

RESEARCH ARTICLE | *Control of Movement*

Automated home cage training of mice in a hold-still center-out reach task

Tejapratap Bollu,¹ Samuel C. Whitehead,² Nikil Prasad,¹ Jackson Walker,¹ Nitin Shyamkumar,¹ Raghav Subramaniam,¹ Brian Kardon,¹ Itai Cohen,² and Jesse H. Goldberg¹

¹Department of Neurobiology and Behavior, Cornell University, Ithaca, New York; and ²Department of Physics, Cornell University, Ithaca, New York

Submitted 1 October 2018; accepted in final form 10 December 2018

Bollu T, Whitehead SC, Prasad N, Walker J, Shyamkumar N, Subramaniam R, Kardon B, Cohen I, Goldberg JH. Automated home cage training of mice in a hold-still center-out reach task. *J Neurophysiol* 121: 500–512, 2019. First published December 12, 2018; doi:10.1152/jn.00667.2018.—An obstacle to understanding neural mechanisms of movement is the complex, distributed nature of the mammalian motor system. Here we present a novel behavioral paradigm for high-throughput dissection of neural circuits underlying mouse forelimb control. Custom touch-sensing joysticks were used to quantify mouse forelimb trajectories with micron-millisecond spatio-temporal resolution. Joysticks were integrated into computer-controlled, rack-mountable home cages, enabling batches of mice to be trained in parallel. Closed loop behavioral analysis enabled online control of reward delivery for automated training. We used this system to show that mice can learn, with no human handling, a direction-specific hold-still center-out reach task in which a mouse first held its right forepaw still before reaching out to learned spatial targets. Stabilogram diffusion analysis of submillimeter-scale micro-movements produced during the hold demonstrate that an active control process, akin to upright balance, was implemented to maintain forepaw stability. Trajectory decomposition methods, previously used in primates, were used to segment hundreds of thousands of forelimb trajectories into millions of constituent kinematic primitives. This system enables rapid dissection of neural circuits for controlling motion primitives from which forelimb sequences are built.

NEW & NOTEWORTHY A novel joystick design resolves mouse forelimb kinematics with micron-millisecond precision. Home cage training is used to train mice in a hold-still center-out reach task. Analytical methods, previously used in primates, are used to decompose mouse forelimb trajectories into kinematic primitives.

joystick; motor primitives; mouse forelimb; reach

INTRODUCTION

An infamous problem in motor control is “the curse of dimensionality”: a hand in motion sweeps through a near-infinite continuum of possible trajectories, making motor control seem intractably high dimensional (Bernstein 1967; Shadmehr and Wise 2005; Woodworth 1899). One solution is to construct movement from a discrete set of elementary building blocks, or motion primitives (Flash and Hochner 2005; Mussa-Ivaldi et al. 1994). For example, when you draw the letter *N* you carve a complex path through space, but *N* can be decom-

posed into three distinct strokes, or kinematic primitives, each with only a few parameters such as direction, speed, and duration (Flash and Hogan 1985; Milner 1992; Viviani and Terzuolo 1982) (Fig. 1, A–C). Motor primitives are abnormally generated and sequenced in disorders such as stroke, dystonia, and Parkinson’s (Desmurget et al. 2004; Inzelberg et al. 1995; Majsak et al. 1998; Rohrer et al. 2004), yet precise circuits that generate primitives, determine their kinematics, and sequence them into a trajectory are poorly understood (Giszter 2015).

One challenge is that kinematic representations and movement initiation signals are not regionally localized but are instead distributed throughout cortical, cerebellar, and basal ganglia circuits (Fortier et al. 1989; Fu et al. 1997; Schwartz 2007; Scott 2003; Shenoy et al. 2013; Turner and Anderson 1997; Wong et al. 2015), complicating the identification of structure-function relationships and motivating causal experiments to provide constraints on theories of motor control (Omran et al. 2017; Pruszynski et al. 2011; Wolff and Ölveczky 2018).

Genetic tools in mice enable temporally precise circuit manipulations (Guo et al. 2014; Tye and Deisseroth 2012), but mechanisms of primitive generation and control remain poorly understood. We set out to resolve forelimb kinematics with sufficient spatiotemporal precision to extract primitives as in Fig. 1, which requires sensitivity to resolve the rapid, tiny details of motion that occur at sharp turns (Gowda et al. 2015; Rohrer and Hogan 2003; Viviani and Flash 1995). We designed ultralow torque touch-sensing joysticks that resolve mouse forelimb kinematics with micron-millisecond spatio-temporal resolution and built an automated home cage system to train mice in a hold-still center-out reach task. To complete the task, mice learned first to actively maintain the joystick in a small center position and then to produce an outward reach to learned spatial targets. The resultant trajectories carved complex paths in space. Algorithms previously used in primates were effective in decomposing forelimb trajectories into kinematic primitives, enabling us to characterize spatiotemporal patterns of mouse forelimb submovements.

MATERIALS AND METHODS

All experiments and procedures were performed according to NIH guidelines and were approved by the Institutional Animal Care and Use Committee of Cornell University.

Address for reprint requests and other correspondence: J. Goldberg, Dept. of Neurobiology and Behavior, Cornell University, Ithaca, NY 14853 (e-mail: jessehgolberg@gmail.com).

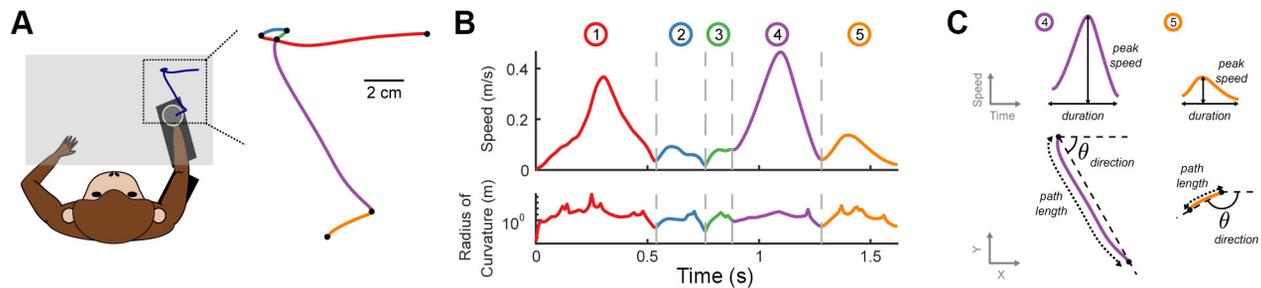


Fig. 1. Primate hand trajectories can be decomposed into kinematic primitives. *A*: schematic of a hand trajectory during a sequential reach task (left) with primate kinematic data from a previous study (Gowda et al. 2015). Black dots denote boundaries separating discrete segments. *B*: speed (top) and radius of curvature (bottom) are plotted as a function of time in the trajectory from *A*. Segment boundaries (gray dashed lines) are detectable as temporally coincident minima. Circled numbers above the plot denote segment number in the movement sequence. *C*: the speed (top) and path (bottom) of the last 2 segments from the trajectory in *A* and *B* (segment numbers 4 and 5). Individual segments are described by kinematic parameters such as peak speed, duration, direction and path length.

Experimental Animals

A total of seven adult male VGAT-ChR2-EYFP line 8 mice [B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP) 8Gfng/J; Jackson Laboratories stock no. 014548] were used. The mice were between 18 and 40 wk old for the duration of the experiment. Under isoflurane induced-anesthesia, we implanted the mice with 400- μm 0.43-NA fiber optic cannulas bilaterally over caudal forelimb area (0.5 mm A/P, \pm 1.5 mm M/L) and rostral forelimb area (2.2 mm A/P, \pm 1.0 mm M/L). All animals were individually housed under a 12-h light-dark cycle for the duration of the study and had continuous access to the task. The animal's daily water intake was monitored using an automated system. If the water intake was less than 1 ml, the system automatically dispensed water to make up the difference.

Principle of Operation and Signal Processing Pipeline

The joystick. The core of our joystick system was a two-axis Hall Sensor (Sentron, 2SA-10G). Moving the joystick changed the angle of the incident magnetic field generated by the magnets on the Hall sensor. The change in incident magnetic field provided x and y voltage (0–5 V) as a linear readout of displacement. The top of the joystick was a conductive ball connected to an active capacitance touch sensor (AT42QT1011) that provided instantaneous read on the joystick being contacted. To calibrate the joystick, i.e., to determine the voltage-versus-distance relationship, we set up the joystick in a custom, precision machined setup that allowed the manipulandum to move along a narrow channel along an outward directed radial axis. Then, using an electronic screw gauge, the joystick was displaced from the center position along the radial axis in increments of 500 μm . At each displacement we recorded the voltage from the joystick's Hall sensor. We then repeated these measurements along a radial axis in a direction nonorthogonal to the first. To quantify spatial resolution we first measured the x , y voltages from the device when it was at rest, and then used this calibration to convert the signal from volts into units of distance (mm). We then calculated the standard deviation of the distance distribution over 2,000 samples (2 s) for each device. For temporal resolution, the Hall sensor is rated to resolve magnetic fields changing at 100 kHz, allowing us to track the change in position with high temporal precision. To quantify stiffness of the joystick, we rotated the joystick by 90°, such that the joystick manipulandum was horizontal to the ground, and applied varying weights to the end of the joystick while measuring the displaced distance. We repeated these measurements by rotating the joystick around the axis of the manipulandum in $\sim 30^\circ$ increments. For construction and design details, see “Instructions for Construction” available at <https://github.com/GoldbergLab/RodentJoystick>.

Fixed post. A fixed post was placed in the home-cage and it also had a conductive ball connected to an active capacitance touch sensor (AT42QT1011).

Nose poke. The nose poke sensor was a modified infrared (IR) diode-phototransistor pair (LTH-301-32) from Lite-on devices. The IR light-emitting diode (LED) and the phototransistor were separated and placed across the nose poke port. When the mouse put its snout through the port, it broke the IR beam and drove the signal high.

Water delivery. The mice received water through a lick spout (H24-01-TB-01, Coulbourn Instruments) connected to a precision solenoid valve (Lee, LHDA2433215H). The valve received water from a reservoir whose level was constantly maintained by closed loop water recirculation, providing stable microliter precision in water delivery over week timescales. In this experiment mice were rewarded with 7 μl of water for a successful trial.

Signal conditioning circuitry. We built signal conditioning circuitry at each home cage to calibrate the sensors and provide noise immunity. Specifically, we 1) used unity gain buffer amplifiers (LM358, Texas Instruments) to protect the joystick signals from noise as they were routed to the computational system; and 2) calibrated the nose poke sensor by biasing the phototransistor at the edge of the linear zone to reduce detection hysteresis, and used a comparator (TL331, Texas Instruments) to convert it to a TTL digital output.

Breakout board and voltage protection system. The digital and analog sensor signals from all the behavioral boxes reached the computational unit (CU) through the breakout board. On the breakout board, for each line we used a thermistor (PRF18BB471QB5RB) to protect the CU from high current surges and Schottky diodes (TBAT54S, Toshiba Electronics) to protect the CU from high voltages. After the voltage/current protection circuitry, analog and digital lines from each behavior box were separated and routed to the right connector.

