Striatal ensembles specify and control granular forelimb actions

Ines Rodrigues-Vaz[1-4]†, Vivek R. Athalye[1,3,4]†, Darcy S. Peterka[1,4,5], and Rui M. Costa[1,3,4,5*]

[1] Zuckerman Mind Brain Behavior Institute, Departments of Neuroscience and Neurology; Columbia University; New York, NY, 10027; USA

[2] Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Lisbon 1400-038, Portugal

[3] Allen Institute, Allen Institute for Neural Dynamics, Seattle, WA, 98109; USA

[4] Aligning Science Across Parkinson's (ASAP) Collaborative Research Network, Chevy Chase, MD20815

[5] Kavli Institute for Brain Science, Columbia University; New York, NY 10027; USA

† These authors contributed equally to this work.

*Corresponding author. rui.costa@alleninstitute.org (RMC)

The ability of the brain to control specific fine actions is crucial for survival. Although striatal dysfunction underlies numerous movement disorders[1–9], its activity has classically been studied in the context of movement reinforcement and invigoration [5,10–13]. Recent studies have revealed that striatal activity encodes specific whole-body [14,15] and forelimb movements[16,17]. However, it is unclear if this activity actually controls specific movements. We developed a closed-loop system to identify and manipulate action-specific neuronal ensembles in real time using holographic optogenetics through a GRIN lens. We designed an isometric force task in which mice performed push or pull actions on an immobile joystick without overt kinematics—actions that engage the same forelimb muscles but with different activation patterns. Using 2-photon calcium imaging, we simultaneously recorded D1- and D2-medium spiny neurons (MSNs) in the dorsolateral striatum - classically viewed as promoting versus inhibiting movement[2,3,18] – and found that both populations equally encoded isometric action identity. Stimulation of action-specific D1- and D2-MSN ensembles increased ongoing force, but only when the stimulated ensemble was congruent with the ongoing action. These results reveal that specific ensembles of D1- and D2-MSNs control specific actions, as granular as distinct muscle activation patterns of the same forelimb. These findings provide a novel framework for understanding how striatal dysfunction can produce highly specific movement impairments in Huntington's disease[2,9] and dystonia[2,6].

The brain must generate signals that enable the execution of appropriate actions, many of which demand granular motor control. Writing, for example, requires precisely coordinated muscle co-contractions in the arm to apply specific forces to a pen. Whereas motor cortices, brainstem, and spinal circuits are classically linked[19] to the online specification and execution of movement kinematics and force, the striatum has been primarily framed as a center for invigoration[5,10–13] and reinforcement[12,13,20–26] rather than as a controller of ongoing action per se. This view has been supported by the fact that Parkinson's disease, where dopamine inputs to striatum are depleted, leads to deficits in action initiation[4], vigor[5], and reinforcement[7,8]. Alternatively, accumulating evidence suggests that the striatum specifies actions[14–17] rather than merely broadcasting a more generic "go/no-go"[27] or vigor signal. This is consistent with

observations that in disorders where striatal neurons are affected, like in Huntington's disease[9] or dystonia[2,6], the execution of particular specific movements is disrupted.

Recent work shows that the activity of specific striatal ensembles encodes particular whole-body [14,15] and forelimb movements[16,17]. However, it remains untested if this specific activity does indeed control the ongoing actions. To address this, we developed a real-time system to identify neuronal ensembles with a statistical model and manipulate them with cellular precision in closed-loop with behavior. Our all-optical system measured neuronal activity with two-photon calcium imaging and stimulated targeted neurons with holographic optogenetics[28–31] deep in the brain through a GRIN lens[30,32].

Using this system, we asked whether the striatum controls ongoing, granular actions. We designed a self-paced isometric force task, in which head-fixed mice executed push or pull forelimb actions on an immobile joystick, activating the same forelimb muscles but with different patterns. This task allows us to test orthogonal aspects to recent studies using cued movements with subtly different kinematics[17], and enables us to test whether the striatum can encode and control different self-paced forelimb actions with similar force and no overt kinematics, just different patterns of activation of the same muscles.

Striatal projection neurons are two distinct cell types, D1- and D2-MSNs, with different dopamine receptors and projection patterns, that have been canonically assigned opponent roles in movement promotion and suppression, respectively [2,3,18,33]. However, they both receive cortical inputs that carry specific movement plans[34–38], both are activated at movement initiation[39,40], and both seem necessary for execution[40,41]. Therefore, we first tested whether D1- and D2-MSNs would encode the identity of specific isometric actions differently or similarly. We found that both the activity of specific D1- and D2-MSN ensembles differentially encoded the identity of the different isometric actions.

Then we tested if activity of these specific ensembles could control particular ongoing movements. We modeled the dimensions of population activity that predict specific actions, identified small, action-specific ensembles within spatially intermingled D1- and D2-MSNs, and stimulated those ensembles precisely during self-paced actions with millisecond latency. We thus tested whether stimulating action-specific ensembles increases the force of the ongoing action, and whether this specificity holds for both D1- and D2-MSNs.

Together, this strategy allowed us to determine that the striatum specifies ongoing actions, with specific ensembles of both D1- and D2-MSNs controlling specific actions, as granular as distinct patterns of forelimb muscles without overt movement.

An isometric force task to study granular forelimb actions

To investigate whether the striatum encodes and controls granular actions of the same body part, we developed a task in which mice performed two forelimb actions (Figures 1A and S4A), consisting of a push or pull force on an immobile joystick in the absence of overt movement (i.e. isometric co-contractions, STAR Methods). Mice initiated actions in a self-paced manner, in the

absence of external cues. The task enabled analysis of push and pull actions with matched continuous parameters including force, duration, and licking (Figure S3A-H, STAR Methods).

Mice performed isometric actions in order to receive sucrose reinforcement. We defined two targets as push and pull that exceed thresholds of force and duration determined for each session (Figure 1B,C, Table 2,3, STAR Methods), and these actions triggered delayed sucrose reward (5 µl per reward). Mice selectively increased performance of whichever action was reinforced, confirming that they performed actions to obtain sucrose (Figures 1D, S2A-C, STAR Methods).

We characterized the simultaneous activity of four forelimb muscles - biceps, triceps, palmaris longus (PL), and the extensor digitorium communis (EDC) - during the performance of these isometric actions (Figure 1E). All four muscles were activated for both isometric push and pull (Figure 1F,G), but the pattern across muscles was action-specific (Fig. 1H). Thus, both isometric push and pull were executed through activation of the same forelimb muscles -- co-contraction of extensor and flexor pairs -- but with action-specific activation patterns.

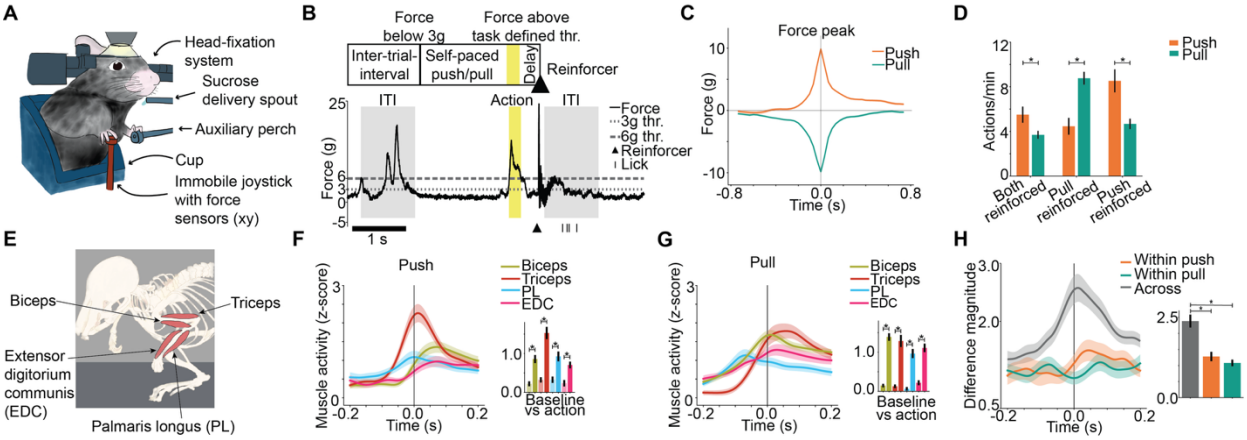

Figure 1: An isometric force task engages the same forelimb muscles with action-specific activation patterns.

(A) Task schematic.

(B) Task events for reinforcement of self-paced actions. Trace of force applied to the joystick on an example trial.

(C) Trial-averaged force trace of 6g cross actions time locked to force peak. By our convention, force in the push direction is positive and pull direction is negative. Data are mean+s.e.m. across 109 sessions from 8 mice.

(D) Rate of 6g cross actions. Mice performed both actions and increased the action which led to sucrose reinforcement. Data are mean+s.e.m. from 8 mice, averaging over sessions in each block (STAR Methods). Mice performed push more than pull when both actions were reinforced (Table S1, ref. 1.1). Mice performed pull more than push when pull was reinforced (Table S1, ref. 1.2). Mice performed push more than pull when push was reinforced (Table S1, ref. 1.3).

(E) Schematic of forelimb muscles simultaneously recorded with EMG.

(F) Trial-averaged z-scored EMG of 6g cross push actions locked to 6g cross, force-matched with pull actions (STAR Methods). Each muscle significantly activated at 6g cross relative to 1 second before (Table S1, ref. 1.4-1.7).

(G) Same as (F) but for pull action, force-matched with push actions. Each muscle significantly activated at 6g cross relative to 1 second before (Table S1, ref. 1.8-1.11).
(H) Vector magnitude of EMG difference at 6g cross. The difference across actions is greater than within each action (Table S1, ref. 1.12,1.13).
(F-H) Data are mean+s.e.m. over n=41 sessions from 4 mice.

Striatal ensembles encode the identity of isometric actions

To test whether the striatum encodes the identity of isometric forelimb actions, we measured the simultaneous activity of D1- and D2-MSNs in the forelimb region of dorsolateral striatum[42,43] (DLS) (Figure 2A,B). To identify D1- and D2-MSNs within the same subject[44], we used transgenic mice that expressed the red label tdTomato either in D1-MSNs (T6 and Ai9xEY217) or D2-MSNs (Ai9xAdora2cre) (Table 1), and we developed an automated method to classify red versus non-red neurons (Figure S1). tdTomato-positive neurons were thus identified as D1-MSNs (or D2-MSNs, depending on the mouse line), and tdTomato-negative neurons were classified as the other subtype; together, D1- and D2-MSNs account for approximately 95% of all striatal neurons, so non-red neurons are overwhelmingly the other MSN subtype[44–46]. We virally expressed GCaMP6f in both MSN populations in DLS using the *Camk2a* promoter[44] and used 2-photon microscopy through a 1mm diameter GRIN lens to image neuronal calcium activity as a proxy for action potentials. We were able to image up to ~200 neurons simultaneously, and the fluorescence signal in both color channels was stable across months of animal training (Figure S3K). We identified ongoing isometric actions that exceeded 3g of force for ~66 ms duration (two imaging frames) and that were temporally isolated (i.e. had no preceding action in the 0.5 seconds before onset). The activity of D1- and D2-MSNs was modulated for both actions (Figure 2B,C), with slightly greater activity in D1-MSNs (Figure 2C).

Many individual neurons modulated their activity only during a specific action (Figure 2D). We therefore asked whether the spatiotemporal activity of MSNs encoded the identity of ongoing actions. We extracted the vector of MSN population activity averaged in 150 ms bins (5 imaging frames) centered at time lags relative to force peak for each trial, and we decoded single-trial population activity into action identity (push or pull) using a linear support vector machine (training in 90% of data, testing in 10% held out data). For time lags before (after) force peak, we analyzed trials with no other action in the 0.5 seconds before (after) force peak. Population activity predicted action identity well above chance with peak accuracy at the time of peak force (Figure 2E,F). In addition, population activity predicted action identity in advance of force peak (Figure 2F). Interestingly, both D1- and D2- MSNs encoded action identity with equal accuracy and timing (Figure 2G).

