins. The stimulus in Figure 5c was the attention sound, indicating a new trial, consisting of a train of short broadband periodic sound bursts with a pitch of 2 kHz. The unit responded to the onset of the first sound burst in the train (mean of N = 233 sound presentations, paired t-test, $p = 8.75 * 10^{-14}$). Figure 5d shows a unit with a response that was loosely locked to a word-like stimulus (mean of N = 106 sound presentations, paired t-test, $p = 1.14 * 10^{-4}$).

Remarkably, in both areas, we found a substantial number of units which did not respond to sounds, but whose activity was strongly modulated by other aspects of the task and the state of the rats. Two of these classes of responses are illustrated here.

In Ins, many recording sites had units that increased their firing rates selectively at specific locations in the arena (Figs. 5e-h). Since the rats spent most of their time close to the walls of the RIFF in this task, the data are summarized as polar plots: the dark gray polar plot represents the average firing rate of the unit in each angular sector, while the light gray line represents the occupation probability of each sector. For example, the unit shown in figure 5e responded essentially only when the rat was at one specific sector (1-way ANOVA, $F(1,35060) = 221$, $p = 7.9 * 10^{-50}$). This sector was not preferred by the rat: it spent much more time in the two flanking sectors. To control for potential confounds, we checked that the firing rate of these units was independent of the time spent in each sector, and that their spatial sensitivity was largely independent of their sensitivity to velocity, head-body angle (see below), and sound presentations (supp. Figs. S5 and S6). For example, to check the independence of spatial sensitivity from velocity, we computed the matrix of mean firing rates as a joint function of spatial position and

velocity. We then computed a rank-1 approximation of this matrix using nonnegative matrix factorization. A good rank-1 approximation to the observed firing rates implies that the two parameters (location and velocity in this example) cause independent changes in the firing rate of the unit, and in particular imply that the location sensitivity is not a side effect of the velocity sensitivity of the firing rates. We therefore verified that this rank-1 approximation explained an appreciable amount of the structure of the joint dependence of firing rates on spatial position and velocity. For that purpose, we used surrogate data to estimate the expected variability of joint firing rate matrices estimated from a rank-1 model (see Methods for details). For the unit depicted in Fig. 5e, 77% of the variance of the observed firing rates was accounted by the rank-1 approximation. This should be compared to the 82% of the variance of the observed firing rates using bootstrapping from the observed firing rates, and 97% of the variance assuming the firing rates had Poisson statistics with the average firing rates given by the rank-1 approximation (for results for all units see supp. Fig. S5c-d). In general, the rank-1 approximation to the observed firing rate accounted for similar amount of variability as expected from the surrogate data, verifying that for the units presented here, spatial sensitivity was indeed independent of velocity. A similar approach was used to test the dependence on head-body angle (Supp. Fig. S5a-b). Responses in the presence and absence of auditory stimulation showed no significant difference (paired t-test, n.s. for each unit separately, see supp. Fig. S6) It is therefore tempting to consider these neurons as ‘place cells’ - encoding the location of the rat in space.

The second type of units illustrated here (Fig. 5i-l) showed sensitivity to the angle between the head and the body of the rat, extracted by the data analysis pipeline from the saved video stream. Such units were found in AC as well as in Ins. To control for confounds, we checked that their sensitivity to the relative head direction was not a consequence of their sound responses, since it was present in the absence of sound. Comparison of the responses in the presence and absence of auditory stimulation showed no significant difference (paired t-test, n.s. for each unit separately, supp. Fig. S7). The head-body angle sensitivity was independent of location in the RIFF, and independent of the absolute angles of the rat body and head in space (see supp. Fig. S8 for full results; for example in the unit depicted in Fig 5i, rank-1 approximation accounted for 93% of the variance in firing rates as a function of head-body angle and spatial position, compared to 96% using bootstrapping and 98% using Poisson statistics). The unit presented in Fig. 5h is the same as that presented in Fig. 5l. It showed sensitivity to both rat location and head-body angle, but those sensitivities were independent of each other.

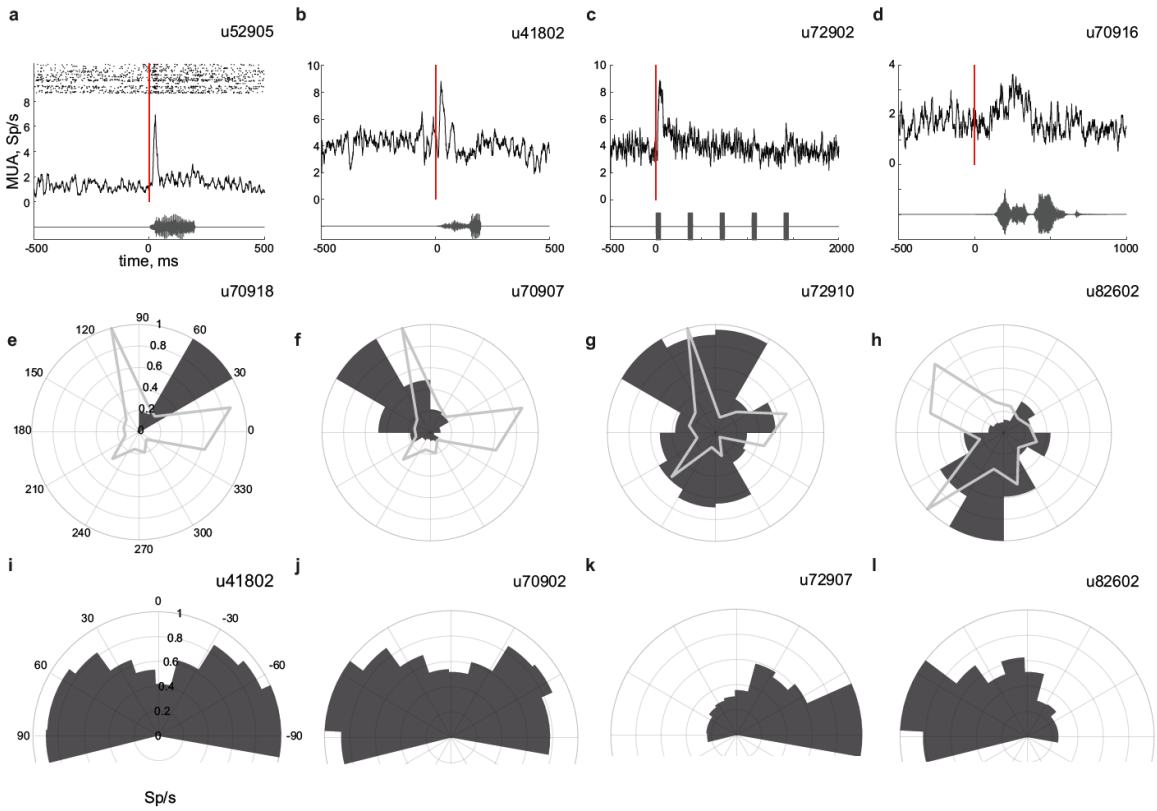

Fig. 5. Neural responses during behavior. **(a-d)** Sound driven responses. Mean over one session of the sound driven responses (top) and the oscillograms of the corresponding auditory stimuli (bottom). Red line indicates sound onset. The dots in (a) show the spike times in the individual trials as a raster display. **(e-h)** Location sensitive units. Each panel shows the mean firing rate of a unit while the animal was in one of 12 sectors of the arena (dark grey), and the time spent in each sector (light grey). **(i-l)** Units sensitive to angle between head and body of the animal. The normalized mean firing rate is plotted when the rat exhibited different angles between head and body (dark gray). The last bin on each side includes the data equal and more extreme than the indicated angles. Unit identifiers are marked above each panel.

6. Discussion

The RIFF is a high-throughput, high-resolution environment for studying the neural basis of behavior in complex tasks. The RIFF interacts with rats using sounds, food and liquid rewards, as well as with punishments such as airpuffs and potentially foot shocks, is fully automatic and can be used to shape behavior without manual interventions. It collects detailed online behavioral information together with synchronized neuronal recordings. We demonstrate here that this system tracks behavior and electrophysiology in a precise manner both in time and in space. Rats rapidly developed typical individual behavioral habits that could be easily documented and interpreted.

We showed that animals learned successful behavioral policies very quickly. This is somewhat surprising, as learning complex tasks can sometimes take weeks or months. Potentially, the large environment and the multiple possible actions available at each moment encouraged more extensive exploration[36,37]. Rats learned to move to reward locations before reward availability was signaled. This finding highlights the importance of continuous position tracking via camera, without which the rats' predictive behavior would have been unobservable. The St+ task described here allows for multiple rewarded behavioral strategies, and we showed that rats find different, individual solutions to the task. More traditional analysis, where only a binary success/failure variable is recorded per trial, would not have allowed for the discovery of inter-individual strategy differences. Moreover, the high-throughput design of the RIFF enabled us to study the stability of individual strategies across days in multiple rats. By combining the time series of rat position and MDP state in the St+ task, we identified the sub-component of behavior where most learning happened. This level of detail in the analysis of behavior would have been impossible in a setup that only tracks reward deliveries, but it can be crucial to guide the analysis of neuronal recordings. We currently use the RIFF to explore the extent to which rats learn more complex reactions for bigger rewards, learn to change fluently between simple and more demanding behavioral routines, and avoid punishments in go/no-go schedules that are embedded within the main tasks. The RIFF serves here as a naturalistic environment providing a set of rules and contingencies that can be approached with different, more or less effective behavioral strategies. The extent to which the habituation procedure, arena size, and the available actions (and possibly other factors) contribute to the speed of learning remains to be studied.

Many high-throughput systems for studying freely moving animals explore learning and plasticity[16], navigation[15,17,18,28,38] exploration [19], social interactions [19–26,39,40] and other behaviors[20,22,41]. Some of these systems may follow the behavior of many individuals simultaneously. However, they are hard to combine with neuronal recordings. A smaller number of published complex behavioral environments include wireless recording of neural activity from single freely moving animals, although with limited behavioral control[29,30] or with a simple behavioral loop that encouraged one specific behavior to obtain a reward[31].

Flexible exploration of animal behavior conjointly with neural responses in complex settings is often addressed by the use of a virtual reality setup[7–15, 42] (and others). Virtual reality systems are designed to solve issues of stability and quality of electrophysiological recording[42], and make it possible to use imaging techniques that are harder to apply in freely moving animals[43]. Still, virtual

reality is only an approximation, and may be associated with unnatural movement patterns[44,45] as well as with neural responses that differ from real environments[14,46].

The RIFF represents a different set of tradeoffs. It was designed for flexibility in the types of behavior that can be studied in it, as exemplified by the two tasks presented here, both of which were run daily (on separate subjects). The 'tight loop' of the main control program makes it possible to use information from many sensors (camera and nose-poke detectors) at short delays (<100 ms) to react to animal actions. At the same time, the RIFF makes it possible to collect high-quality electrophysiological signals (in the current implementation, with up to 64 electrodes at 32 kHz each).

The high-resolution data collected in the RIFF allow for good experimental control in spite of the substantial freedom of behavior that we allowed the animals to have. In consequence, we identified a number of unexpected and interesting connections between behavior and neural activity. We present here examples of neurons in auditory areas, in primary as well as high-order secondary fields, that show sensitivity to non-auditory features of the tasks. The highly detailed data collected from each animal allowed us to control these findings for a large number of confounds, such as concomitantly presented sounds, head and body directions in space, and animal velocity.

In particular, we provide here for the first time evidence for spatially-tuned neurons in an auditory cortical field. The spatially-tuned neurons we describe here lie in the auditory field of insular cortex (Ins). While this field is definitely spatially separate from the primary auditory cortex in rats[47], it receives direct projections from the primary division of the auditory thalamus, the ventral medial geniculate body[48], and has neurons with frequency-tuned short-latency responses[49]. Thus, the spatially-tuned neurons described here are intermixed with neurons that do not differ much in their response properties from neurons in the primary auditory cortex. This evidence provided here adds to other recent studies that described spatially tuned neurons in many structures outside the hippocampal system[50,51], including visual cortex[52], somatosensory cortex[53], retrosplenial cortex[54], and premotor cortex[55]. The role of these neurons in auditory processing, as well as in spatial navigation, requires further study.

Importantly, the RIFF is more than its current implementation: it represents a concept that can be easily extended by adding more cameras and other sensors for documentation of behavior (for example a motion capture setup [56, 57]), more actuators for interacting with the animals, and higher channel count electrodes for recording electrophysiological signals. The RIFF shifts the emphasis

in the design of an experiment and the analysis of behavior from single, simplistic measures of behavior such as success and failure into the documentation of a high rate, on-going process that is better adapted, and therefore easier to link, to the dynamics of the mammalian brain.

Materials and methods

Behavioral Setup

The arena, 160 cm in diameter, is made of 18 modular sections (28 cm width, 50 cm height - aluminum skeleton and 5 mm thick black matte Perspex sheets; one section is illustrated in Fig. 2b). The floor is made of two sheets, tightly connected by screws. The arena is mounted on a 5 cm high supporting structure, of the same size and shape of the floor sheets with depressions for the connection between the two sheets as well as for the insertion of gridded floors (for possible foot shock punishment). The arena is electrically grounded with a cable running around the outer perimeter near the bottom of the arena.

Each of the 18 sections contains three panels where additional equipment can be inserted. In the RIFF, the outer panels of every third section contain two nose-poke ports each (Fig 2; a total of 12 ports in 6 sections, [SPECIAL.090-SE v1.0, LaFayette-Campden Instruments, Loughborough, UK). One of the two ports in each section is connected to a liquid pump and the other to a pellet dispenser. In addition, all 12 ports are connected to an airpuff valve (80210raas, LaFayette-Campden Instruments, Loughborough, UK). At the top of each panel with a port, a free field speaker (MF1, Tucker Davis Technologies, Alachua, FL, USA) is mounted in a custom made holder with adjustable vertical angle. Such section, with its two ports and two loudspeakers, is termed an ‘Interaction area’ (IA) in the paper.

The ports are controlled by a combination of commercial software (ABET II, LaFayette-Campden Instruments, Loughborough, UK) and custom-written MATLAB program (The MathWorks, Inc., Natick, MA, USA). Signals to and from the ports interface with a digital I/O card (NI PCIe-6509 96 ch, National instruments, Austin, TX, USA) through specially designed hardware. The digital IO signals are monitored and changed by the custom-written MATLAB program that implements the MDP in the main computer.

Auditory Stimulation

In both tasks described here (St+ task and L/D task, described below), a small set of sounds was used. These sounds were synthesized ahead of time and stored in computer files. Pure tone stimuli were synthesized in MATLAB. Word-like stimuli were created using the following procedure. Words were selected from free online recordings (http://shtooka.net). They were then processed using the STRAIGHT Vocoder[58] and MATLAB. For sounds in the St+ task, the spectrotemporal envelope was extracted by STRAIGHT and shifted above 1 kHz. For the sounds in the L/D task, the spectrotemporal envelope extracted by STRAIGHT[58] was shifted and stretched along the frequency scale, so that the envelope levels at frequencies 0.1 kHz and 20 kHz of the original sound were shifted to frequencies 1 kHz and 40 kHz (with linear interpolation on a logarithmic scale in between). The resulting spectrotemporal envelope of all sounds was used to resynthesize sounds using the original pitch contour.