Computational System.

Hardware. The core of the computational system was a single-board real-time input/output (sbRIO) system (sbRIO-9636, National Instruments). The sbRIO had both a field-programmable gate array (FPGA) and a real-time (RT) processor on the same board. It had 16 analog inputs, 4 analog outputs, and 28 digital input/output ports. The FPGA is programmed through LabVIEW. See “Instructions for Construction” at <https://github.com/GoldbergLab/RodentJoystick> for construction and design details.

Overview of the software infrastructure. There were four major components of the software architecture: 1) Every millisecond the FPGA code processed both the digital inputs (nose poke sensors, touch sensors) and analog inputs (joystick x,y) to determine whether the trial was “live” and whether the “hold” and reach-out contingency was met. The FPGA code also directly interacts with the digital outputs. It controlled the solenoid valve for the water reward, and it could also be used, in future optogenetic experiments, to control the masking light used during photoinhibition and pattern the optogenetic output. The FPGA took $\sim 37 \mu\text{s}$ to process the contingency requirements. At the end of each millisecond, the FPGA wrote all the input

sensor data and the generated outputs to a first in-first out (FIFO) register on the FPGA. Using an onboard FIFO prevented sampling errors and discontinuities. 2) Every second, the code on the RT processor read from the FIFO register using a direct memory access method and packaged it into a single array element. The sbRIO had limited onboard memory and could not hold the vast behavioral data being generated; it required a secondary storage solution. Importantly, the data transfer needed to be fast, in real time, and have temporal continuity. We used Network Stream Objects to set up communications between the sbRIO, and the server-grade acquisition computer. We developed a simple data protocol containing the data, the timestamp, and a frame ID. The data from the RT processor was packaged using this protocol and sent over the network to the acquisition computer. In addition to packaging and transferring the data, the code on the RT processor also had the programmatic access to all the parameters of the task, i.e., hold time, reach angle, water dispense time, optogenetic power, and optogenetic stimulation time and had the ability to deliver rewards to mouse boxes with a tunable Poisson distribution. 3) The code on the acquisition computer (“PC code”) was set up to interface with the appropriate sbRIO and continuously monitored the network stream for any new data packets. If a data packet was detected, the PC code unpacked the data, split it into separate channels, identified the sequence of the data, created filenames for each data file, and then wrote the behavioral data onto the Redundant Array of Independent Disks server with generated filenames. Each filename was coded to provide information on the mouse it came from and the settings of the behavioral paradigm. 4) During training, the behavioral data was automatically analyzed at the end of the day (~11 PM) to determine the new training parameters. The code combined and extracted all valid trials from the day, the number of rewards, distributions of hold times and reach directions. If the animal hadn’t had enough successful trials to meet its daily water requirements, the system communicated with the RT code to automatically dispense water to make up the difference and the contingency of the task remained unchanged. If the animal met its daily requirement of water intake, the contingency was updated to reward only 15% of the trials based on the current distributions of the hold time and reach direction.

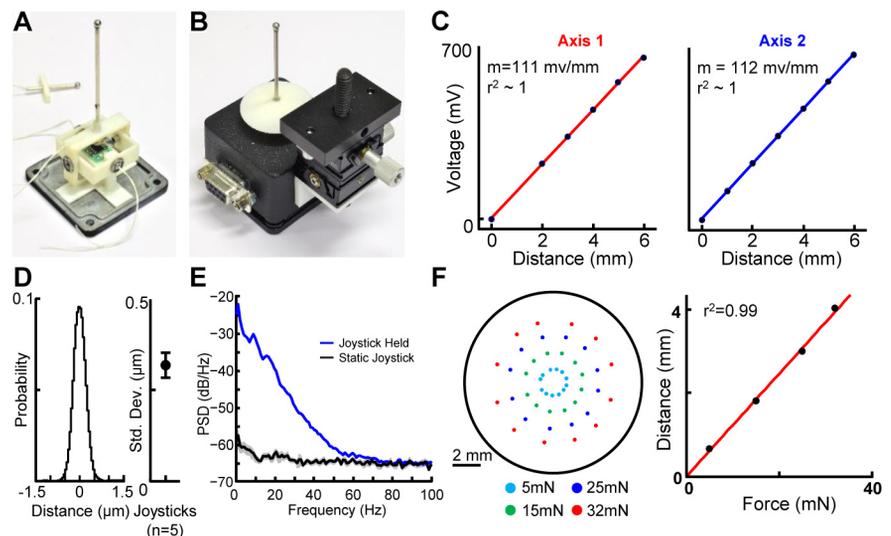
Behavioral shaping. We used a sequence of reward contingencies to train mice to perform “hold-still” + direction-specific center-out reaches. 1) On *day 1* water-restricted mice (~6 h) were placed into joystick-endowed home cages. Water was automatically dispensed with a Poisson distribution (mean 150 s). If the mice simultaneously did a nose poke (np) and contacted both joystick and the fixed post for more than 50 ms they were rewarded with water. 2) On *day 2-day 3* after mice started interacting with the joystick, we changed the reward contingency such that the mice only get rewarded if they contact the

joystick after both the nose poke and fixed post. 3) After the mice learned to contact the joystick last in sequence, we changed the reward contingency such that mice had to move the joystick out to at least 4 mm from the center to get rewarded. 4) Once the mice started reaching out to 4 mm, we gradually shaped the “hold time” of the trajectories. Hold time was defined as the amount of time the trajectory spent within 2-mm radius of the center. We shaped the trajectories by setting the required hold time for reward each day as greater than the 75th percentile of hold time distribution of the previous day. This was repeated till the 75th percentile of the hold time was greater than 100 ms, at which point, the required hold time was set to 100 ms. 5) We then characterized the reach-angle distributions of the mice. The reach angle was defined as the angle at which the trajectories crossed the 4-mm reward threshold. 6) To shape an animal’s reach in a given direction we chose a “reward zone” defined as a 60° arc at the outer threshold (e.g., Figure 7B, center, right). To drive learned changes in reach direction, we rotated this arc daily, such that the new position of the arc would reward 15% of the animal’s trials. In this way, an animal was always “chasing a moving target.” For clockwise (CW) rotation, the arc’s counterclockwise (CCW) edge was placed at the 15th percentile on the CW side of the reach distribution. For CCW rotation, the arc’s CW edge was placed at the 15th percentile on the CW side of the reach angle distribution. The rotations were done till the mean of distribution was different by at least 60° from the initial mean reach direction.

Quantification and Statistical Analyses

Trajectory level analysis. Trajectories were acquired at 1 kHz and low-pass filtered at 50 Hz with an 8-pole Butterworth filter in software. The 50-Hz cutoff was determined by analyzing the spectral power of the joystick trajectories (Fig. 2E) and identifying the frequency where the power in the joystick movements was still distinguishable from the background sensor noise. Hold time was defined as the amount of time the trajectory was within the inner radius (2 mm) starting from joystick contact. The reach direction was defined as the angle at which the outer radius was transected by the trajectory. Instantaneous velocity was calculated as a one-sample difference of the position vector. Note that, in subsequent sections, we refer to “instantaneous velocity” as “velocity.” To make sure that we were specifically analyzing trajectories that were attributable to right forelimb movement with the animal in a consistent posture, only the trajectories that were contacted after the nose poke and fixed-post contact were considered valid trials eligible for reward. If the mouse exited the nose-poke port or lost contact with the fixed post or the joystick, the trial was immediately failed. The mouse also failed the trial if it exited the inner radius earlier than its hold-still requirement

Fig. 2. A joystick for mice. A–E: characterization of the joystick. A and B: the joystick was a manipulator capped with a capacitive touch sensor and mounted on a modified Gimbal assembly equipped with a 2-axis Hall sensor. C–E: Hall sensor voltage measured as a function of distance across 2 nonorthogonal radial axes. D: spatial resolution characterization: measured probability distribution of a single joystick displacement for a single joystick (*left*) and mean \pm SE standard deviation of the distribution across 5 joysticks (*right*). E: example of force-displacement characterization for a single joystick. Each dot indicates the measured displacement by a fixed force, color coded at *bottom* (*left*). Average joystick displacement as a function of force.



or if it reached in the wrong direction. There were no cues for failed or successful rewards (except the water delivery apparatus). However, a blue LED masking light roughly at eye level to the mouse was turned on at joystick contact. The mouse had to recontact the joystick to start a new trial after both successful and failed trials.

Minimum jerk decomposition and trajectory segmentation. To decompose complex trajectories into primitives, we employed two methods: the first was to fit trajectory kinematics to a linear combination of minimum jerk basis functions (called submovements) (Gowda et al. 2015; Rohrer and Hogan 2003; Viviani and Flash 1995), and the second was to break up trajectory kinematics into segments, where segment boundaries were defined by temporally coincident minima of velocity and radius of curvature (Milner 1992; Viviani and Terzuolo 1982). To implement the former, we followed the method from Gowda et al. (2015), using the MATLAB package described therein. In short, the code performed a least-squares fit of the six-dimensional description of joystick kinematics (position, velocity, and acceleration for both the x and y coordinates) to a series of minimum-jerk basis, whose velocity profile took the functional form

$$v(t; t_0, D, A) = (A/D) \times [30\tau^4 - 60\tau^3 + 30\tau^2]$$

where τ gives the normalized time, defined as $\tau = (t - t_0)/D$, t_0 is the initiation time of the submovement, D is the duration of the submovement, and A is the amplitude. Importantly, the function is defined to be zero outside the bounds $[t_0, t_0 + D]$. The fit procedure did not require strong assumptions about the number of submovements before running the fit, but rather iteratively added submovements to improve the overall fit until a cost function threshold was reached.

The second decomposition algorithm, which used coincident velocity and radius of curvature minima and is the primary decomposition method discussed in the main text, did not rely on any optimization or assumptions regarding the parameterization of primitives. To determine these segment boundaries, we identified the velocity and radius of curvature minima for each trajectory and assigned a segment boundary whenever these minima were temporally separated by no more than 1 ms. This method acted as a sharp turn detector and identified temporal boundaries within a trajectory that must have arisen from a new force applied to/via the forelimb.

Segment and submovement analysis. We performed analyses on primitives—obtained using the methods described in the previous section—to gain a finer-scale description of joystick trajectories. Because we are primarily interested in task-related movements, we excluded from our analyses the first primitive produced during each trial. This ensured that our data set did not include motion related to grasping the joystick, but instead only included controlled motion. Furthermore, we excluded from our analyses any primitive with duration < 10 ms. This cutoff was applied to all trajectory segments (whose boundaries are determined by coincident minima in velocity and radius of curvature). Choice of 10 ms for this cutoff was informed by the filtering frequency. The 50-Hz low-pass filter permits detection of 20-ms period sinusoids, each of which will contain two extrema, with 10-ms duration associated with each extremum. From the primitives we obtained through decomposition/segmentation, we extracted various kinematic parameters to describe said primitive. In the case of minimum jerk basis functions, each submovement was uniquely defined by four parameters: the start time, duration, amplitude (path length), and direction of the submovement. Each of these was calculated directly from the code in Gowda et al. (2015). Peak speed for submovements was obtained by taking the ratio of amplitude to duration, multiplied by the constant 1.875 (reflecting the characteristic bell-shape of the submovement basis function).

