Supplemental Figure 1

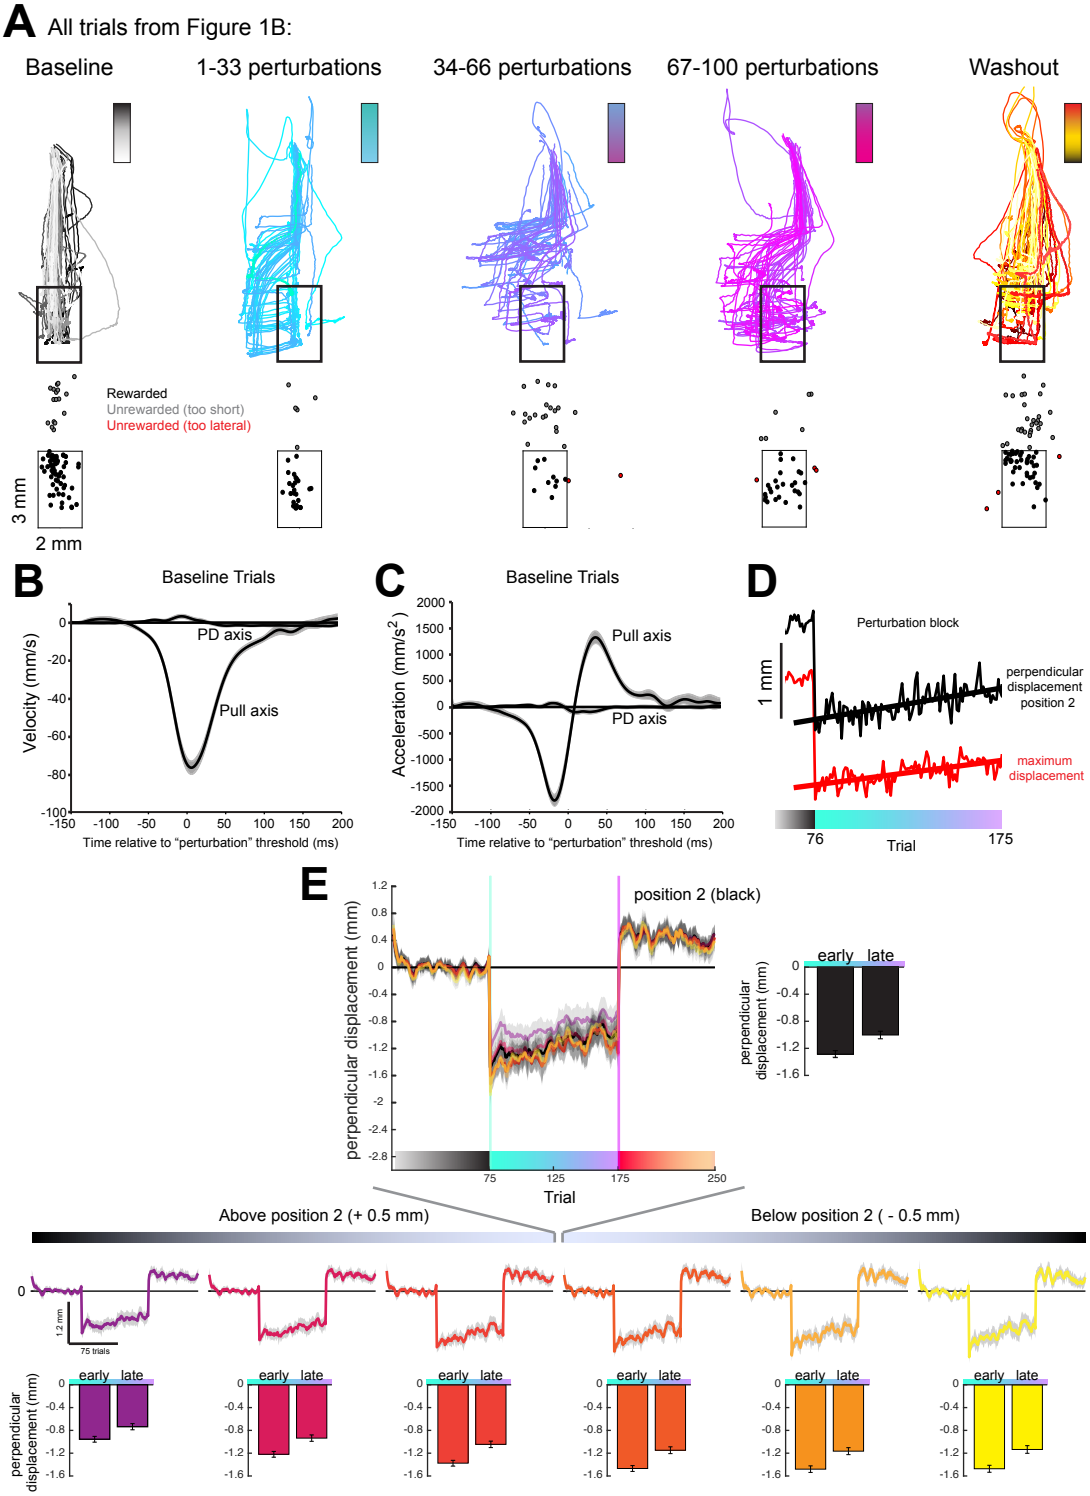

Supplemental Figure 1. Characterization of kinematics and adaptation.
Related to Figure 1.

(A) Top: These are all the trials in the session shown in Figure 1B (left). They are color coded within block, as denoted in the panel. The black box denotes the reward zone. Bottom: endpoints in relation to reward box. Note that during baseline most unrewarded trials are too short. Later unrewarded trials are also either too short, or too laterally deviated.
(B) Velocity profile of baseline pulls (neg. values denote movements towards the target box), aligned to the point where a perturbation onset would occur.
(C) Acceleration profile of baseline pulls, aligned as in B.
(D) Quantification of the perturbation block trials by either position 2 (black; mean is shown, as in Figure 1, and see STAR Methods), or by taking the maximum displacement in every trial (red; mean is shown). Thicker lines indicate a linear fit of the average data, demonstrating quantification of adaptation is comparable if using maximum or computing the position in a specific spatial location during the pull.
(E) Top: Perpendicular displacement as shown in Figure 1 (in black), overlaid with 6 slices spaced from 0.5 above to 0.5 mm below position 2. Bottom: Each slice that is overlaid above is shown independently, from dark purple to yellow denoting slices from the top of the pull towards the end point (but all slices are taken within the force field onset). Quantification of early perturbation trials (average ± sem of trials 76-86) and late perturbation trials (164-175) in each slice is shown below the respective slice.

