Cortico-subthalamic projections send brief stop signals to halt visually-guided locomotion

Elie M. Adam[1*], Taylor Johns[1] and Mriganka Sur[1*]

[1]Picower Institute for Learning and Memory,
Massachusetts Institute of Technology
Cambridge, MA 02139

[*]Corresponding Authors (eadam@mit.edu, msur@mit.edu)

Summary

Goal-directed locomotion necessitates control signals that propagate from higher-order areas to regulate spinal mechanisms. The cortico-subthalamic hyperdirect pathway offers a short route for cortical information to reach locomotor centers in the brainstem. We developed a task where head-fixed mice run to a visual landmark, then stop and wait to collect reward, and examined the role of secondary motor cortex (M2) projections to the subthalamic nucleus (STN) in controlling locomotion. Our modeled behavioral strategy indicates a switching point in behavior, suggesting a critical neuronal control signal at stop locations. Optogenetic activation of M2 axons in STN leads the animal to stop prematurely. By imaging M2 neurons projecting to STN, we find neurons that are active at the onset of stops, when executed at the landmark but not spontaneously elsewhere. Our results suggest that the M2-STN pathway can be recruited during visually-guided locomotion to rapidly and precisely control the mesencephalic locomotor region through the basal ganglia.

Keywords

Subthalamic nucleus, secondary motor cortex, hyperdirect pathway, locomotion, goal-directed, visually-guided, stopping, landmark, controller.

Introduction

Coordinated movement, and in particular locomotion, is enabled by distributed spinal and brain circuits. Although the executive mechanisms for locomotion are implemented in the spinal cord (Goulding, 2009; Grillner, 2003; Kiehn, 2016, 2006), locomotion is importantly regulated by supraspinal circuitry (Ferreira-Pinto et al., 2018; Kim et al., 2017; Ryczko and Dubuc, 2013). Brainstem circuits can induce locomotion upon targeted stimulation (Caggiano et al., 2018; Capelli et al., 2017; Josset et al., 2018; Roseberry et al., 2016); however, goal-directed locomotion, especially when guided by sensory information, requires control signals from higher-order areas directed towards spinal mechanisms (Arber and Costa, 2018; Drew et al., 2004; Grillner et al., 2008). The repertoire of higher-order signals required to regulate locomotion, and the neural circuitry that enables such signaling, are particularly unclear in behavioral settings (Arber and Costa, 2018). The control principles and substrates governing action can further suggest principles for cognitive control, particularly if the same neural substrate supports multiple functions.

While the regulation of locomotion is mainly implemented in the brainstem, planning is considered to arise at the level of the cortex (Churchland et al., 2010; Economo et al., 2018; Shenoy et al., 2013; Svoboda and Li, 2018; Wong et al., 2015). The basal ganglia are situated between these two centers and importantly regulate voluntary movement (Bolam et al., 2000; Graybiel, 2000). The direct and indirect pathway, initiated from the striatum, bidirectionally control locomotion (Roseberry et al., 2016) through the substantia nigra pars reticulata (SNr) and its control over the mesencephalic locomotor region (MLR) (Freeze et al., 2013; Liu et al., 2018). The basal ganglia additionally admit a cortical input straight to the subthalamic nucleus (STN) (Nambu et al., 2002, 2000) which projects to SNr (Hamani et al., 2004). This pathway has been termed the hyperdirect pathway, and evidence across species, notably humans, has shown stop activity at the source of this pathway in reactive stop-signal or go/no-go tasks (Aron et al., 2016; Aron and Poldrack, 2006; Eagle et al., 2008; Wessel and Aron, 2017). The STN itself is considered pivotal in stopping movement (Hamani et al., 2004; Schmidt et al., 2013). In rodents, reducing STN's excitatory output by a partial knockout of vGlut2 in STN induces hyperlocomotion (Schweizer et al., 2014), and lesions of STN induce impulsive responding (Baunez and Robbins, 1997; Eagle et al., 2008; Uslaner and Robinson, 2006). More recently, in mice, optogenetic areal activation of STN excitatory cells disrupts self-initiated bouts of licking

(Fife et al., 2017). Thus, the hyperdirect pathway stands as an ideal short-latency cortico-brainstem route for fast control of locomotion.

Goal-directed locomotion implies a proactive locomotor plan that is implemented to achieve a needed goal. From an engineering standpoint, we can enforce a desired goal trajectory in a system through feedback control (Aström and Murray, 2010). Feedback control is based on an error signal: a measured discrepancy between what we would like the system to do (reference) and what it is actually doing (output). By processing such a signal through a controller and feeding it to the system (plant) we intend to control, we ensure adequate performance (Dahleh et al., 2004; Oppenheim et al., 1996). If such a principle is implemented in neural circuits, then we would expect surges in neural signals upon sudden changes in planned locomotion trajectories, and that such signals would drive movement corrections to ensure fast control.