Similarly, in the case of segments defined by coincident minima of velocity and radius of curvature, we extracted various kinematic parameters, including the duration, path length, peak speed, and direction. Here we defined segment direction as the angle between the initial position of the segment and the position of the segment at its maximum speed.

Finally, for many analyses we separately analyzed the primitives involved in the hold-still and reach portions of the task. The hold primitives for each trial were defined as the primitives that occur before the first inner threshold crossing of the trajectory—this corresponded directly to the definition of the hold period in the task structure. The “reach” primitive for each trajectory was defined as the first primitive to cross the outer threshold, regardless of whether the hold period was completed successfully. Thus each trajectory contained at most one reach segment, but potentially multiple overlapping reach submovements.

Stabilogram diffusion analysis of the hold-still. The stabilogram diffusion function (SDF) gives the mean squared displacement over a specified time window. The SDF was calculated on hold-still part of the trajectories by only analyzing the trajectory up to the onset of the segment that exited the inner radius. The mean squared displacement

$$\langle \Delta r^2 \rangle \text{ for a time interval } \Delta t \text{ is given by } \langle \Delta r^2 \rangle_{\Delta t} = \frac{\sum_{i=1}^N \Delta r_i^2}{N},$$

where N is the total number of points for the time interval across all the trajectories for that condition and Δr^2 is the squared displacement for each of those points. The slope and the offset of the log-log plot of the SDF is estimated by linear fit to the first 10 ms of the SDF. The inflection point is defined as the intersection of the line estimated by the linear fit of the first 10 ms with the tangent drawn at the peak of the SDF in the first 100 ms.

Statistical analyses. Statistical analyses were performed using standard tests in MATLAB, including two-sample t -tests and Wilcoxon signed-rank tests, with Bonferroni corrections for multiple comparisons. To avoid the effects of repeated measures on a single mouse, statistical tests were performed on data aggregated at the mouse level (i.e., the data on which we ran the test consisted of the medians for each mouse across multiple sessions).

RESULTS

A Novel Sensor Quantifies Mouse Forelimb Kinematics with Micron-Millisecond Spatiotemporal Precision

Joysticks can be used in rodents to resolve forelimb kinematics during reach tasks (Francis and Chapin 2004; Kimura et al. 2012; Mathis et al. 2017; Miri et al. 2017; Morandell and Huber 2017; Panigrahi et al. 2015; Slutzky et al. 2010; Wagner et al. 2017). We implemented a novel joystick design to increase spatial precision, reduce displacement force and ensure an isotropic force profile. We designed a capacitive touch-sensing joystick that used contactless magnetic field sensing to detect motion, endowing it with micron-scale resolution (Fig. 2, A–D, average spatial resolution: 320 ± 35 nm, $n = 5$ joysticks). We also replaced the standard two axis spring recentering mechanism with a single pair of magnets, resulting in a stable, uniform force-displacement relationship [Fig. 2F, stiffness: 8.11 mN/mm (or 0.82 gf/mm), $r^2 = 0.99$; see MATERIALS AND METHODS]. These modifications enabled us to resolve tiny details of mouse forelimb motion with minimal load.

To stabilize body posture, the joystick was integrated into a narrow “reward port” consisting of five parts: 1) the joystick, which detects right paw contact and x – y movements; 2) a touch-sensing fixed post positioned for the left paw; 3) two side-walls that constrain the animal’s body position and orientation; 4) an IR-sensing nose poke; and 5) a solenoid-controlled water dispensation spout within tongue’s reach of the nose poke sensor. The requirement that animals engage the joystick, fixed post, and nose poke contacts constrained the animal’s posture and ensured joystick manipulation with the right forepaw (Fig. 3A, Supplemental Movie S1; Supplemental Material for this article is available online at the Journal website.). By

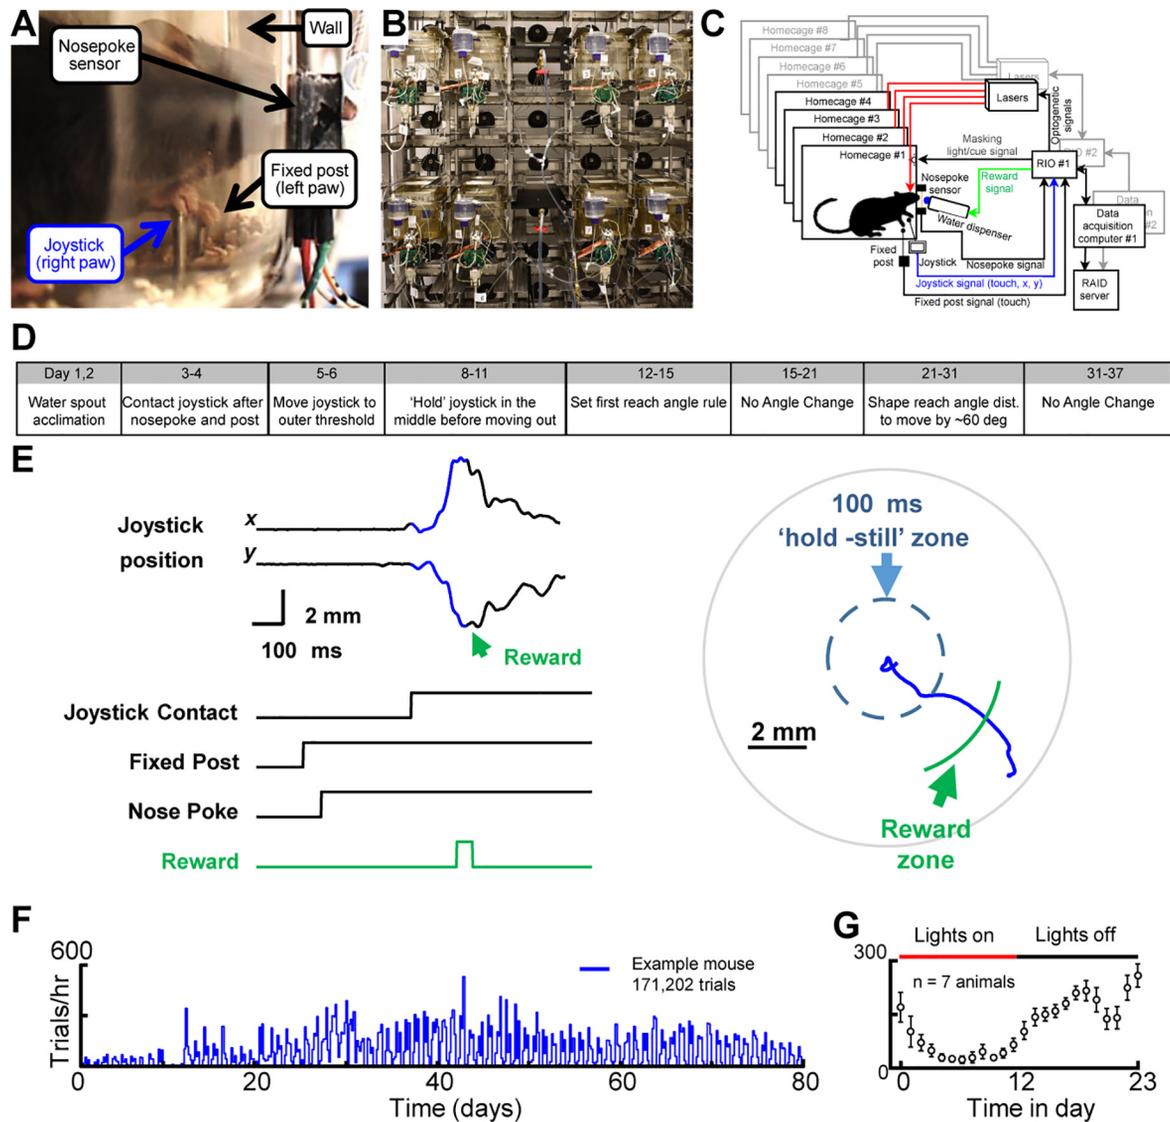


Fig. 3. Automated home cage training of mice in a hold-still center-out reach task. *A*: still frame of mouse engaged in the task. *B*: photograph of home cage training room in mouse facility. *C*: schematic of signal processing pipeline for each home cage. RIO, real-time input/output processor. *D*: timeline of behavioral shaping. *E*: example trajectory and sensor data from a single trial. *F*: trials per hour exhibited by an example animal for 80 days. *G*: trials per hour as a function of time in day (mean \pm SE, $n = 7$ animals).

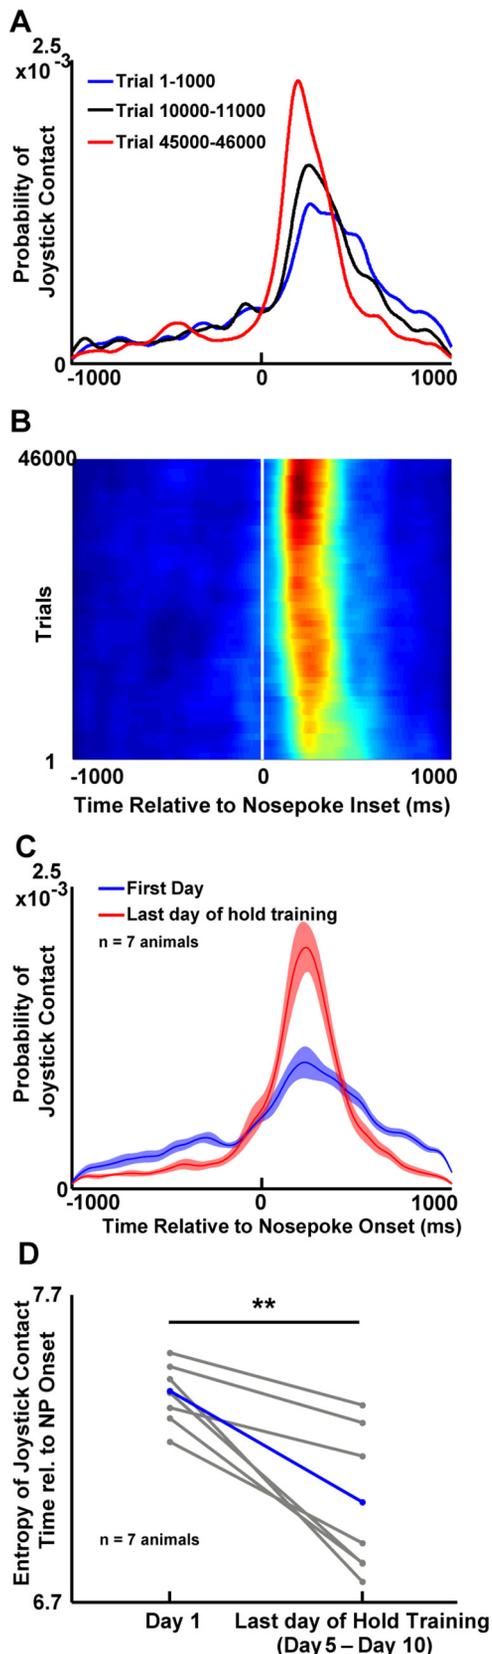
design, each animal was required to manipulate the joystick with its right forelimb. The fixed post and joystick positions could be switched in future studies to further examine potential handedness in mice (Signore et al. 1991).