During the experiment, all sounds were loaded to memory. When a sound had to be presented, it was transduced to voltage signals by a sound card (M-16 and MADIface USB, RME Audio Interfaces, Munich, Germany), attenuated (PA5, Tucker Davis Technologies, Alachua, FL, USA), and then played through one or more of the 12 speakers, driven by stereo power amplifiers (SA1, TDT) .

The acoustics of the arena were studied in detail[59]. In short, the arena produces two major reflections, one from the floor and the other from the wall behind the animal. The dominant reflection from the floor creates spectral modulation with a period of 5 kHz and depth of about 5 dB.

Behavioral training

The experiments were carried out in accordance with the regulations of the ethics committee of The Hebrew University of Jerusalem. The Hebrew University of Jerusalem is an Association for Assessment and Accreditation of Laboratory Animal Care (AAALAC) accredited institution.

Seven adult female Sabra and Sprague Dawley rats were used (Envigo LTD, Israel). All efforts were taken to create a low-stress, rat-friendly living environment enabling experimental animals to freely express their innate behaviors. Upon arrival, animals were housed in groups (2-3) or individually depending on experiment in the same room in which the RIFF is situated and where all experiments took place. For the first 7-10 days rats were habituated to the experimenter. During

this period, rats were handled frequently and introduced to new types of foods and flavored water which were later used as rewards in the experimental arena. The procedure was aimed to decrease stress and prevent neophobia towards rewards given during behavioral training. Rewards consisted of palatable 45 mg food tablets in six flavors: basic company flavor, banana, bacon, chocolate, peanut butter and fruit punch (diet AIN-76A: RodTab45MG; RodTabBan45MG; RodTabBcn45MG; RodTabChoc45MG; RodTabPbtr45MG; RodTabFrtP45MG5TUL, TestDiet, Richmond, IN, USA). Rats received also 4 types of fluids: mineral water, 4% sucrose in mineral water, 0.1% saccharine in mineral water, and 1:1 mixture of 4% sucrose and 0.1 % saccharine solutions (sucrose and saccharin, Sigma-Aldrich, St. Louis, MO, USA).

In order to motivate animals to work for food, rats were food-restricted up to 85% of their ad libitum body weight. Rats were subjected also to mild water restriction before behavioral sessions (4-12h). Experiments were carried out 5 days a week and during the weekend animals had free access to standard rodent food and unsweetened mineral water.

Multiple strategy task (St+)

In this task, rats had to move to specific locations at defined times in order to obtain rewards. The experiment consisted of a repeated iteration over three events: Attention (ATT) - Target (TGT) - Feedback (FDB). Each event was associated with a specific sound that was played whenever the event occurred. The exact paradigm is described in Fig. S1a. The ATT event consisted of the presentation of an ATT sound - a 50 ms broadband 2 KHz periodic sound that was repeated 5 times at intervals of 0.3 s. It was played from all 12 speakers simultaneously. The position of the rat was determined 2.5 s after the ATT event, and one of the IAs was selected as the target location. The selection was done according to the current and previous position of the rat (see below for details). Then a TGT sound was presented from the target location. Each IA was associated with a specific TGT sound. The TGT sounds consisted of a sequence of six 50 ms long narrowband sounds with inter-sound intervals of 250 ms covering a range of 1-16 kHz. Six different permutations of those sequences were assigned as target sounds to each IA for the whole experiment (target area 1: 1, 1.7, 3, 5.2, 9.2, 16 kHz; target area 2: 1.7, 16, 1, 9.2, 5.2, 3 kHz; target area 3: 3, 1.7, 1, 16, 9.2, 5.2 kHz; target area 4: 5.2, 3, 1.7, 1, 16, 9.2 kHz; target area 5: 9.2, 16, 5.2, 1, 3, 1.7 kHz; target area 6: 16, 9.2, 5.2, 3, 1.7, 1 kHz). After the TGT event, the rat had to nose-poke into one of the two ports at the target IA. Nose-poking within 20 seconds after the TGT event was rewarded by food pellets or water, depending on which of the two ports of the IA were poked. Reward was provided only if the first nose-poke was in the correct IA, and

only to that first nose-poke. The first nose-poke (in a correct or wrong location) elicited the FDB event, consisting of a sound indicating whether the location was correct or not. A timeout (20 seconds without poking after a TGT event) was also followed by an error FDB sound. After each FDB event, an inter-trial period of 3 seconds led to the next ATT event.

The target locations were selected in the following way. The arena was subdivided (Fig. S1a) into A zones, in the immediate vicinity of the IAs; B zones, adjacent to the A zones but 17 cm hubward; C zones, separating the A/B zones of nearby IAs; and a large D zone in the center of the arena. The A zones were slightly wider (33 degrees) than C zones (27 degrees). When the rat was in the D zone, nothing happened. If the rat was in one of the A/B/C zones, an ATT sound was presented 3 s following the last FDB event. The rat had to make a choice where to go within the next 2.5 s. The rat could stay in the same A zone in which it received the last reward, in which case the target was presented from the next A zone in the counterclockwise direction. It could move to another A zone, which was then selected as the target. The rat could therefore move from one A zone to the next, harvesting reward from each IA in turn. It could move to a B zone, in which case the next target consisted of the adjacent IA. The rat could therefore cycle between adjacent A and B zones, continuously harvesting reward from the same port. The rat could move to a C zone, in which case a random target was selected. In the data shown here, rats learned an 'A' strategy - moving around the arena from one IA to its neighbor, circling the arena. In fact, the rats rapidly learned the timing structure of the task, and often moved to the next A zone before the ATT sound was presented (Fig. 4d).

The rats were trained in three phases. During phase 1 rats were trained to poke for reward. Multiple food and water restricted rats were placed together in the arena for about 12 hours for five consecutive days. The sounds were not presented at this point. The rewards were delivered in a fixed ratio (FR) schedule. The number of pokes needed to get reward was increased every day as follows: FR2, FR7, FR10, FR12, FR15.

In phase 2, rats established instrumental and pavlovian associations with relevant sounds and their sequences. Rats were trained individually, each during a 70 min long session performed once a day. The training was carried out for 30 days. The sound level was increased gradually to avoid suppression of behavioral responses and to introduce animals gradually to the relatively high SPL of the sounds during the main experiment. Each day, only 1 IA was accessible for a rat during the session. Interactive ports in all other 5 areas were closed by custom blinds. Consecutive sessions were carried out in areas 1 to 6 and such cycle was repeated 5 times.

There were 3 types of sounds presented in this phase. The rat started the trial by poking into the accessible port. Each poke (photo beam break) evoked the ATT sound. After five successful pokes with the ATT sound, the sixth poke triggered reward delivery (1 unit) and immediately the target sound (TGT) was played from the two speakers above the accessible IA. The TGT sound was followed by 1 s of silence and the word-like stimulus derived from the word "correct" played from two speakers above the accessible IA. An obligatory 2 s interval followed before the RIFF reacted again to nose-pokes of the rat. Thus, during that stage, while sounds were associated with reward, it was the nose-poke that triggered the sounds rather than the reverse (as in the main task).

Phase 3 was the main experiment as described above. Rats performed a 70 min long experimental session once every day.

The behavioral data in Fig. 4 was collected during the first two days of Phase 3 of the task. However, the electrophysiological data shown in Fig. 5 was collected after one additional refinement of the task (phase 4). In randomly selected 10% of the trials, a warning sound was presented (the phrase, composed of word-like stimuli: "do not go'', played from all speakers). The rats had to avoid accessing the ports in order to avoid air-puff, starting 1.5 s after the warning sound (to let the rat retract its nose from the port within in case it was poking when the warning sound was presented). An airpuff was triggered by a nose poke for pokes that occurred during the next 3.5 s. If the rat successfully avoided poking any of the ports, a safety sound (the phrase, composed of word-like stimuli: "it's safe") was presented from all speakers. Otherwise, an airpuff was delivered and the phrase "you failed" (composed of word-like stimuli) was presented from all speakers. After safety or punishment sounds, a new trial was initiated with ATT sound.

Localization/discrimination task (L/D)

In this task, each of the IAs was associated with a sound. The sounds were 200 ms excerpts of 6 different words (the word "here" in 6 languages: English, French, German, Italian, Polish and Russian) in a female voice, that were processed using STRAIGHT vocoder to fit the hearing range of the rats (see Auditory stimulation section for details). The rats had to go to a central circular area to initiate a new trial (Fig. S1b). When a crossing into the central area was detected, one of the 6 IAs was randomly selected by the program and its associated sound was played once every 2 seconds from both speakers for up to 10 times (20 s in total), or until the rat poked in any port. A poke in one of the two ports of the IA from which the sound was played resulted in a reward

(food or water, according to the poked port). A wrong poke or a timeout (20 s without a poke) resulted in the termination of the trial.

The rats were habituated to the RIFF during 2 nights (24 hours in total). During habituation, The RIFF was divided into 6 equal sectors centered on the IAs (supp. Fig. S1b). Rats were placed in the RIFF and were able to explore freely. Whenever the rat entered an active sector, the associated sound was played from the speakers of the IA in the sector. The sound repeated every 2 s until the rat poked in one of the ports of the IA and got a reward, with a timeout after 2 minutes. After a nose poke, the current sector and its two neighbors became inactive and the three other sectors became active, so that the rat had to cross to one of the three sectors on the opposite side of the RIFF in order to initiate a new interaction.

Following the habituation, the rats were exposed to increasingly stricter versions of the main task in 12 hour overnight daily sessions. For the first 6 training sessions the rats were allowed to poke in more than one port before trial termination (10 pokes in the first and second sessions, decreased to 6 and 4 for one session each, 2 pokes for 2 sessions, then down to 1 for the rest of the experiment). In addition, the requirements for initiating a trial became stricter: on the first training session, the central area had a radius of 50 cm, on the 2nd session of 40 cm, and from the 3rd session on the central area had a radius of 30 cm. Before electrode implantation rats were switched from 12 hour overnight sessions to 3 hour morning sessions.

Electrodes

Rats were chronically implanted with 32-channel silicon probes (ASSY-116_E-2, Cambridge Neurotech, UK). Prior to implantation thin flexible ground wire was soldered to the electrode's PCB ground contacts. The electrodes were aligned and glued to the microdrive shuttle with epoxy (Nano Drive, Ronal Tool Company, Inc., PA, USA). The microdrive with the electrodes and connector were held by a custom holder attached later to a stereotaxic apparatus for implantation. The moveable parts of the microdrive were covered with paraffin oil to prevent possible leak of dental cement between them. Before implantation, the electrodes were cleaned by washing them with a 4% tergazyme solution in purified DDW water (Alconox, Jersey City, NJ, USA) and afterwards carefully washed in purified DDW to remove all residues of tergazyme. Before implantation the electrodes and the holder were disinfected in a UV sterilizer.

Surgery

The implantation of the silicon probes was performed in two stages: (1) preparation of the base for the implant and (2) implantation of microelectrodes into the brain tissue.

1. Preparation of the base

Rats were initially anesthetized in an induction chamber with sevoflurane (8% in oxygen, Piramal Critical Care Inc., Bethlehem, PA, USA). The head was shaved and they were placed in a stereotaxic instrument with a mask for gas anesthesia (David Kopf Instruments, CA, USA). Sevoflurane concentration was slowly adjusted to a level of 2–2.5% and maintained at this level throughout the surgery. A surgical level of anesthesia was verified by the lack of a pedal-withdrawal reflex and slow, regular breathing rate. Body temperature was controlled by a closed loop heating system with a rectal probe (Homeothermic Monitoring System, Harvard Apparatus, MA, USA). The eyes were protected with sterile eye drops for dry eyes (Viscotears Liquid Gel, Carbomer: polyacrylic acid 2 mg/g, Berlin, Germany) and the skin on the head was disinfected with a povidone-iodine solution (10%, equivalent to 1% iodine, Rekah Pharm. Ind. Ltd., Holon, Israel). To prevent postoperative pain, rats received during the surgery subcutaneous injection of Carprofen 50 mg/ml (5% W/V) in a dose of about 12 mg/kg (Norocarp, Norbrook Laboratories Limited, Newry, Co. Down, Northern Ireland).

A 1.5–2 cm longitudinal cut of the skin on the head was made and the dorsal surface of the skull was exposed. The opened skin was stretched and the eyes were closed. The left temporal muscle was pulled away to expose also the lateral surface of parietal and temporal bones. The connective tissue covering the bones was removed and bones were treated with a 15% hydrogen peroxide solution (Sigma Aldrich Inc., St. Louis, MO, USA) which was washed off with sterile saline after approximately 10-20 s. When the surface of the skull was clean and dry, a reference point for the entry point of the recording electrodes was marked (Insular cortex (InsC): AP = -1.0 mm, ML = -6.1 mm; Primary auditory cortex (A1): AP = -5.1 mm, ML = guided by landmarks on the lateral surface of parietal and temporal bones). Subsequently, 7 small holes for supporting screws were drilled and screws were tightly screwed into the frontal, parietal and occipital bones. Ground wire, soldered previously to one of the screws, was placed in the left frontal bone. The screws were fixed together and to the bone first with resin and then acrylic dental cements (Super-bond C&B, Sun Medical, Moriyama, Shiga, Japan; Coral-fix, Tel Aviv, Israel) forming a base of the implant. At the selected electrode implantation site, a thin polyimide tube was placed on the skull vertically

and cemented with the rest of the implant base. The tube served as a guide to the implantation site for the next surgery. The free end of the ground wire was twisted and covered with a polyethylene cap cemented to the rest of the implant.

The wounds were cleaned and treated in situ with antibiotic ointment (synthomycine, chloramfenicol 3%, Rekah Pharm. Ind. Ltd., Holon, Israel). The skin was sutured in the anterior part of the implant with one or two sutures (Nylon, Assut sutures, Corgémont, Switzerland) to stretch the skin around the base of the implant. The skin around the wound was cleaned and covered with a povidone-iodine solution (10%). The rats received intraperitoneal injection of the antibiotic enrofloxacin 50mg/ml (5% W/V) in a dose of 15 mg/kg diluted with saline to 1 ml (Baytril, Bayer Animal Health GmbH, Leverkusen, Germany). After surgery, animals were housed individually to prevent them from chewing the implants. Carprofen or other similar NSAID dissolved in palatable wet food was provided at the home cage for the first few days after surgery. The rats were allowed at least 1 week of recovery post-surgery before restarting behavioral training.

2. Implantation of the microelectrodes

When the wounds were completely healed following the first surgery (14-28 days), the recording electrodes were implanted. To minimize tissue damage, a small craniotomy was made, drilling solely through the base of the implant, and thus leaving the healed skin intact to accelerate recovery and reduce the pain.