Supplemental Figure 2

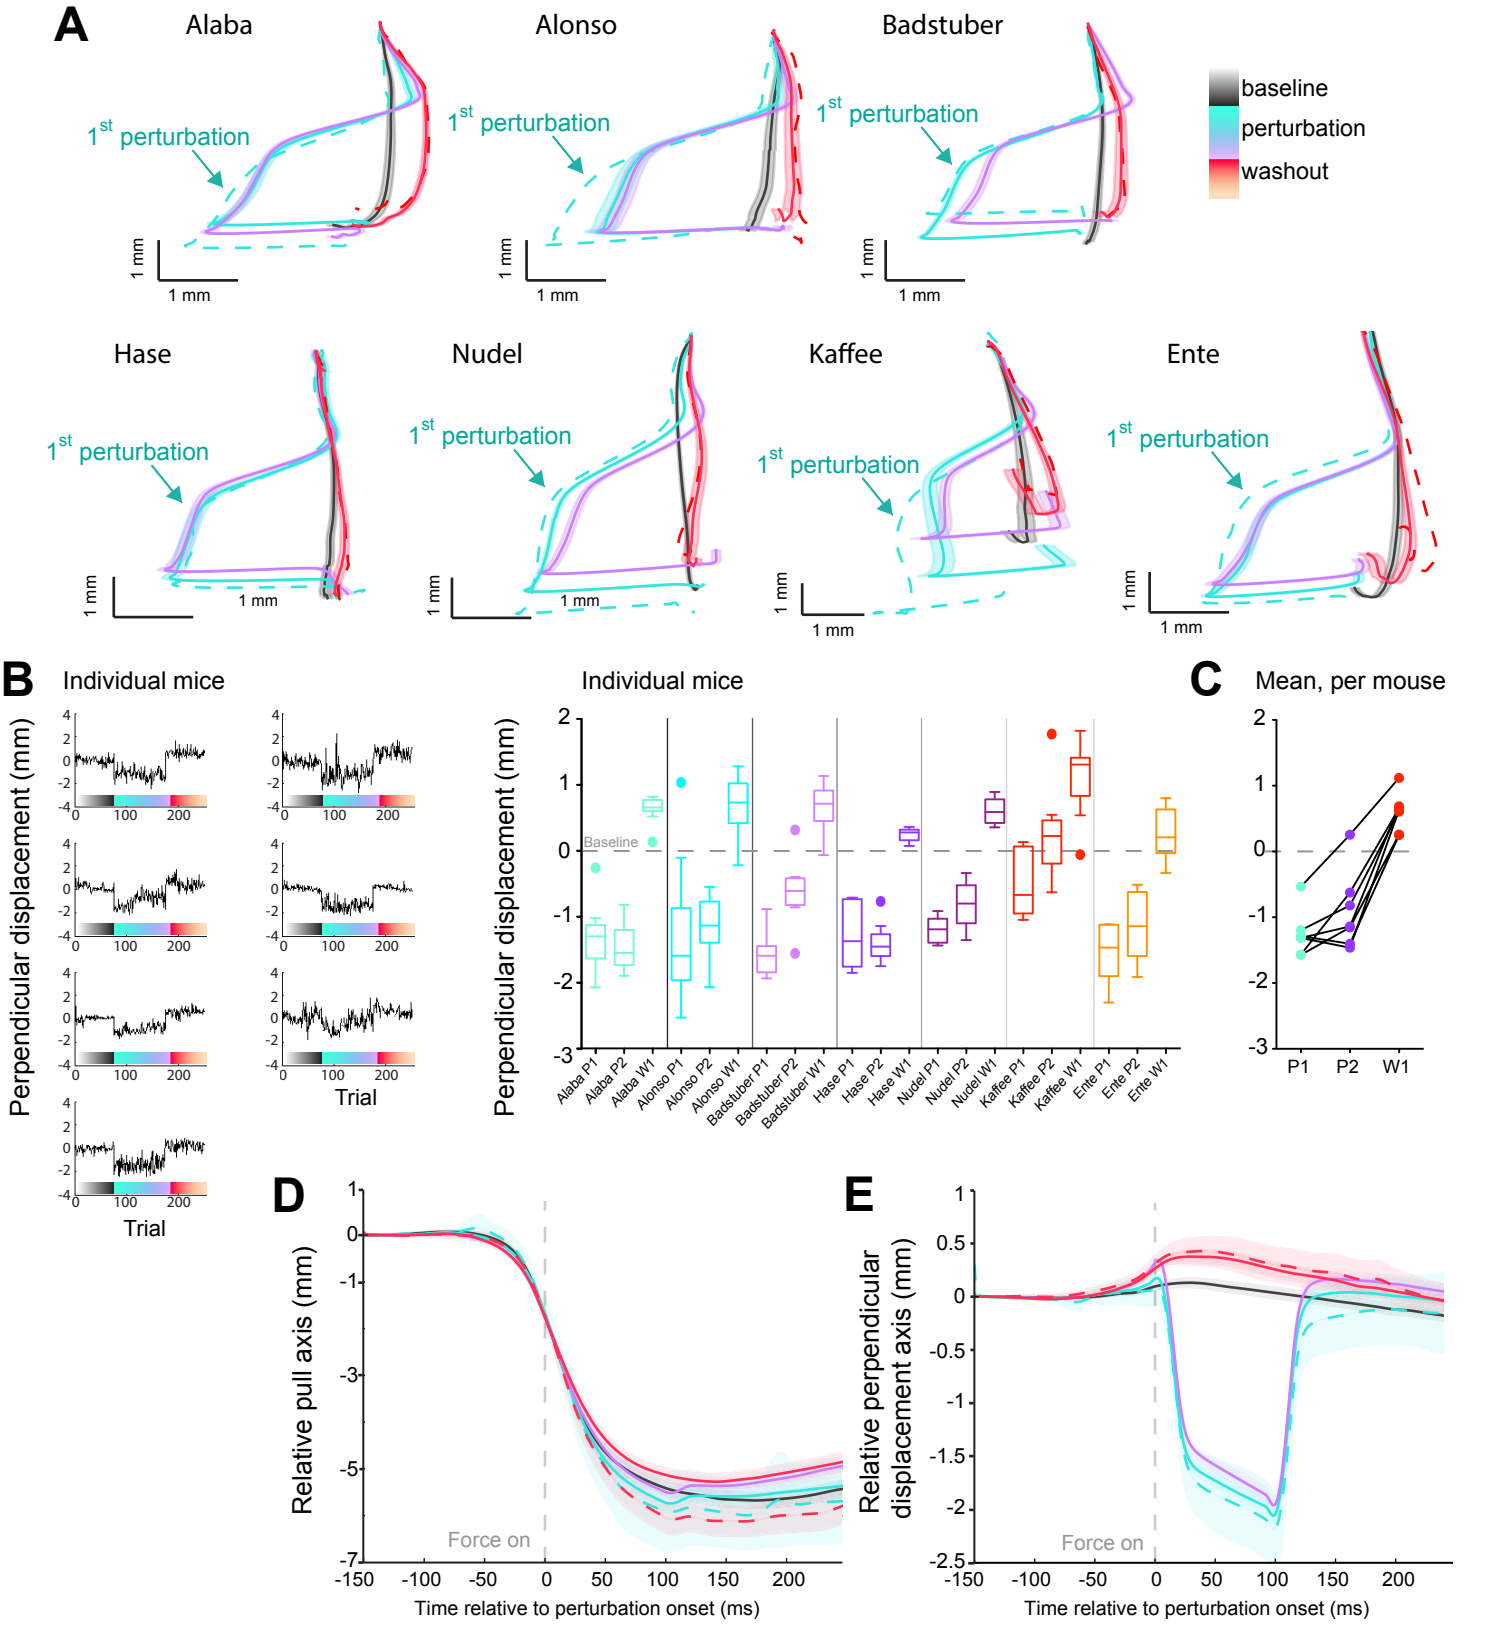

Supplemental Figure 2. Adaptation in individual mice.
Related to Figure 1.