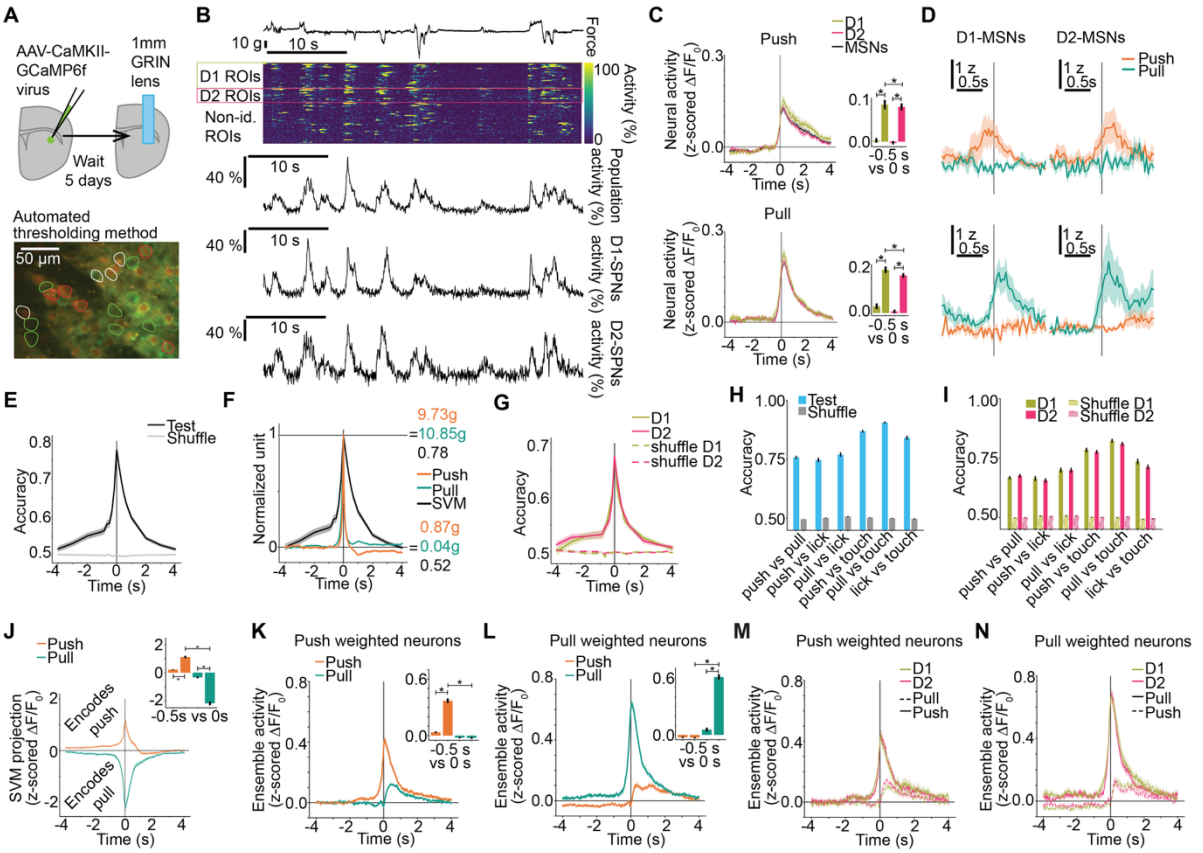

Figure 2: D1- and D2-MSNs both encode identity of isometric forelimb actions.

(A) Schematic of experimental procedures. Top: GCaMP6f virus injection and lens implant in transgenic mice to image D1- and D2-MSNs simultaneously. Bottom: Schematic of automated method for identifying D1- and D2-MSNs based on transgenic expression of tdTomato (STAR Methods).

(B) Example trace of force, individual neuronal activity for all ROIs in this field of view, population activity averaged across neurons, averaged activity for D1-MSNs (n=21 neurons) and averaged activity for D2-MSNs (n=21 neurons). Neural activity is $\Delta F/F_0$.

(C) Population activity locked to peak force of push and pull actions, averaged across D1-, D2-, and all MSNs. Data are mean+s.e.m. across 86 sessions from 8 mice. Activity increased at push force peak ("0 s") relative to "-0.5 s" for D1-MSNs (Table S1, ref. 2.1) and D2-MSNs (Table S1, ref. 2.2). Activity increased at pull force peak ("0 s") relative to "-0.5 s" for D1-MSNs (Table S1, ref. 2.3) and D2-MSNs (Table S1, ref. 2.4). D1-MSNs were more active than D2-MSNs at force peak (Table S1, ref. 2.5,2.6). Bar plots were $\Delta F/F_0$ averaged across 3 event-centered imaging frames.

(D) Example traces of individual D1- or D2-MSN activity locked to force peak for push and pull actions.

(E) Decoding all MSN activity to predict action identity on single trials with a model fit at each time lag relative to force peak. Accuracy at force peak was greater than a model ("shuffle") trained with shuffled action identity (Table S1, ref. 2.7). Accuracy was greater than shuffle when averaged over time points in the 4 seconds before force peak (Table S1, ref. 2.8) and in the 0.5

seconds before force peak during which there was no other action (Table S1, ref. 2.9). Data are mean+s.e.m. across n=87 sessions from 8 mice.

(F) Decoding all MSN activity overlayed on the average force of actions (normalized to 0 at -4s and 1 at peak of 0s).

(G) Decoding D1- or D2- MSN activity separately. Accuracy at force peak was greater than shuffle for D1-MSNs (Table S1, ref. 2.10) and D2-MSNs (Table S1, ref. 2.11), with no significant difference between D1- and D2-MSNs (Table S1, ref. 2.12). Accuracy was greater than shuffle when averaged over time points in the 4 seconds before force peak for D1-MSNs (Table S1, ref. 2.13) and D2-MSNs (Table S1, ref. 2.14) and in the 0.5 seconds before force peak for D1-MSNs (Table S1, ref. 2.15) and D2-MSNs (Table S1, ref. 2.16). Data are mean+s.e.m. across n=87 sessions from 8 mice.

(H-L) Data are mean+s.e.m. across 86 sessions from 8 mice.

(H) Decoding MSN activity at force peak for push and pull, predicting action identity between pairs of actions including push, pull, touch, and licking. Accuracy was greater than shuffle for all action pairs (Table S1, ref. 2.17-2.22).

(I) Same as (G), decoding D1- and D2-MSNs activity. Decoding push versus pull was not significantly different between D1- and D2-MSNs (Table S1, ref. 2.23). Accuracy was greater than shuffle for all action pairs for D1-MSNs (Table S1, ref. 2.24-2.29) and D2-MSNs (Table S1, ref. 2.30-2.35).

(J) Projection of population activity (z-scored $\Delta F/F_0$) on the SVM dimension locked to force peak. Inset: average SVM projection at 500 ms before force peak (“-0.5 s”) and at force peak (“0 s”). For the push action, the SVM projection at force peak (“0 s”) increased relative to “-0.5 s” (Table S1, ref. 2.36). For the pull action, the SVM projection at force peak (“0 s”) decreased relative to “-0.5 s” (Table S1, ref. 2.37). The SVM projection at force peak was greater for push than for pull (Table S1, ref. 2.38).

(K) Ensemble activity (z-scored $\Delta F/F_0$) averaged across push weighted neurons, locked to force peak for push and pull actions. Inset: average activity at -0.5 s before force peak and at force peak for each action. The activity of push weighted neurons increased at push force peak (“0 s”) relative to “-0.5 s” (Table S1, ref. 2.39). Activity at push force peak was greater than at pull force peak (Table S1, ref. 2.40).

(L) Ensemble activity (z-scored $\Delta F/F_0$) averaged across pull weighted neurons, locked to force peak for push and pull actions. Inset: average activity at -0.5 s before force peak and at force peak for each action. The activity of pull weighted neurons increased at pull force peak (“0 s”) relative to “-0.5 s” (Table S1, ref. 2.41). Activity at pull force peak was greater than at push force peak (Table S1, ref. 2.42).

(M) Same as (K) for D1- and D2-MSNs. Activity of push weighted D1- and D2-MSNs around push force peak (-0.5 s to 0.5 s) is not significantly different (Table S1, ref. 2.43). Data are mean+s.e.m. across 63 sessions from 8 mice.

(N) Same as (L) for D1- and D2-MSNs. Activity of pull weighted D1- and D2-MSNs around pull force peak (-0.5 s to 0.5 s) is not significantly different (Table S1, ref. 2.44). Data are mean+s.e.m. across 73 sessions from 8 mice.

Population activity at force peak predicted action identity even for trials when other actions were allowed in the 0.5s temporal window before and after (Figure 2H). D1- and D2-MSNs encoded these actions with equal accuracy (Figure 2I). Given that the striatum encodes continuous kinematics[16,17], we verified action identity was encoded for trials with matched continuous parameters including force, duration, and licking (Figure S3I,J).

To understand the structure of this encoding, we examined how action identity was embedded within the high dimensional space of population activity where each neuron's instantaneous activity is one axis. The SVM identifies a single dimension that best predicts the animal's action at the time of the force peak. It follows that the projection of population activity on that SVM dimension (a weighted sum over neurons) showed differential modulation at force peak for push and pull actions (Figure 2J). The SVM's weights partition neurons into action-specific groups whose activation biases the decoder to predict push or pull. The push-weighted neurons increased activity at push force peak, but not for pull force peak (Figure 2K). Similarly, the pull weighted neurons increased activity at pull force peak but not for push force peak (Figure 2L). D1- and D2-MSNs showed no difference in SVM weights and action-specific modulation of push- and pull-weighted neurons (Figure 2M,N).

The striatum has a classical role in the reinforcement of actions, with D1- and D2-MSNs thought to have opponent[13,25] or different[26] roles. Therefore, we examined if encoding of action identity would change whether the action is reinforced or not. We found that both types of MSNs encode ongoing action identity regardless of whether both or one of the particular actions were reinforced (Figure S2D-G). Further, we tested whether striatal plasticity is necessary for action encoding. Notably, we found that striatal activity encoded the identity of specific actions even in mice with impaired striatal plasticity[47] who could not learn which actions were reinforced[48–50] (Figure S2H-J). These results reveal that although the striatum has a clear role in reinforcement and learning, it can encode ongoing actions independently of these roles.

Together, these results demonstrate that striatal population activity robustly encodes the identity of ongoing forelimb actions, with equal contributions from D1- and D2-MSNs.

Specific stimulation of striatal neurons encoding specific actions

We next addressed our central question: whether these ensembles causally control the specific actions they encode. One critical technical challenge was that we needed to target perturbations to specific and spatially-intermingled D1- and D2- MSNs that encode action identity. For this, we developed an approach to stimulate individual MSNs with 2-photon optogenetics through a GRIN lens[30] in mice co-expressing GCaMP6f and red-shifted activating opsin ChRmine[29] (Figure 3A-C). We used a spatial light modulator to holographically target a group of user-defined MSNs for simultaneous stimulation[28]. Another technical challenge was that action duration was very short, and we needed to trigger stimulation upon self-initiated actions with low latency, as force onsets rapidly over milliseconds (Figure 1C). For this, we developed an experimental platform that allowed us to flexibly select the trigger action and stimulation pattern and begin the stimulation with ~2ms total latency (Figures 3D and S4B).