As previously, rats were initially anesthetized in an induction chamber with sevoflurane (8% in oxygen). After induction, rats were transferred to a stereotaxic instrument with a mask for gas anesthesia. Sevoflurane concentration was slowly adjusted to the level of 2–2.5% and maintained at this level throughout the procedure. The eyes were protected with sterile saline and body temperature was controlled by a closed loop heating system with a rectal probe.

The dental cement above the implantation site marked by polyimide tube was removed gradually using a dental drill until the skull was exposed. The craniotomy was performed by drilling, and a 0.4-0.8 mm long slit in the dura was gently resected. The electrodes were slowly inserted into the brain tissue using a single axis micromanipulator (MO-10, Narishige, Tokyo, Japan). The craniotomy was sealed with paraffin oil and elastic silicone polymer (Duragel, Cambridge Neurotech, UK). The microdrive and connector were fixed to the base of the implant with acrylic dental cement. A ground wire was soldered between the base and the electrodes connector and

covered with acrylic dental cement. To mechanically stabilize the connection between the implant and the recording device, an additional supporting connector was cemented to the implant (6 pins of 853 Interconnect Socket, MILL-MAX MFG. CORP., New York, USA). A custom plastic enclosure with a screw cap was cemented for implant protection. The robust design of the implant protected the microdrive and the electrodes against mechanical stress while the rats were moving, improving recording stability (supp. Fig. S9d-f). The weight of the whole construct didn't exceed 11.5 g.

At the end of the surgery, to prevent postoperative pain, rats received subcutaneous injection of Carprofen 50mg/ml (12 mg/kg). Rats received intraperitoneal injection of enrofloxacin 50mg/ml (15 mg/kg) diluted with saline to 1 ml. Rats were allowed at least 3 days of recovery post-implantation before recordings.

Wireless electrophysiology in freely moving rats

Two wireless recording systems were used: (1) modular 64 channel neural logger (RatLog-64, Deuteron Technologies, Jerusalem, Israel) and (2) 64-channel wireless transmission system (TBSI W64, Triangle BioSystems International, Durham, NC, USA).

1. Neural logger

For the experiments described here we used a single processor board and a single amplifier board to record from 32-channel silicon probes at a sampling rate of 32 kHz. The data were saved on a 64 GB micro SD card, and copied to a computer after each session. The logger was equipped with an audio microphone and 9-axis motion sensor. An electrically insulated, compact enclosure for the logger, with a separate compartment for a battery, protected the logger from mechanical shocks (supp. Fig. S9). The enclosure included an interconnector, to protect the board connector against mechanical damage caused by attaching and releasing over hundreds of recording sessions (supp. Fig. S9a). The interconnector was also used to rotate the device about 34 degrees backwards from the vertical axis, to allow the rats free movement and free access to the ports in the IAs (supp. Fig. S9b). The enclosure of the logger had a supporting connector (6 pins of 852 + 853 Interconnect Socket, MILL-MAX MFG. CORP., New York, USA) that matched the supporting connector attached to the implant, in addition to a 36-pin omnetics connector (A79029-001, Omnetics Connector Corporation, Minneapolis, MN, USA) for the electrical signals from the

electrodes. The total weight of all elements was approximately 16 g. The load was balanced, to avoid pulling the rat head in any direction. Before every recording session, the device was additionally secured with autoclave sticky tape (Sigma-Aldrich, St. Louis, MO, USA) to prevent loosening of the connectors during the 2-3 hours duration of the recording sessions. Maximal recording time using a 300 mAh battery was about 3 hours (LiPo battery 582030).

2. Wireless analog transmission system

The system consisted of an analog transmitter and receiver (TBSI W64, Triangle BioSystems, Durham, NC, USA). The output signal of the receiver was routed to a data acquisition system (AlphaLab SnR, Alpha Omega) and digitized at a sampling rate of 44 kHz. The recordings were made from 32-channel silicon probes as above. A small interconnector with a battery holder was prepared in the lab. The headstage was placed horizontally with a battery holder on its left side enabling animals to freely access to all locations in the experimental arena, and in particular to the IAs. The total weight of all elements was approximately 13.5 grams of well-balanced load. Before each recording session the device was additionally secured with autoclave sticky tape (Sigma-Aldrich, St. Louis, MO, USA) to prevent loosening of connectors contact during long active behavioral sessions (up to 12 hours). Maximal recording time using a 260 mAh battery was about 11.5 hours (LiPo battery 601240).

Recordings

Neural signals were recorded in reference to a ground placed in the frontal bone. For the logger, the analog bandpass filter was set to 10-7000 Hz or 300-7000 Hz depending on the experiment. For the wireless system, the low frequency of the bandpass filter was set to 0.07 Hz. Recording sessions took place five days a week and implants were checked daily. Spiking activity was screened immediately after each recording session was finished. The electrodes were kept in the same position as long as spiking activity was detected on many contacts. When signals deteriorated, the animals were briefly sedated with sevoflurane and the electrodes were lowered in steps 25, 50 or 100 µm into the brain tissue. Electrodes were moved typically every 1-7 days. Electrodes were never moved up.

The video imaging system

A custom video system tracked the trajectory of the animal in real-time at 30 frames/second (supp. Movie S1). The localization latency was less than 30 ms and a compact storage format resulted

in data volume of 4.3 GB/h. The system robustly tracked rats of different sizes and was unaffected by the presence of a head implant.

1. Hardware

A monochrome camera (DMK 33G445 GigE, TheImagingSource, Bremen, Germany) was mounted on the ceiling above the center of the arena. A wide-field lens (T3Z3510CS, Computar, North Carolina, US) covered the whole arena. A LED ring (50 cm in diameter, CN_R64, NanGuang, Chenghai, China) placed around the camera provided a uniform diffuse illumination (figure s1.a). The luminance of the ring was set to the minimal level that was sufficient for exposure time of 25 ms.

The camera was connected to a computer running the Windows operating system with a GigE cable that supplied power and transmitted the grayscale 1280x960 pixel images. The triggering pulses were sent from the camera by a Hirose cable to the digital input section of a multifunction device (RX8, Tucker-Davis Technologies). The RX8 sub-sampled the incoming trigger train from 30 Hz to a 1 Hz pulse train and output it to the main recording unit (AlphaLab SnR, AlphaOmega). These 1 Hz triggers also powered a LED that was placed outside the arena but inside the field of view of the camera (supp. Fig. S10b). The LED state was used to synchronize the images with the rest of the recorded data during the experiment.

2. The real-time module – image acquisition and rat localization

A custom graphical user interface (programmed in MATLAB) was used to launch the tracking program and monitor its activity in real time (supp. Fig. S10c). A single background image of the empty arena was acquired during the initialization of the tracking program. The real-time tracking routine executed an acquisition-localization tight loop: Every time a new image was acquired from the camera, the rat was located by subtracting an image of the empty arena from the current image of the arena with the rat inside. The difference image highlighted the area where the rat was located. The location of the rat was estimated by fitting an ellipse to the rat pixels and computing its center of mass. The x and y coordinates of the estimated rat center of mass were sent to the computer that controlled the experiment over an Ethernet connection and are also stored as metadata.

To reduce data storage, the video stream was saved in a compact form. Each rat image was cropped around the rat center of mass, since only this area was expected to change from one

frame to the next. The original image could be approximately reconstructed by placing the cropped rat image on top of the background image at the location indicated by the coordinates of the center of the rat, stored in the metadata file. In addition, one of the corners of the bounding box was replaced by a small rectangle that showed the LED (supp. Fig. S10d). The images and the metadata were stored in a compact database at the rate of 4.3 GB/h. The processing of a single video frame required less than 30 ms.

Post-experiment data processing

After the conclusion of each experiment, data were processed by a custom, automated pipeline written in MATLAB (tested with version 2020b). Data processing started with the extraction, alignment, and merging of the time stamps from all data acquisition devices. Time stamps were generated for changes in internal states, the acquisition of video images, behavioral events, sound presentations, and neural recordings. The digital lines of an AlphaLab SnR system (Alpha Omega) were used as the main synchronization hub for aligning the triggers coming from all devices. Each device that generated timing triggers had at least one of those triggers also channeled in parallel to a digital line on the main synchronization hub, and the timing events that were co-registered on both devices were used to align all other events from that device on a common timeline. Thus, the output files of the SnR system were used as the main timing files during post experimental data processing. All behavioral events (nose-pokes, food/water delivery, foot shocks) as well as the video images were parsed, and the extracted information stored in MATLAB files. Finally, the neural recordings were processed, detecting and sorting spikes (see next section).

Spike Sorting

We developed a pipeline that processes an experimental session of 3 hours in about 10 minutes, based on the Kilosort2 spike sorting program (https://github.com/MouseLand/Kilosort2). The processing is predominantly autonomous, requiring only a few minutes of manual curation for every session.

Spike sorting began by removing common modes from the neural recordings - the mean activity of all spatially proximal recording channels, such as all channels on a single shank of a silicon probe (supp. Fig. S3a.1-2). Once the common mode was removed, large artifacts were identified by computing the standard deviation over electrodes at each point in time, and locating peaks in

the resulting time series. Signal sections around each of these peaks were zeroed by multiplication with a 10 ms window with 2 ms rising and falling edges (supp. Fig. S3a.3). Such artifacts were mostly generated by rat movements. The total recording duration that was zeroed during a typical experimental session was about 5%.

Spikes were detected and clustered by the Kilosort2 program. In order to process large amounts of data, we developed a fast and almost automatic pipeline, which can be complemented by detailed manual curation where necessary. Kilosort2 was used to produce the initial set of candidate spike clusters. Automatic rejection of noise clusters was performed according to two criteria (supp. Fig. S3b): The dominant channel of the cluster template has exactly one minimum and one maximum in its voltage trace, and on each electrode the voltage at the template endpoints had to be close to 0 (in consequence, confining expected spike duration to <1.8 ms). For each unit that passed this test, an inter-spike-interval histogram was computed alongside the amplitude and firing rate histograms, heat map of raw spike waveforms, and smoothed time series of spike amplitude and firing rate (supp. Fig. S3c). All those were printed to an image file that summarizes the quality of that unit.

Lastly, a manual labeling interface allowed the experimenter to sort the units into well isolated ones, multi-unit clusters or noise. This interface displayed the images with the quality evaluation information and stored the user-inserted labels with the data. A typical recording session of 1-2 hours with 32 electrodes produced typically a few tens of non-noise units, so that the manual labeling could be completed in several minutes.

Feature extraction from rat images – convolutional neural network

Pose information was extracted from the saved video stream, offline, in two steps: First, the nose, the base of the neck and the base of the tail were located on the image by a convolutional neural network (sup. Fig. S2a, supp. Movie S2). Then, the relevant angles were computed by geometric calculations (supp. Movie S3, green arrow represents the body direction). The neural network was trained de-novo, including data labeling, model training and time optimizations of the inference process.

1. Network architecture

The input to the model is a square grayscale image of the rat, as produced by the rat tracking module. The output consists of three pairs of Cartesian coordinates, representing the three points of nose, neck and the base of the tail on the rat image. The model is thus a multivariate regressor that predicts 6 continuous variables.

We used a custom convolutional neural network architecture, following the recent advances in data-driven computer vision (supp. Fig. S3b). Six convolutional layers are stacked in a feed-forward manner, each accompanied by a batch-normalization layer and a rectifier non-linear activation function. A max-pool layer operates after every second convolutional layer. The last layer is a linear readout from the previous activation map stack. The network has 0.3 million trainable parameters.

2. Database creation

We used a custom programmed graphical user interface (written in MATLAB) to manually mark the locations of the nose, neck and base of the tail on each image of the training set. The final training database included 1500 labeled images of 4 rats of different ages, with and without a head implant.

3. Network training and inference enhancements

The training database was augmented to produce additional training examples and to emphasize the irrelevant degrees of freedom in the image space. The augmentation included rotation of images by multiples of 90 degrees, horizontal mirror flips, adding small amounts of noise to pixels or labels, and rigid translations of images by a few pixels along the horizontal and vertical axes. The labels were re-adjusted for each augmentation type to correctly represent the updated body part location. This augmentation procedure increased the original dataset by a factor of 64, from 1500 to 96,000 training examples.

The network was initialized with the default Kaiming distribution[60], trained with the Adam optimizer[61] on 256 images in each batch. The input images were subsampled to 50x50 pixels. The network was trained for 2 hours on a NVIDIA GTX 1080Ti GPU (NVIDIA, California, US).

After the training was completed, the network could be used to predict the labels of previously unseen images. Prediction used image augmentations that resemble those used during the training: An input image was processed 4 times, rotated by 0, 90, 180 or 270 degrees. The predicted 4 labels were corrected for the corresponding image rotation and averaged to obtain the final label for the original image. The high temporal correlations of the image stream were leveraged to increase the accuracy by smoothing the labels with a uniform window of 7 frames, removing large, biologically implausible jumps.

An hour of a typical experiment produced ~10^5 images, which were labeled by the trained model in ~20 seconds.

4. Computation of allo- and ego- centric directions

Once the locations of the nose, neck and the base of the tail are predicted by the network, they are used to calculate the direction of various body parts. The main body direction is defined as the vector from the base of the tail to the base of the neck. The head direction is defined as the vector from neck to the nose point. The head turn angle is calculated as the difference between the head direction and the body direction, with 0 corresponding to the three points being along a single line. The range of head turn angles was typically between -30 and 30 degrees.

Final database formats

The data processing pipeline parsed each experiment into a set of data files, metadata descriptors and statistical visualizations. This format is uniform and allows an abstraction of the experiment parameters such as the type of neural recording system, the stimulus types, and so on. All behavioral and neural events are stored in a lossless manner, allowing maximal temporal resolution. The storage footprint of this format is about 12GB/hour.

For rapid analysis and visualization, a temporally coarse-grained version is generated by discretizing the timeline into a 10 Hz grid, and aggregating the separate data files into a single data table (stored in one file). This format uses about 1% of the storage of the raw storage (<0.1GB/hour). Each data row in the resulting data table has a time-stamp that indicates its start time on the time-axis of the original data. These time-stamps can be used as indexing variables for tracing the data back to the high-resolution files. The 10 Hz resolution is high enough to

faithfully represent rat behavior and many of the firing rate fluctuations that are related to it, but is too coarse for detecting well-timed sound responses. The analysis of such responses requires accessing the data at the original temporal resolution.

The experiment visualizer

All processed data from an experiment can be loaded into an interactive user interface that presents the recorded experiment as an interactive movie (at a rate of 10 Hz, see fig. 3.f). The raw neural data is loaded along with the detailed time information of the sorted spikes, making it possible to view both the firing rate of each unit and the times of individual spikes with a 1 ms resolution.

The visualizer is used for multiple purposes. During the development of the experiment, replay of the experiment with 100 ms resolution makes it possible to verify that the experimental logic is correctly implemented and executed by the hardware, and the rat indeed performs the required sequence of actions. The visualizer is also used for verifying the collected data; the temporal resolution of the visualizer makes it possible to detect potential data collection malfunctions such as synchronization errors. During data analysis, the visualizer can be used as a tool for understanding fine details of rat behavior or for observing the progress of the learning process. Lastly, the visualizer can plot the raw neural traces at any time point of the experiment, making it possible to observe the relationships between spiking activity and other events in the RIFF. Figures 3.c, 3.f show images produced by the visualizer. The visualizer was used to generate supp. Movies S4 and S5.