Automated Home Cage Training of Mice in a Hold-Still Center-Out Reach Task

Automated training facilitates high-throughput experimentation on rodents in sophisticated learning tasks (Erlich et al. 2015; Murphy et al. 2016; Poddar et al. 2013; Woodard et al. 2017). We incorporated the joystick-reward port into rack-mountable, fully automated mouse home cages (Fig. 3, *B–F*). Mice enjoyed continuous, ad libitum access to the joystick for their daily water, resulting in thousands of trajectories per animal per day ($2,904 \pm 531$ trials per day, mean \pm SE, $n = 7$ animals). We built a three-stage signal processing system pipeline to automate training and data acquisition: 1) an FPGA implemented millisecond timescale RT analysis of all sensors,

including joystick position, for closed-loop control of reward dispensation; 2) an RT processor implemented second-time-scale analysis for selective acquisition of trajectory and sensor data associated with eligible trials; and 3) a host PC implemented day-timescale analysis of recent joystick manipulation patterns for automated contingency updates underlying training (Fig. 3*C*; see MATERIALS AND METHODS and “Instructions for Construction” available at <https://github.com/GoldbergLab/RodentJoystick>).

A sequence of fully automated reward contingency updates shaped right forelimb trajectories in a direction-specific center-out reach task (Fig. 3*D*; see *Behavioral shaping* under MATERIALS AND METHODS). First, mice were trained to contact the joystick after the nose poke and fixed-post to ensure that joystick movement was attributable to right paw (Fig. 4, *A* and *B*). The timing of joystick contact became stereotyped with experience (Fig. 4, *C* and *D*; joystick contact onset: 184 ± 14.5 ms after nose poke; entropy of joystick contact time relative to



nose poke onset: 7.37 ± 0.03 (day 1) versus 7.014 ± 0.084 (criterion day), mean \pm SE, $P < 0.01$, paired t -test). Next mice were trained to hold the joystick within an inner radius of 2 mm for 100 ms before reaching past an outer radius of 4 mm (Fig. 3E). This “hold period,” defined as the latency from joystick contact to the moment of inner-radius transection, was implemented to study neural basis of maintaining stability.

Mice Learn To Hold Still Before Reaching Out

Due to the joystick’s low stiffness, the hold-in-center requirement approximates an inverted pendulum or stick balancing problem, in which a control policy is required to produce corrective micromovements to counter deviations from the center position (Anderson 1989; Bhounsule et al. 2015; Cabrera and Milton 2002; Funato et al. 2017). Consistent with this, whereas early in training trajectories exhibited rapid displacements from the center position, later in training all mice produced a clearly resolved sequence of “micromovements” that maintained the joystick within the inner radius for longer periods of time (Fig. 5) [mean hold time: 54 ± 7 ms (first 100 trials) vs. 97 ± 15 ms (at criterion), $P < 0.01$, $n = 7$ animals].

To test whether micromovements produced during the hold reflected an active control policy for maintaining position, we computed SDFs, which plot mean square displacement for all pairs of points in a trajectory as a function of time interval Δt (Collins and De Luca 1993; Peterka 2000) (MATERIALS AND METHODS; Fig. 5H). The analysis effectively distinguishes between different types of motion. For example, the SDF of a classic random walk exhibits a slope of 1. The SDF of a periodic tremor exhibits multiple peaks (Fig. 5G). Persistent motion biased toward increased time-dependent displacement exhibits a slope greater than 1. Finally, antipersistent motion biased toward restoring position exhibits a slope less than 1 (Fig. 5G). In balance tasks such as human standing, SDFs exhibit a characteristic shape: at low intervals, persistent motion is observed. An inflection point is observed at the transition to antipersistent motion thought to mark the onset of control process that corrects for displacements and returns to center position (Peterka 2002).

SDFs of trajectories acquired during the hold period strongly resembled SDFs observed during human standing. At short latencies, the time-dependent displacement adhered to a power law (mean slope of log-log: 1.957 ± 0.0028 , $n = 7$ animals), reflecting persistent motion consistent with the absence of a corrective process. A second mode appeared after a brief delay (mean delay to inflection point: 27.59 ± 1.56 ms, $n = 7$ animals), reflecting the onset of a distinct process that reversed time-dependent displacement (Fig. 5G). These findings suggest that an active policy, potentially similar to one used during the

Fig. 4. Mice learn to contact the joystick after the nose poke. *A*: probability distribution of time of joystick contact relative to nose poke onset in an example mouse at early (trials 1–1,000, red), intermediate (trials 10,000–11,000, black) and late stages of training (trials 45,000–46,000, blue). *B*: evolution of stereotypy in joystick contact time. Distribution of joystick contact times relative to nose poke onset as a function of trial number. *C*: time of joystick contact relative nose poke across animals ($n = 7$ animals, mean \pm SE). *D*: entropy of time of the joystick contact relative to nose poke on day 1 in the home cage and the last day of hold training (range: 5–10 days later across mice).

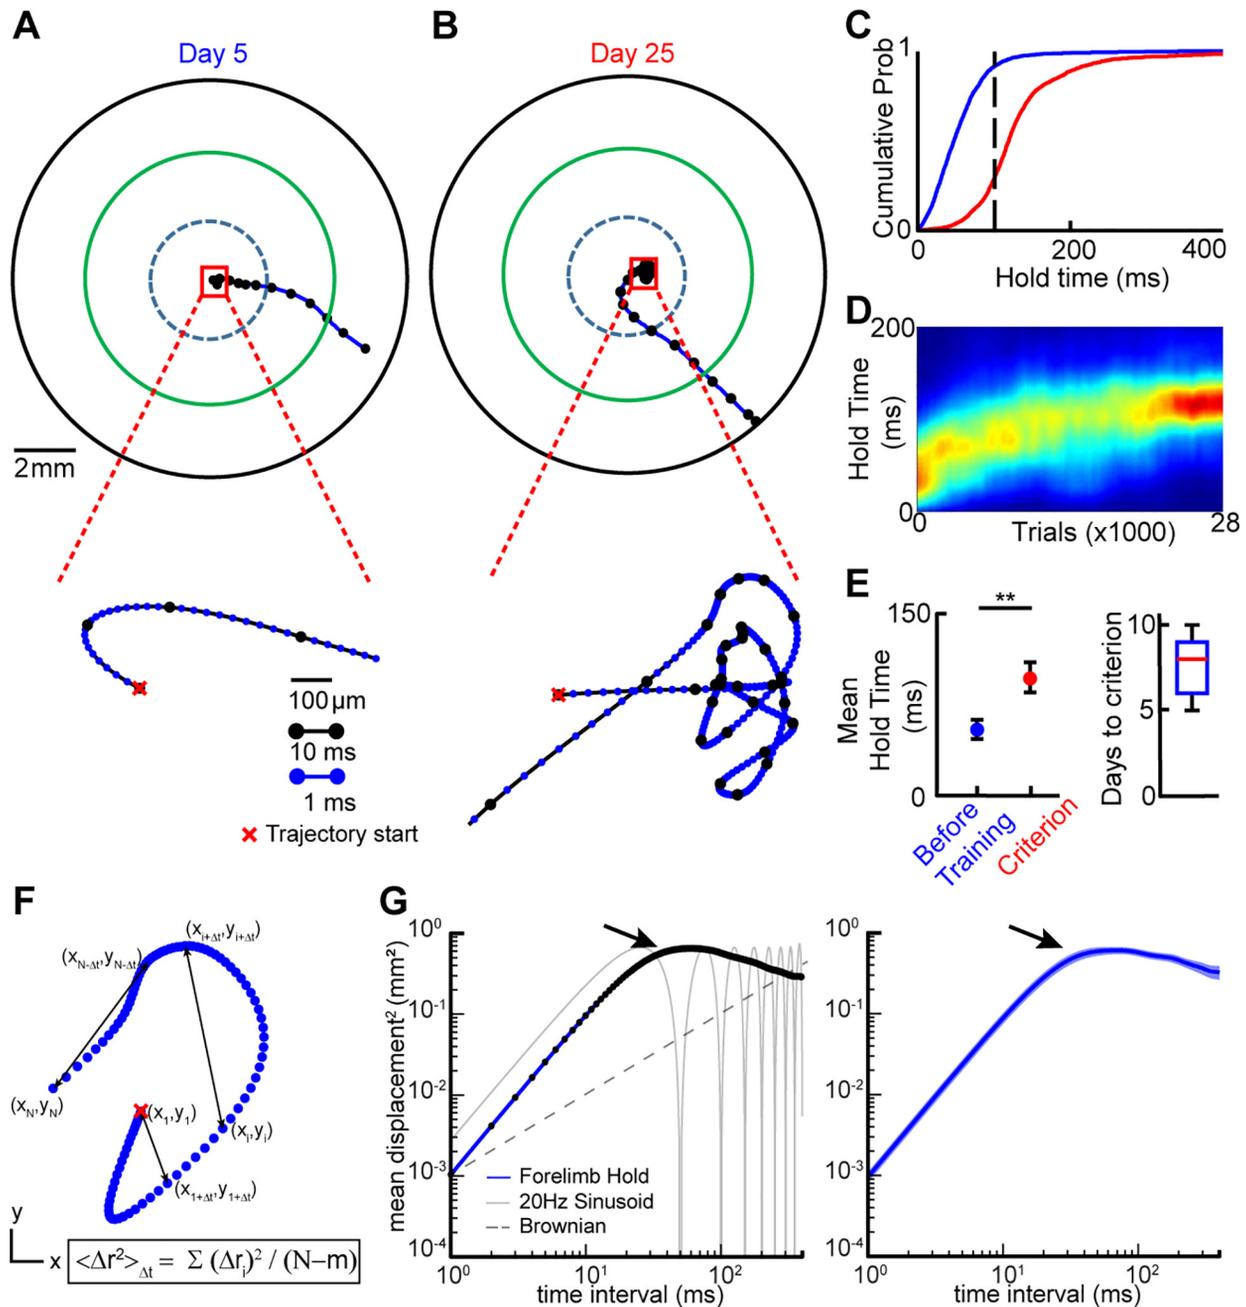


Fig. 5. Mice learn hold-in-center requirement by producing decamicon-scale micromovements. *A*: example trajectory produced at *day 5* since home cage introduction. Blue dashed circle indicates inner “hold” region; green circle indicates reward zone; black circle indicates joystick boundaries. Black dots denote 10-ms intervals in the trajectory. Note distinct displacement scale bar for the expanded view of movement inside the hold region (*bottom*). *B*: data plotted as in *A* for a trajectory produced by the same mouse 20 days later. Expanded view at *bottom* plots details of micromovements produced to maintain hold requirement. *C*: cumulative probability (Prob) of the hold time of trajectories before (blue) and after (red) training for the example mouse. *D*: Heat map showing the probability density plot of hold times as a function of trial number for the mouse shown in *A–C*. *E*: 75th percentile hold time before training and at criterion for all mice (*left*) and number of days to criterion (*right*) (data are mean \pm SE across 7 animals). *F*: schematic describing calculation of the stabilogram diffusion function (SDF) of a short trajectory segment *G*: example SDF from a single animal (*left*, blue line) and across all animals (*right*; $n = 7$, line and shading represent mean \pm SE of SDF). Example SDFs of a purely Brownian motion (*left*, dashed gray) and a pure sinusoid (*left*, solid gray) are shown for comparison.

distinct inverted pendulum problem of maintaining of upright balance, was implemented to achieve the hold-still component of the task. Thus the hold-still part of the task provides an opportunity to study feedback control processes underlying balance.