(A) Temporally averaged trajectories as computed and colored in Figure 1B, per mouse.

(B) Left: mean perpendicular displacement (PD) for each mouse at position 2. Right: PD of the first ten (P1) and last ten (P2) perturbation trials, and the first ten (W1) washout trials, averaged for each mouse, demonstrating all mice displayed adaptation (reduction in perpendicular displacement and/or an aftereffect). Baseline (corrected within mouse) is the dashed grey line. Whiskers: Tukey.

(C) Mean, per mouse, of perpendicular displacement: first ten, last ten, and first ten washout trials.

(D) Relative distance versus time in the pull axis, i.e. towards the mouse (see Methods); baseline (black), first perturbation (dashed green), early perturbation (green), late perturbation (purple), first washout (dashed red), early washout (red). All pulls were aligned to perturbation onset (corresponds to 0 ms in time axis). Starting location is zeroed per pull (see STAR Methods)

(E) Relative distance in the perpendicular axis, same coloring as in D.

Supplemental Figure 3

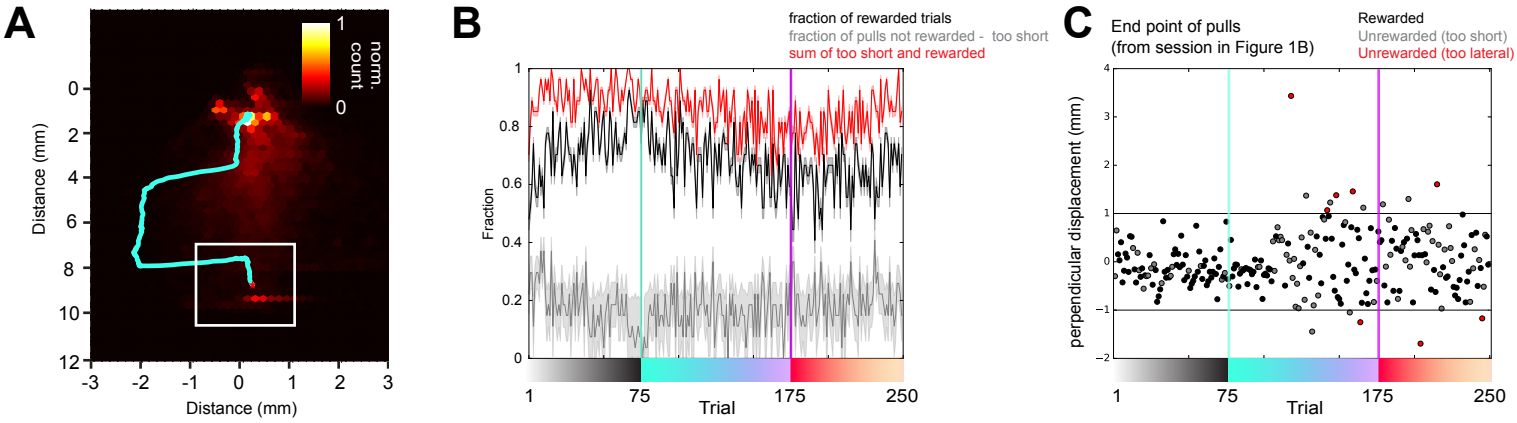

Supplemental Figure 3. Reward rate and first-perturbation illustration.
Related to Figure 2.
(A) Individual mouse example; Heat map of baseline-training days to demonstrate the typical area covered by the paw. The counts per hexagonal tile were transformed to log-scale and normalized. Sea green line is the mouse's first ever perturbation trial, white box is target box, and the red star is the reward location (paw position at opening of water valve).
(B) Average reward rate of the adaptation task (shading is s.e.m.). Fraction of too short trials is indicated in grey, this type of error is constant throughout session. Other errors are too lateral or too quick to leave reward box, as indicated in the main text
(C) Example end point errors from the session shown in Figure 1B and Figure S1A.