Statistical analysis

Exploratory analysis of the dependence of spike trains on any of the measured parameters was conducted using linear mixed effects models (MATLAB function fitlme). Explanatory variables included location and kinematic parameters, body and head direction parameters, and sounds presentations. The models were fitted to the firing rates of the individual units collected in all experimental sessions of each rat. To check the effect of any of these parameters on the neuronal responses, it was used as a fixed effect (in order to account for a non-zero mean across the neuronal population) as well as a random factor depending on each recorded unit. These random factors are highly regularized (they are estimated under the assumption that they are instances of a Gaussian variable with a data-dependent covariance matrix[62]), and are therefore

conservative estimates of the dependence of the firing rate of each specific unit on the parameter of interest.

We selected for further study units with prominent random effects. Prominent random effects were defined as those whose magnitude was greater than one standard deviation of the corresponding fixed effect. The logic behind this choice is based on the observation that for a unit that has a prominent random effect, the dependence of its firing rate on the parameter of interest is sufficiently different from that of the population mean. This method identified sound-sensitive neurons, but in addition it identified rat location, velocity, and head-body angle as variables that affected the activity of many units in all rats.

Specifically in this paper we highlight location sensitive units and head-body angle sensitive units. While units could show sensitivity to multiple variables, in the units highlighted here we verified that the sensitivity to location, for example, did not depend on other variables of interest (and therefore was not the result of a spurious correlation due to another primary dependence). Thus, location sensitive units were individually tested for velocity and head-body angle sensitivity using the following method. For each unit, the average firing rate as a joint function of location and one additional variable (for example, velocity) was computed, usually using a 10-by-10 grid of bins. In order to check the independence of the tuning to the two variables, a rank-1 approximation was computed for each of these matrices using nonnegative matrix factorization (Matlab function nnmf, see supp. Fig. S5). We then compared the matrix containing the average firing rates with its rank-1 approximation. To measure this, the squared L2 norm of the residual was subtracted from the L2 norm of the original matrix, and normalized by the L2 norm of the original matrix. The resulting number is the fraction of the data variability explained by the rank-1 approximation.

Since the original matrix was estimated from a finite sample, it is expected to have a high rank even if its ideal structure has rank 1. Therefore, two methods were used to estimate the expected fraction of data variability explained by a rank-1 approximation to a noisy data matrix whose underlying structure is also of rank 1. In the first method the original data was used for bootstrapping: the firing rates observed within each location bin were resampled, and assigned to random velocity bins according to their observed probability. The mean for each one of the matrix entries was calculated. The resulting 10-by-10 matrix was processed in the same way as the original data. This process was repeated for 100 times, the fraction of explained variability averaged and compared to that of the original matrix. In the second method, the rank-1 approximation was used to generate surrogate data with Poisson distribution, using the same

number of counts in each bin of the original matrix. This matrix was processed in the same way as the original data, and the fraction of explained data variability from this process was compared to that of the original data.

Units sensitive to head-body angle were individually tested using the same method for joint sensitivity to either absolute head angle, absolute body angle, or location (see supp. Fig. S8).

All the units were also tested for the effects of stimulus driven responses (supp. Figs S6 and S7). The mean responses with and without sound presentations were computed as a function of the relevant parameter (in each location bin for the location-sensitive units, for each head-body angle bin for the head-body sensitive units). The mean firing rates in the presence and absence of sound were then compared using a paired t-test.

References

1. Schnupp, J., Nelken, I. & King, A. J. *Auditory Neuroscience: Making Sense of Sound*. (MIT Press, Cambridge, Massachusetts, 2010).
2. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* **33**, 89–108 (2010).
3. Carandini, M. & Churchland, A. K. Probing perceptual decisions in rodents. *Nat. Neurosci.* **16**, 824–831 (2013).
4. Poo, M.-M. *et al.* What is memory? The present state of the engram. *BMC Biol.* **14**, 40 (2016).
5. Saraf-Sinik, I., Assa, E. & Ahissar, E. Motion makes sense: an adaptive motor-sensory strategy underlies the perception of object location in rats. *J. Neurosci. Off. J. Soc. Neurosci.* **35**, 8777–8789 (2015).
6. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
7. Sofroniew, N. J., Cohen, J. D., Lee, A. K. & Svoboda, K. Natural whisker-guided behavior by head-fixed mice in tactile virtual reality. *J. Neurosci. Off. J. Soc. Neurosci.* **34**, 9537–9550 (2014).
8. Keller, G. B., Bonhoeffer, T. & Hübener, M. Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. *Neuron* **74**, 809–815 (2012).
9. Radvansky, B. A. & Dombeck, D. A. An olfactory virtual reality system for mice. *Nat. Commun.* **9**, 839 (2018).
10. Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L. & Tank, D. W. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat. Neurosci.* **13**, 1433–1440 (2010).
11. Dombeck, D. A., Khabbaz, A. N., Collman, F., Adelman, T. L. & Tank, D. W. Imaging Large-Scale Neural Activity with Cellular Resolution in Awake, Mobile Mice. *Neuron* **56**, 43–57 (2007).
12. Harvey, C. D., Collman, F., Dombeck, D. A. & Tank, D. W. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* **461**, 941–946 (2009).
13. Hölscher, C., Schnee, A., Dahmen, H., Setia, L. & Mallot, H. A. Rats are able to navigate in virtual environments. *J. Exp. Biol.* **208**, 561–569 (2005).
14. Aghajan, Z. M. *et al.* Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. *Nat. Neurosci.* **18**, 121–128 (2015).
15. Go, M. A. *et al.* Place Cells in Head-Fixed Mice Navigating a Floating Real-World Environment. *Front. Cell. Neurosci.* **15**, (2021).

16. Keating, P., Dahmen, J. C. & King, A. J. Context-Specific Reweighting of Auditory Spatial Cues following Altered Experience during Development. *Curr. Biol.* **23**, 1291–1299 (2013).
17. Tsoar, A. *et al.* Large-scale navigational map in a mammal. *Proc. Natl. Acad. Sci.* **108**, E718–E724 (2011).
18. Khan, Z., Herman, R. A., Wallen, K. & Balch, T. An outdoor 3-D visual tracking system for the study of spatial navigation and memory in rhesus monkeys. *Behav. Res. Methods* **37**, 453–463 (2005).
19. Ballesta, S., Reymond, G., Pozzobon, M. & Duhamel, J.-R. A real-time 3D video tracking system for monitoring primate groups. *J. Neurosci. Methods* **234**, 147–152 (2014).
20. Matsumoto, J. *et al.* A 3D-Video-Based Computerized Analysis of Social and Sexual Interactions in Rats. *PLOS ONE* **8**, e78460 (2013).
21. de Chaumont, F. *et al.* Real-time analysis of the behaviour of groups of mice via a depth-sensing camera and machine learning. *Nat. Biomed. Eng.* **3**, 930–942 (2019).
22. de Chaumont, F. *et al.* Computerized video analysis of social interactions in mice. *Nat. Methods* **9**, 410–417 (2012).
23. Weissbrod, A. *et al.* Automated long-term tracking and social behavioural phenotyping of animal colonies within a semi-natural environment. *Nat. Commun.* **4**, 2018 (2013).
24. Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S. & de Polavieja, G. G. idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. Methods* **11**, 743–748 (2014).
25. Romero-Ferrero, F., Bergomi, M. G., Hinz, R. C., Heras, F. J. H. & de Polavieja, G. G. idtracker.ai: tracking all individuals in small or large collectives of unmarked animals. *Nat. Methods* **16**, 179–182 (2019).
26. Straw, A. D., Branson, K., Neumann, T. R. & Dickinson, M. H. Multi-camera real-time three-dimensional tracking of multiple flying animals. *J. R. Soc. Interface* **8**, 395–409 (2011).
27. Fry, S. N., Rohrseitz, N., Straw, A. D. & Dickinson, M. H. TrackFly: Virtual reality for a behavioral system analysis in free-flying fruit flies. *J. Neurosci. Methods* **171**, 110–117 (2008).
28. Fry, S. N., Bichsel, M., Müller, P. & Robert, D. Tracking of flying insects using pan-tilt cameras. *J. Neurosci. Methods* **101**, 59–67 (2000).
29. Finkelstein, A. *et al.* Three-dimensional head-direction coding in the bat brain. *Nature* **517**, 159–164 (2015).
30. Yartsev, M. M. & Ulanovsky, N. Representation of Three-Dimensional Space in the Hippocampus of Flying Bats. *Science* **340**, 367–372 (2013).

31. Nourizonoz, A. *et al.* EthoLoop: automated closed-loop neuroethology in naturalistic environments. *Nat. Methods* **17**, 1052–1059 (2020).
32. Sutton, R. S. & Barto, A. G. *Reinforcement learning: an introduction*. (MIT Press, 20).
33. Mallory, C. S. *et al.* Mouse entorhinal cortex encodes a diverse repertoire of self-motion signals. *Nat. Commun.* **12**, 671 (2021).
34. Rodgers, K. M., Benison, A. M., Klein, A. & Barth, D. S. Auditory, somatosensory, and multisensory insular cortex in the rat. *Cereb. Cortex N. Y. N 1991* **18**, 2941–2951 (2008).
35. Kimura, A., Imbe, H. & Donishi, T. Efferent connections of an auditory area in the caudal insular cortex of the rat: anatomical nodes for cortical streams of auditory processing and cross-modal sensory interactions. *Neuroscience* **166**, 1140–1157 (2010).
36. de Hoz, L. & Nelken, I. Frequency tuning in the behaving mouse: different bandwidths for discrimination and generalization. *PloS One* **9**, e91676 (2014).
37. Rosenberg, M., Zhang, T., Perona, P. & Meister, M. Mice in a labyrinth: Rapid learning, sudden insight, and efficient exploration. *bioRxiv* 2021.01.14.426746 (2021) doi:10.1101/2021.01.14.426746.
38. Lim, J. & Celikel, T. Real-time contextual feedback for close-loop control of navigation. *J. Neural Eng.* **16**, 065001 (2019).
39. Hong, W. *et al.* Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. *Proc. Natl. Acad. Sci.* **112**, E5351–E5360 (2015).
40. Shemesh, Y. *et al.* High-order social interactions in groups of mice. *eLife* **2**, e00759 (2013).
41. Minderer, M., Harvey, C. D., Donato, F. & Moser, E. I. Virtual reality explored. *Nature* **533**, 324–325 (2016).
42. Bjerre, A.-S. & Palmer, L. M. Probing Cortical Activity During Head-Fixed Behavior. *Front. Mol. Neurosci.* **13**, (2020).
43. Huang, K.-H. *et al.* A virtual reality system to analyze neural activity and behavior in adult zebrafish. *Nat. Methods* **17**, 343–351 (2020).
44. Whishaw, I. Q. *et al.* Organization of the reach and grasp in head-fixed vs freely-moving mice provides support for multiple motor channel theory of neocortical organization. *Exp. Brain Res.* **235**, 1919–1932 (2017).
45. Ravassard, P. *et al.* Multisensory Control of Hippocampal Spatiotemporal Selectivity. *Science* **340**, 1342–1346 (2013).
46. The International Brain Laboratory *et al.* Standardized and reproducible measurement of decision-making in mice. *eLife* **10**, e63711 (2021).

47. Rodgers, K. M., Benison, A. M., Klein, A. & Barth, D. S. Auditory, Somatosensory, and Multisensory Insular Cortex in the Rat. *Cereb. Cortex* **18**, 2941–2951 (2008).

48. Takemoto, M., Hasegawa, K., Nishimura, M. & Song, W.-J. The insular auditory field receives input from the lemniscal subdivision of the auditory thalamus in mice: Auditory thalamocortical connections. *J. Comp. Neurol.* **522**, 1373–1389 (2014).

49. Kimura, A., Imbe, H. & Donishi, T. Efferent connections of an auditory area in the caudal insular cortex of the rat: anatomical nodes for cortical streams of auditory processing and cross-modal sensory interactions. *Neuroscience* **166**, 1140–1157 (2010).

50. Jankowski, M. M. *et al.* Evidence for spatially-responsive neurons in the rostral thalamus. *Front. Behav. Neurosci.* **9**, (2015).

51. Jankowski, M. M. & O'Mara, S. M. Dynamics of place, boundary and object encoding in rat anterior claustrum. *Front. Behav. Neurosci.* **9**, (2015).

52. Long, X., Deng, B., Cai, J., Chen, Z. S. & Zhang, S.-J. *A compact spatial map in V2 visual cortex*. http://biorxiv.org/lookup/doi/10.1101/2021.02.11.430687 (2021) doi:10.1101/2021.02.11.430687.

53. Long, X. & Zhang, S.-J. A novel somatosensory spatial navigation system outside the hippocampal formation. *Cell Res.* (2021) doi:10.1038/s41422-020-00448-8.

54. Mao, D., Kandler, S., McNaughton, B. L. & Bonin, V. Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nat. Commun.* **8**, 243 (2017).

55. Yin, A., Tseng, P. H., Rajangam, S., Lebedev, M. A. & Nicolelis, M. A. L. Place Cell-Like Activity in the Primary Sensorimotor and Premotor Cortex During Monkey Whole-Body Navigation. *Sci. Rep.* **8**, 9184 (2018).

56. Gardner, R. J. *et al.* Toroidal topology of population activity in grid cells. *bioRxiv* 2021.02.25.432776 (2021) doi:10.1101/2021.02.25.432776.

57. Marshall, J. D. *et al.* Continuous Whole-Body 3D Kinematic Recordings across the Rodent Behavioral Repertoire. *Neuron* **109**, 420-437.e8 (2021).

58. Kawahara, H. *et al.* Tandem-STRAIGHT: A temporally stable power spectral representation for periodic signals and applications to interference-free spectrum, F0, and aperiodicity estimation. in *2008 IEEE International Conference on Acoustics, Speech and Signal Processing* 3933–3936 (2008). doi:10.1109/ICASSP.2008.4518514.

59. Kazakov, A. & Nelken, I. Acoustic calibration in an echoic environment. *J. Neurosci. Methods* **309**, 60–70 (2018).

60. He, K., Zhang, X., Ren, S. & Sun, J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. *ArXiv150201852 Cs* (2015).

61. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. *ArXiv14126980 Cs* (2017).

62. A Model for Longitudinal Data. in *Linear Mixed Models for Longitudinal Data* (eds. Verbeke, G. & Molenberghs, G.) 19–29 (Springer, 2000). doi:10.1007/978-0-387-22775-7_3.