We thus developed a task where head-fixed mice run to a visual landmark, then stop and wait to collect reward, and examined the projections from secondary motor cortex (M2) to STN. We hypothesized that these projections convey rapid signals that halt locomotion, indicating that the hyperdirect pathway acts as a controller onto the mesencephalic locomotor region in the midbrain. Here, we report the existence of such signals sent from M2 to STN that halt visually-guided locomotion.

Results

We trained mice to run, stop and wait at landmarks to collect reward.

We developed a task that allows us to examine proactive visually-guided locomotion stops (**Figure 1-A**). A head-fixed mouse was positioned on a self-propelled treadmill, in a virtual runway flanked on both sides by a continuous streak of LEDs. At the start of a trial, the animal was presented with a visual landmark, consisting of a lit contiguous subset of LEDs, at a variable position from the animal. The movement of the treadmill was coupled to the movement of the landmark. As the mouse rotated the treadmill to move forward, the landmark approached the mouse. The mouse was then required to run and stop at the landmark, holding its position for 1.5 seconds to collect reward. If the mouse waited at the landmark for the required time, then it received a reward tone and a water reward simultaneously. If the mouse either ran to the end of the runway, bypassing the landmark, or failed to stop at the landmark within 30s, then it

received a miss tone. After the reward or miss tone, all the LEDs were turned off and a new trial started after 1s, with a landmark appearing.

We ensured that the distance over which the landmark position randomly varied is greater than the width of the landmark. As such, the animal could not have only relied on internally keeping track of the distance, and indeed the task is visually-guided (**Figure 1-B**). The task elicited an ON-OFF locomotion pattern (**Figure 1-C**), of which we carefully examined the stops.

The behavior suggests a sudden switch in locomotion state.

We then modeled the behavior of the animal in a single trial as an optimal-control problem. The trained mice were water-regulated, and we assumed that a trained mouse strives to collect the maximum amount of reward possible during the session, while its motivation is sustained. As the trials are of variable duration contingent on the animal's movement, maximizing reward is equivalent to minimizing time in a trial to collect reward. From an optimal-control perspective, we stated the problem as follows. Starting from an initial position away from the landmark, the mouse is tasked to pick a locomotor plan u_t (control policy) that dictates its locomotion pattern so as to minimize time to collect reward. The task was based only on positive reinforcement; the animal was not punished for miss trials. Furthermore, the animal was allowed to consider any locomotion trajectory to collect a reward, as long as it held its position at the landmark for 1.5 seconds. As such, there are no major constraints on the locomotor plan, excepts (a) how it affects the locomotion speed of the animal, and (b) that there is a maximal speed that the animal can achieve. For (a) we considered the simplest possible dynamics as a first-order ordinary differential equation between the speed of the animal v_t and the locomotor plan u_t, parametrized by a time-constant τ. For (b) we let the locomotor plan u_t be bounded between two values, u_{min} and u_{max}, and therefore had the speed v_t also be bounded. We further added boundary conditions, and the optimal-control problem was then set up as follows: (**Methods**)

$$
\begin{aligned}
\arg\min_{u_t} \quad & \mathrm{T} \\
\text{subject to:} \quad & \frac{dv_t}{dt} = -\frac{1}{\tau} v_t + u_t && (1) \\
& u_{\min} \leq u_t \leq u_{\max} && (2) \\
& v_T = 0 \\
& d_T = 0 \\
& d_0 = \text{initial position, with the landmark as origin}
\end{aligned}
$$

We then solved this problem to obtain the optimal solution. Using Lagrangian methods (Bertsekas et al., 1995), we solved for the optimal solution **(Methods)** and found it to be a bang-bang control policy: the optimal solution consists of having the locomotor plan be at u_{max} from the start of the trial up to a switching time, where it abruptly changes to u_{min}.

Briefly, to collect reward as soon as possible, the animal should accelerate as much as possible up to a switching point, which likely occurs before the animal arrives at the landmark, then suddenly brake as much as possible to arrive to a full halt at the landmark. This model gave us two features. First, this optimal solution suggests that there is an essential switching point in behavior. If the brain generates a signal to stop, it should occur around this switching point. Second, by fixing u_{min} and u_{max}, the stopping behavior can be parameterized by τ, a time constant, which reflects the speed of stopping.

From each session, we recovered time windows around all the stops at the landmark, which we termed landmark-stop windows. In each time window, we let the switching point correspond to the time point where the speed of the animal last peaked before stopping (**Methods**). We then aligned all the speed trajectories in the time windows along their switching point (**Figure 2-B**). The model fits well to the average trajectory and we can deduce an average time constant of about 100ms (**Figure 2-C**) (**Methods**). Importantly, having a small enough time constant is essential for the animal to stop in a timely manner at the landmark. For instance, if τ was equal to 1s, the animal will not stop in time and will miss the landmark (**Figure 2-C**)

Activating M2 axons in STN leads to stopping.