Supplemental Figure 4

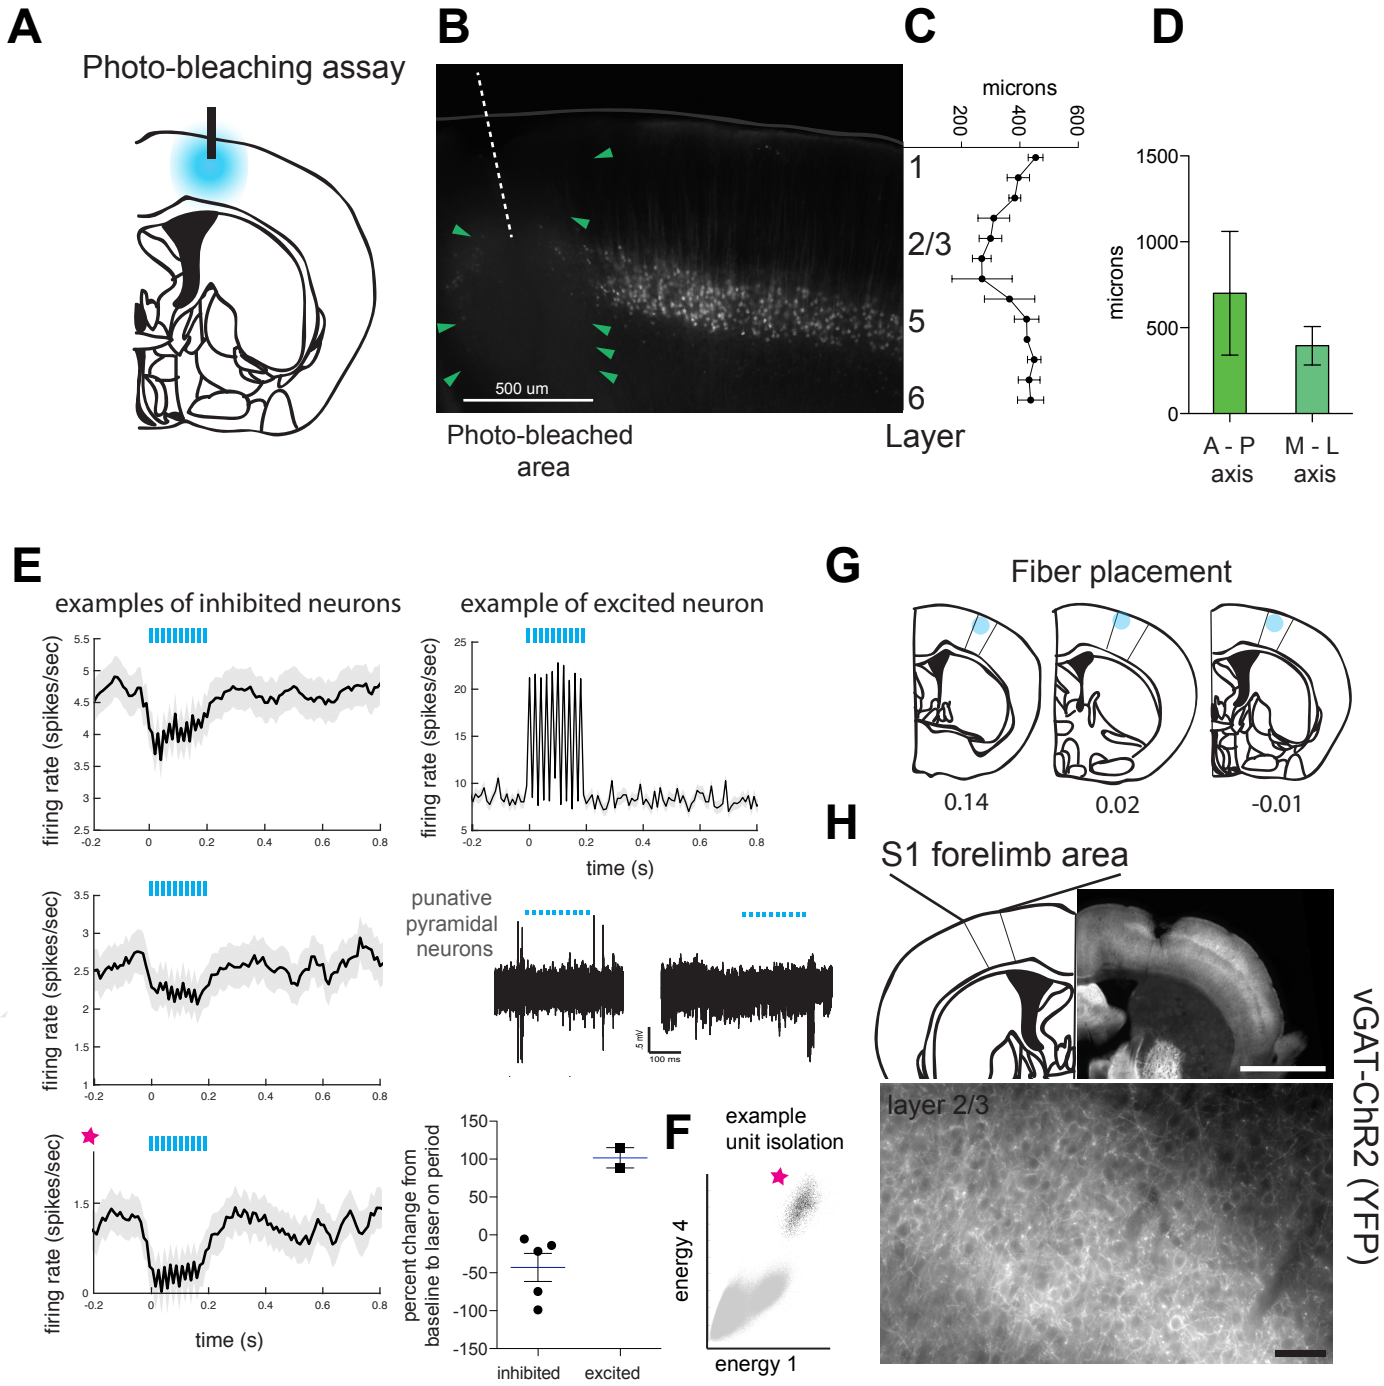

Supplemental Figure 4. Photoinhibition and light spread characterization.
Related to Figure 3.

(A) Diagram of the implantation of the optical fiber. Thy1-YFPH mice were implanted with an optical fiber (see STAR Methods) then 20 mW of 473 nm light was applied for 600 seconds (n= 3 mice; as in (Guo et al., 2014)). The mice were immediately perfused and post-fixed for 24 hours before quantification.

(B) Example image of the photo-bleached area in the medial-lateral axis; green arrows highlight the boundaries. White dashed line shows fiber track position. Grey line denotes the pia.

(C) Quantification of layer-specific bleaching across, measured as average distance across the boundaries of bleaching, in 13 spatial bins (spanning layer 1 to 6). As Thy1 marks predominantly layer 5, apical dendrites, somas and axons were carefully taken into account to measure the extent of bleaching.

(D) Brains were cut into 100 um sections. Quantification of maximum bleaching areas defined across the anterior to posterior (A-P) axis and the medial to lateral (M-P) axis (n = 3 mice, mean ± standard deviation). As the light should radiate in a spherical way, averaging across the A-P and M-L axis yields a light spread of 421 ± 180 µm (i.e. ± ~300 µm in each direction from the tip of the fiber).

(E) PSTH of example neurons in S1 during photoinhibition testing, raw data examples from tetrode-based recordings, and quantification (n = 5) neurons; -50.2 ± 22% activity during photoinhibition compared to baseline) and one example presumptive GABA neuron (spike width < 0.35) showing time-locked excitation to laser application.

(F) Example of a well isolated unit (corresponding to the pink start in E); L-ratio = 0.03.

(G) Atlas schematic showing location of optical fiber placement in each of the VGAT-ChR2 mice (n = 3 mice). Values indicate distance from Bregma.

(H) Top: Example histology image (vGAT-ChR2). Scale bar is 2 mm. Bottom: Higher magnification of VGAT-ChR2; scale bar is 50 µm.