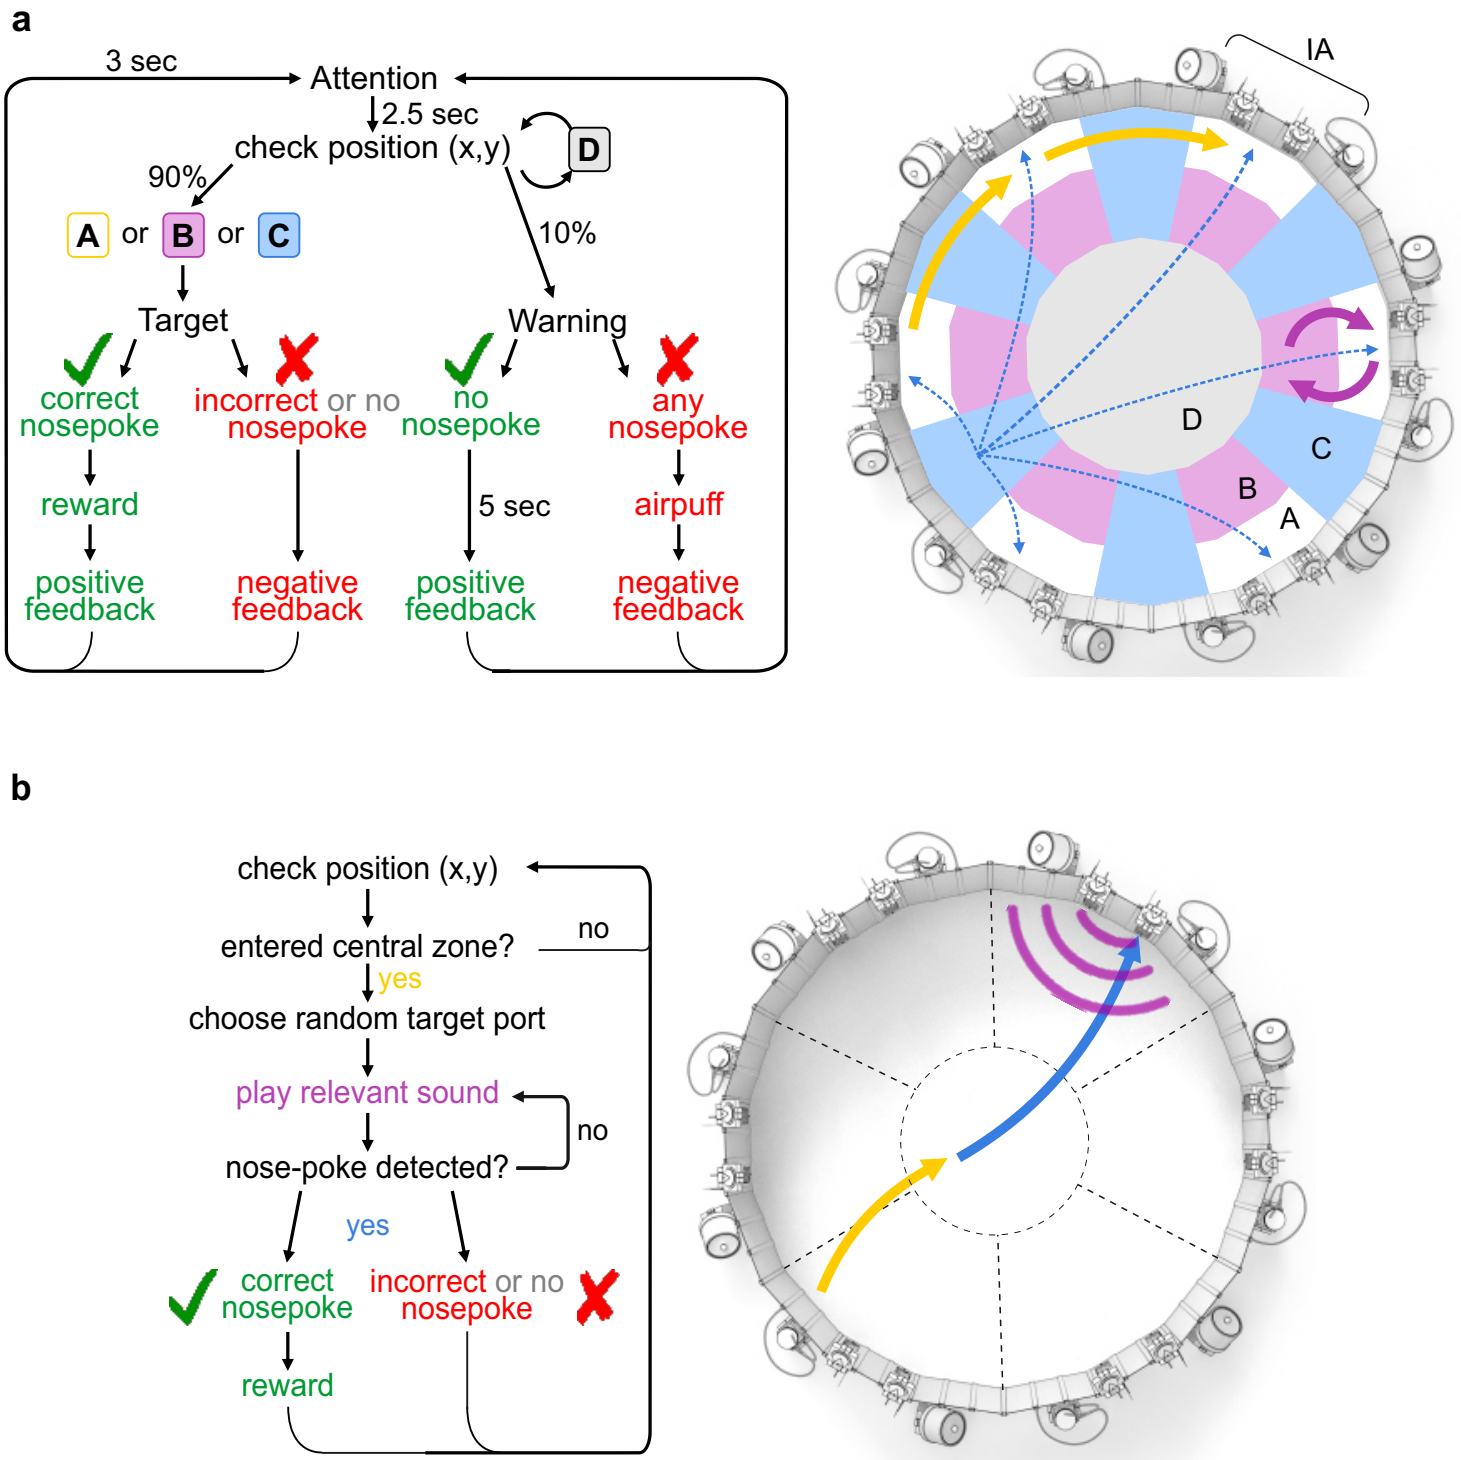

Fig. S1. Structure of the two tasks described in the paper. (a) The multiple strategies task (St+). On the left, a flowchart of the experiment is displayed as a real-time loop. At the check position stage, three strategies were available to the rat, marked as A, B and C. These strategies are illustrated in the diagram of the arena on the right, using the same color code. The A strategy consisted of moving from one interaction area to another (usually a neighboring area), which was then selected as the next target. The B strategy consisted of cycling from an A area to the associated B area, which led to the selection of the same A area as the next target. The C strategy consisted of moving to a C area, in which case a random port was selected as the next target. **(b)** Diagram of the localization/discrimination task (L/D). Flowchart of the experiment real-time loop (left) and the corresponding events in a diagram of the arena (right), plotted with the same color code.

a.

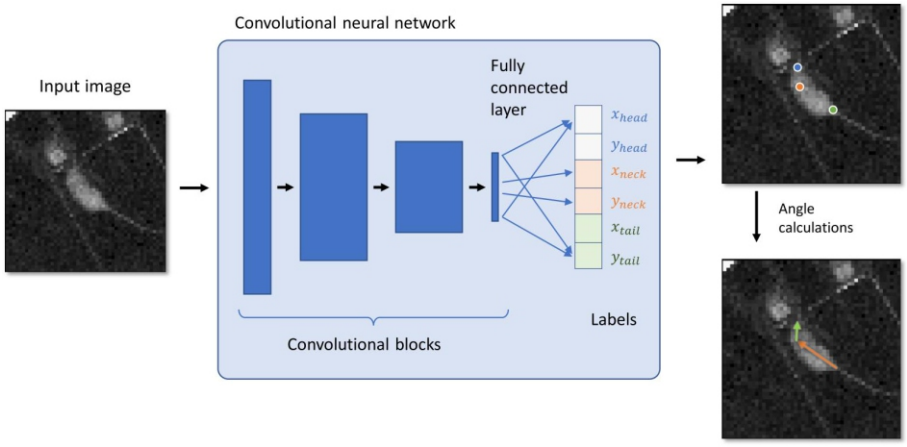

b.

Convolution neural network layers

Type	Input channels	Kernels	Kernel size	Other
Input	1			Input res. 50x50
Convolution	1	32	3	
Batch Norm	32			
ReLU				
Convolution	32	32	3	
Batch Norm	32			
ReLU				
Max Pooling	32		2	Stride = 2
Convolution	32	64		
Batch Norm	32			
ReLU				
Convolution	64	64		
Batch Norm	64			
ReLU				
Max Pooling			2	Stride = 2
Convolution	64	128		
Batch Norm	128			
ReLU				
Convolution	128	128		
Batch Norm	128			
ReLU				
Max Pool	128		2	Stride = 2
Flatten				
Linear	4608	6		

Fig. S2. Feature extraction from the video images. (a) A feed-forward convolutional neural network estimates the locations of the head, the base of the neck and the base of the tail for each input image. These three markers are then used for the calculation of the body and head angles (bottom right image). **(b)** The custom architecture of the neural network is optimized for the grayscale rectangular input, reducing the volume of the trained model and decreasing the inference times.

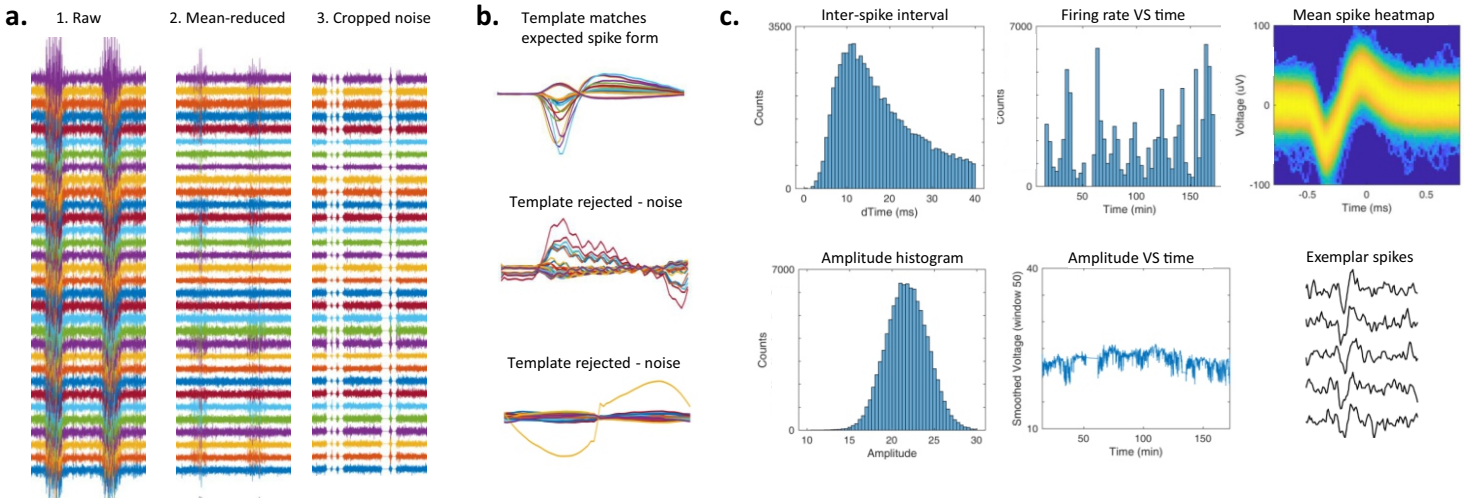

Fig. S3. Neural data processing and spike sorting. (a) Extracellular neural recordings in freely behaving rats (32 simultaneously recorded channels) include periods of high noise (left panel). Noise components that are common to all channels can be largely removed by subtracting the average waveform (middle panel). Remaining noise segments are identified by their amplitude and by their high variance across channels, and are silenced (right panel). The resulting neural data is then processed by the Kilosort2 program. **(b)** The clusters detected by the Kilosort2 program are automatically characterized for implausible spike shapes. The top example here was automatically identified as a true spike, while the bottom two examples were determined to be noise. **(c)** Statistics of the neural activity are produced for each cluster and used for a final manual curation by the experimenter.