Mice Learn To Reach in Different Rewarded Directions

Once the contact, hold-still, and reach sequence was learned, all outward reach directions were rewarded, enabling each

animal’s natural reach direction and variability to be quantified (Fig. 6). Mice learned to execute the contact, hold-still, and reach sequence approximately 1 wk after being placed into a home cage (Fig. 5, *C–E*, 8.3 ± 0.75 days, mean \pm SE $n = 7$ animals; see MATERIALS AND METHODS). Reach directions, defined as the angle at which the outer radius was transected, were next rotated clockwise (CW) or counterclockwise (CCW) by contingency updates that rewarded the 15th (for CW) or

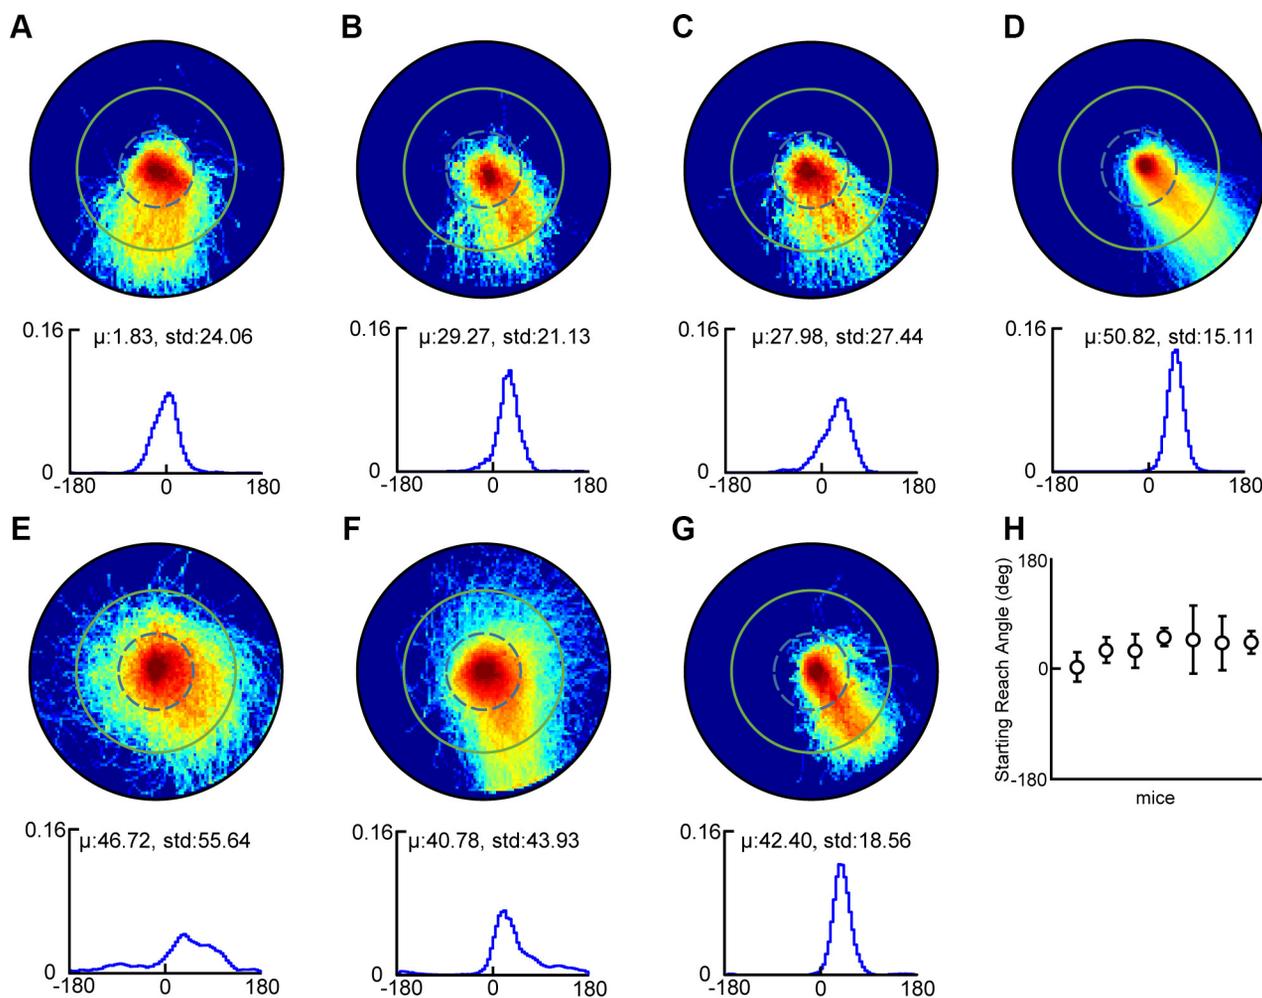


Fig. 6. Trajectory patterns exhibited by mice during all-directions-rewarded stage of training. *A, top*: probability density function of all trajectories produced by a single mouse during a single day, immediately after hold training, when all directions were rewarded (green circle indicates reward zone). *Bottom*: corresponding reach direction distribution (down corresponds to 0°). *B–G*: data plotted as in *A* for all mice in the study. *H*: reach angles (mean \pm SD) produced by each mouse in the study at all-directions-rewarded stage shown in *A–G*.

85th (for CCW) percentile tails of their current reach direction distribution (Fig. 7, *A–C*, MATERIALS AND METHODS). All mice learned to change their reach directions commensurate with rewarded contingency changes and changed their reach direction at a rate of 6.8 ± 2.3 (CW) and 4.6 ± 1.46 (CCW) $^{\circ}$ /day (mean \pm SE) (Fig. 7*C–D*).

Quantification of Primitive Kinematics with Trajectory Decomposition

The joystick's high spatiotemporal precision allowed us to analyze kinematic patterns of movement beyond the relatively coarse trajectory-level metrics described above. Trajectories exhibited complex, multip peaked speed profiles, suggesting that they were composed of a sequence of discrete segments (Figs. 1, *A–C* and 8, *A–C*). We tested the utility of two distinct decomposition algorithms commonly used in primate studies. One algorithm assumes that trajectories are derived from minimum-jerk basis functions with only three parameters: peak speed, duration, and the time in the sequence at which it is generated (see MATERIALS AND METHODS) (Gowda et al. 2015; Rohrer and Hogan 2003; Viviani and Flash 1995). Using this method we were able to decompose mouse trajectories into superpositions of bell shaped velocity curves (Fig. 8, *A–C*),

similar to what was previously observed in mice (Mathis et al. 2017; Panigrahi et al. 2015). Following convention in primates we termed this class of kinematic primitives "submovements." We also used a second decomposition algorithm that imposes segment boundaries at movement discontinuities revealed by temporally coincident minima in the radius of curvature and speed (Figs. 1 and 8, *A* and *B*) (Viviani and Terzuolo 1982). This method essentially functions as a "sharp turn detector" and identifies segments bounded by moments in a trajectory when a new force was imposed on the forelimb and, in a task that lacks external perturbations, also defines moments when an efferent neural command signal must have been generated by the central nervous system (Milner 1992). We term this class of kinematic primitives "segments." Thus we define two types of kinematic primitives: 1) submovements, computed as minimum jerk basis functions; and 2) segments, computed as parts of a the trajectory bound by coincident minima in radius of curvature and speed.

Trajectories were tortuous and complex. Decomposition enabled each trajectory to be analyzed as a sequence of discrete primitives, each of which differed in only a few kinematic parameters such as duration, complexity, peak speed, path length, and direction (Fig. 1*C*). We used trajectory decompo-