Supplemental Figure 5

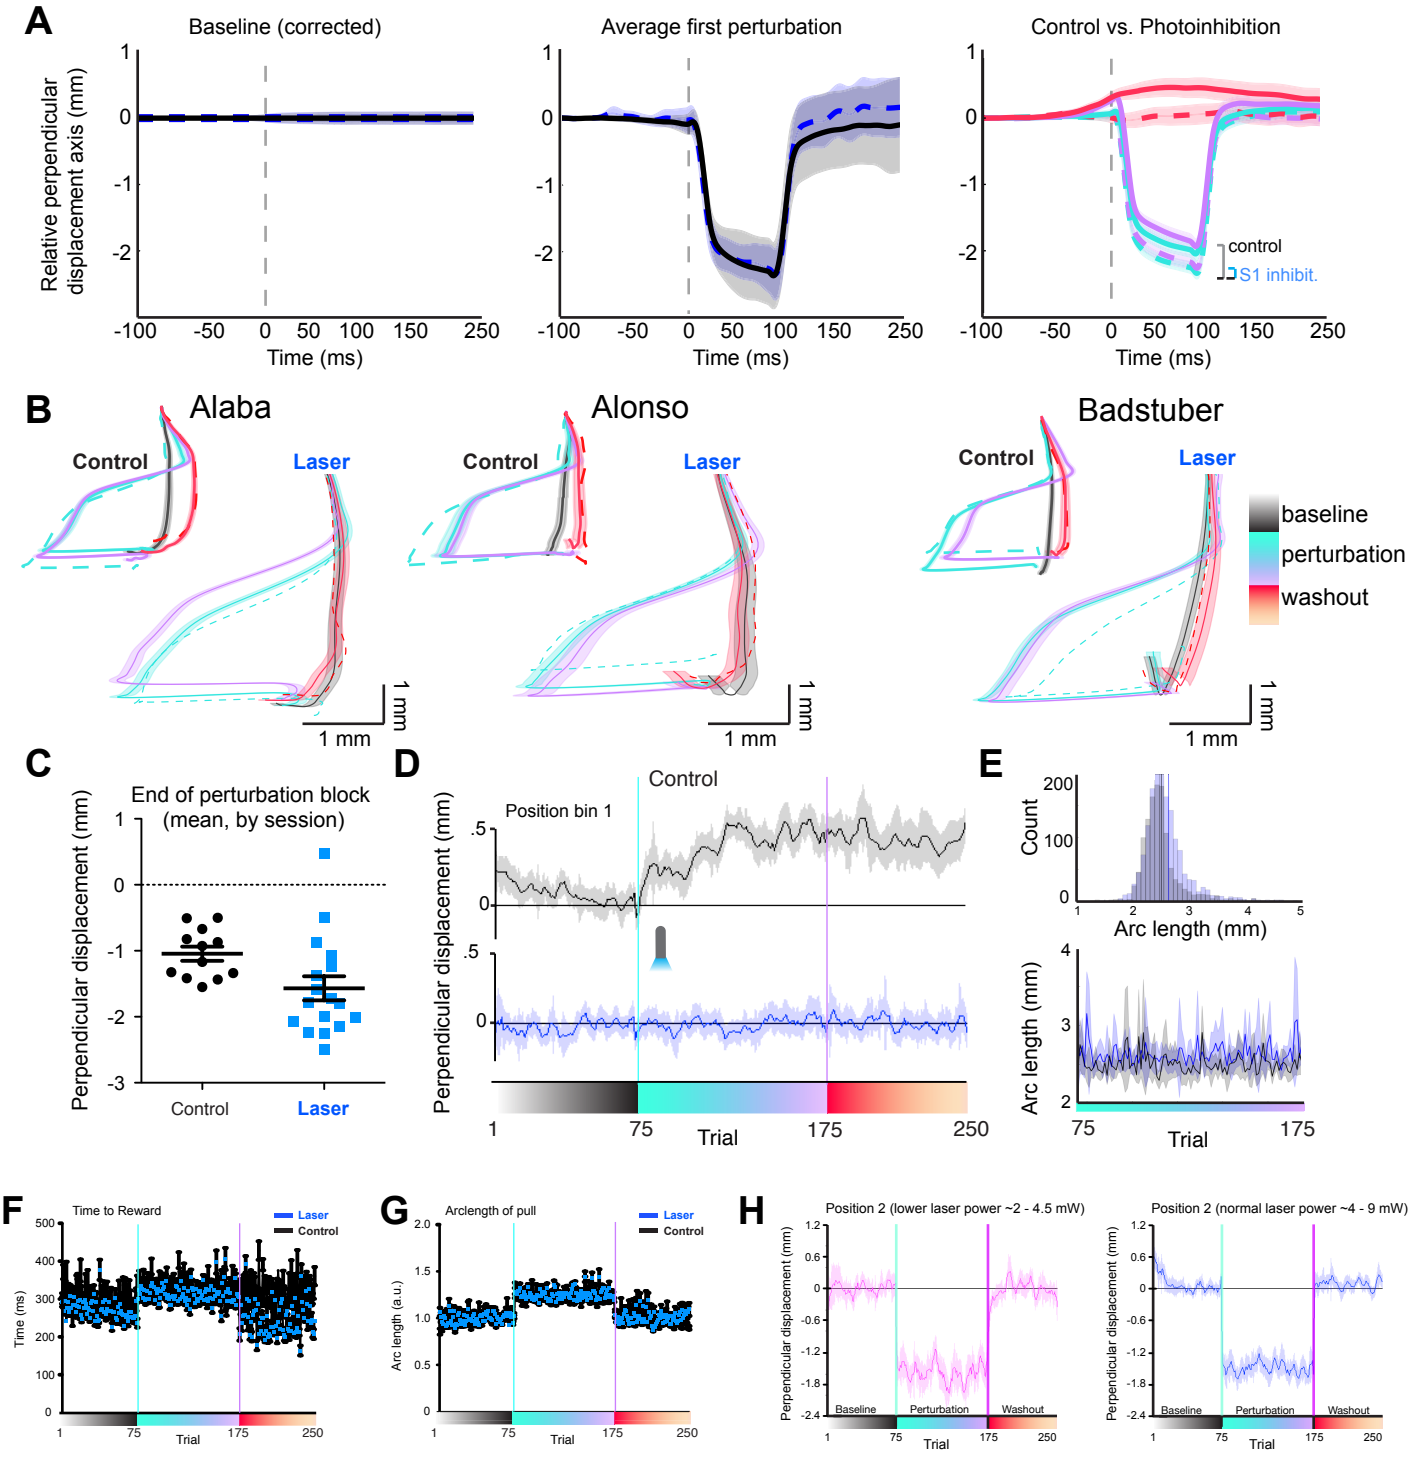

Supplemental Figure 5. Photoinhibition abolishes adaptation.
Related to Figure 3.

(A) Relative perpendicular displacement in the perpendicular axis (see STAR Methods) during baseline (left panel), the first perturbation trial (middle panel), and the perturbation and washout blocks (right panel): baseline (black), first perturbation (dashed green), early perturbation (green), late perturbation (purple), first washout (dashed red), and early washout (red). Dashed lines are for photoinhibition, solid for control sessions. The brackets highlight the change from first perturbation to the late perturbation average, per condition – the horizontal black bar is the first perturbation level (see middle panel for comparison). Note that initial perpendicular deviation is similar for both conditions (middle panel).
(B) Individual mice averaged trajectories during photoinhibition sessions, with corresponding control sessions (as shown in Figure S2).
(C) Averages of trials 165-175 for different sessions (n =12 (control) and n =17 (S1 inhibited)) per condition.
(D) Perpendicular displacement across blocks in position bin 1 (above the force onset), demonstrating predictive steering in control but not photoinhibition sessions.
(E) Arc length of "L" turn post force off is slightly longer during photoinhibition sessions. Control: 2.50 ± 0.02 mm vs. photoinhibition: 2.67 ± 0.01 mm, Mann-Whitney test, $P = 6.4 \times 10^{-28}$. Bottom panel shows session average ± s.e.m. of arc length by trial.
(F) Time to reward (from the spatial point of force onset to water valve onset).
(G) Arc length of pull (from the spatial point of "force on" to reward onset), normalized by the average arc length during baseline.
(H) Averaged perpendicular displacement at position 2 for additional sessions with normal light levels (4 - 9 mW, as was applied in Figure 3), vs. light at half-power (2 - 4.5 mW). 'Normal' power level is shown in blue (n = 12 sessions), and low level is shown in pink (n = 12 sessions).