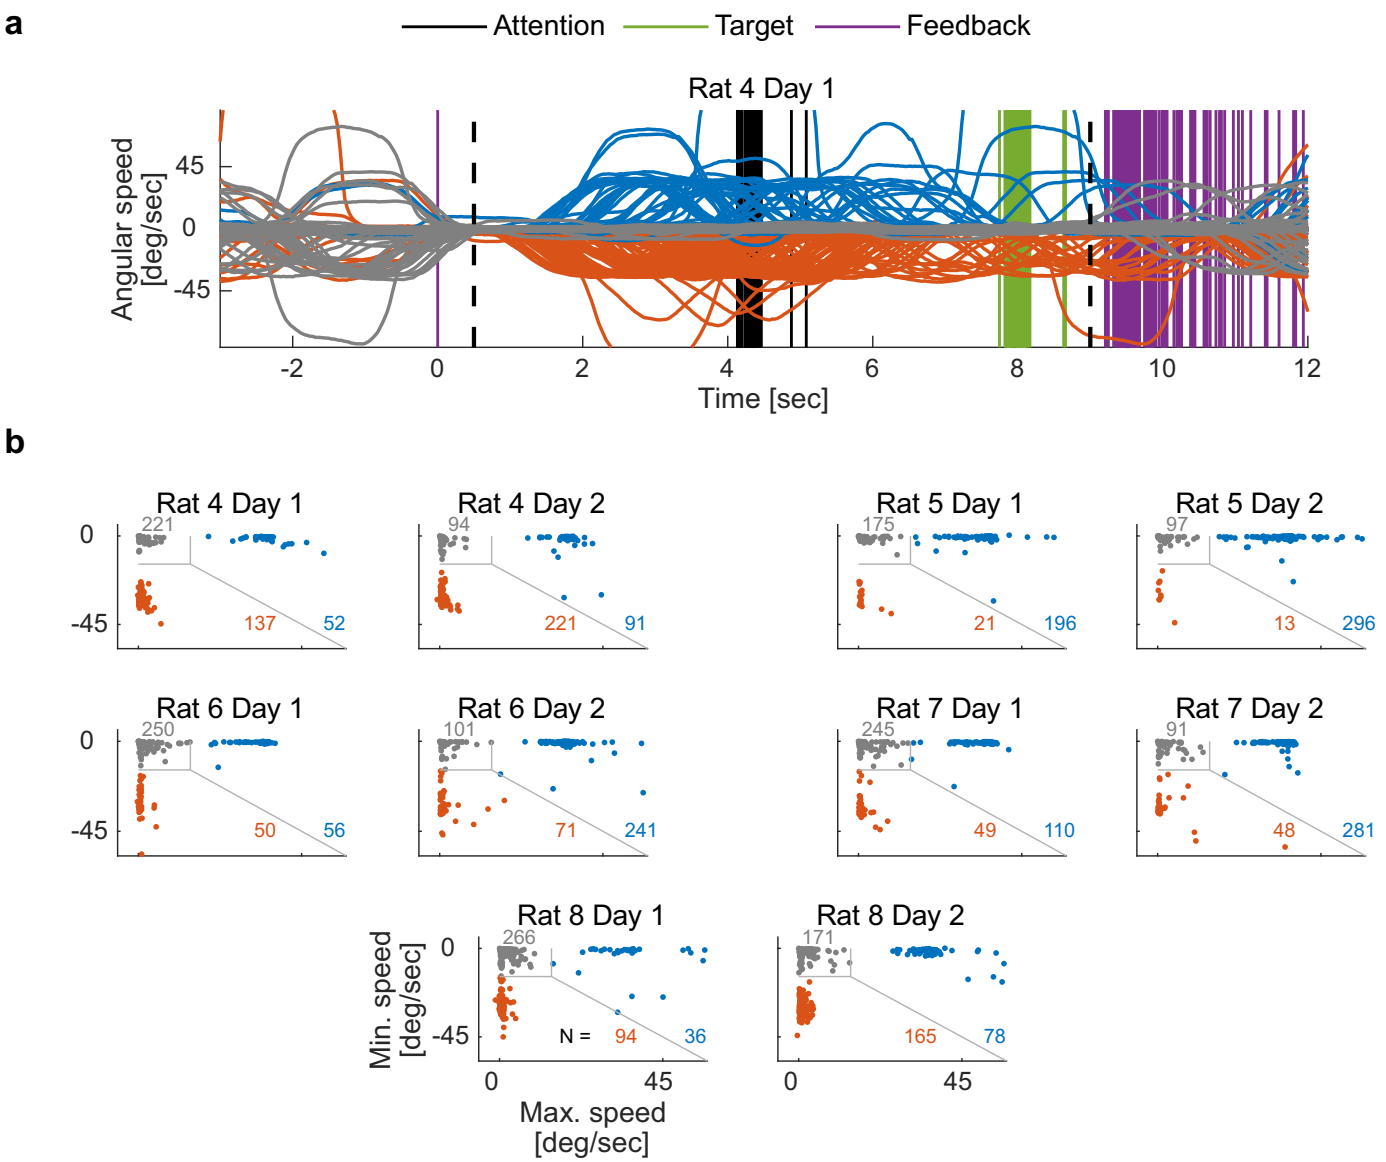

Fig. S4. Classification of trials into three types. (a) Angular running speed for all trials performed by rat 4 on day 1. For each trial, the maximal angular speed in the clockwise and counterclockwise direction was extracted in a time window lasting from 0.5 seconds to 9 seconds following the feedback sound of the previous trial. A trial was classified as "Sit" (grey lines) if the absolute value of the angular speed never exceeded 0.25 radians/s (14.3 degrees/s), otherwise as "Run clockwise" (blue lines) or "Run counterclockwise" (red lines), according to the direction with the higher maximal speed. **(b)** Trial clusters were clearly separated in all rats. Scatter plots show each trial of each rat and each day according to the maximal angular speed in the clockwise and counterclockwise directions. Colors as in (a). The number of trials of each trial cluster are indicated.

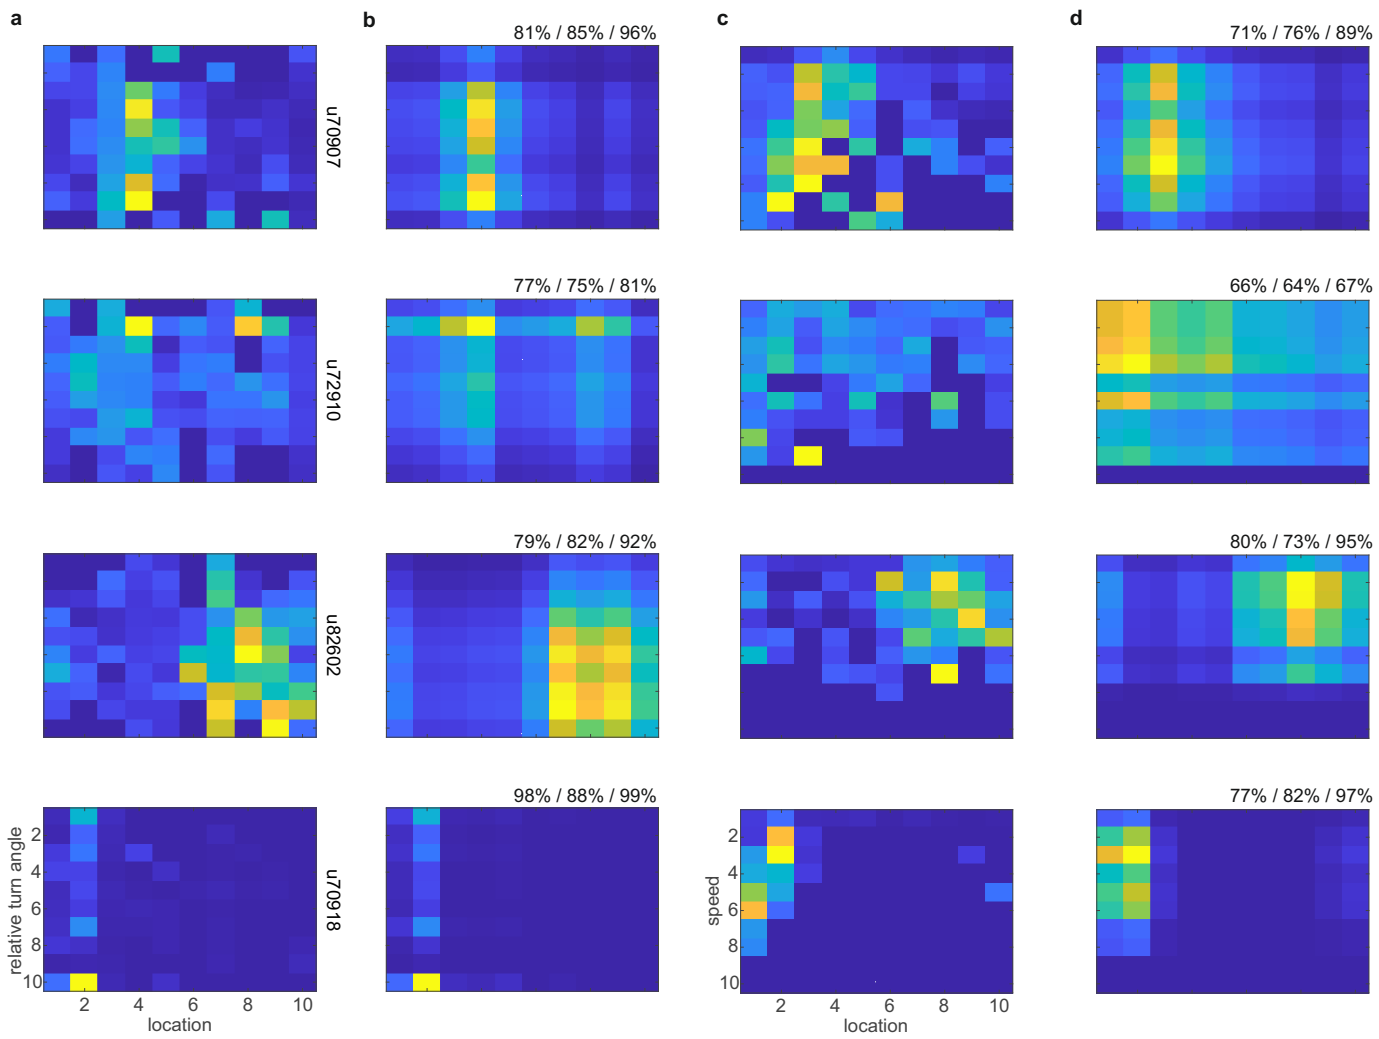

Fig. S5. Joint sensitivity of the units in Figs. 5e-h to location, head-body angle, and velocity. (a) Mean firing rates for location (abscissa) and relative head-body angle (ordinate). **(b)** Non-negative rank-1 matrix approximation of the matrices in (a). Top right corner: the fraction of the data variability explained by the rank 1 approximation for the original data / bootstrapping method / Poisson distribution approximation **(c)** Mean firing rates for location (abscissa) and velocity (ordinate). **(d)** Same as (b) for the matrices in (c).

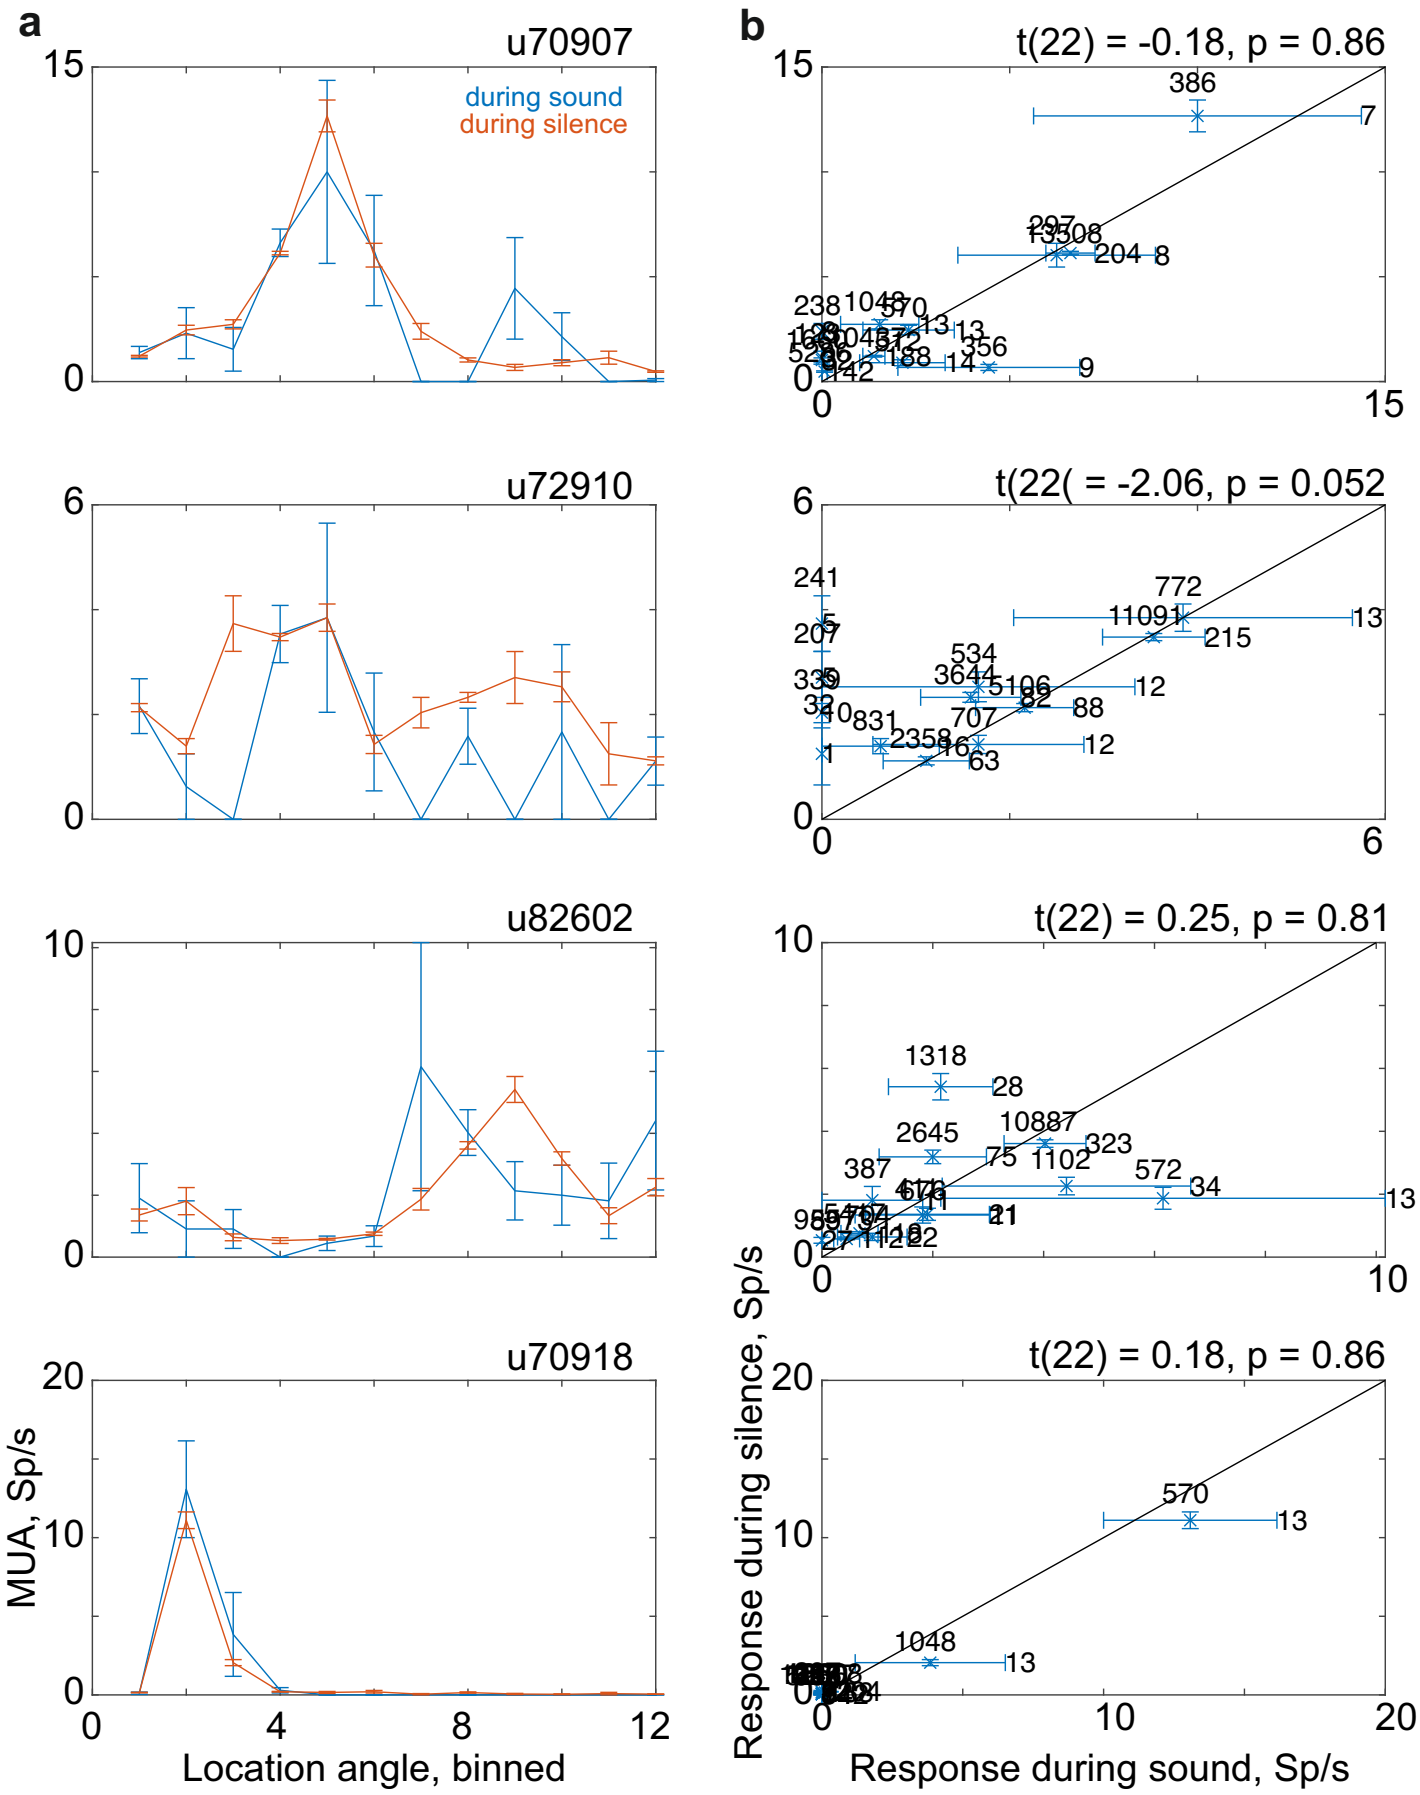

Fig. S6. Joint sensitivity of the units in Figs. 5e-h to location and sound. (a) Mean responses In the presence (blue) and absence (red) of sound. Error bars are s.e.m. **(b)** Scatter plots of the data as in (a). Horizontal error bar indicates s.e.m during sound presentation, vertical error bar indicates s.e.m during silence. The number of instances from which the mean and s.e.m were derived are indicated on the right of each point (during sound presentation), and on top (during silence).