Our behavioral model indicates a switching point in behavior, and suggests that the brain generates a signal at that time. As we expect the signaling to be the result of an instantiated locomotor plan, frontal associative regions are good candidates to contain such plans. In fact, activity emanating from the right inferior frontal cortex and the pre-supplementary motor area in humans has generally preceded stops in stop-signal reaction tasks and go/no-go tasks (Aron and Poldrack, 2006; Eagle et al., 2008; Nachev et al., 2008; Swann et al., 2012). The pre-supplementary motor area is of particular interest. Although its homologous structure in mice is unclear, some of its properties have been often considered to coincide with the medial part of M2 (Barthas and Kwan, 2017). Additionally, the task is visually-guided and necessitates visuo-spatial processing. Medial M2 has been considered to be part of the visual-subnetwork (Barthas and Kwan, 2017; Zhang et al., 2016) in the mouse brain. This is supported by the dense projection of medial M2 to visual areas (Attinger et al., 2017; Zhang et al., 2016, 2014), and

reciprocal connections with visual cortex, retrosplenial cortex (Yamawaki et al., 2016) and the lateral posterior nucleus of the thalamus, among other brain structures, for inflow and outflow of visuo-spatial information. The BG also admit a functional subdivision—typically defined into motor, associative and limbic—which is preserved in the STN and the hyperdirect pathway (Alexander et al., 1986; Bolam et al., 2000; Hamani et al., 2004; Hintiryan et al., 2016; Hooks et al., 2018; Mandelbaum et al., 2019). STN receives projections from most of frontal cortex, and the medial M2 area then corresponds to the associative subdivision of the BG and STN, ideal for visuomotor transformation and integration.

The M2-STN pathway then offers a very quick route for information to achieve rapid stops. We hypothesized that if signals are sent, they can be sent from M2 to STN to achieve a fast control of locomotion. We thus asked: if activity is sent along the M2-STN pathway, does it lead the animal to stop locomoting?

To address this question, we injected an adeno-associated virus (AAV) expressing ChR2 under the CaMKII promoter in M2, and implanted ipsilaterally an optic-fiber above STN to target the M2 axons there (**Figure 3-A,B**). On a random subset of trials, we delivered a brief burst of blue (473nm) light (at 20Hz for 500ms) into the optic-fiber to activate the axons once the animal crossed the middle of the runway, while running towards the landmark (**Figure 3-A**). We found that whenever we activate the axons, the animal stops prematurely (**Figure 3-C,D**). This suggests that if the brain transmits a signal along the hyperdirect pathway right before stopping, it will causally trigger the animal to stop.

Stop activity is seen in M2-STN neurons on landmark stops but not mid-stops.

We next asked: is there activity at the onset of stopping in the M2-STN pathway? To address this, we expressed GCaMP6f in M2 neurons projecting to STN and imaged their calcium activity during behavior using two-photon microscopy **(Methods)**. As done for the behavioral model, for each session, we defined landmark-stop windows all aligned to the switching point. For each imaged neuron, we retrieved the activity during the landmark-stop windows and averaged all the activity. Hierarchical clustering of the average responses revealed a heterogeneity in the types of neural responses (**Figure 4-A**).

For a more principled clustering, we performed low-rank approximation and defined the low dimensional subspace explaining more than 80% of the energy in the neural population. Each neuronal response then corresponded to a weighted combination of the basis functions

(**Figure 4-B**) with an additional noise orthogonal to the defined subspace. These basis functions reflect a symmetry around the switching, suggesting a transition in neuronal responses around it. To then decompose our neuronal responses into responses related to the various temporal epochs around the stop, we performed a change of basis, by defining three types of 'ideal' neurons and projecting them onto the subspace **(Methods)**. The ideal pre-stop, stop and post-stop neurons correspond to having activity in only one epoch: before, on and after the switching (**Figure 4-C**). The average response of each neuron can then be expressed as a weighted combination of these three types, plus some additional noise. We then used the weights to cluster the neurons into three groups, keeping only the neurons where their response energy is more than 85% explained. We then recovered three groups of neurons, revealing a fraction that are active during stops (**Figure 4-D**).

Our analysis had, thus far, investigated locomotion stops that are executed at the landmark. Upon inspecting the behavior of the animal, we found spontaneous stops as the animal runs towards the landmark. We then collected these stops by defining spontaneous-stop windows (as done for landmark-stop windows) and aligning them to the switching point there. (**Figure 4-E**) We then retrieved the neural activity and computed the average responses for the stop neurons. We found that the M2 neurons are active only at stops at the landmark that are visually-guided but not at spontaneous stops (**Figure 4-F,G**).

Discussion

Our results suggest a signal that is sent down along cortico-subthalamic projections to halt locomotion. This signaling likely controls the tonic inhibition of SNr onto MLR activity. The mechanisms by which these signals can control MLR activity in a timely manner remains unresolved. STN activity can reach MLR through SNr, but also through direct STN-MLR projections. The interaction of these two pathways can prove to be key in enabling a quick control of locomotion.