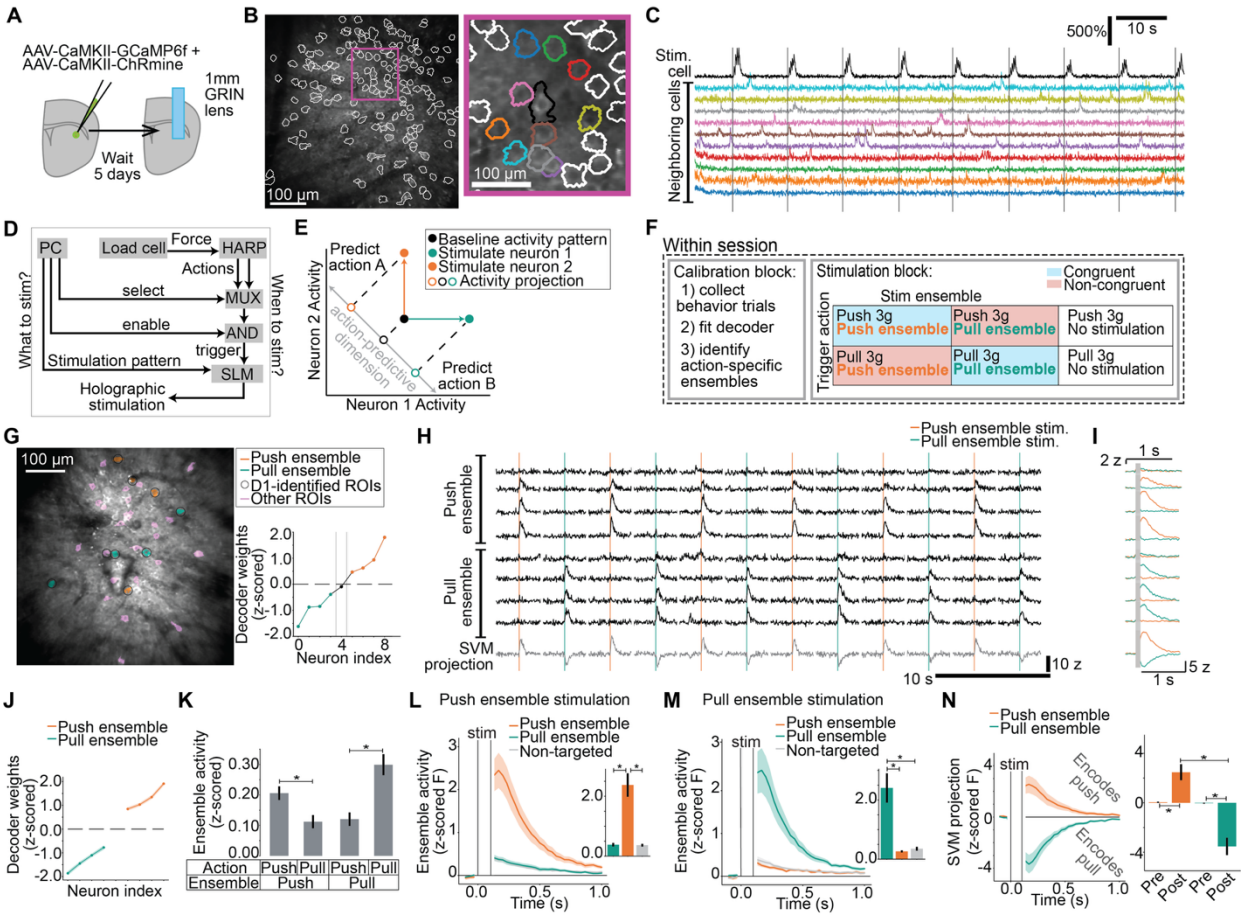

Figure 3: Closed-loop modeling and manipulation of neuronal ensembles that encode action identity.

(A) Schematic of GCaMP6f and ChRmine viral injection and lens implant in transgenic mice to image and photostimulate D1- and D2-MSNs.

(B) Example field of view with a single neuron targeted for stimulation. Inset: zoom in onto targeted neuron in black.

(C) Example activity traces of neurons in (B) inset. Targeted neuron (in black) and neighboring non-targeted neurons' activity during 2-photon optogenetic stimulation.

(D) Schematic of closed-loop system to select the pattern of holographic stimulation and the action that triggers stimulation (Figure S4B for more details).

(E) Cartoon depicting how activating specific neurons can manipulate a dimension of population activity that encodes and predicts specific actions.

(F) Schematic for conducting the modeling and manipulation experiment.

(G) Left: example field of view with regions of interest labeled based on their decoder weights and identification as D1-MSNs. Right: Decoder weights for identified D1-MSNs, colored by assignment into action-specific ensembles.

(H) Z-scored fluorescence of each ROI member of each action-specific ensemble, and the projection of population activity on the SVM decoder dimension. Four consecutive stimulation trials concatenated from the beginning, middle, and end of the example session.

(I) Trial-averaged traces of targeted neurons and projection of population activity on the action-predictive SVM dimension.

(J-N) Data are mean+s.e.m. across n=21 sessions and 10 mice.
(J) Average decoder weights of action-specific ensembles. The top four neurons with the largest weights are shown.
(K) Average activity (z-scored fluorescence) of action-specific ensembles in 100 ms before force peak during trials in the calibration block. Push ensemble activity was larger for push than pull (Table S1, ref. 3.1). Pull ensemble activity was larger for pull than push (Table S1, ref. 3.2).
(L-M) Neural activity averaged across neurons (z-scored fluorescence) within each ensemble, locked to stimulation of each ensemble. Inset: average ensemble activity in 100 ms after stimulation.
(L) After push ensemble stimulation, activity of the push ensemble was greater than the pull ensemble (Table S1, ref. 3.3) and non-targeted neurons (Table S1, ref. 3.4). There was no significant difference between the pull ensemble and non-targeted neurons (Table S1, ref. 3.5).
(M) After pull ensemble stimulation, activity of the pull ensemble was greater than the push ensemble (Table S1, ref. 3.6) and non-targeted neurons (Table S1, ref. 3.7). There was no significant difference between the push ensemble and non-targeted neurons (Table S1, ref. 3.8).
(N) Projection of population activity on the SVM dimension locked to ensemble stimulation. Inset: Average SVM projection in the 100 ms before stimulation ("pre") and after stimulation ("post"). Stimulating the push ensemble increased the "post" SVM projection relative to "pre" (Table S1, ref. 3.9). Stimulating the pull ensemble decreased the "post" SVM projection relative to "pre" (Table S1, ref. 3.10). The "post" SVM projection was greater after push ensemble stimulation than pull ensemble stimulation (Table S1, ref. 3.11).

Within a single session, we identified action-specific MSN ensembles and stimulated them in closed-loop during specific, self-paced actions (Figure 3E,F). In a calibration block, we collected behavior trials and fit a decoder to predict action identity from MSN activity. As before (Figure 2K-N), we identified action-specific ensembles (Figure 3G) using the weights of the force peak SVM. We defined the action-specific ensembles as equally sized groups of 4-11 neurons with the largest decoder weights (mean ensemble size was 6.7 neurons over 21 sessions). Then in the stimulation block, specific self-paced actions triggered stimulation of action-specific ensembles (Figure 3H,I). We interleaved 6 conditions: two actions that triggered stimulation (3g push, 3g pull) and three stimulation patterns for each action (push ensemble, pull ensemble, no stimulation) (Figure 3F).

The identified ensembles had similar magnitude decoder weights (Figure 3J) and modulated specifically for their corresponding actions (Figure 3K). Stimulation specifically activated the targeted ensemble (Figure 3L,M). We verified that this stimulation of a small number of action-specific neurons bi-directionally drove the action-specific SVM dimension of population activity (Figure 3N). Thus, selective stimulation of small, action-specific striatal ensembles was sufficient to bias population activity toward encoding the corresponding action (Figure 3E).

Specific striatal ensembles control specific actions

We used this approach to test whether action-specific striatal ensembles control specific actions or simply invigorate any ongoing action. If ensembles control specific actions, stimulating a pull

ensemble should selectively increase pull force, with no effect or an opposing effect during push. If ensembles merely invigorate movement, any ensemble stimulation should increase force regardless of action identity. We refer to the pairing of an ongoing action with ensemble stimulation specific to that action as "congruent" (e.g. pull ensemble stimulation during pull action), and the pairing of an ongoing action with ensemble stimulation specific to the other action as "non-congruent" (e.g. push ensemble stimulation during pull action) (Figure 3F). This also allowed us to test whether D1-MSNs promote and D2-MSNs suppress action, or both promote action.

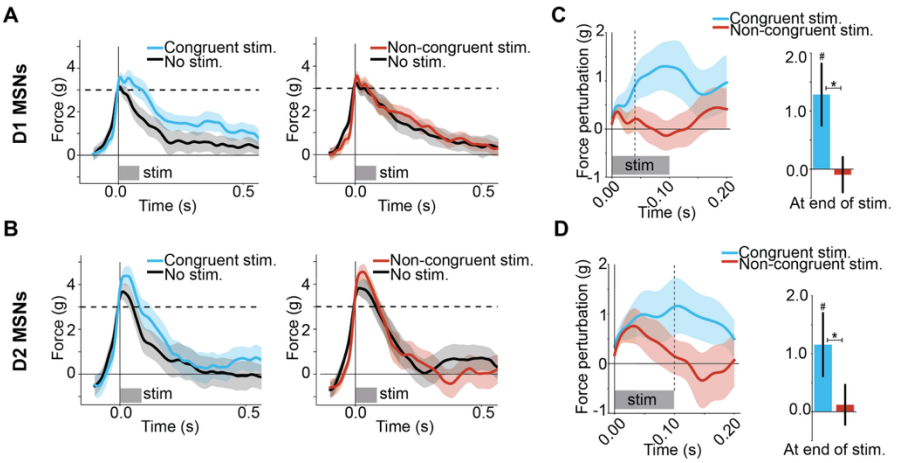

Figure 4: Action-specific ensembles of D1- and D2- MSNs control congruent actions.

(A-B) Trial-averaged force traces locked to 3g cross that triggers holographic stimulation. In this plotting convention, force is positive for the action executed at time 0 (push or pull) and negative for the other action.

(A) Stimulation of D1-MSN ensembles. Data are mean+s.e.m. across n=12 sessions and 9 mice.

(B) Stimulation of D2-MSN ensembles. Data are mean+s.e.m. across n=9 sessions and 5 mice.

(C-D) Force perturbation, i.e. the trial-averaged force trace on stimulation trials minus non-stimulation trials. Right panel is the force perturbation at the end of stimulation, averaged from 100ms to 110ms.

(C) Force perturbation from stimulation of D1-MSN ensembles. The perturbation was greater than 0 with congruent stimulation (Table S1, ref. 4.12) and was greater than non-congruent stimulation (Table S1, ref. 4.11) at the end of stimulation. The perturbation was not significantly different than 0 with non-congruent stimulation (Table S1, ref. 4.13). The perturbation first became significantly different between congruent and non-congruent stimulation at 40ms after stimulation onset (Table S1, ref. 4.5).

(D) Force perturbation from stimulation of D2-MSN ensembles. The perturbation was greater than 0 with congruent stimulation (Table S1, ref. 4.25) and was greater than non-congruent stimulation (Table S1, ref. 4.24) at the end of stimulation. The perturbation was not significantly different than 0 with non-congruent stimulation (Table S1, ref. 4.26). The perturbation first became significantly different between congruent and non-congruent stimulation at 100ms after stimulation onset (Table S1, ref. 4.24).

Stimulating specific D1-MSN ensembles congruent with the ongoing action increased ongoing force (Figure 4A left). In contrast, no change in ongoing force was observed after stimulating the non-congruent ensemble (Figure 4A right). Notably, stimulating specific D2-MSN ensembles congruent with the ongoing action also increased ongoing force (Figure 4B left), whereas stimulation of the non-congruent ensemble did not (Figure 4B right). In other words, stimulation of the same ensemble during different ongoing actions had different effects (e.g. stimulating the pull ensemble increased force during pull, but not during push).

We quantified the force perturbation as the difference between force traces in stimulation trials and non-stimulation trials. Stimulation of action-specific ensembles increased congruent but not non-congruent force for both D1- and D2-MSNs (Figure 4C,D). Congruent stimulation of D1-MSN ensembles increased perturbation relative to non-congruent stimulation as early as 40 ms after start of stimulation, whereas D2-MSNs ensemble stimulation increased perturbation later, around 100 ms after start of stimulation. Thus, the dynamics of the stimulation of D1- vs D2-MSNs were slightly different, suggesting that D1- and D2 populations do not have a redundant role. Together, these results demonstrate that specific ensembles of both D1- and D2-MSNs causally control specific ongoing actions, increasing force selectively for the action they encode.

Discussion

Here we show that striatal ensembles that encode specific ongoing actions can indeed control them. Further, this control operates at the granularity of distinct muscle activation patterns of the same forelimb. We identified and stimulated small, action-specific ensembles of D1- and D2-MSNs during self-paced isometric actions and found that stimulation selectively increased force for the congruent action. Despite the striatum's indirect access to spinal circuits, these results establish it as an online controller of specific actions and forces — a role beyond its classical contributions to invigoration and reinforcement. To achieve this, we developed a model-guided, all-optical[28–30,32] closed-loop system to identify and manipulate action-specific ensembles in real time, even deep in the brain — addressing a critical challenge in motor neuroscience: linking the encoding properties of specific neural ensembles to their role in the execution of movement.