Supplemental Figure 6

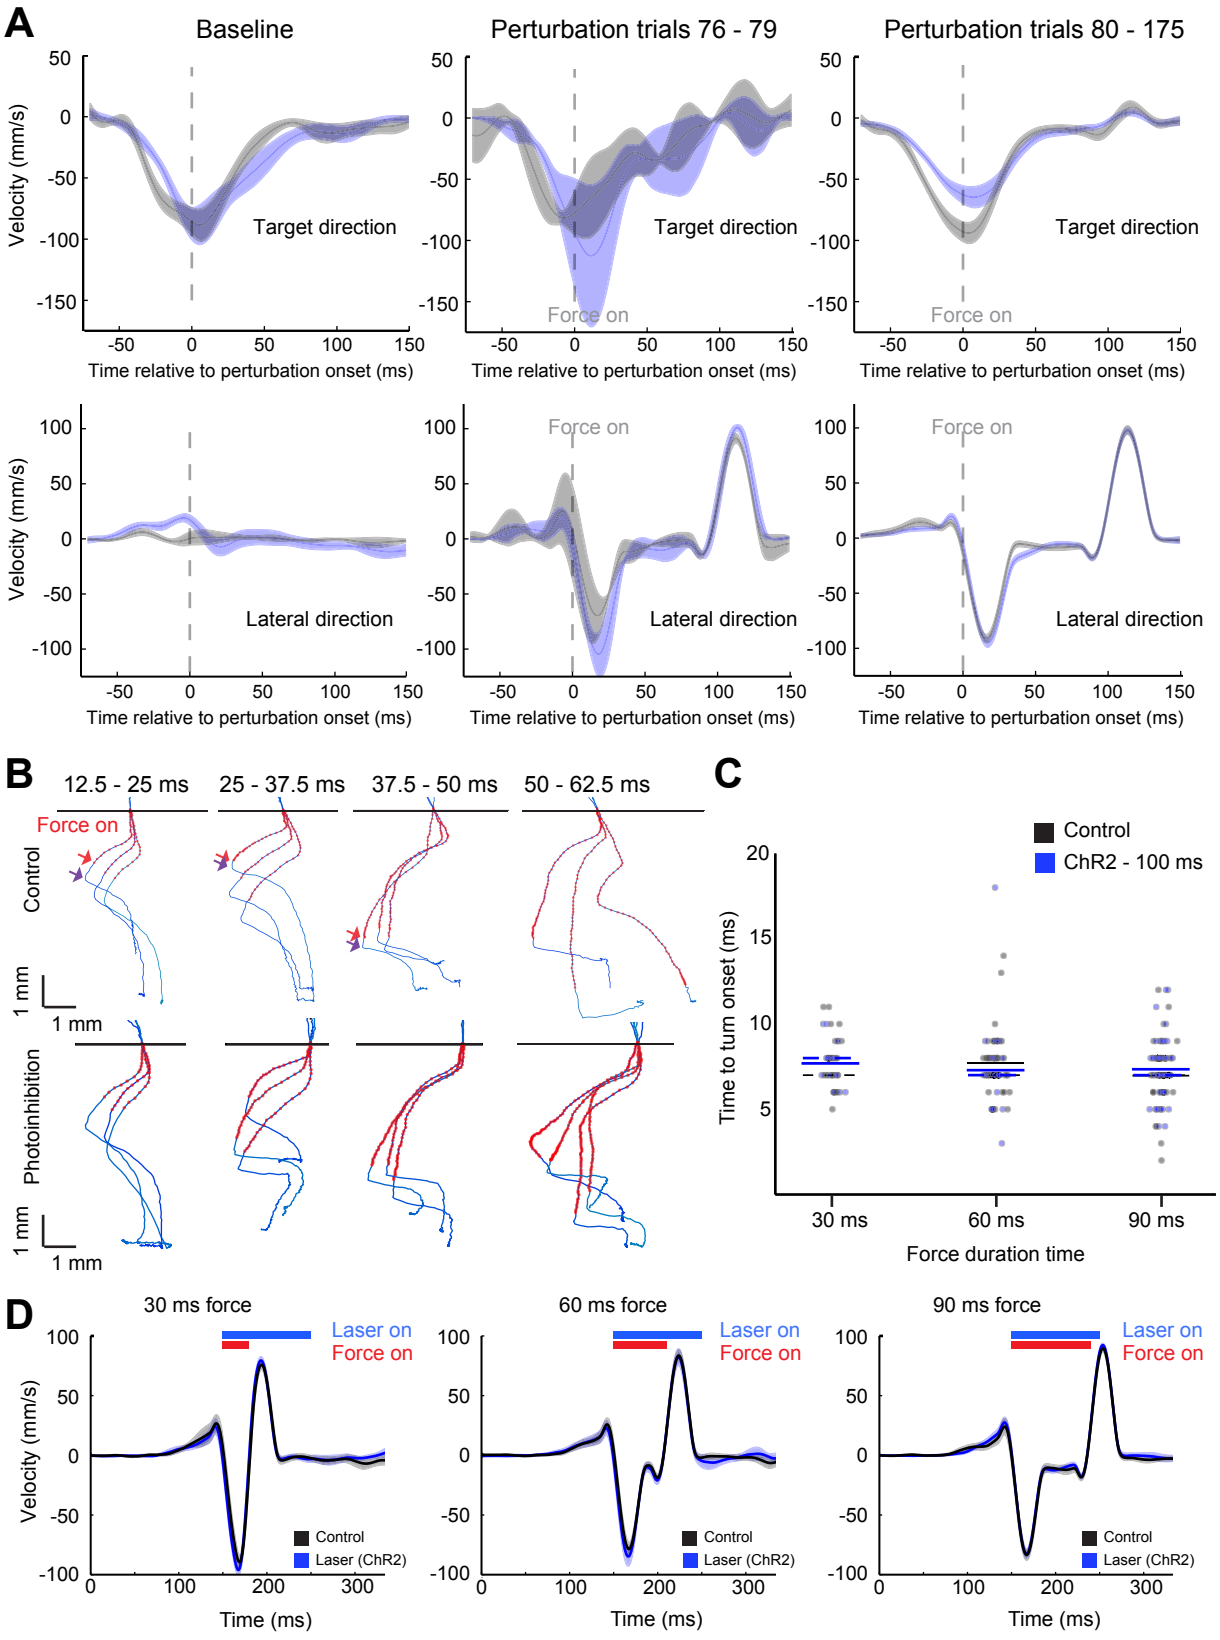

Supplemental Figure 6. Kinematics and motor control during photoinhibition.
Related to Figure 3.

(A) Velocity profiles during a control or photoinhibition sessions (from the same mouse). Notice there is an adaptation-related difference between control and photoinhibition that emerges after several perturbation trials.

(B) Example pulls during force fields of various durations without (top row) and with (bottom) photoinhibition for 100ms duration. Arrows denote the force off and turn point (i.e. first inflection point post force off). We quantified this as time to turn in C.

(C) 'Time to turn' (from force onset) during the motor control task with 100 ms with photoinhibition (blue) or without (black), paired with 30, 60, or 90 ms force fields, demonstrating no significant difference between conditions.

(D) Velocity profiles during force fields of 30, 60, or 90 ms durations, with and without photoinhibition (mean with CI shown) for the sessions with 10% randomly selected perturbation trials (to avoid adaptation).