Populations of neurons in the brainstem, notably in the medulla, have been found to stop ongoing locomotion (Bouvier et al., 2015; Capelli et al., 2017; Grätsch et al., 2019; Juvin et al., 2016). While such reticulo-spinal cells can be expected to be recruited in the context of our task, the nature of the signal is different. These impinge directly onto spinal circuits, and offer a low-level direct control of locomotion. However, the cortico-subthalamic signaling necessitates further processing to reach locomotor circuits, and the processing should be designed so that it

ensures fast locomotion halts. From an engineering perspective, feedback control can be leveraged to steer a system's trajectory as needed and ensured desired operation, which in our case consists of additionally ensuring fast responses. This principle is often realized by computing an error signal—reflecting the discrepancy between a reference trajectory and the current trajectory of the system—and employing it to oppose the deviation of the system from a desired operation (Aström and Murray, 2010). Our behavioral model implicates a locomotion plan as a reference signal, from which such error signals can potentially be derived. Error signals have been reported widely in the brain, and particularly throughout expectation perturbations such as visuomotor mismatch (Attinger et al., 2017; Heindorf et al., 2018; Marple-Horvat et al., 1993), as signaling prediction errors. The signals we elucidate could also be of a similar 'error' nature. The mechanisms by which they arise and how they are employed to ensure fast control are subject of further research.

Stopping in such a task is typically considered to recruit proactive inhibition as opposed to reactive inhibition. Indeed, the animal can see the landmark as it approaches, and can prepare to stop. The situation then lends itself to a potential locomotor plan that is implemented without being interrupted. However, the role of the hyperdirect pathway has been extensively studied in reactive settings, in which it is found to be engaged (Aron et al., 2003; Eagle et al., 2008; Nachev et al., 2008, 2007). In reactive stop-signal reaction tasks or go/no-go tasks, participants are signaled to immediately halt an ongoing (or to be initiated) action. Proactive inhibition is considered to be heavily mediated by the indirect pathway, but there is certainly evidence of the hyperdirect pathway, and more generally the same 'stopping network', as equally being involved (Aron, 2011; Meyer and Bucci, 2016).

In addition to the indirect pathway in the basal ganglia, alternative routes for locomotion halts signaling might include a visual tectal pathway, consisting of direct projections from the superior colliculus to the mesencephalic locomotor regions (Roseberry et al., 2016). We expect such a signal to be more engaged upon sudden flashes of the landmark, likely in a more reactive setting, perhaps engaging circuits typically used for startle responses (Liang et al., 2015). Differently, glycinergic neurons in the pontine reticular formation (PRF) project to the intralaminar thalamic nuclei (IL), and stimulation of their axons in IL produced behavioral arrest (Giber et al., 2015). M2 directly projects to PRF, and the projections might directly drive such glycinergic cells to achieve fast locomotion halts. Aside from the described additional pathways, various others inducing behavioral arrest can be found (Klemm, 2001; Roseberry and Kreitzer, 2017). These are generally recruited via different mechanisms, though some can overlap with ours. Further tasks enhancements are needed to elucidate their contributions.

Acknowledgment

This work was supported by: National Institute of Neurological Disease and Stroke, U01 NS090473 (M.S.), National Eye Institute, R01 EY007023 (M.S.) and a JPB Foundation Picower Fellowship (E.A.).

Author Contributions

E.A. and M.S. conceived the project and developed the concepts presented. E.A. designed the experiments and performed them with assistance from T.J.. E.A. analyzed the data and developed the theory. E.A. and M.S. wrote the manuscript with input from T.J..

Competing Interests

The authors declare no competing interests.