These results offer new views on the canonical model that D1- and D2-MSNs respectively promote and suppress movement[3,9,18,33] via the striatonigral and striatopallidal pathways. While bulk manipulations have supported opposing effects[13,18], subsequent manipulations[40,41] suggested that their coordinated activity are necessary for movement execution. Here we identified and stimulated action-specific ensembles of D1- and D2-MSNs and found that both increase the force of congruent actions. Thus, D2-MSNs can not only promote movement, but also control specific actions in real time, online. This finding suggests that D2-MSN projections to GPe can have effects that are independent of canonical inhibition of GPi/SNr, possibly through different anatomical circuits linking GPe directly to thalamus, cortex and brainstem[51]. It is also important to note that we focused our recordings and manipulations on forelimb region of dorsolateral striatum, and understanding the function of D1- and D2-MSNs would likely depend on the striatal subregion[16,33,52], each receiving distinct topographic inputs from cortex[42,43] and thalamus, and also projecting to specific outputs. Finally, the faster onset of perturbations

evoked by D1-MSN congruent stimulation suggests different rather than opposing roles for D1- and D2-MSNs in the control of movement.

This study positions the striatum as a center for controlling granular ongoing actions, despite its indirect connections to the spinal cord. Although basal ganglia outputs tonically suppress downstream motor centers[27,53–55], recent work reveals highly specific output dynamics during forelimb behavior[56]. Our work provides a foundation for future work on a fundamental problem: to understand how action-specific ensembles in the center of the brain can have such remarkably specific control over downstream motor centers. Conceptually, we hypothesize that the striatum guides motor centers into specific subspaces of neural population dynamics which implement action plans and execution[57], consistent with evidence that striatal activity[58] and corticostriatal plasticity[48] are required for the brain to learn and control low-dimensional cortical dynamics[59–61] that directly operate brain-machine interfaces. More generally, our experimental approach can be applied to motor circuits across the brain to test not only what neurons encode but also what their activity causally produces[62] — both in terms of downstream population dynamics and movement.

The discovery that striatal ensembles control granular actions offers a framework for understanding movement disorders with striatal dysfunction characterized by impaired, highly specific motor patterns—such as task-specific dystonias[63] (e.g. writer's cramp, musician's dystonia) and the choreic disruptions of selective movements observed in Huntington's disease[2,9].

Resource Availability

Lead contact

Further information and requests should be directed to the lead contact, Rui M. Costa.

Materials availability

No unique reagents were generated in this study.

Acknowledgements

We thank A. Vaz, M. Correia and G. Martins for mouse colony management. We thank the members of the Scientific Hardware Platform of the Champalimaud Institute for help with hardware and software. We thank H. Rodrigues for help with experimental setups. We thank T. Tabachnik and R. Hormigo for technical advice and hardware support. We thank L. Hammond and the Zukerman Institute's Cellular Imaging Platform for guidance in imaging and analysis. We thank A.M. Vicente for blinding the experimenter and discussions about mouse training. We thank Bruker and M.J. Fox for help developing the Bruker Software API. We thank K. Deisseroth, Y.S. Kim, and C. Ramakrishnan for providing the ChRmine virus for holographic stimulation. This work was supported by Fundacao para Ciencia e Tecnologia (FCT PD/BD/105950/2014) to I.R.-V. This work was supported by BRAIN Initiative National Institute of Mental Health postdoctoral fellowship (1F32MH118714-01) and NIH NINDS Pathway to Independence Award (1K99NS128250-01) to V.R.A. This work was supported by NS (NS104649), the Aligning Science Across Parkinson's (ASAP-020551) through the Michael J.

Fox Foundation for Parkinson's Research (MJFF) and Simons-Emory International Consortium (Simons 717104) to R.M.C, and by the National Institutes of Health (NIH, U19: 5U19NS104649) to D.S.P and R.M.C.

Author contributions

I.R.-V., V.R.A., D.S.P. and R.M.C. designed the study, interpreted results and wrote the manuscript. I.R.-V. performed experiments. I.R.-V. and V.R.A. analyzed data and developed the model-guided system for closed-loop stimulation.

Declaration of interests

The authors declare no competing interests.

STAR Methods

Experimental model and study participant details

Animals

All experiments and procedures were performed according to National Institutes of Health (NIH) guidelines and approved by the Institutional Animal Care and Use Committee (IACUC) of Columbia University. Adult mice, aged 2–6 months, were used for these experiments. For labeling D1-MSNs and D2-MSNs we used the following transgenic lines (Table 1): Tg(Drd1-cre)EY217Gsat (RRID:MGI:4366805), Tg(Drd1-cre)EY217Gsat crossed with Ai9(RCL-tdT) (RRID:IMSR_JAX:007905), Tg(Adora2a-cre)KG139Gsat (RRID:MMRRC_036158-UCD), Tg(Adora2a-cre)KG139Gsat crossed with Ai9(RCL-tdT) and Drd1a-tdTomato line 6 (RRID:IMSR_JAX:007905). For experiments testing the necessity of plasticity in the striatum we used the following lines (Table 1): RGS9-cre (RRID:IMSR_JAX:020550) crossed with NMDAR1-loxP (RRID:IMSR_JAX:036352) [47]. For comparison, we used the negative Cre littermates as controls. Mice used for experiments were individually housed and kept under a 12-h light-dark cycle.

Line name	Extended name	Source	Available at	Code	Ref.
Tg(Drd1-cre)EY217Gsat	STOCK Tg(Drd1-cre)EY217Gsat/Mmucd	Gerfen lab*	MMRRC	030778	Gong et al, 2007
Ai9(RCL-tdT)	B6;129S6-Gt(ROSA)26Sortm9(CAG-tdTomato)Hze/J	Jackson Laboratories	Jackson Laboratories	007905	Madisen et al, 2010
Drd1a-tdTomato line 6	B6.Cg-Tg(Drd1a-tdTomato)6Calak/J	Jackson Laboratories	Jackson Laboratories	016204	Ade et al, 2011
Tg(Adora2a-cre)KG139Gsat	B6.FVB(Cg)-Tg(Adora2a-cre)KG139Gsat/Mmucd	MMRRC*	MMRRC	036158	Gong et al, 2007
RGS9-cre	B6;129S-Rgs9$^{tm1.1(cre)Yql}$/J	NIH from Yuqing Li*	Jackson Laboratories	020550	Dang et al, 2006
NMDAR1-loxP	B6.129S2(Cg)-*Grin1*^{tm1Yql}/NkzaJ	NIH from Yuqing Li*	Jackson Laboratories	036352	Dang et al, 2006

* These lines were kept in house and backcrossed for more than 10 generations.

Table 1 – Mouse lines used in experiments.

Randomization and blinding

The experimenter was blinded to the experimental groups for plasticity experiments.

All animals were run on each day in a fixed random order.

Holographic manipulations were performed according to a pre-defined protocol and controlled by Matlab in closed-loop with the mouse's behavior - not manually by the experimenter. Sessions for stimulating D1- and D2-MSNs were interleaved as much as possible.

Method Details

Stereotaxic Surgery

Surgical procedures for imaging of neurons in the dorsolateral striatum (DLS) were performed in two steps consisting of viral injection at least 5 days prior to lens implant, following protocols described in protocols.io ("Viral injection" and "Lens implant").

Before each surgery, analgesia was administered to ensure 72 hours of pain relief. The mouse was placed in a stereotaxic frame under isoflurane anesthesia on a heated pad.

For viral injection, capillary tubes (Drummond 3.5" #3-000-203-G/X) were pulled using a Pipette puller (Sutter P-2000; Heat 360, FIL 5, Vel 25, Del 10, Pul 50) and backfilled with mineral oil. A craniotomy was made using a dental drill, and the capillary tube was lowered to the target coordinates. For imaging, 300 nl of AAV5.CaMKII.GCaMP6f (titer: 1×10^{12}–1×10^{13} GC/ml) was injected into the left DLS (AP: +0.75 mm, ML: 2.5 mm, DV: 2.4 mm) at a rate of ~5 nl/pulse. For holographic optogenetics experiments, an additional 300 nl of AAV8.CaMKII.ChRmine (titer: ~1×10^{11} GC/ml) was co-injected with AAV5.CaMKII.GCaMP6f at a 1:1 ratio into the left DLS at the same coordinates. Following injection, the capillary tube was left in place for at least 5 minutes before slow withdrawal, and the skin was sealed with sutures.

At least 5 days after viral injection, a 1 mm GRIN lens implant surgery was performed. Following craniotomy and cortex removal – using vacuum aspiration with 1 mm blunt needle – the GRIN lens was implanted at the same coordinates as the viral injection at depth of DV: 2.0 mm. A head bar was secured and a protective cap was built around the lens, which was optionally covered with silicone sealant (Kwick-Sil or Kwick-Cast) for protection.

For experiments requiring electromyographic (EMG) recordings of forelimb muscle activity, we fabricated EMG electrodes in-house as previously described[64] and in protocols.io ("EMG fabrication"). Briefly, two pieces of insulated braided stainless-steel wire were cut, knotted together, and secured to a working surface. Under a microscope, 0.5 mm portions of insulation were stripped from each strand just below the knot, leaving exposed contact sites separated by 0.5 mm. The wire ends were inserted into hypodermic needles and crimped in place to facilitate muscle insertion, and the opposing ends were soldered into corresponding sites on a miniature Omnetics connector.

EMG electrodes were then implanted into forelimb muscles (biceps, triceps, palmaris longus, and extensor digitorum communis) following lens implant[36], as described in protocols.io ("EMG implant"). Prior to surgery, mice were habituated to an Elizabethan collar and gel food for 4 days to minimize post-surgical disruption. On the day of surgery, analgesia was administered as described above. Electrodes were inserted into each target muscle using the hypodermic needle as a guide and secured in place by knotting the wire at the exit point to fix the recording sites within the muscle. Each incision was closed with sutures and covered with antibiotic ointment. The EMG connector was attached to the back of the skull or existing head cap using dental cement (Metabond), and any neck incision was similarly closed with sutures and antibiotic ointment. After surgery, the Elizabethan collar was placed on the mouse for at least 5 days. Food gel and soften pellets were provided for at least 5 days post-surgery.

Histology

Histology was performed as described in protocols.io (Histology) to verify lens coordinates and viral expression. Cardiac perfusion used 4% PFA solution to ensure fixation of the brain tissue. At least 24 hours after the cardiac perfusion, the brain was moved into 0.05% Thimerosal-PBS and stored in the fridge until histological processing.

Brains were sliced coronally at 50 µm using a vibratome (Leica VT1000) and collected into 1X PBS. Immunohistochemistry was performed as follows:

- To localize the imaging site, sections were incubated with an Alexa Fluor 488-conjugated GFP antibody at 1:1000 in 0.4% Triton-PBS overnight at room temperature.
- To identify tdTomato somatic expression, sections were incubated with a primary anti-RFP antibody at 1:2000 in 0.4% Triton-PBS overnight at 4°C, followed by a secondary anti-Rabbit Alexa Fluor 647 antibody at 1:1000 in 0.4% Triton-PBS for 2 hours at room temperature.

All sections were counterstained with DAPI at 1:1000 in PBS for 15 minutes and then mounted on glass slides with Mowiol solution, allowed to dry overnight, and sealed with nail polish.

Sections were imaged using a custom-built automated slide scanner equipped with a 4x 0.4-NA Plan Apo objective (Nikon Instruments Inc.) and P200 slide loader (Prior Scientific), controlled by NIS-Elements using custom acquisition scripts, from the Zuckerman Institute's Cellular Imaging Platform. For detection of DAPI, Alexa Fluor 488, and Alexa Fluor 647, the following excitation wavelengths and filters were used, respectively: 405 nm (filter 435/26 nm), 488 nm (filter 525/30 nm), and 647 nm (filter 705/72 nm). Image processing and analysis were performed using Brain J (RRID:SCR_027061).

Two-action isometric task

Setup

Mice were head-fixed and placed in a 3D-printed cup[65] with a sucrose spout placed within tongue reach. The right forelimb interacted with an immobile joystick (3 mm screw) mounted on 3D-printed frames attached to two load cell sensors (Phidgets 3139_0), one measuring push/pull and the other measuring left/right forces. The left forelimb rested on an auxiliary pole equipped with two additional load sensors measuring push/pull and vertical forces. The cup holding the mouse was mounted on a fifth load sensor to capture overall body movement.

Signals from all load sensors were recorded at 1 kHz using the Champalimaud Foundation Scientific (CFS) “Scientific Board v1.3” with “Quad Load Cell v1.2” module, which also detected threshold crossings based on experimenter-defined parameters. Data were visualized and stored using the CFS “Load Cells Visualizer” software.