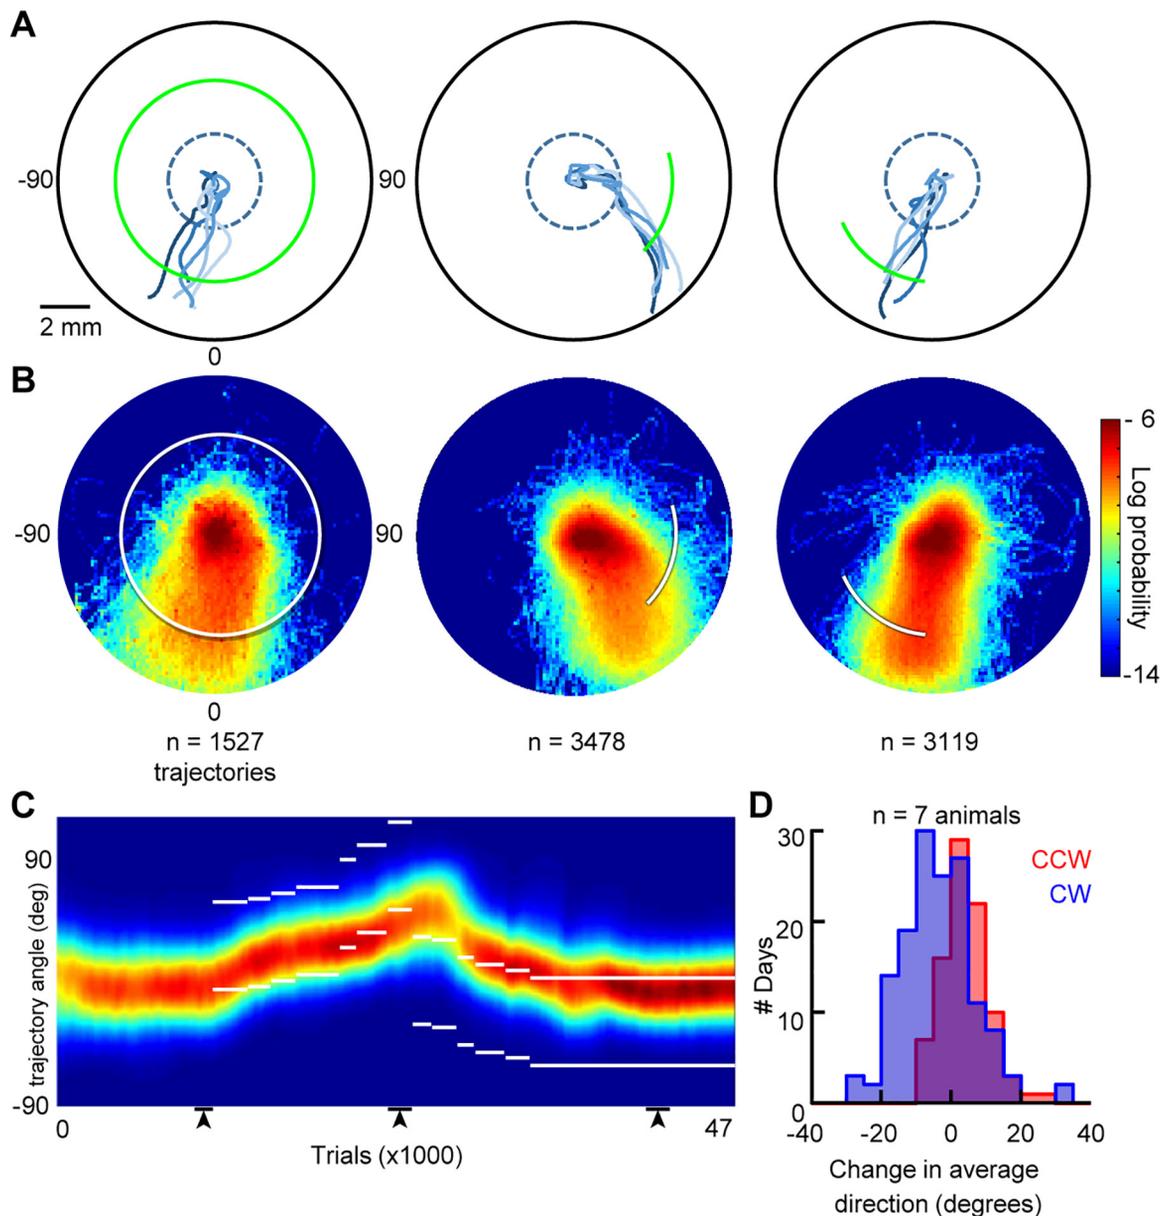


Fig. 7. Mice learn to reach to spatial targets. *A*: example trajectories from 3 nonconsecutive days (5 trajectories plotted per day). Green line indicates the reward zone. *B*: 2D-probability distributions of all trajectories for the mouse on days from which trajectories were sampled from in *A*; white bar indicates the reward zone. Number of trajectories denoted at *bottom*. *C*: probability distribution of reach direction as a function of trial number for an example mouse; white bars indicate the rewarded zone boundaries. Black bars at *bottom* indicate data from the 3 nonconsecutive days plotted in *A* and *B* (left to right order preserved). *D*: histogram showing change in reach direction per day as reward zones were rotated clockwise (CW; negative angle change, blue) or counterclockwise (CCW; positive angle change, red).

sition to test how movements differed during hold-still and reach components of the task. During hold-still periods, durations, peak speeds, and path lengths of primitives produced were significantly smaller than during reaches ($P < 0.05$, Wilcoxon signed-rank test with Bonferroni correction; Fig. 8, Table 1). Yet, interestingly, the difference in segment durations between hold and reach was subtle compared with differences in segment speed and path length (effect size defined as $(\langle \text{reach} \rangle - \langle \text{hold} \rangle) / \langle \text{hold} \rangle$ for duration: 0.85 ± 0.09 ; peak speed: 1.52 ± 0.24 ; path length: 3.22 ± 0.66 , Fig. 8, *G–I*). In fact, segment path length was more strongly predicted by its peak speed ($R^2 = 0.90 \pm 0.02$) than by its duration (0.64 ± 0.09 , grand mean \pm SE across $n = 7$ animals, Fig. 9). Thus mouse forelimb trajectories roughly adhere to the isochrony principle, previously

observed in primate reach and human handwriting tasks, in which a segment's path length is more strongly predicted by its peak speed than by its duration (Viviani and Flash 1995; Viviani and McCollum 1983).

Results on primitive duration, speed, path length, and direction were independently replicated when we instead decomposed trajectories into submovements using a minimum jerk model (Fig. 10, Table 1).

DISCUSSION

We developed a touch-sensing, low-torque joystick that resolves mouse forelimb kinematics with micron-millisecond spatiotemporal precision. We built joysticks into a computer-

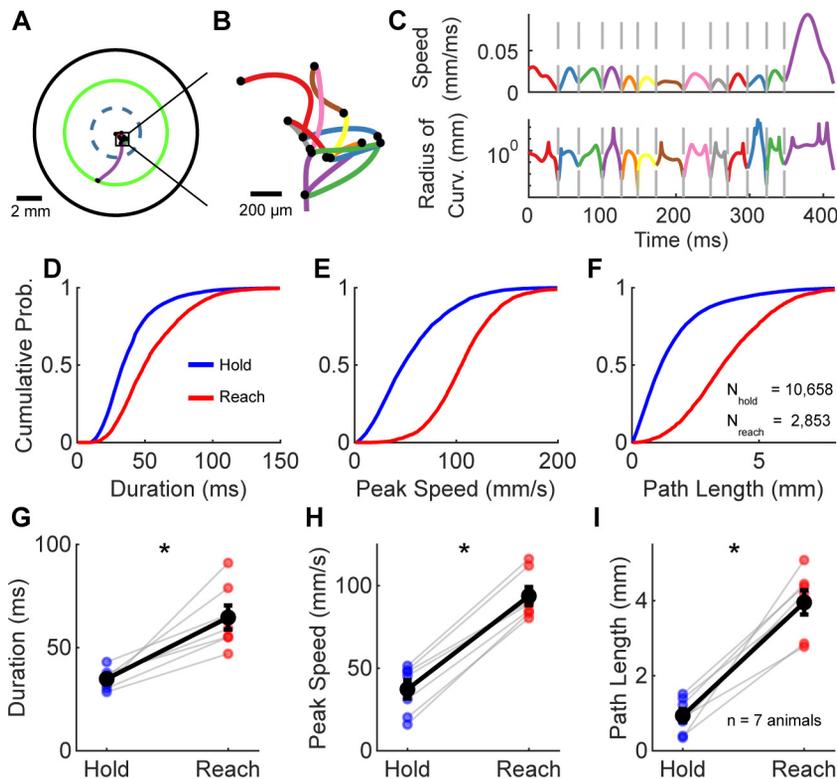


Fig. 8. Primitive kinematics of segments during the hold and the reach. *A–C*: data plotted as in Fig. 1, *A* and *B* for a mouse forelimb trajectory. *A*: example mouse forelimb trajectory with black dots denoting boundaries separating decomposed segments. *B*: expanded view zooming in on micromovements produced during the hold period of the trajectory; different segments are color coded. *C*: speed (top) and radius of curvature (bottom) are plotted as a function of time in the trajectory from *A* and *B*. Segment boundaries (gray dashed lines) are detectable as temporally coincident minima in speed and radius of curvature (see MATERIALS AND METHODS). *D–F*: cumulative probability distributions of durations (*D*), peak speeds (*E*) and path lengths (*F*) of segments from a single session (1 mouse on 1 day). Blue and red traces indicate segments produced in the hold zone and the single segment in a trajectory that transected the outer radius (i.e., the reach segment), respectively. *G–I*: comparison of median values for duration (*G*), peak speed (*H*), and path length (*I*) across different animals for hold (blue) and reach (red) segments ($n = 7$ mice, $*P < 0.05$ Wilcoxon signed-rank test; see Table 1). Black data points show grand mean \pm SE across all 7 mice.

ized home cage system that automatically trains mice to produce complex, directed center-out forelimb trajectories. We also used trajectory decomposition to test hypotheses about how kinematic parameters of primitives are controlled during hold-still and reach components of the task. This automated system presents a platform for testing how genetically tractable neural circuits in mice maintain positional stability and direct the limb to spatial targets.

Previous studies have examined joystick, lever or trackball manipulation in response to cues (Kimura et al. 2012; Miri et al. 2017; Morandell and Huber 2017; Sanders and Kepecs 2012; Wagner et al. 2017), to overcome external force fields (Mathis et al. 2017), or to directly measure force (Becker et al. 2016; Fowler et al. 1994; Francis and Chapin 2004; Hays et al. 2013; Pasquini et al. 2018; Vigarar et al. 2013). Importantly, some of these systems can also be used to aid learning and rehabilitation due to their ability to partially or completely assist in task performance (Pasquini et al. 2018; Spalletti et al. 2014). We developed a task in which freely moving animals, without cues, visual guidance, or externally-induced sensory

Table 1. Kinematics of segments and submovements

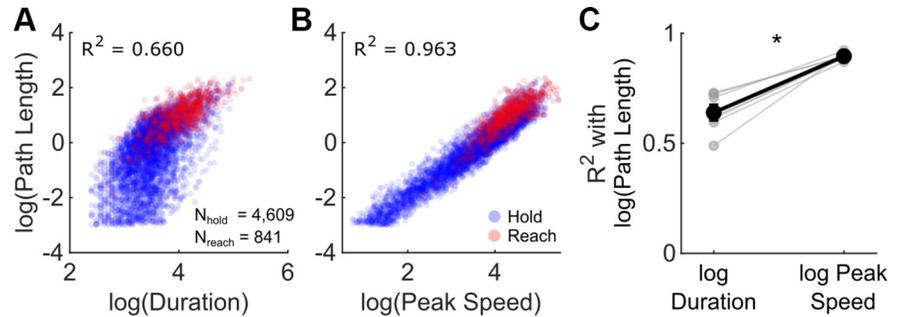
Parameter	Hold	Reach	<i>P</i> -Value (Wilcoxon signed rank test)
<i>Segments</i>			
Duration, ms	34.8214 \pm 1.8754	64.5893 \pm 5.8180	0.0234
Peak speed, mm/s	37.2505 \pm 5.5466	93.7205 \pm 5.4379	0.0234
Path length, mm	0.9366 \pm 0.1752	3.9537 \pm 0.3218	0.0234
<i>Submovements</i>			
Duration, ms	30.6609 \pm 0.9434	47.2012 \pm 2.0409	0.0234
Peak speed, mm/s	25.0533 \pm 3.5159	61.0414 \pm 3.8092	0.0234
Path length, mm	0.4209 \pm 0.0641	1.5523 \pm 0.0834	0.0234

Hold and reach values are means \pm SE.

prediction errors, produce complex trajectories that are directed to spatial targets. Several extensions of our system are possible. For example, our system can easily be adapted to include cues or dynamic external force fields, and group housing with radio frequency identification-tagged animals can allow for a four- to fivefold increase in training capacity for the same amount of hardware (Silasi et al. 2018). The system could also be rendered more compatible with electrophysiology by adding electronics to mitigate potential cross-talk between the capacitive touch sensor and neurophysiological signals (Totten et al. 2018).

A novel feature of our task was the requirement to hold still. Static maintenance of limb posture during appetitive tasks may require specific patterns of activity to actively suppress movement and/or actively maintain current limb position (Ebbesen and Brecht 2017; Kaufman et al. 2014; Shadmehr 2017). Given the joystick's compliance, holding still while applying downward force (i.e., during gripping) created an inverted pendulum, or stick-balancing, problem, in which corrective micromovements are required to maintain position (Anderson 1989; Bhounsule et al. 2015; Cabrera and Milton 2002). To succeed in this task, any drift from center position must be counteracted, which requires proprioception of limb displacement and movement in the opposite direction of ongoing motion. All mice gradually learned to hold their right forelimb still in a small 2-mm "hold zone," and all mice implemented this hold by producing clearly resolved micromovements with speeds, path lengths, and durations significantly less than during reaches. Stabilogram diffusion analysis suggested that these micromovements were not randomly produced but were consistent with an active control policy to maintain center position that, interestingly, strongly resembled one used for maintenance of upright posture in humans (Peterka 2002). It would be interesting to test how cortical, striatal, and cerebellar activity

Fig. 9. Segment path length is better predicted by its peak speed than its duration. Joint distributions of segment path length and duration (A) and segment path length and peak speed in “hold” (blue) and “reach” (red) segments (B) for an example mouse. C: coefficients of determination for path length predicted from segment duration or peak speed for each mouse (gray) with grand mean \pm SE (black). * $P < 0.05$ Wilcoxon signed-rank test; $n = 7$ mice.



manipulations affect performance in this balancelike task for the forelimb. The ability to hold still before outward reaching may also enable, in future studies, joystick contact-triggered optical manipulation to control activity in target circuits before the generation of the outward-directed reach.