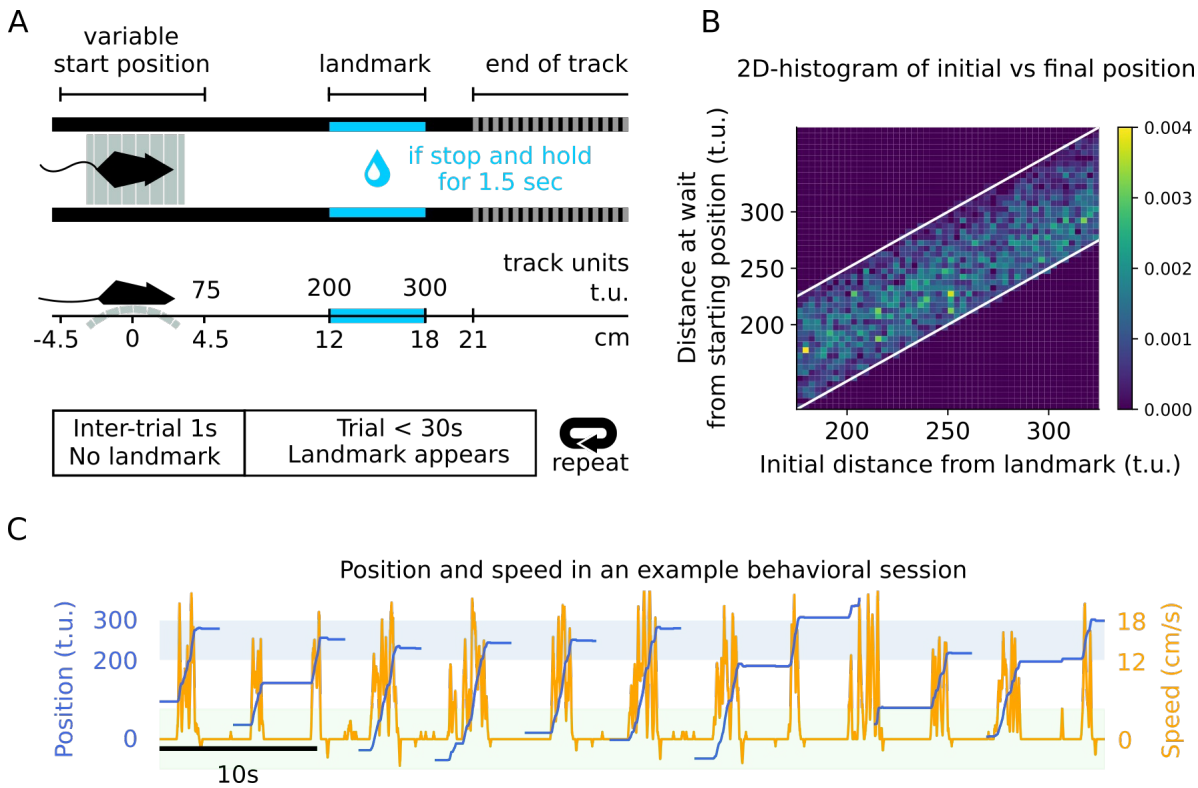

Figure 1: We trained mice to run, stop and wait at visual landmarks to collect reward.

(A) Head-fixed mice were trained to run towards a landmark and wait there to collect reward. A trial started by the appearance of a landmark at a variable distance. If the mouse waited at the landmark for 1.5s within 30s of the start of the trial, the mouse received an auditory reward tone followed by a drop of water. If the mouse failed to wait at the landmark for 1.5s, for example, by reaching the end of the track, waiting after the landmark, or never reaching the landmark in time, the mouse received an auditory white noise. Position is defined in terms of track units (t.u.) with 200t.u. corresponding to 12cm. **(B)** The distance of the landmark from the initial position does not affect the animals' final stop position. The animal is thus using the visual cues, instead of relying on other mechanisms. (N=10 mice, 3 sessions each) **(C)** The task enables an ON-OFF locomotor pattern, where both initiation and termination appear to be sudden.