A pyControl “Lickometer v1.0” board detected licking and contact with the joystick and auxiliary pole. Task control and sucrose delivery were managed by a pyControl[66] “Breakout 1.2” board, which logged events at 1 kHz and opened solenoid to dispense the reinforcer - 5 µl of 10% sucrose solution. Sucrose flowed via gravity through tubing and was delivered through a blunt 16G needle spout positioned near the animal’s mouth. Solenoid timing was calibrated before each session, and the sucrose solution was prepared fresh weekly. Hardware schematics are shown in Figure S4a.

Training

Mice were habituated to head fixation for 15 min on a running wheel over four days and then food-restricted to 80–85% of initial body weight. Training for the two-action isometric task was performed as

described in protocols.io ("Behavior training"). Mice pushed or pulled an immobile, pressure-sensitive joystick in a self-paced manner to receive reinforcer. A task action was defined as a force exceeding a task-defined threshold for longer than the task-defined duration. Thresholds, durations, and reinforcement delays were progressively adjusted during training – these are referred to as "Task Parameters" (Table 2). Reinforcement followed the end of a valid action after a delay. A 3 g quiescence threshold was required for trial initiation.

Reinforcement block	*Parameters label*	*Task action threshold (g)*	*Task action duration (ms)*	*Reinforcement delay (ms)*
Both	1	4	10	100
Both	2	4	20	200
Both	3	6	20	300
A or B	1	4	20	300
A or B	2	6	20	300
A or B	3	6	50	300

Table 2 – Changes in task parameters for different reinforcement blocks.

Each trial consisted of: (1) a 1 s inter-trial interval (ITI); if forces were below 3 g at 1 s, a new trial began, otherwise the ITI reset; (2) a trial period, during which a task action (push or pull exceeding the threshold for the task duration) triggered delayed sucrose delivery. Trials terminated upon reinforcement or when force fell below task-defined threshold after an action that did not meet task-defined duration.

Training comprised three sequential blocks with distinct reinforcement contingencies. In Block Both, push and pull actions were both reinforced, up to 50 reinforcements per action per session. On the final session of Block Both, the less frequent action was designated action A and the more frequent as action B. In Block A, only action A was reinforced – starting after last session of Block Both; in Block B, only action B was reinforced – starting after last session of Block A. Block Both sessions ended after 50 reinforcements of each action or 30 min. Blocks A and B ended after 200 reinforced actions or 30 min. Task parameters were adjusted during training, advancing according to Table 2 when mice executed >50 reinforced task actions in a session. Block Both concluded when mice exceeded 50 reinforced actions in two consecutive sessions at Parameters 3. Blocks A and B each concluded when the reinforced action comprised ≥80% of total actions - defined as reinforced task actions divided by total action A and action B executions - over three consecutive sessions at Parameters 3. Conclusion of Block B marked the end of training, after which animals were returned to ad libitum feeding.

For Rgs9-LCre::Grin1tm1Yql homozygous mice (Striatal NR1-KO) and Cre-negative control littermates, the same training protocol was used, with two modifications. Each block was capped at 1.2× the number of sessions required by the third fastest mouse in the cohort to reach Parameters 3 (Table 2). Animals were also advanced to the next block if they met give-up criteria, defined as fewer than ten 3g threshold crossings (≥50 ms duration) averaged over three consecutive sessions.

Selected sessions for analysis of behavioral training

Because animals were trained to criterion, the number of sessions per block varied across individuals. To enable cross-animal comparisons, specific sessions were selected for analysis. In Block Both, four sessions were chosen, and in Blocks A and B, five sessions were chosen, following the selection rules detailed in Table 3. These sessions were used in Figures S2 and S3.

Reinforcement block	*Selected session*	*Selection Rules*
Both	1	1st day at Parameters 1 for Block Both
Both	2	1st day at Parameters 2 for Block Both
Both	3	1st day at Parameters 3 from Block Both
Both	4	Last day before change in reinforcement schedule
A or B	1	1st day at Parameters 1 for Block A or B
A or B	2	1st day at Parameters 2 or Block A or B
A or B	3	1st day at Parameters 3 or Block A or B
A or B	4	Middle day between selected session 3 and 5
A or B	5	Last day of current reinforcement Block

Table 3 – Selected sessions for comparison across animals for all reinforcement blocks.

Adaptive Isometric task

Setup

Mice were trained on this task to test causal striatal control of specific ongoing actions. To do that, we developed a modified version of the two-action forelimb task that promoted balanced performance of both actions. Reinforcement probability was adjusted online based on recent performance: the proportion of push and pull actions was tracked over the last 10 trials. When both actions occurred in fewer than 70% of recent trials, both were reinforced; if one action exceeded 70%, that action was unrewarded.

The setup was otherwise identical to the two-action isometric task (Figure S4A), with additional hardware for closed-loop holographic stimulation (Figure S4B). This system required flexibility to: (1) select, on a single-trial basis, the threshold crossing to trigger stimulation; (2) trigger laser with millisecond latency; and (3) specify the holographic 2-photon pattern, on a single-trial basis. Features 1 and 2 required an improved hardware system for processing force sensors, while feature 3 required a software update to the two-photon microscope's stimulation interface.

Features 1 and 2 were accomplished using an updated load-cell system ("HARP Load Cells Acquisition v1.1" and "HARP Load Cell Interface v1.1," CFS Hardware Platform), which evaluated in parallel when the force exerted on the joystick crossed six experimenter-defined thresholds. In the holographic experiments, six thresholds were defined: 3 g, 4 g, and 6 g for push and pull actions. To select which threshold would trigger stimulation on a given trial, and to restrict the trigger signal to specific time windows within each trial, a custom circuit board consisting of a MUX and an AND gate was built to route the appropriate trigger signal to the microscope while minimizing the delay between threshold crossing and stimulation onset. The MUX and AND gate was controlled by custom MATLAB scripts running on the microscope control computer, which determined which force threshold should serve as the trigger on each trial and when to enable it. Digital communication between the task boards, the computer, and the MUX and AND gate was implemented using a NIDAQ DIO board (USB-6001).

Feature 3 was accomplished through a newly implemented function in the PrairieLink API of the Bruker microscope, which allowed two-photon stimulation patterns to be specified dynamically during an ongoing imaging experiment. Schematics of the full setup, including both the behavioral and stimulation hardware, are presented in Figures S4.

Training

This task was used for holographic stimulation experiments and followed the two-action forelimb task with modified reinforcement to promote balanced performance of both actions. Mice were trained for 3 weeks

with adaptive reinforcement and no stimulation. Training was conducted partly in behavior boxes in the absence of imaging, prior to head-fixed sessions under the microscope.

2-photon microscope

The optical setup comprised two femtosecond pulse lasers coupled to a custom-modified two-photon laser scanning microscope (Ultima In Vivo, Bruker Corporation, Billerica, Massachusetts) equipped with an 8 kHz resonant scanner. The laser source for imaging was a pulsed Ti:sapphire laser (Chameleon Vision-S, Coherent, Inc., Saxonburg, Pennsylvania) tuned to 920 nm for GCaMP6f imaging, with power controlled by a Pockels cell (350–105-BK, 302RM controller, Conoptics, Inc., Danbury, Connecticut). The laser beam was expanded by a 1:2 telescope (Thorlabs GBE02-B) and further scaled in the microscope by a 1:1.33 telescope before being coupled into a scan lens (f = 75 mm), tube lens (f = 180 mm), and objective lens (16X 0.8NA N16XLWD-PF, Nikon Corporation, Tokyo, Japan). The laser could also be directed to a non-resonant scanning path where both X and Y scanning were controlled by galvanometric mirrors. The fluorescence signal was collected through the objective lens, split from the IR beam with a 670 nm long-pass dichroic (ZT670rdc, Chroma Technology Corp., Bellows Falls, Vermont), and coupled into a liquid light guide assembly. The light was split by a 565 nm long-pass dichroic mirror (HQ565dcxr, Chroma Technology Corp.), with the green and red channels each directed to GaAsP PMTs through 510/20 and 607/45 bandpass filters, respectively (Chroma Technology Corp.).

The photostimulation path was largely independent from the imaging path. The laser source for photostimulation was a low repetition rate (1 MHz) amplified laser (Monaco 1035–40, Coherent, Inc., Saxonburg, Pennsylvania) operating at 1035 nm, with power controlled by an integrated acousto-optic modulator. The beam was expanded by a 1:2 telescope (Thorlabs GBE02-B) and directed into the microscope, where it could be steered either by a dedicated pair of galvanometers or coupled to a customized light path including a spatial light modulator (Bruker Neurolight 3D, Billerica, Massachusetts) for holographic stimulation. The photostimulation beam was combined and made coaxial with the imaging beam just before the scan lens via a 1030 nm short-pass dichroic (T1030SP, Chroma Technology Corp.).

Imaging and photostimulation were controlled by PrairieView (Bruker Corporation, Billerica, Massachusetts). Custom software defined stimulation patterns and triggered closed-loop photostimulation based on behavioral events - MATLAB (The Mathworks, Inc., Natick, Massachusetts) and Python, interfaced with PrairieView via the PrairieLink API.

2-photon imaging

Before each experimental cohort, the imaging and stimulation laser powers were measured at the output of 16X Nikon objective and confirmed stable across months. A custom computerized goniometer was used to precisely and reproducibly angle the head of the mouse such that the imaging lens was orthogonal to the beam path.

2-photon imaging was performed as described in protocols.io (“2-photon imaging”). Prior to each session, laser alignment was confirmed at the objective output. The mouse was head-fixed in the behavioral setup under the microscope, and the top of the GRIN lens was located and cleaned if necessary. Imaging power was continuously monitored using a beam pickoff (Thorlabs BSF10-B) and power meter (Thorlabs PDA100A2), with the Pockels cell bias adjusted accordingly throughout the session.

The typical imaging power was <50 mW, and could be up to 80 mW for imaging deeper than ~250 μm below the terminus of the GRIN lens. Images were acquired at 30 Hz over a 512 × 512 pixel field of view. For Figures 1 and 2, a 1.5x zoom was used over 496.1 μm × 496.1 μm; for Figures 3 and 4, a 1.5x zoom over 550.4 μm × 550.4 μm and a 2x zoom over 412.7 μm × 412.7 μm were used.

For cell type identification, 1000 frames of functional and structural images were acquired using green and red PMTs respectively. These data were then processed as described below (STAR methods – Distinguishing D1- and D2-MSNs). Structural images and z-stacks were acquired at the end of each session using both green and red channels to confirm ROI identity (Figure S1).

Most mice were imaged for the entire duration of training, with the same field of view tracked across sessions. ROIs identified in previous sessions were used to help locate the same field of view, and custom code was written to track the same ROIs across days.

Protocol to image the same field of view and neurons across sessions

For holographic stimulation experiments, the same field of view and neurons were imaged across paired sessions. After the first session, imaging data were processed with Suite2p[67] to extract active ROIs. In the paired session, PrairieView's "Brightness over Time" (BOT) feature was used to annotate these neurons in the GUI. XYZ coordinates were adjusted to best match the previous session's field of view.

Calibration of 2-photon stimulation spatial targeting and power

When doing holographic stimulation, the spatial targeting of 2-photon stimulation was calibrated before each session (protocols.io, "2-photon stimulation calibration"). Using PrairieView's "burn spots" routine, we iteratively aligned burned spots on a fluorescent sample to target locations, allowing the software to calibrate galvos' control for accurate stimulation. Calibration was stable across sessions, requiring only minor adjustments. This calibration was done using the same objective used for the experimental recordings.

Before experiments, we calibrated stimulation power and duration for each ROI. Using custom Python scripts with PrairieView's "BOT" feature, we visualized ROI responses while sweeping power (3–7 mW per target) and duration (50-100 ms) to identify the minimum parameters that reliably activated the targeted neurons.

Holographic optogenetics experiments

We performed holographic optogenetics experiments to test the effect of stimulating specific neuronal ensembles on self-paced forelimb actions during the Adaptive Isometric task.

Experiments relied on pairs of sessions in which we imaged the same field of view and neurons. On the first session, we performed two-photon imaging as mice performed the task, used Suite2p to extract functional ROIs, and automatically labelled ROIs as D1- or D2-MSNs. On the second session, we performed the stimulation experiment, consisting of a calibration block followed by a stimulation block (Figure 3F). For each stimulation session, cell type choice (D1- or D2-MSNs) was determined by balancing the number of sessions per cell type and the number of neurons available for stimulation.