Supplemental Figure 7

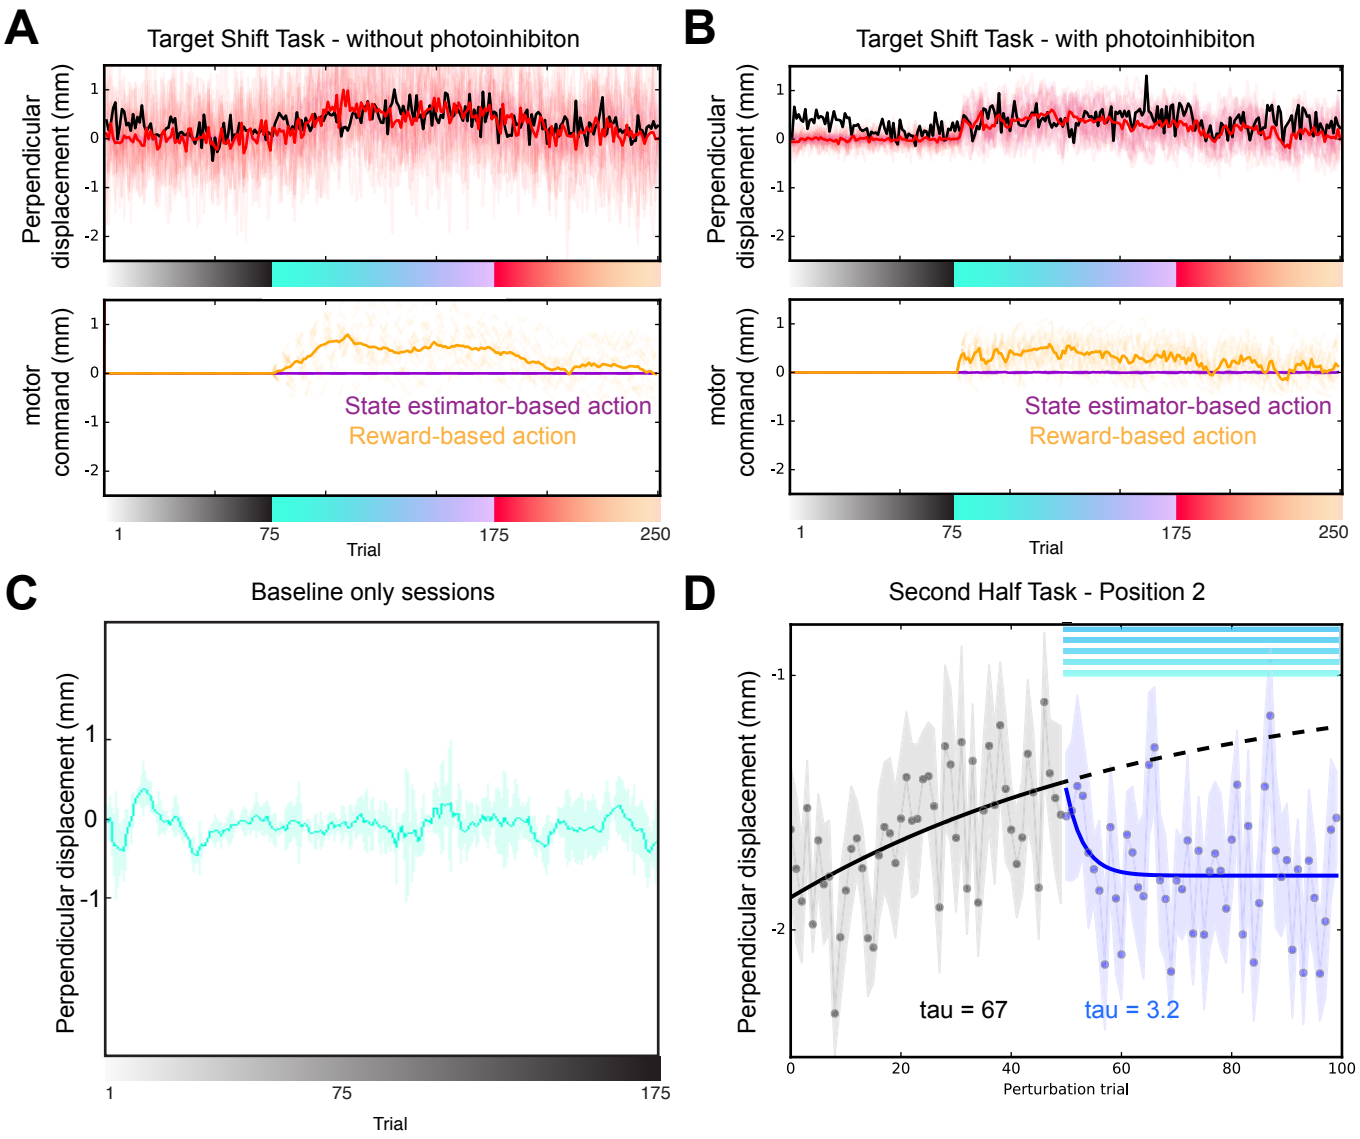

Supplemental Figure 7. The role of S1 in updating the memory of the perturbation.
Related to Figure 4.
(A) Best fit of hybrid model to control data of target-shift task (also see Figure 3G). The motor command arising from the actor-critic reinforcement-learning model drives the learning.
(B) Best fit of hybrid model to photoinhibited data of target-shift task (also see Figure 3G).
(C) Average of perpendicular deviation at position 2 during 'baseline only' sessions (one randomly selected day from each of the vGAT-ChR2 mice) showing no systematic drift without target shift.
(D) Perpendicular deviation at position 2 during the perturbation block of the "second half task" with fitted exponential increase and decrease to illustrate the time scales of adaptation and forgetting (without updating).