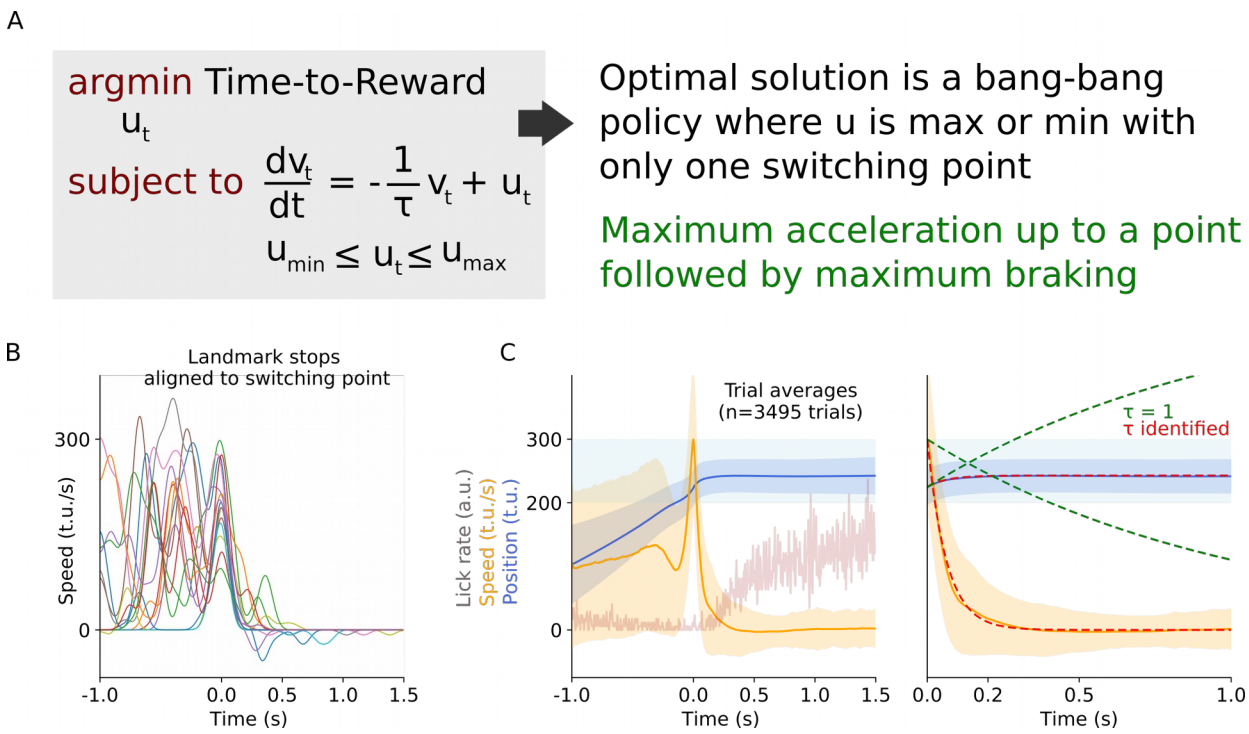

Figure 2: The behavior suggests a sudden switch in locomotion state.

(A) We cast the behavior as an optimal-control problem. The mouse is tasked to pick a locomotor plan (u_t) that minimizes the time required to collect reward. This plan was mainly subject to two constraints. First, the locomotor plan dictates the speed of the animal, and the relation is governed by a first-order ordinary differential equation, parametrized by a time-constant τ. Second, the locomotor plan, and therefore the speed, are bounded and cannot be infinite. The optimal solution for any given initial condition (i.e., velocity and position of the animal) is a bang-bang control policy: the locomotor plan u_t at any time is either equal to u_{min} or u_{max}, and can admit at most one switching point in value. The optimal locomotor plan corresponds to enforcing maximum acceleration up to a switching point after which it enforces maximum braking to halt in time. **(B)** During a trial, whenever the animal stops at the landmark, we corresponded the switching time point to the time point of the last major velocity peak before stopping. From the behavioral traces, we then collected time-windows around the switching points, 1s before and 1.5s after, that we termed landmark-stop windows. The graph shows examples of speed traces in landmark-stop windows aligned to the switching point. **(C)** The graph shows the average speed and position around a switching point. The shaded regions correspond to the standard-deviation of the sampled distribution. The licking rate shows that animals began licking as soon they stopped, likely in anticipation of the upcoming reward. We

can identify the time constant τ to fit the data. Importantly, τ needs to be small enough to ensure a quick response. For example, if τ is set to 1s, then the animal will not stop fast enough and will miss the landmark. (N=10 mice, 3 sessions each)

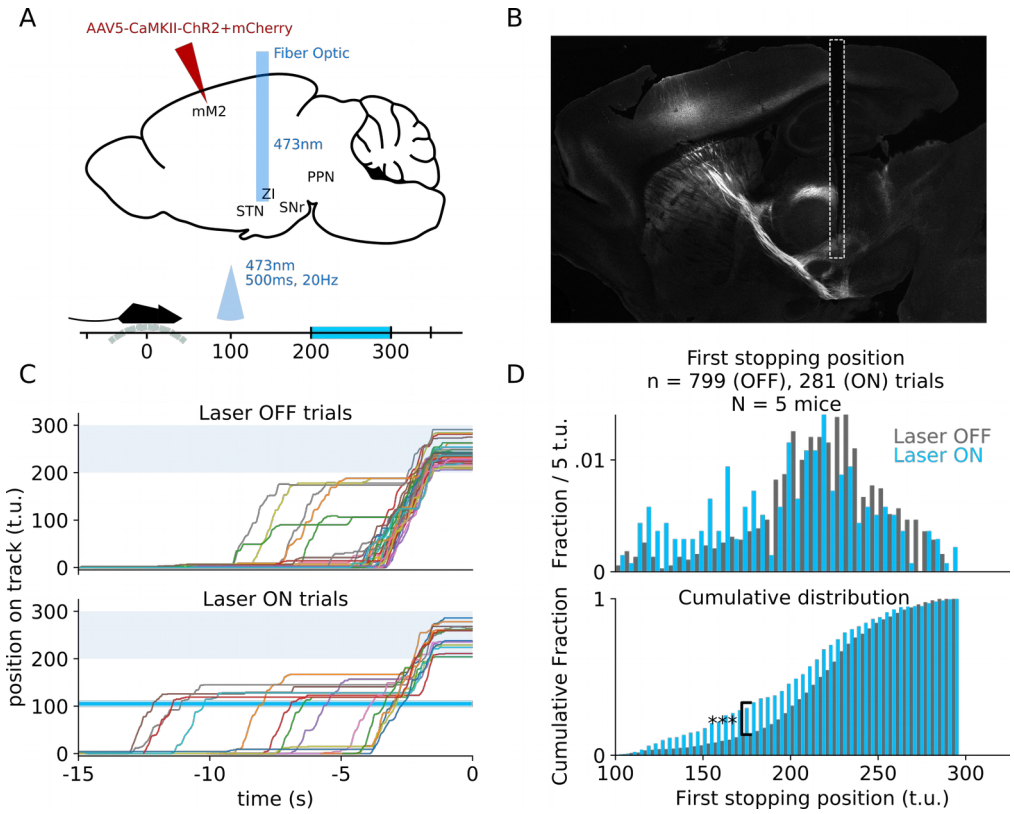

Figure 3: Activating M2 axons in STN leads to stopping.