After executing the hold, all mice reached out in different directions to spatial targets. We found that mice learned to reach in whatever direction was associated with water reward. Importantly, the rate at which mice learned to reach in different directions was relatively slow ($\sim 4,000$ of trials per 10°). In other studies in head-fixed mice, animals learn to reach in different directions relatively quickly, even on a trial-by-trial basis depending on the position of a pellet, water droplet, or joystick target zone (Galinañes et al. 2018; Kimura et al. 2012; Miri et al. 2017; Morandell and Huber 2017; Sanders and Kepecs 2012). In uncued, freely moving tasks where behavior is incrementally shaped, learning can be much slower. The number of trials required to change reach direction was on the same order of magnitude as the number of trials for rats to increase intertap intervals (Kawai et al. 2015) or for songbirds to change the pitch or note intervals of their songs (Ali et al. 2013; Andalman and Fee 2009; Tumer and Brainard 2007).

To parse the large, complex data sets generated by our experiments, we employed two analysis methods for extracting

movement “building blocks” from forelimb trajectories: 1) segmentation based on coincident minima of velocity and radius of curvature, and 2) submovement decomposition based on fitting trajectories to minimum jerk basis functions. Both methods represent distinct, testable hypotheses regarding how complex movements are generated by the nervous system. In the case of minimum jerk submovements, forelimb trajectories are decomposed into a series of parameterized basis functions, all of which have the same velocity profile by definition; the form for this minimum jerk basis function is the result previous studies showing this form of movement across different animals and experimental conditions (Gowda et al. 2015). By contrast, segments defined by coincident minima in velocity and radius of curvature are not parameterized and thus make fewer assumptions about the shape of movement primitives (Milner 1992; Viviani and Terzuolo 1982). As a result, segment shapes are more variable. Future experiments that allow simultaneous monitoring of neural activity during our joystick task may enable identification of possible neural correlates of these two frameworks for motion primitives. Moreover, such experiments could allow us to test for signatures of nonprimitive-based frameworks, particularly those that take an approach based on optimal feedback control to explain motor planning and execution (Todorov and Jordan 2002).

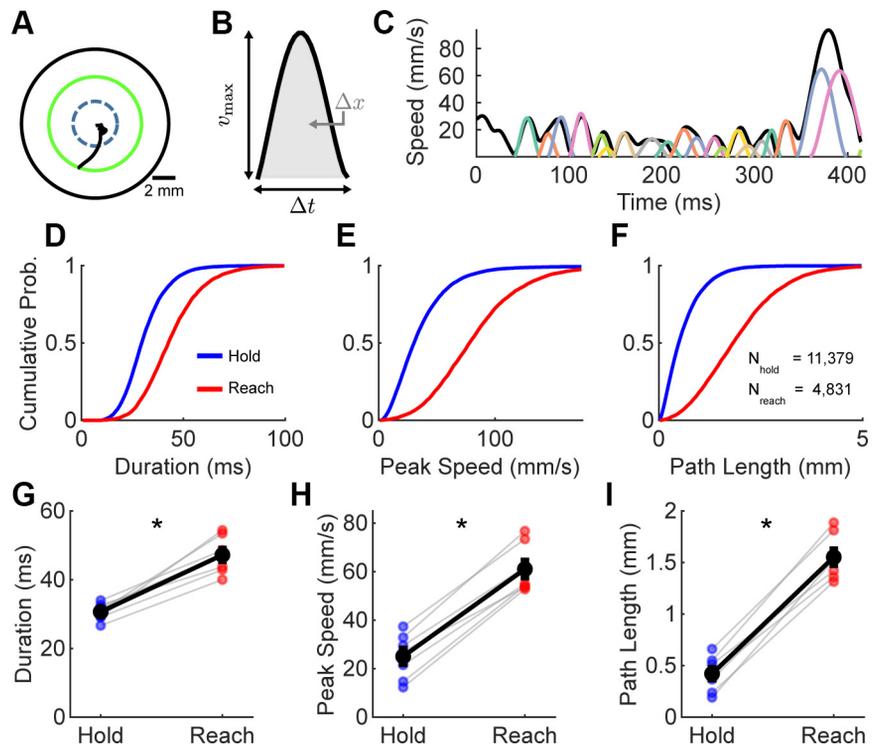


Fig. 10. Primitive kinematics of submovements during the hold and the reach. A: example mouse forelimb trajectory. B: example of a minimum-jerk basis function. C: example decomposition of the trajectory in A into minimum jerk basis functions, or submovements, using algorithms previously used in primate studies (Gowda et al. 2015; Rohrer and Hogan 2003) (see MATERIALS AND METHODS). D–F: cumulative probability distributions of durations (D), peak speeds (E), and path lengths (F) of submovements. Blue and red traces indicate submovements produced in the hold zone and the single submovement in a trajectory that transected the outer radius (i.e., the reach submovement), respectively. G–I: comparison of median values for duration (G), peak speed (H), and path length (I) across different animals for hold (blue) and reach (red) submovements ($n = 7$ mice, * $P < 0.05$ Wilcoxon signed-rank test; see Table 1). Black data points show grand mean \pm SE across all 7 mice.

Our system offers several advantages for the study of mouse forelimb control. First, automated home cage training will enable high-throughput experimentation. Second, the system operates closed loop and is easily compatible with optogenetic experiments that enable the activity of specific cell types and/or projections to be manipulated at precise task events, such as joystick contact or forepaw position. Third, the joystick's relatively low compliance makes holding-still a difficult, yet achievable task that opens up opportunities for studying neural mechanisms of postural stability. Finally, the joystick provides an increased spatial precision that resolves submillimeter details of movements, even at sharp turns. This resolution enables trajectory decomposition methods, previously used in primates, to segment complex trajectories into kinematic primitives.

ACKNOWLEDGMENTS

We thank Joe Fetcho, Melissa Warden, Andy Ruina, Lena Ting, and Goldberg laboratory members for helpful discussions and providing comments on the manuscript.

GRANTS

J. H. Goldberg was supported by NIH (DP2HD087952), the Pew Charitable Trusts, and the Esther A. and Joseph Klingenstein Fund. T. Bollu and S. C. Whitehead were supported by the Cornell Mong Neurotech Fellowship.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

T.B. and J.H.G. conceived and designed study; T.B., N.P., J.W., N.S., and R.S. performed experiments; T.B. and S.C.W. analyzed data; T.B., S.C.W., I.C., and J.H.G. interpreted results of experiments; T.B., S.C.W., and J.H.G. prepared figures; T.B., S.C.W., B.K., and J.H.G. drafted manuscript; T.B., S.C.W., and J.H.G. edited and revised manuscript; T.B., S.C.W., B.K., and J.H.G. approved final version of manuscript.

ENDNOTE

At the request of the authors, readers are herein alerted to the fact that additional materials related to this manuscript [software (trajectory analysis, real time control software, automation analysis) and the hardware design schematics, including part files and building instructions] may be found at the institutional website of one of the authors, which at the time of publication they indicate is: <https://www.github.com/GoldbergLab/RodentJoystick>. These materials are not a part of this manuscript, and have not undergone peer review by the American Physiological Society (APS). APS and the journal editors take no responsibility for these materials, for the website address, or for any links to or from it.

REFERENCES

- Ali F, Otchy TM, Pehlevan C, Fantana AL, Burak Y, Ölveczky BP. The basal ganglia is necessary for learning spectral, but not temporal, features of birdsong. *Neuron* 80: 494–506, 2013. doi:10.1016/j.neuron.2013.07.049.
- Andalman AS, Fee MS. A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors. *Proc Natl Acad Sci USA* 106: 12518–12523, 2009. doi:10.1073/pnas.0903214106.
- Anderson CW. Learning to control an inverted pendulum using neural networks. *IEEE Contr Syst Mag* 9: 31–37, 1989. doi:10.1109/37.24809.
- Becker AM, Meyers E, Sloan A, Rennaker R, Kilgard M, Goldberg MP. An automated task for the training and assessment of distal forelimb function in a mouse model of ischemic stroke. *J Neurosci Methods* 258: 16–23, 2016. doi:10.1016/j.jneumeth.2015.10.004.
- Bernstein NA. *The Coordination and Regulation of Movements*. Oxford, UK: Pergamon, 1967.
- Bhounsule PA, Ruina A, Stiesberg G. Discrete-decision continuous-actuation control: balance of an inverted pendulum and pumping a pendulum swing. *J Dyn Syst Meas Control* 137: 051012, 2015. doi:10.1115/1.4028851.
- Cabrera JL, Milton JG. On-off intermittency in a human balancing task. *Phys Rev Lett* 89: 158702, 2002. doi:10.1103/PhysRevLett.89.158702.
- Collins JJ, De Luca CJ. Open-loop and closed-loop control of posture: a random-walk analysis of center-of-pressure trajectories. *Exp Brain Res* 95: 308–318, 1993. doi:10.1007/BF00229788.
- Desmurget M, Gaveau V, Vindras P, Turner RS, Broussolle E, Thobois S. On-line motor control in patients with Parkinson's disease. *Brain* 127: 1755–1773, 2004. doi:10.1093/brain/awh206.
- Ebbesen CL, Brecht M. Motor cortex - to act or not to act? *Nat Rev Neurosci* 18: 694–705, 2017. doi:10.1038/nrn.2017.119.
- Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* 4: e05457, 2015. doi:10.7554/eLife.05457.
- Flash T, Hochner B. Motor primitives in vertebrates and invertebrates. *Curr Opin Neurobiol* 15: 660–666, 2005. doi:10.1016/j.conb.2005.10.011.
- Flash T, Hogan N. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5: 1688–1703, 1985. doi:10.1523/JNEUROSCI.05-07-01688.1985.
- Fortier PA, Kalaska JF, Smith AM. Cerebellar neuronal activity related to whole-arm reaching movements in the monkey. *J Neurophysiol* 62: 198–211, 1989. doi:10.1152/jn.1989.62.1.198.
- Fowler SC, Davison KH, Stanford JA. Unlike haloperidol, clozapine slows and dampens rats' forelimb force oscillations and decreases force output in a press-while-licking behavioral task. *Psychopharmacology (Berl)* 116: 19–25, 1994. doi:10.1007/BF02244866.
- Francis JT, Chapin JK. Force field apparatus for investigating movement control in small animals. *IEEE Trans Biomed Eng* 51: 963–965, 2004. doi:10.1109/TBME.2004.827463.
- Fu Q-G, Flament D, Coltz JD, Ebner TJ. Relationship of cerebellar Purkinje cell simple spike discharge to movement kinematics in the monkey. *J Neurophysiol* 78: 478–491, 1997. doi:10.1152/jn.1997.78.1.478.
- Funato T, Sato Y, Fujiki S, Sato Y, Aoi S, Tsuchiya K, Yanagihara D. Postural control during quiet bipedal standing in rats. *PLoS One* 12: e0189248, 2017. doi:10.1371/journal.pone.0189248.
- Galiñanes GL, Bonardi C, Huber D. Directional reaching for water as a cortex-dependent behavioral framework for mice. *Cell Reports* 22: 2767–2783, 2018. doi:10.1016/j.celrep.2018.02.042.
- Giszter SF. Motor primitives—new data and future questions. *Curr Opin Neurobiol* 33: 156–165, 2015. doi:10.1016/j.conb.2015.04.004.
- Gowda S, Overduin SA, Chen M, Chang YH, Tomlin CJ, Carmena JM. Accelerating submovement decomposition with search-space reduction heuristics. *IEEE Trans Biomed Eng* 62: 2508–2515, 2015. doi:10.1109/TBME.2015.2434595.
- Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. Flow of cortical activity underlying a tactile decision in mice. *Neuron* 81: 179–194, 2014. doi:10.1016/j.neuron.2013.10.020.
- Hays SA, Khodaparast N, Sloan AM, Hulsey DR, Pantoja M, Ruiz AD, Kilgard MP, Rennaker RL II. The isometric pull task: a novel automated method for quantifying forelimb force generation in rats. *J Neurosci Methods* 212: 329–337, 2013. doi:10.1016/j.jneumeth.2012.11.007.
- Inzelberg R, Flash T, Schechtman E, Korczyn AD. Kinematic properties of upper limb trajectories in idiopathic torsion dystonia. *J Neurol Neurosurg Psychiatry* 58: 312–319, 1995. doi:10.1136/jnnp.58.3.312.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* 17: 440–448, 2014. doi:10.1038/nn.3643.
- Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP. Motor cortex is required for learning but not for executing a motor skill. *Neuron* 86: 800–812, 2015. doi:10.1016/j.neuron.2015.03.024.
- Kimura R, Saiki A, Fujiwara-Tsukamoto Y, Ohkubo F, Kitamura K, Matsuzaki M, Sakai Y, Isomura Y. Reinforcing operandum: rapid and reliable learning of skilled forelimb movements by head-fixed rodents. *J Neurophysiol* 108: 1781–1792, 2012. doi:10.1152/jn.00356.2012.
- Majsak MJ, Kaminski T, Gentile AM, Flanagan JR. The reaching movements of patients with Parkinson's disease under self-determined maximal speed and visually cued conditions. *Brain* 121: 755–766, 1998. doi:10.1093/brain/121.4.755.
- Mathis MW, Mathis A, Uchida N. Somatosensory cortex plays an essential role in forelimb motor adaptation in mice. *Neuron* 93: 1493–1503.e6, 2017.