Within the calibration block, we performed two-photon imaging as mice performed the adaptive isometric task for up to 30 minutes. Using PrairieView's "BOT" feature, we pre-defined ROIs identified from the previous session and rapidly extracted their fluorescence, averaged over the pixels within a circle defined by each Suite2p mask's center and radius. We extracted behavioral data and identified force threshold crossings of 3 g (hereafter "3g cross actions"). For either D1- or D2-MSNs, we z-scored each ROI's fluorescence across the recording and fit an SVM to predict push versus pull using single trials of neural activity averaged over the 5 imaging frames centered at force peak. If SVM accuracy on test trials exceeded 55%, we used the SVM weights to identify push- and pull-weighted ensembles with equal numbers of neurons, and plotted their activity time-locked to the force peak of 3g cross actions. If the

push (pull) ensemble showed greater activity for push (pull) actions, we proceeded to holographic stimulation.

We then defined a software command via the PrairieLink API to specify holographic stimulation for each ensemble. Stimulation was targeted to the center of each ROI in the ensemble, with total power calculated as the product of the number of neurons per ensemble and the desired average power per neuron. Stimulation duration was typically 100 ms - sufficient to reliably activate neurons expressing ChRmine. To cover each ROI's circular extent, dedicated galvos traced 5 spirals per stimulation with a diameter matching the average ROI diameter. Stimulation was configured to be delivered upon a trigger signal corresponding to a force threshold crossing.

Within the stimulation block, we performed two-photon imaging and closed-loop holographic stimulation as mice performed the adaptive isometric task for 30 minutes. We pre-defined a schedule interleaving conditions based on which action triggered stimulation (3g cross push or pull) and which stimulation pattern was delivered (push ensemble, pull ensemble, or no stimulation). Trials were organized into alternating blocks of 6, triggered by one action (either 3g cross push or 3g cross pull), with the following stimulation sequence per block: (1) no stimulation, (2) push ensemble, (3) no stimulation, (4) no stimulation, (5) pull ensemble, (6) no stimulation. A 2-second inter-trial interval was imposed between trials.

Our custom setup (Figures 3 and S4B) enabled selection of which force threshold crossing served as the trigger on interleaved trials and stimulation delivery within milliseconds of threshold crossing — critical given the rapid timescale of force changes. As described in the setup section above, the microscope computer ran custom MATLAB code to step through the scheduled trials, using the PrairieLink API to send stimulation commands and a NIDAQ DIO board (USB-6001) to select and enable force thresholds via the MUX and AND gate. The trigger signal output from the MUX was also received by the computer to track stimulation delivery and advance to the next inter-trial interval and trial.

EMG recordings

EMG recordings were performed as described in protocols.io ("EMG recordings"). Custom-made electrodes were implanted in forelimb muscles and connected to an Omnetics connector. Prior to each session, mice were head-fixed in the behavioral setup under the microscope and the 16-channel ZC16 headstage (Tucker-Davis Technologies – TDT) was connected to the implanted connector. EMG recordings were set to true differential mode in Synapse recording software, and the signal quality was confirmed when the mouse moved. Signal quality was optimized by grounding the mouse to the air table. Signals were acquired via the ZC16 headstage connected to a TDT RZ5D Bioamplifier and PZ5-32 preamplifier, and recorded at 24 kHz. For analysis, EMG signals were downsampled to 1 kHz, high-pass filtered at 40 Hz, rectified, and smoothed with a Gaussian kernel (SD = 25 ms)[68]. Data included 41 recording sessions from 4 mice performing the two-action isometric task (Figure 1).

Quantification and Statistical Analysis

Analysis was performed using custom scripts in Python.

Data extraction and synchronization

All data streams—task events (1 kHz), joystick forces (1 kHz), EMG (24 kHz), 2-photon imaging/voltage (30 Hz/1 kHz) - were synchronized via the ITI signal, providing a unified timestamp across modalities for precise alignment of behavior, forces, EMG and imaging.

Calibrated force measurement

The joystick voltage signal was converted to force using a linear model:

$$\text{Force (g)} = (\text{Voltage} - \text{baseline}) \times \text{conversion factor}$$

Baseline voltage was measured before each session and defined as the average voltage with no weight on the sensors. The conversion factor was determined by applying known weights (0, 2, 5, 20 g) and fitting a linear model. Conversion was stable and re-calibrated weekly to correct drift.

Definition of actions

For analysis, we defined several actions from force and touch signals. Task actions were push and pull events identified by the hardware from 1 kHz force traces, with intervals from threshold rise to fall. 3 g cross actions were pushes or pulls exceeding 3 g for ≥66 ms with joystick touch ≥100 ms before crossing to ensure isometric behavior; force traces were downsampled to 30 Hz. 6 g cross actions were a subset of 3 g crosses peaking above 6 g, the final training threshold. Touch was defined as joystick contact <2 g for >1 s, and licks as ≥5 licks with <0.5 s inter-lick interval, excluding overlap with 3 g crosses. We also analyzed special subsets: pre-isolated 3 g crosses had no other 3 g crosses in the 0.5 s before, post-isolated 3 g crosses had none in the 0.5 s after, and matched actions were push/pull trials matched for continuous-valued features (STAR Methods - Matching trials across actions).

Action events

For push and pull actions, analysis was aligned to two key events within the action interval: the moment the force threshold was crossed (rising edge) and the moment of peak force. For touch and lick actions, analysis was aligned to the midpoint of the action interval.

Matching trials across actions

We tested whether muscle and neural activity encoded push and pull actions using trials matched for average features (Figures 1F-H and S3). Tolerances were: average force ≤0.1 g (Figure 1F-H); peak force ≤0.1 g, two-axis force magnitude at peak ≤0.2 g, action duration ≤30 ms, and lick rate/probability during actions and at peak ≤0.5 Hz and ≤0.1, respectively (Figure S3).

Trials were iteratively removed to achieve matching while minimizing loss of trials. Features were z-scored, then for each iteration: (1) define action 1 as the one with fewer trials, action 2 as the other; (2) compute the target vector as the difference in mean features across actions; (3) zero in-tolerance features (i.e. set the difference to zero for each feature that is within-tolerance); (4) compute dot products of action 2 trials with the target; (5) remove the trial with the smallest dot product; (6) repeat until all features met tolerance.

For Figure S3, trials were pre-selected such that the time between solenoid opening and force threshold crossing was at least 4 seconds, **ensuring that analyzed actions did not occur during sucrose consumption**.

Time course of action force

Trial-averaged joystick force was computed as peri-event time histograms (PETHs) locked to action events: 6 g cross push/pull actions were aligned to force peak (Figure 1C). In Figure 2F, -4–0 s shows pre-isolated and 0–4 s shows post-isolated 3 g cross force, illustrating trials used to predict action identity from neural activity.

Quantification of actions across reinforcement

Action performance was quantified across the reinforcement schedule. For task actions, we quantified action rate for each session (Figure S2B,C). For 6 g cross actions, we quantified action rate pooled over sessions in each reinforcement block (Figures 1D). For trials matched across 3g cross actions, we

quantified number of actions, peak and average force, two-dimensional force magnitude, average lick probability, action duration, action rate, and action proportion (Figure S3A-H).

For Striatal NR1-KO mice, we quantified 6g cross actions with the overall action rate for each session (Figure S2H), and the action rate pooled over sessions in each reinforcement block (Figure S2I).

EMG analysis

We analyzed pre-processed EMG (STAR Methods – EMG Recordings) from four forelimb muscles for 6 g cross actions with force matched within 0.1 g. EMG channels were z-scored across each session and aligned to the 6 g cross event. We compared the vector magnitude of EMG difference across actions to the difference within each individual action. Concretely, for each session, we split trials for each action into two halves and calculated the EMG PETH for each half. At each time point locked to 6g cross, we calculated the vector magnitude of difference across halves within an action, as well as the difference across actions using the PETH from one half. We compared across-action and within-action difference in the 100ms centered at 6g cross (Figure 1H).

Neuron detection and activity extraction with Suite2p

2-photon imaging data were processed with Suite2p[67] for motion correction, neuron detection, and fluorescence extraction; ROIs were classified using a classifier trained on sparse striatal MSN activity, neuropil-corrected ($F_{corrected} = F - 0.7 \times F_{neu}$), and neural activity was reported as z-scored $\Delta F/F_0$ with F_0 as the 1-minute running median. Figure 3H,I,K-M shows z-scored, raw fluorescence averaged over ROI pixels from PrairieView's "BOT".

Distinguishing D1- and D2-MSNs

We developed an automated thresholding method to identify D1- and D2-MSNs based on fluorophore expression (Figure S1). Functional imaging (green channel) was used to detect ROIs with Suite2p, while structural imaging (red channel) determined co-localization of cell-type–specific fluorophores. The overall workflow is outlined in Figure S1A. The method involved three main steps: (1) aligning and processing structural and functional data, (2) labeling ROIs manually (ground truth) and automatically using red-channel thresholds, and (3) comparing manual and automatic labeling.

1) **Structural data processing:**
   Functional data were motion-corrected and ROIs extracted using Suite2p (Figure S1B). As fluorophore expression is soma-restricted, only pixels within a 12-pixel diameter circle of each ROI center were used. Structural images (red channel, 1000 frames) were motion-corrected against the functional template, then background-subtracted using a 25-pixel rolling ball filter in ImageJ (Figure S1C,D).
2) **ROI labeling:**
   ROIs near the edge of the GRIN lens were excluded to avoid aberrations; the center of the lens was estimated from the red channel by computing the median coordinates of pixels above the 20th percentile intensity. Valid ROIs fell within a circle around this center.
   a. *Manual labeling:*
      Manual labels were assigned on average, motion-corrected, background-subtracted structural images overlaid with functional ROIs (Figure S1E). Labels were "red," "non-red," or "no-ID" if uncertain. Each mouse had at least two sessions labeled three times (folds); only ROIs consistently labeled across all folds were retained (Figure S1F,G). Thresholds of 0.7 (red) and 0.3 (non-red) were used for comparisons in Figure S1.
   b. *Automatic labeling:*

Structural images were processed with Ilastik[69] (v1.3.3) to estimate per-pixel probability of being red (Figure S1H). ROIs with >0.7 red pixels were labeled “red,” <0.3 as “non-red,” and 0.3–0.7 as “no-ID” (Figure S1I).

3) **Manual vs. automatic comparison:**
Label accuracy, swapping errors, and single-method errors were quantified across threshold pairs to optimize performance (Figure S1J-M).

Thresholds of 0.3/0.7 (non-red/red) were used for general characterization (Figure 2), and 0.45/0.55 thresholds were used for holographic stimulation experiments (Figures 3 and 4).

Time course of neuronal activity

Neuronal activity was analyzed time-locked to action events. An example trace of activity averaged over D1-, D2-, and all MSNs is shown alongside force (Figure 2B). We analyzed session-averaged PETHs aligned to the force peak of pre-isolated 3 g cross actions (Figure 2C). We specifically analyzed activity at force peak and 0.5 seconds before force peak, computed as z-scored F/F0 averaged over 3 frames centered at each time point. Averaging over neurons of each cell type, we compared activity at force peak versus 0.5 seconds before force peak (paired t-test). We also compared population-averaged activity across actions at force peak (paired t-test). We plotted PETHs of example individual D1- and D2-MSNs (Figure 2D). Only sessions with ≥20 trials per action and ≥10 neurons were analyzed (same as in STAR Methods – Support Vector Machine (SVM) decoder model).

Support Vector Machine (SVM) decoder model

To test whether MSNs encoded specific actions, we trained a linear SVM (sklearn.svm.LinearSVC) to predict action identity (e.g. push versus pull) from single-trial population activity. Each trial vector comprised the $\Delta F/F_0$ of each neuron averaged over 5 frames (~167 ms) centered on the action event (i.e. the frame at the event and 2 frames before and after) (Figures 2H,I and S2). Only sessions with ≥20 trials per action and ≥10 neurons were analyzed, and trials were weighted by 1/(number of trials per action class) to prevent bias. Trials were randomly split 90/10 for training/testing, repeated 100 times. The regularization parameter (C) was optimized over a range of 1×10^{-5} to 100 via 10-fold cross-validation on the training set (with 80% of training data used for sub-training and 20% for sub-testing in each sub-fold). Accuracy was computed for each action separately and overall, and compared to a shuffle control **in which the model was trained on randomized action labels but tested on non-randomized labels**.