(A) An AAV virus expressing ChR2 was unilaterally injected in M2 of wild-type mice (N=5) and an optic fiber implanted over STN (ipsilateral to the injection site) to optogenetically target M2 axons in STN. On pseudorandom trials, we shine blue light (473nm) for 500ms at 20Hz as soon as the animal crossed track position 100. **(B)** M2 projects directly to STN via the hyperdirect pathway. The image shows the projections and a fiber-optic placement. **(C)** On laser trials, the animal stopped soon after triggering the laser, before arriving to the landmark. **(D)** The plots show the distribution of the first position the animal stops at, after position 100. We observe a shift in the distribution, towards position 100. This indicates that on laser trials, the animal is stopping prematurely (Kolmogorov–Smirnov test, $p<0.001$).

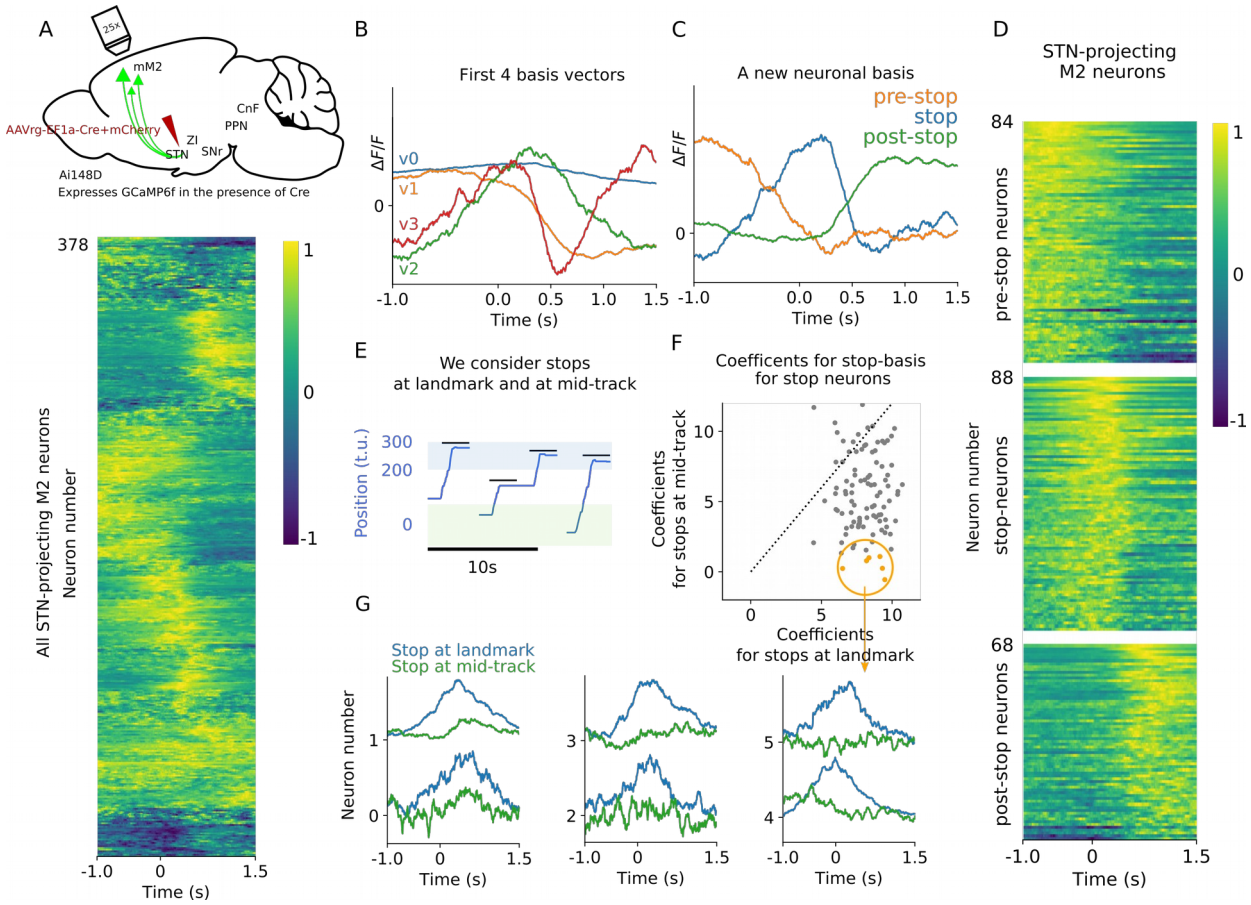

Figure 4: Stop activity is seen in STN-projecting M2 neurons at landmark stops but not at mid-stops.