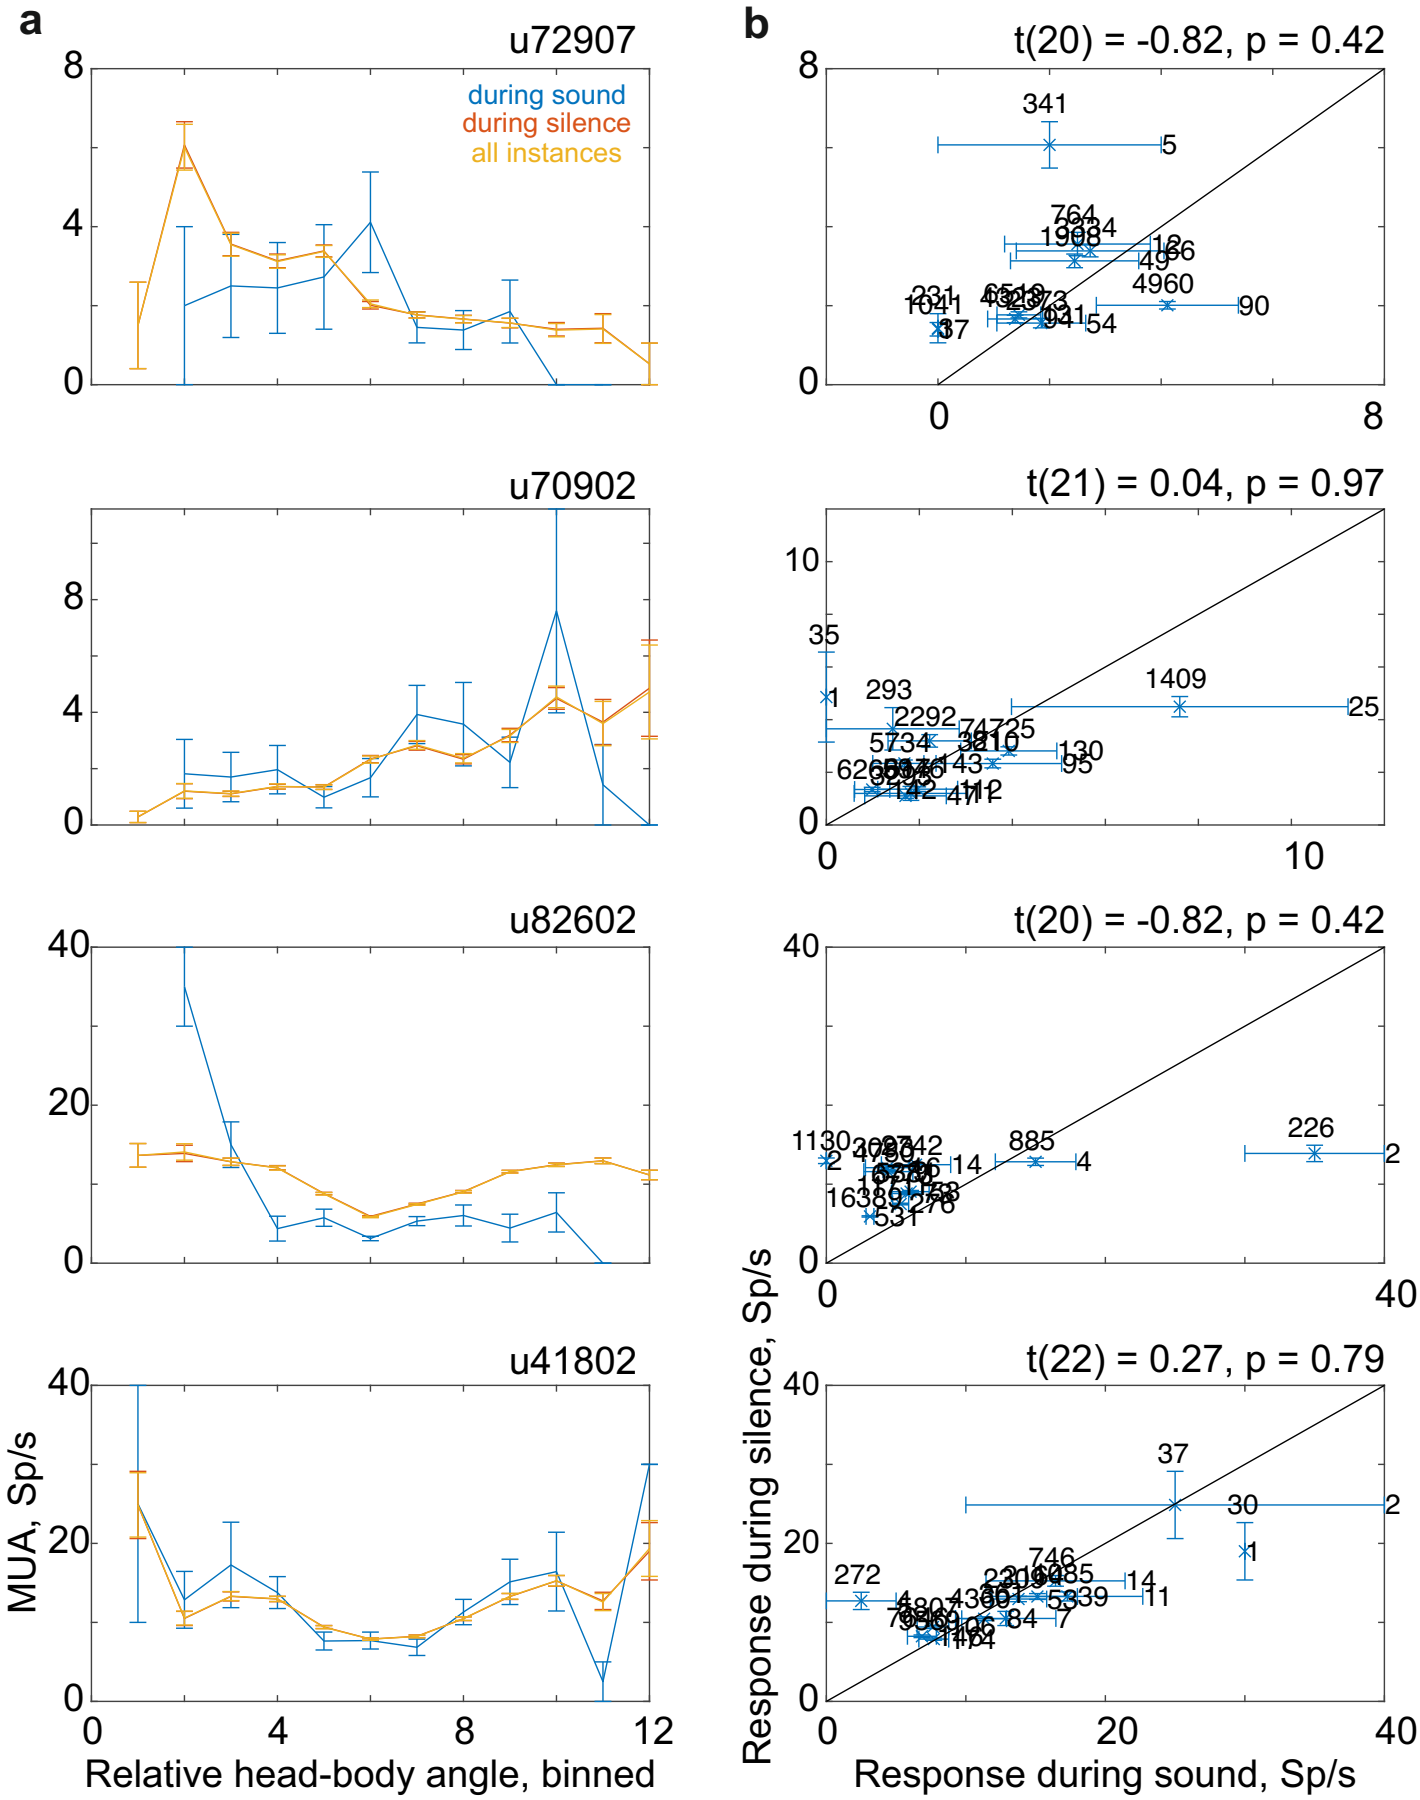

Fig. S7. Joint sensitivity of the units in Figs. 5i-l to location and sound. Same format as Fig. S6.