Matching number of D1- and D2-MSNs for SVM decoding

To compare D1- and D2-MSNs, we equalized their numbers for SVM decoding (Figure S3K). For the cell type with more neurons, a random subset was selected in each training/testing fold to match the number of neurons in the other cell type.

SVM decoding across sessions

We reported SVM decoding accuracy for individual sessions and pooled across sessions, under different reinforcement conditions.

SVM accuracy across each pair of actions, pooling the selected sessions across all reinforcement blocks:

- Figure 2H decoded the population activity vector including all MSNs. For push and pull, 3g cross actions were used, and neural activity was centered at force peak.
- Figure 2I is the same as Figure 2H, but decoding population activity of a cell type (D1- or D2-).
- Figure S3I is the same as Figure 2H, and Figure S3J is the same as Figure 2I, except for matched 3g cross actions and separating by action A and B rather than push and pull.

SVM accuracy for each selected session across reinforcement:

- Figure S2D decoded all MSNs at force peak for 3g cross actions, predicting push versus pull.

- Figure S2E (left) is the same as Figure S2D, but decoding population activity of a cell type.
- Figure S2F (left) is the same as Figure S2D, except with one SVM predicting action A versus touch and another SVM predicting action B versus touch.

SVM accuracy pooling selected sessions within each reinforcement block:

- Figure S2E (right) decoded population activity of a cell type at force peak for 3g cross actions, predicting push versus pull.
- Figure S2F (right) decoded all MSNs at force for 3g cross actions, with one SVM predicting action A versus touch and another SVM predicting action B versus touch.
- Figure S2G is the same as Figure S2F (right), but decoding population activity of a cell type.
- Figure S2J is the same as Figure S2E (right), except decoding all MSNs for NR1-KO and littermate control mice.

Time course of action decoding

We used the SVM to analyze the time course of action decoding. We extracted and decoded neural activity centered at specific time points relative to action events. In Figure 2E-G, we analyzed non-overlapping 5 frame bins locked to force peak. The bin at time 0 was centered at force peak. From -4 seconds to 0 seconds, the SVM was applied on trials from pre-isolated 3g cross actions. From after 0 seconds to 4 seconds, the SVM was applied on trials from post-isolated 3g cross actions. The SVM used all MSNs in Figure 2EF and matched numbers of D1- or D2-MSNs in Figure 2G.

Overlay of decoder accuracy and force

Decoder accuracy and force were visualized in normalized units (Figure 2F) by centering data at -4 seconds from force peak (i.e. subtracting the force at -4 seconds from all time points) and scaling by the window's (-4 s to +4 s) maximum trial-averaged value (i.e. dividing by the maximum).

SVM weights and dimension

The SVM assigned positive (negative) weights to neurons whose activation biased the SVM to predict push (pull). For analysis of the weights directly, we z-scored the weights for each decoder (Figure 3G,J). The weights define a dimension of population activity that best discriminates action identity. We extract this dimension by using the decoder's weight vector and dividing by its L2 norm. This is a vector $w \in R^N$, where N is the number of neurons and the vector magnitude is 1 ($\|w\|_2 = 1$).

Time course of SVM projection

We visualized the time course of action-specific population activity by projecting neural activity onto the SVM dimension. At each time point, the vector of population activity $x \in R^N$ (where each entry is one neuron's activity) was projected onto the SVM dimension: $w^T x \in R^1$, resulting in a scalar that represents the weighted sum of neuronal activity. In Figure 2J, we used the SVM trained at force peak and projected activity from pre-isolated 3g cross actions time locked to force peak. In Figure 3H,I and N, we used the SVM trained at force peak and projected activity locked to holographic stimulation.

SVM-weighted neural ensembles

We defined neuronal ensembles that biased the SVM toward push or pull. For each action, we selected the top $k *$ ensemble_frac neurons with the largest positive (push) or negative (pull) weights, where k is the smaller of the number of push- or pull-weighted neurons. We used ensemble_frac $= 0.25$ to visualize the most strongly weighted neurons (Figure 2K-N) and ensemble_frac $= 1$ to select as many neurons as possible for holographic optogenetic stimulation (Figure 3G-N).

Time course of SVM-weighted neural ensembles

We analyzed the activity of push-weighted and pull-weighted neural ensembles that biased the SVM to predict push or pull.

- We averaged the activity over the neurons in the push-weighted and pull-weighted ensembles time-locked to force peak of 3g cross actions (Figure 2K,L).
- We separately analyzed the average activity of the identified D1- and D2-MSNs within these ensembles (Figure 2M,N).
- In Figure 3K, we averaged ensemble activity in the 100ms before force peak of 3g cross actions.
- In Figure 3L,M, we time-locked ensemble activity to holographic stimulation.

Effect of holographic stimulation on actions

We compared force in stimulation trials versus no-stimulation trials that immediately preceded stimulation trials. No-stimulation trials that occurred closer in time to the previous stimulation trial than the next stimulation trial were excluded. We analyzed force sampled at 1 kHz and low-pass filtered at 30 Hz (10th-order Butterworth filter), providing both high resolution visualization and de-noising.

Given that stimulation occurred with a random (but small) latency from 3g cross, we performed an alignment procedure to match the force of no-stimulation trials to stimulation trials in the time leading up to stimulation. We identified the time in the no-stimulation trial that best matched the following stimulation trial on 1) the force at time of stimulation ("target instantaneous force") and 2) the average force in the 100 ms window preceding time of stimulation ("target preceding force"). Concretely, we identified the time in each no-stimulation trial that minimized a cost, which was the sum of two terms: 1) the absolute difference between force at the time in the no-stimulation trial and the "target instantaneous force", and 2) the absolute difference between the average force in the 100 ms preceding the time in the no-stimulation trial and the "target preceding force." This alignment was performed on the raw 1 kHz force trace before filtering at 30 Hz.

We analyzed the difference between stimulation trials locked to stimulation time and no-stimulation trials locked to this aligned time. Trials were pooled into "congruent stimulation" and "non-congruent stimulation" conditions (Figure 4). Congruent stimulation consisted of 1) stimulation of the push ensemble triggered on push 3g cross, and 2) stimulation of the pull ensemble triggered on pull 3g cross. Non-congruent stimulation consisted of 1) stimulation of the push ensemble triggered on pull 3g cross, and 2) stimulation of the pull ensemble triggered on push 3g cross. Accordingly, the number of data points for statistics (Figure 4C,D) was 2×num_sessions.

Statistical analysis

Statistical analysis was performed using custom scripts in Python. We used the ‘scipy.stats’ package to perform paired t-tests (‘ttest_rel’ method), unpaired t-tests (‘ttest_ind’ method), and linear regression (‘linregress’ method). Linear mixed effects models were performed using the ‘statsmodels’ package and ‘mixedlm’ method. Significance was set at P=0.05. Statistical information in Tables S1 and S2.

References

1. Donaldson, I., Marsden, C.D., Schneider, S., and Bhatia, K. (2012). Marsden’s Book of Movement Disorders (Oxford University Press) https://doi.org/10.1093/med/9780192619112.001.0001.
2. Albin, R.L., Young, A.B., and Penney, J.B. (1989). The functional anatomy of basal ganglia disorders. Trends Neurosci. *12*, 366–375. https://doi.org/10.1016/0166-2236(89)90074-X.

3. DeLong, M.R. (1990). Primate models of movement disorders of basal ganglia origin. Trends Neurosci. *13*, 281–285. https://doi.org/10.1016/0166-2236(90)90110-V.
4. Giladi, N., McMahon, D., Przedborski, S., Flaster, E., Guillory, S., Kostic, V., and Fahn, S. (1992). Motor blocks in Parkinson's disease. Neurology *42*, 333–333. https://doi.org/10.1212/WNL.42.2.333.
5. Mazzoni, P., Hristova, A., and Krakauer, J.W. (2007). Why Don't We Move Faster? Parkinson's Disease, Movement Vigor, and Implicit Motivation. J. Neurosci. *27*, 7105. https://doi.org/10.1523/JNEUROSCI.0264-07.2007.
6. Simonyan, K., Cho, H., Hamzehei Sichani, A., Rubien-Thomas, E., and Hallett, M. (2017). The direct basal ganglia pathway is hyperfunctional in focal dystonia. Brain *140*, 3179–3190. https://doi.org/10.1093/brain/awx263.
7. Shohamy, D., Myers, C.E., Grossman, S., Sage, J., Gluck, M.A., and Poldrack, R.A. (2004). Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. Brain J. Neurol. *127*, 851–859. https://doi.org/10.1093/brain/awh100.
8. Knowlton, B.J., Mangels, J.A., and Squire, L.R. (1996). A Neostriatal Habit Learning System in Humans. Science *273*, 1399–1402. https://doi.org/10.1126/science.273.5280.1399.
9. Albin, R.L., Reiner, A., Anderson, K.D., Penney, J.B., and Young, A.B. (1990). Striatal and nigral neuron subpopulations in rigid Huntington's disease: Implications for the functional anatomy of chorea and rigidity-akinesia. Ann. Neurol. *27*, 357–365. https://doi.org/10.1002/ana.410270403.
10. Panigrahi, B., Martin, K.A., Li, Y., Graves, A.R., Vollmer, A., Olson, L., Mensh, B.D., Karpova, A.Y., and Dudman, J.T. (2015). Dopamine Is Required for the Neural Representation and Control of Movement Vigor. Cell *162*, 1418–1430. https://doi.org/10.1016/j.cell.2015.08.014.
11. Dudman, J.T., and Krakauer, J.W. (2016). The basal ganglia: from motor commands to the control of vigor. Neurobiol. Cogn. Behav. *37*, 158–166. https://doi.org/10.1016/j.conb.2016.02.005.
12. Turner, R.S., and Desmurget, M. (2010). Basal ganglia contributions to motor control: a vigorous tutor. Mot. Syst. – Neurobiol. Behav. *20*, 704–716. https://doi.org/10.1016/j.conb.2010.08.022.
13. Yttri, E.A., and Dudman, J.T. (2016). Opponent and bidirectional control of movement velocity in the basal ganglia. Nature *533*, 402–406. https://doi.org/10.1038/nature17639.
14. Klaus, A., Martins, G.J., Paixao, V.B., Zhou, P., Paninski, L., and Costa, R.M. (2017). The Spatiotemporal Organization of the Striatum Encodes Action Space. Neuron *95*, 1171-1180.e7. https://doi.org/10.1016/j.neuron.2017.08.015.
15. Markowitz, J.E., Gillis, W.F., Beron, C.C., Neufeld, S.Q., Robertson, K., Bhagat, N.D., Peterson, R.E., Peterson, E., Hyun, M., Linderman, S.W., et al. (2018). The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. Cell *174*, 44-58.e17. https://doi.org/10.1016/j.cell.2018.04.019.