- Milner TE.** A model for the generation of movements requiring endpoint precision. *Neuroscience* 49: 487–496, 1992. doi:10.1016/0306-4522(92)90113-G.
- Miri A, Warriner CL, Seely JS, Elsayed GF, Cunningham JP, Churchland MM, Jessell TM.** Behaviorally selective engagement of short-latency effector pathways by motor cortex. *Neuron* 95: 683–696.e11, 2017.
- Morandell K, Huber D.** The role of forelimb motor cortex areas in goal directed action in mice. *Sci Rep* 7: 15759, 2017. doi:10.1038/s41598-017-15835-2.
- Murphy TH, Boyd JD, Bolaños F, Vanni MP, Silasi G, Haupt D, LeDue JM.** High-throughput automated home-cage mesoscopic functional imaging of mouse cortex. *Nat Commun* 7: 11611, 2016. doi:10.1038/ncomms11611.
- Mussa-Ivaldi FA, Giszter SF, Bizzi E.** Linear combinations of primitives in vertebrate motor control. *Proc Natl Acad Sci USA* 91: 7534–7538, 1994. doi:10.1073/pnas.91.16.7534.
- Omrani M, Kaufman MT, Hatsopoulos NG, Cheney PD.** Perspectives on classical controversies about the motor cortex. *J Neurophysiol* 118: 1828–1848, 2017. doi:10.1152/jn.00795.2016.
- Panigrahi B, Martin KA, Li Y, Graves AR, Vollmer A, Olson L, Mensh BD, Karpova AY, Dudman JT.** dopamine is required for the neural representation and control of movement vigor. *Cell* 162: 1418–1430, 2015. doi:10.1016/j.cell.2015.08.014.
- Pasquini M, Lai S, Spalletti C, Cracchiolo M, Conti S, Panarese A, Caleo M, Micera S.** A robotic system for adaptive training and function assessment of forelimb retraction in mice. *IEEE Trans Neural Syst Rehabil Eng* 26: 1803–1812, 2018. doi:10.1109/TNSRE.2018.2864279.
- Peterka RJ.** Postural control model interpretation of stabilogram diffusion analysis. *Biol Cybern* 82: 335–343, 2000. doi:10.1007/s004220050587.
- Peterka RJ.** Sensorimotor integration in human postural control. *J Neurophysiol* 88: 1097–1118, 2002. doi:10.1152/jn.2002.88.3.1097.
- Poddar R, Kawai R, Ólveczky BP.** A fully automated high-throughput training system for rodents. *PLoS One* 8: e83171, 2013. doi:10.1371/journal.pone.0083171.
- Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH.** Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* 478: 387–390, 2011. doi:10.1038/nature10436.
- Rohrer B, Fasoli S, Krebs HL, Volpe B, Frontera WR, Stein J, Hogan N.** Submovements grow larger, fewer, and more blended during stroke recovery. *Mot Contr* 8: 472–483, 2004. doi:10.1123/mcj.8.4.472.
- Rohrer B, Hogan N.** Avoiding spurious submovement decompositions: a globally optimal algorithm. *Biol Cybern* 89: 190–199, 2003. doi:10.1007/s00422-003-0428-4.
- Sanders JJ, Kepecs A.** Choice ball: a response interface for two-choice psychometric discrimination in head-fixed mice. *J Neurophysiol* 108: 3416–3423, 2012. doi:10.1152/jn.00669.2012.
- Schwartz AB.** Useful signals from motor cortex. *J Physiol* 579: 581–601, 2007. doi:10.1113/jphysiol.2006.126698.
- Scott SH.** The role of primary motor cortex in goal-directed movements: insights from neurophysiological studies on non-human primates. *Curr Opin Neurobiol* 13: 671–677, 2003. doi:10.1016/j.conb.2003.10.012.
- Shadmehr R.** Distinct neural circuits for control of movement vs. holding still. *J Neurophysiol* 117: 1431–1460, 2017. doi:10.1152/jn.00840.2016.
- Shadmehr R, Wise SP.** *The Computational Neurobiology of Reaching and Pointing: A Foundation for Motor Learning.* Cambridge, MA: MIT Press, 2005.
- Shenoy KV, Sahani M, Churchland MM.** Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* 36: 337–359, 2013. doi:10.1146/annurev-neuro-062111-150509.
- Signore P, Nosten-Bertrand M, Chaoui M, Roubertoux PL, Marchaland C, Perez-Diaz F.** An assessment of handedness in mice. *Physiol Behav* 49: 701–704, 1991. doi:10.1016/0031-9384(91)90305-8.
- Silasi G, Boyd JD, Bolaños F, LeDue JM, Scott SH, Murphy TH.** Individualized tracking of self-directed motor learning in group-housed mice performing a skilled lever positioning task in the home cage. *J Neurophysiol* 119: 337–346, 2018. doi:10.1152/jn.00115.2017.
- Stutzky MW, Jordan LR, Bauman MJ, Miller LE.** A new rodent behavioral paradigm for studying forelimb movement. *J Neurosci Methods* 192: 228–232, 2010. doi:10.1016/j.jneumeth.2010.07.040.
- Spalletti C, Lai S, Mainardi M, Panarese A, Ghionzoli A, Alia C, Gianfranceschi L, Chisari C, Micera S, Caleo M.** A robotic system for quantitative assessment and poststroke training of forelimb retraction in mice. *Neurorehabil Neural Repair* 28: 188–196, 2014. doi:10.1177/1545968313506520.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002. doi:10.1038/nn963.
- Totten D, Novik L, Christe K, Lemoy MJ, Roberts J, Carmena JM, Morecraft RJ, Ganguly K.** An automated behavioral apparatus to assess distal forelimb function in non-human primates (Preprint). bioRxiv 396572, 2018. doi:10.1101/396572.
- Tumer EC, Brainard MS.** Performance variability enables adaptive plasticity of ‘crystallized’ adult birdsong. *Nature* 450: 1240–1244, 2007. doi:10.1038/nature06390.
- Turner RS, Anderson ME.** Pallidal discharge related to the kinematics of reaching movements in two dimensions. *J Neurophysiol* 77: 1051–1074, 1997. doi:10.1152/jn.1997.77.3.1051.
- Tye KM, Deisseroth K.** Optogenetic investigation of neural circuits underlying brain disease in animal models. *Nat Rev Neurosci* 13: 251–266, 2012. doi:10.1038/nrn3171.
- Vigaru BC, Lambercy O, Schubring-Giese M, Hosp JA, Schneider M, Osei-Atiemo C, Luft A, Gassert R.** A robotic platform to assess, guide and perturb rat forelimb movements. *IEEE Trans Neural Syst Rehabil Eng* 21: 796–805, 2013. doi:10.1109/TNSRE.2013.2240014.
- Viviani P, Flash T.** Minimum-jerk, two-thirds power law, and isochrony: converging approaches to movement planning. *J Exp Psychol Hum Percept Perform* 21: 32–53, 1995. doi:10.1037/0096-1523.21.1.32.
- Viviani P, McCollum G.** The relation between linear extent and velocity in drawing movements. *Neuroscience* 10: 211–218, 1983. doi:10.1016/0306-4522(83)90094-5.
- Viviani P, Terzuolo C.** Trajectory determines movement dynamics. *Neuroscience* 7: 431–437, 1982. doi:10.1016/0306-4522(82)90277-9.
- Wagner MJ, Kim TH, Savall J, Schnitzer MJ, Luo L.** Cerebellar granule cells encode the expectation of reward. *Nature* 544: 96–100, 2017. doi:10.1038/nature21726.
- Wolff SB, Ólveczky BP.** The promise and perils of causal circuit manipulations. *Curr Opin Neurobiol* 49: 84–94, 2018. doi:10.1016/j.conb.2018.01.004.
- Wong AL, Haith AM, Krakauer JW.** Motor planning. *Neuroscientist* 21: 385–398, 2015. doi:10.1177/1073858414541484.
- Woodard CL, Bolaños F, Boyd JD, Silasi G, Murphy TH, Raymond LA.** An automated home-cage system to assess learning and performance of a skilled motor task in a mouse model of Huntington’s disease. *eNeuro* 4: ENEURO.0141-17.2017, 2017. doi:10.1523/ENEURO.0141-17.2017.
- Woodworth RS.** Accuracy of voluntary movement. *The Psychological Review: Monograph Supplements* 3(3): i, 1899.