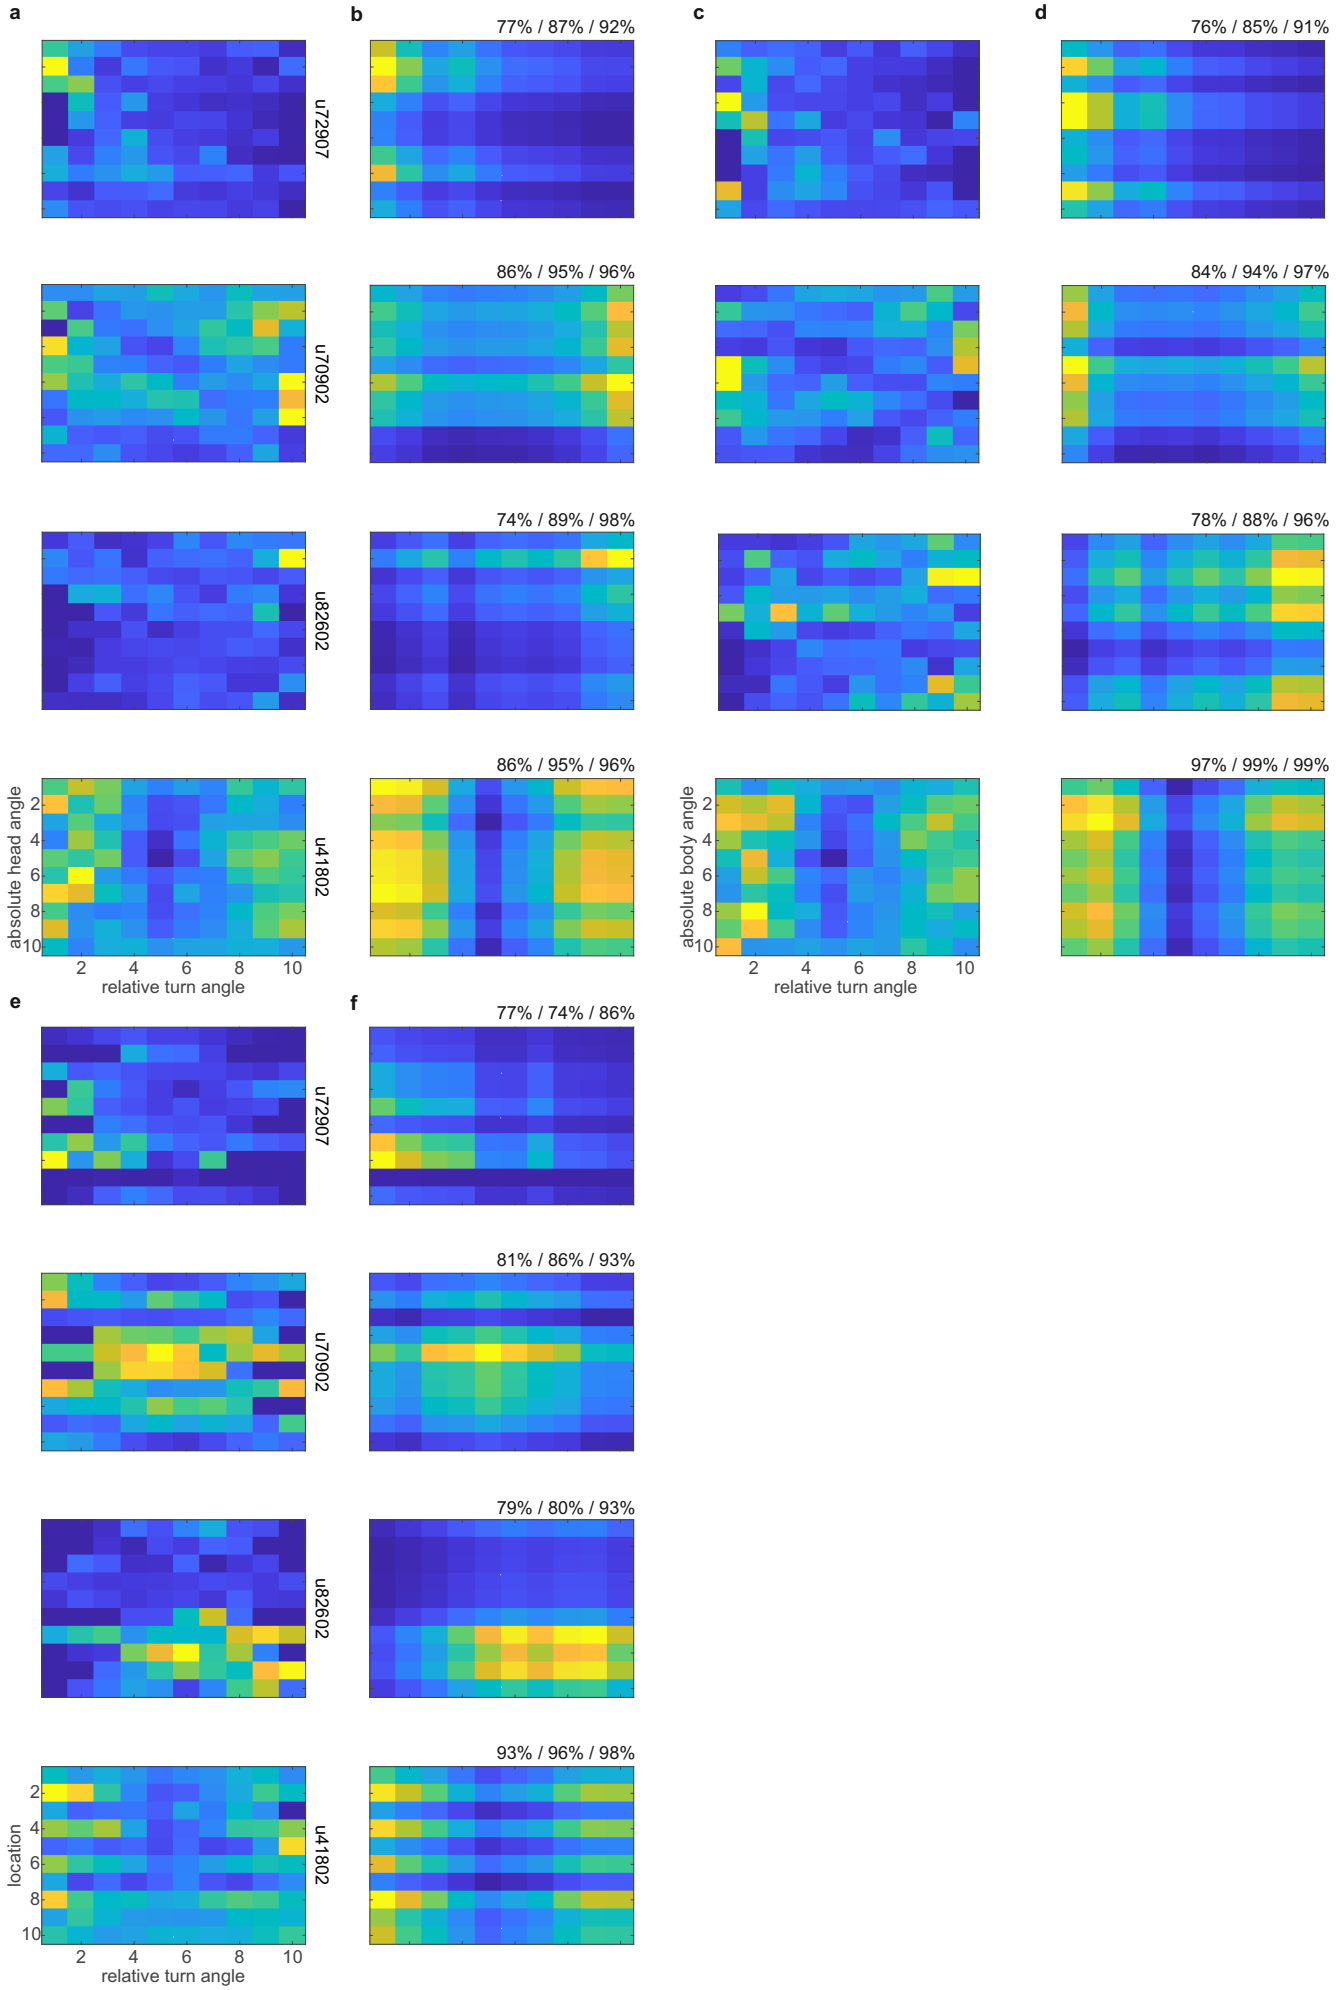

Fig. S8. Joint sensitivity of the units in Figs. 5i-l to head-body angle, absolute head angle, absolute body angle, and location. Same representations as in Fig. S4. The abscissa represents head-body angles, while the ordinate represents absolute head angle ((a) and (b)), absolute body angle ((c) and (d)), and location ((e) and (f)).

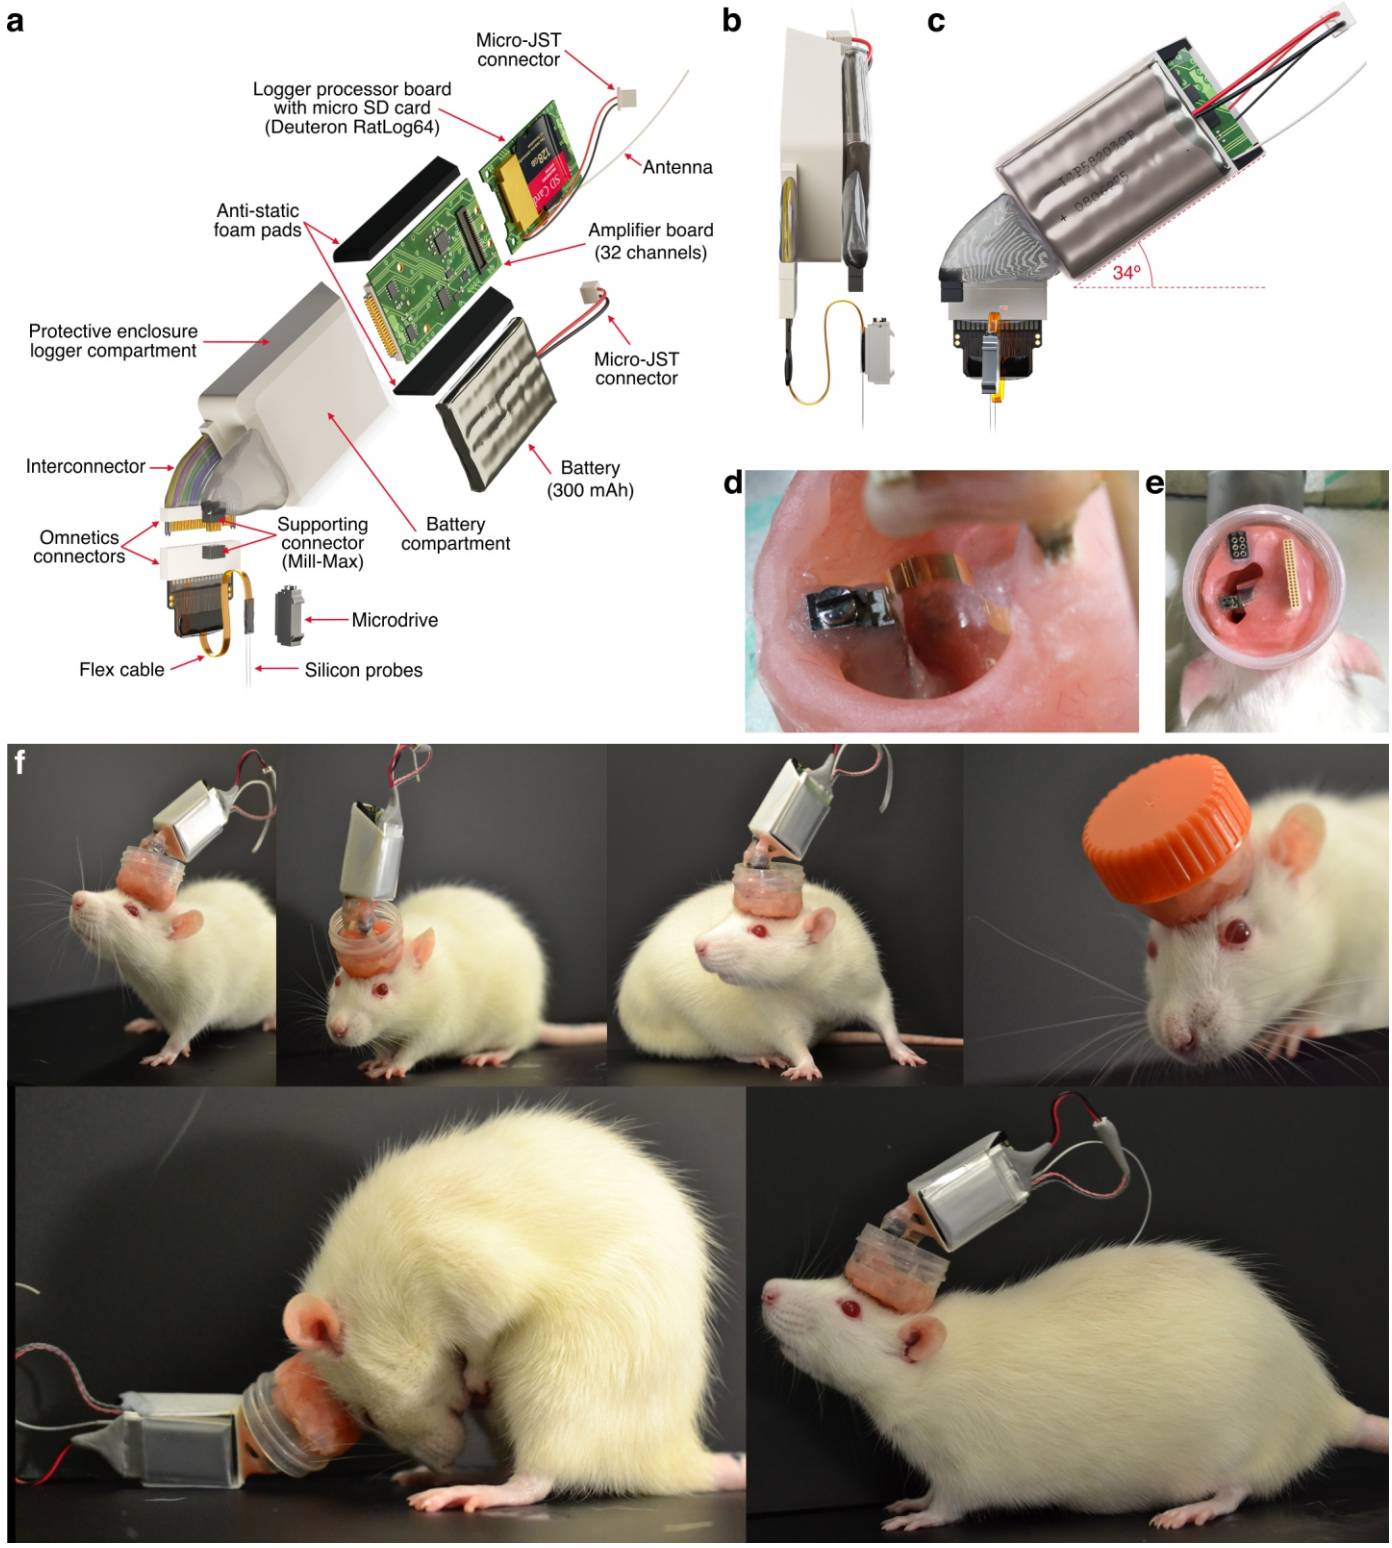

Fig. S9. An approach for chronic wireless recordings in rats. (a) Complete recording set. The neural logger and battery are in the protective plastic case that can be attached to the electrode's connector. The 300 mAh battery is placed in a separate compartment, connects to the logger through a micro-JST connector, and can be easily changed at any moment during the experiment. Anti-static foam pads are placed on the sides of the logger components (amplifier and processor boards) to protect the logger against mechanical shocks. The protective case has a 36-pin omnetics connector matching that on the 32-channel silicon probe, as well as a small Mill-Max connector which mechanically stabilizes the case during recording sessions. The silicon probe is mounted on the Microdrive (Cambridge Neurotech). **(b and c)** Side and front views of the recording set. The device is inclined to the back in order to allow rats natural movements and undisturbed access to interactive ports. **(d)** Photo of the implant. The silicon probe is mounted on the Microdrive cemented to the skull. The moveable parts of the implant are covered with paraffin oil. The flex cable of the probe is bent to provide a long travel distance for the electrodes. **(e)** Finished implant with added plastic protection for flex cable and the enclosure of the whole implant. **(f)** Photos of female rats with a 32-channel moveable silicon probe implant and the wireless data logger in the case with a battery. The recording set enables natural movements, is easily carried by the rats, and is well protected against mechanical shocks. The enclosure can be closed with a plastic cap (orange) to protect the implant in the home cage. The big size of the cap prevents the implant from getting stuck between parts of the home cage.

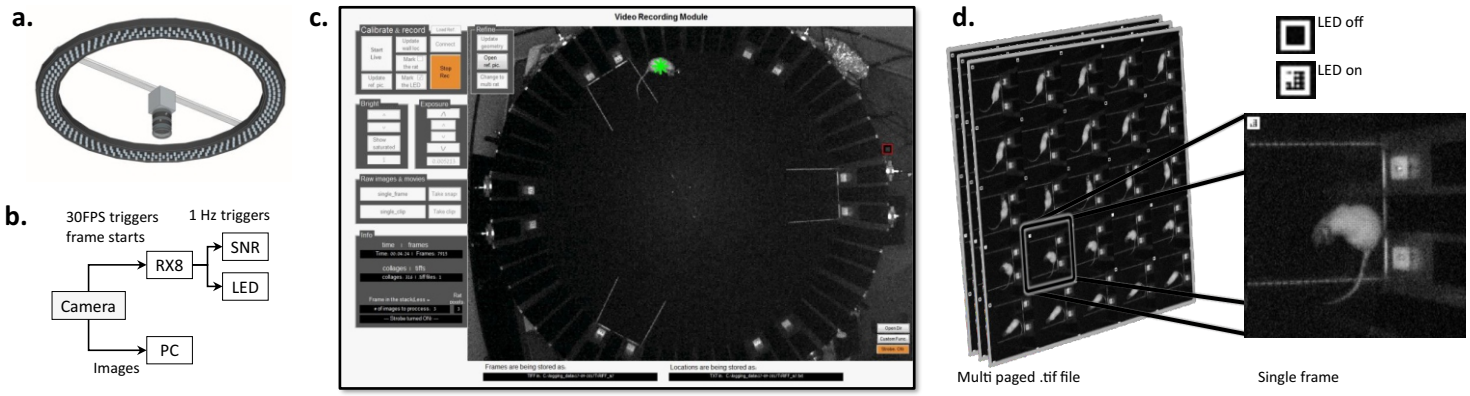

Fig. S10. The real-time imaging module. (a) Diagram of the camera (DMK 33G445 GigE, TheImagingSource) and the light ring, mounted on the ceiling above the arena's center. **(b)** Synchronization diagram of the image stream. Triggers that indicate frame acquisition were sent to a digital processor that sub-sampled them from 30 Hz to 1 Hz. The 1 Hz triggers were then simultaneously recorded on the common synchronization hardware, and also powered a LED in the field of view of the camera. **(c)** Graphical user interface of the real-time imaging module. The LED is marked by a small red square on the right side of the arena. The rat center of mass is marked by a green asterisk. **(d)** For efficient data storage, rat images were cropped around its center of mass and stored in a multi page .tif file. A cropped image of the LED was stored in the upper left corner of each image, allowing for time synchronization during the post-processing steps.