16. Dhawale, A.K., Wolff, S.B.E., Ko, R., and Ölveczky, B.P. (2021). The basal ganglia control the detailed kinematics of learned motor skills. Nat. Neurosci. *24*, 1256–1269. https://doi.org/10.1038/s41593-021-00889-3.
17. Park, J., Polidoro, P., Fortunato, C., Arnold, J., Mensh, B., Gallego, J.A., and Dudman, J.T. (2025). Conjoint specification of action by neocortex and striatum. Neuron *113*, 620-636.e6. https://doi.org/10.1016/j.neuron.2024.12.024.
18. Kravitz, A.V., Freeze, B.S., Parker, P.R.L., Kay, K., Thwin, M.T., Deisseroth, K., and Kreitzer, A.C. (2010). Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. Nature *466*, 622–626. https://doi.org/10.1038/nature09159.
19. Arber, S., and Costa, R.M. (2018). Connecting neuronal circuits for movement. Science *360*, 1403–1404. https://doi.org/10.1126/science.aat5994.
20. Joel, D., Niv, Y., and Ruppin, E. (2002). Actor–critic models of the basal ganglia: new anatomical and computational perspectives. Neural Netw. *15*, 535–547. https://doi.org/10.1016/S0893-6080(02)00047-3.
21. Niv, Y. (2009). Reinforcement learning in the brain. J. Math. Psychol. *53*, 139–154. https://doi.org/10.1016/j.jmp.2008.12.005.
22. Ito, M., and Doya, K. (2011). Multiple representations and algorithms for reinforcement learning in the cortico-basal ganglia circuit. Curr. Opin. Neurobiol. *21*, 368–373. https://doi.org/10.1016/j.conb.2011.04.001.
23. Samejima, K., Ueda, Y., Doya, K., and Kimura, M. (2005). Representation of Action-Specific Reward Values in the Striatum. Science *310*, 1337–1340. https://doi.org/10.1126/science.1115270.
24. Reynolds, J.N.J., Hyland, B.I., and Wickens, J.R. (2001). A cellular mechanism of reward-related learning. Nature *413*, 67–70. https://doi.org/10.1038/35092560.
25. Kravitz, A.V., Tye, L.D., and Kreitzer, A.C. (2012). Distinct roles for direct and indirect pathway striatal neurons in reinforcement. Nat. Neurosci. *15*, 816–818. https://doi.org/10.1038/nn.3100.
26. Vicente, A.M., Galvão-Ferreira, P., Tecuapetla, F., and Costa, R.M. (2016). Direct and indirect dorsolateral striatum pathways reinforce different action strategies. Curr. Biol. *26*, R267–R269. https://doi.org/10.1016/j.cub.2016.02.036.
27. Hikosaka, O., Takikawa, Y., and Kawagoe, R. (2000). Role of the Basal Ganglia in the Control of Purposive Saccadic Eye Movements. Physiol. Rev. *80*, 953–978. https://doi.org/10.1152/physrev.2000.80.3.953.
28. Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D.S., and Yuste, R. (2018). Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions. eLife *7*, 1–21. https://doi.org/10.7554/eLife.32671.
29. Marshel, J.H., Kim, Y.S., Machado, T.A., Quirin, S., Benson, B., Kadmon, J., Raja, C., Chibukhchyan, A., Ramakrishnan, C., Inoue, M., et al. (2019). Cortical layer–specific critical dynamics triggering perception. Science *365*, eaaw5202–eaaw5202. https://doi.org/10.1126/science.aaw5202.
30. Piantadosi, S.C., Zhou, Z.C., Pizzano, C., Pedersen, C.E., Nguyen, T.K., Thai, S., Stuber, G.D., and Bruchas, M.R. (2024). Holographic stimulation of opposing amygdala ensembles bidirectionally modulates valence-specific behavior via mutual inhibition. Neuron *112*, 593-610.e5. https://doi.org/10.1016/j.neuron.2023.11.007.

31. Gill, J.V., Lerman, G.M., Zhao, H., Stetler, B.J., Rinberg, D., and Shoham, S. (2020). Precise Holographic Manipulation of Olfactory Circuits Reveals Coding Features Determining Perceptual Detection. Neuron *108*, 382-393.e5. https://doi.org/10.1016/j.neuron.2020.07.034.
32. Vinograd, A., Nair, A., Kim, J.H., Linderman, S.W., and Anderson, D.J. (2024). Causal evidence of a line attractor encoding an affective state. Nature *634*, 910–918. https://doi.org/10.1038/s41586-024-07915-x.
33. Cruz, B.F., Guiomar, G., Soares, S., Motiwala, A., Machens, C.K., and Paton, J.J. (2022). Action suppression reveals opponent parallel control via striatal circuits. Nature *607*, 521–526. https://doi.org/10.1038/s41586-022-04894-9.
34. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448. https://doi.org/10.1038/nn.3643.
35. Inagaki, H.K., Chen, S., Daie, K., Finkelstein, A., Fontolan, L., Romani, S., and Svoboda, K. (2022). Neural Algorithms and Circuits for Motor Planning. Annu. Rev. Neurosci. *45*, 249–271. https://doi.org/10.1146/annurev-neuro-092021-121730.
36. Nelson, A., Abdelmesih, B., and Costa, R.M. (2021). Corticospinal populations broadcast complex motor signals to coordinated spinal and striatal circuits. Nat. Neurosci. *24*, 1721–1732. https://doi.org/10.1038/s41593-021-00939-w.
37. Economo, M.N., Viswanathan, S., Tasic, B., Bas, E., Winnubst, J., Menon, V., Graybuck, L.T., Nguyen, T.N., Smith, K.A., Yao, Z., et al. (2018). Distinct descending motor cortex pathways and their roles in movement. Nature *563*, 79–84. https://doi.org/10.1038/s41586-018-0642-9.
38. Churchland, M.M., and Shenoy, K.V. (2024). Preparatory activity and the expansive null-space. Nat. Rev. Neurosci. *25*, 213–236. https://doi.org/10.1038/s41583-024-00796-z.
39. Cui, G., Jun, S.B., Jin, X., Pham, M.D., Vogel, S.S., Lovinger, D.M., and Costa, R.M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. Nature *494*, 238–242. https://doi.org/10.1038/nature11846.
40. Sheng, M., Lu, D., Shen, Z., and Poo, M. (2019). Emergence of stable striatal D1R and D2R neuronal ensembles with distinct firing sequence during motor learning. Proc. Natl. Acad. Sci. *116*, 11038–11047. https://doi.org/10.1073/pnas.1901712116.
41. Tecuapetla, F., Jin, X., Lima, S.Q., and Costa, R.M. (2016). Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution. Cell *166*, 703–715. https://doi.org/10.1016/j.cell.2016.06.032.
42. Foster, N.N., Barry, J., Korobkova, L., Garcia, L., Gao, L., Becerra, M., Sherafat, Y., Peng, B., Li, X., Choi, J.-H., et al. (2021). The mouse cortico–basal ganglia–thalamic network. Nature *598*, 188–194. https://doi.org/10.1038/s41586-021-03993-3.
43. Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., et al. (2016). The mouse cortico-striatal projectome. Nat. Neurosci. *19*, 1100–1114. https://doi.org/10.1038/nn.4332.
44. Parker, J.G., Marshall, J.D., Ahanonu, B., Wu, Y.-W., Kim, T.H., Grewe, B.F., Zhang, Y., Li, J.Z., Ding, J.B., Ehlers, M.D., et al. (2018). Diametric neural ensemble dynamics in parkinsonian and dyskinetic states. Nature *557*, 177–182. https://doi.org/10.1038/s41586-018-0090-6.

45. Gerfen, C.R., and Surmeier, D.J. (2011). Modulation of Striatal Projection Systems by Dopamine. Annu. Rev. Neurosci. *34*, 441–466. https://doi.org/https://doi.org/10.1146/annurev-neuro-061010-113641.
46. Maltese, M., March, J.R., Bashaw, A.G., and Tritsch, N.X. (2021). Dopamine differentially modulates the size of projection neuron ensembles in the intact and dopamine-depleted striatum. eLife *10*, e68041. https://doi.org/10.7554/eLife.68041.
47. Dang, M.T., Yokoi, F., Yin, H.H., Lovinger, D.M., Wang, Y., and Li, Y. (2006). Disrupted motor learning and long-term synaptic plasticity in mice lacking NMDAR1 in the striatum. Proc. Natl. Acad. Sci. *103*, 15254–15259. https://doi.org/10.1073/pnas.0601758103.
48. Koralek, A.C., Jin, X., Long II, J.D., Costa, R.M., and Carmena, J.M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. Nature *483*, 331–335. https://doi.org/10.1038/nature10845.
49. Jin, X., and Costa, R.M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. Nature *466*, 457–462. https://doi.org/10.1038/nature09263.
50. Santos, F.J., Oliveira, R.F., Jin, X., and Costa, R.M. (2015). Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. eLife *4*, e09423. https://doi.org/10.7554/eLife.09423.
51. Dong, J., Hawes, S., Wu, J., Le, W., and Cai, H. (2021). Connectivity and Functionality of the Globus Pallidus Externa Under Normal Conditions and Parkinson's Disease. Front. Neural Circuits *Volume 15-2021*. https://doi.org/10.3389/fncir.2021.645287.
52. Yin, H.H., Mulcare, S.P., Hilário, M.R.F., Clouse, E., Holloway, T., Davis, M.I., Hansson, A.C., Lovinger, D.M., and Costa, R.M. (2009). Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. Nat. Neurosci. *12*, 333–341. https://doi.org/10.1038/nn.2261.
53. Arber, S., and Costa, R.M. (2022). Networking brainstem and basal ganglia circuits for movement. Nat. Rev. Neurosci. *23*, 342–360. https://doi.org/10.1038/s41583-022-00581-w.
54. Redgrave, P., Prescott, T.J., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? Neuroscience *89*, 1009–1023. https://doi.org/10.1016/S0306-4522(98)00319-4.
55. Grillner, S., Robertson, B., and Stephenson-Jones, M. (2013). The evolutionary origin of the vertebrate basal ganglia and its role in action selection. J. Physiol. *591*, 5425–5431. https://doi.org/10.1113/jphysiol.2012.246660.
56. Falasconi, A., Kanodia, H., and Arber, S. (2025). Dynamic basal ganglia output signals license and suppress forelimb movements. Nature *644*, 749–758. https://doi.org/10.1038/s41586-025-09066-z.
57. Athalye, V.R., Carmena, J.M., and Costa, R.M. (2020). Neural reinforcement: re-entering and refining neural dynamics leading to desirable outcomes. Neurobiol. Behav. *60*, 145–154. https://doi.org/10.1016/j.conb.2019.11.023.
58. Neely, R.M., Koralek, A.C., Athalye, V.R., Costa, R.M., and Carmena, J.M. (2018). Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. Neuron *97*, 1356-1368.e4. https://doi.org/10.1016/j.neuron.2018.01.051.

59. Athalye, V.R., Ganguly, K., Costa, R.M., and Carmena, J.M. (2017). Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. Neuron *93*, 955-970.e5. https://doi.org/10.1016/j.neuron.2017.01.016.
60. Athalye, V.R., Khanna, P., Gowda, S., Orsborn, A.L., Costa, R.M., and Carmena, J.M. (2023). Invariant neural dynamics drive commands to control different movements. Curr. Biol. *33*, 2962-2976.e15. https://doi.org/10.1016/j.cub.2023.06.027.
61. Athalye, V.R., Santos, F.J., Carmena, J.M., and Costa, R.M. (2018). Evidence for a neural law of effect. Science *359*, 1024–1029. https://doi.org/10.1126/science.aao6058.
62. Kao, T.-C., and Hennequin, G. (2019). Neuroscience out of control: control-theoretic perspectives on neural circuit dynamics. Comput. Neurosci. *58*, 122–129. https://doi.org/10.1016/j.conb.2019.09.001.
63. Stahl, C.M., and Frucht, S.J. (2017). Focal task specific dystonia: a review and update. J. Neurol. *264*, 1536–1541. https://doi.org/10.1007/s00415-016-8373-z.
64. Akay, T., Acharya, H.J., Fouad, K., and Pearson, K.G. (2006). Behavioral and Electromyographic Characterization of Mice Lacking EphA4 Receptors. J. Neurophysiol. *96*, 642–651. https://doi.org/10.1152/jn.00174.2006.
65. Mosberger, A.C., Sibener, L.J., Chen, T.X., Rodrigues, H.F.M., Hormigo, R., Ingram, J.N., Athalye, V.R., Tabachnik, T., Wolpert, D.M., Murray, J.M., et al. (2024). Exploration biases forelimb reaching strategies. Cell Rep. *43*. https://doi.org/10.1016/j.celrep.2024.113958.
66. Akam, T., Lustig, A., Rowland, J.M., Kapanaiah, S.K., Esteve-Agraz, J., Panniello, M., Márquez, C., Kohl, M.M., Kätzel, D., Costa, R.M., et al. (2022). Open-source, Python-based, hardware and software for controlling behavioural neuroscience experiments. eLife *11*, e67846. https://doi.org/10.7554/eLife.67846.
67. Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L.F., Dalgleish, H., Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. bioRxiv, 061507. https://doi.org/10.1101/061507.
68. Miri, A., Warriner, C.L., Seely, J.S., Elsayed, G.F., Cunningham, J.P., Churchland, M.M., and Jessell, T.M. (2017). Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. Neuron *95*, 683-696.e11. https://doi.org/10.1016/j.neuron.2017.06.042.
69. Sommer, C., Straehle, C., Köthe, U., and Hamprecht, F.A. (2011). Ilastik: Interactive learning and segmentation toolkit. In 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 230–233. https://doi.org/10.1109/ISBI.2011.5872394.