Supplemental Figure 8

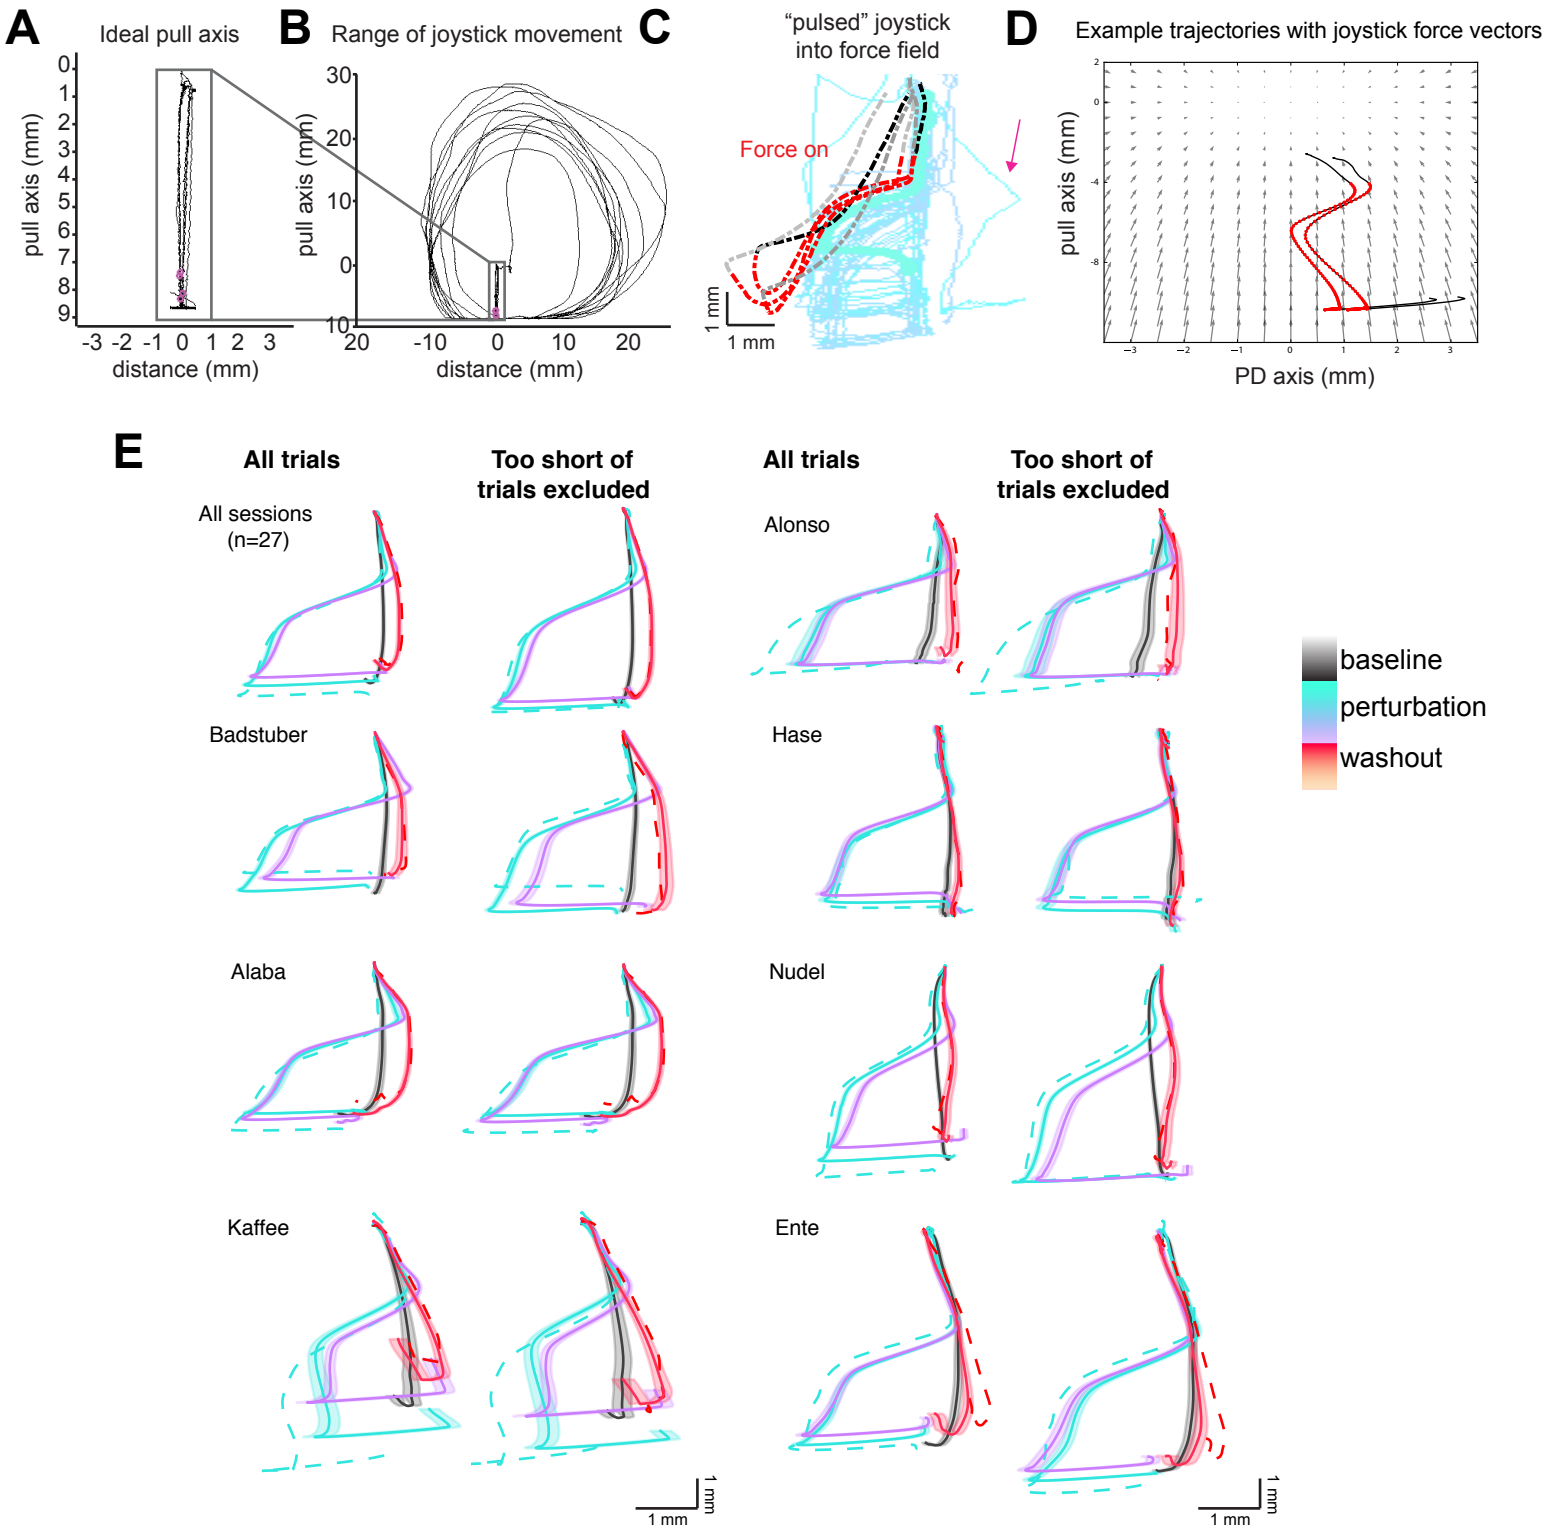

Supplemental Figure 8.
Related to STAR Methods.
(A) Ideal pull axis (joystick placed in a guide).
(B) The joystick had 2 dimensions of freedom, allowing for 360 degree movements.
(C) The return trajectories when the joystick (without a mouse) was pushed along the pull axis until it interacted with the magnet (with mouse-imposed trajectories overlaid). Red denotes force on time. Pink arrow highlights a pull from a mouse where the force offset does not return to midline.
(D) Example L-turns that are away from the magnet even on the positive perpendicular displacement side of the midline pull axis (half-plane), where the spring's lateral force component aligns with the magnet's force field direction.
(E) "All trials" are average trajectories without excluding any trials as shown in Fig. S2A (see STAR Methods). To the right of each plot are the average trajectory plots when pulls that do not reach the upper edge of the target box are excluded (this affects ~10% of the trials, see Methods). The dashed sea-green line shows the average first perturbation trial, and the dashed red line the average of the first five washout trials. The solid lines are average paw paths ± s.e.m. The solid green line shows the average over the first 25 trials during the perturbation epoch, the solid purple over the last 25 trials of the perturbation epoch, and the solid red line over the first 25 washout trials.

Supplemental Materials, Movie Legends

Movie 1. Related to Figure 1.
example baseline trial, slowed by a factor of 10.

Movie 2. Related to Figure 3.
example perturbation trial, slowed by a factor of 10.

Movie 3. Related to Figure 3.
example perturbation trials with photoinhibition to forelimb S1, slowed by a factor of 10.