(A) We injected an AVV in STN of Ai148D mice (N=4) to retrogradely express Cre, and thereby GCaMP6f in neurons projecting there. The plot shows the normalized average calcium activity (DFF) of the labeled M2 neurons, within landmark-stop windows aligned to the switching point. Applying hierarchical clustering reveals a heterogeneity in the dynamics, with neurons selectively active at different epochs around the stop. **(B)** The mean DFF of all the neurons were then pooled together into one data matrix. A low-rank decomposition yields a basis for a four-dimensional subspace that captures most of the energy in the responses. The basis functions show a symmetry, reflecting a potential switch in dynamics. **(C)** To decompose the calcium signals into physiological responses related to various epochs of the stop, we performed a change of basis, and derived a three-dimensional subspace spanned by 3 templates of ideal neurons, that are active before stops (pre-stop), during stops (stop) and after stops (post-stop). Each M2 neuron can then be written as a weighted combination of these 3 templates, with an

additional noise component orthogonal to the subspace spanned by the templates. **(D)** We then clustered all the neurons with respect to those templates using their respective weights, keeping only the ones that have most of their energy in the subspace. We find population of neurons that are active pre-stop, at stop and post-stop with somewhat similar fractions. **(E)** In addition to visually-guided stops, constituting the landmark-stop windows, we investigated spontaneous stops. We then defined spontaneous-stop windows to be 1s before and 1.5s after a switching point (as defined for landmark-stop windows) that occur in the middle of the track and not near the landmark. **(F)** For each neuron, we recovered the DFF during the spontaneous-stop windows, aligned them to the switching point and averaged their activity. The scatter plot presents the activity weight that the stop neuron template contributes to landmark-stops versus spontaneous-stops activity. We find that the landmark stops have higher coefficients, indicating a higher activity in M2 at stops that are performed at the landmark than stops that are performed spontaneously in the middle of the track. **(G)** The graphs show examples of the DFF activity of six M2 neurons projecting to STN. The response for spontaneous-stops is normalized to the maximum value of the corresponding response for landmark-stops. We find that they are highly active on landmark-stops but are quiet on spontaneous-stops. This suggests that the activity sent from M2 to STN is a visually-guided activity.

References

Alexander, G.E., DeLong, M.R., Strick, P.L., 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annu. Rev. Neurosci. 9, 357–381. https://doi.org/10.1146/annurev.ne.09.030186.002041

Arber, S., Costa, R.M., 2018. Connecting neuronal circuits for movement. Science 360, 1403–1404. https://doi.org/10.1126/science.aat5994

Aron, A.R., 2011. From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. Biol. Psychiatry 69, e55-68. https://doi.org/10.1016/j.biopsych.2010.07.024

Aron, A.R., Fletcher, P.C., Bullmore, E.T., Sahakian, B.J., Robbins, T.W., 2003. Stop-signal inhibition disrupted by damage to right inferior frontal gyrus in humans. Nat. Neurosci. 6, 115–116. https://doi.org/10.1038/nn1003

Aron, A.R., Herz, D.M., Brown, P., Forstmann, B.U., Zaghloul, K., 2016. Frontosubthalamic Circuits for Control of Action and Cognition. J. Neurosci. 36, 11489–11495. https://doi.org/10.1523/JNEUROSCI.2348-16.2016

Aron, A.R., Poldrack, R.A., 2006. Cortical and Subcortical Contributions to Stop Signal Response Inhibition: Role of the Subthalamic Nucleus. J. Neurosci. 26, 2424–2433. https://doi.org/10.1523/JNEUROSCI.4682-05.2006

Aström, K.J., Murray, R.M., 2010. Feedback systems: an introduction for scientists and engineers. Princeton university press.

Attinger, A., Wang, B., Keller, G.B., 2017. Visuomotor Coupling Shapes the Functional Development of Mouse Visual Cortex. Cell 169, 1291–1302.e14. https://doi.org/10.1016/j.cell.2017.05.023

Barthas, F., Kwan, A.C., 2017. Secondary Motor Cortex: Where ‘Sensory’ Meets ‘Motor’ in the Rodent Frontal Cortex. Trends Neurosci. 40, 181–193. https://doi.org/10.1016/j.tins.2016.11.006

Baunez, C., Robbins, T.W., 1997. Bilateral Lesions of the Subthalamic Nucleus Induce Multiple Deficits in an Attentional Task in Rats. Eur. J. Neurosci. 9, 2086–2099. https://doi.org/10.1111/j.1460-9568.1997.tb01376.x

Bertsekas, D.P., Bertsekas, D.P., Bertsekas, D.P., Bertsekas, D.P., 1995. Dynamic programming and optimal control. Athena scientific Belmont, MA.

Bolam, J.P., Hanley, J.J., Booth, P. a. C., Bevan, M.D., 2000. Synaptic organisation of the basal ganglia. J. Anat. 196, 527–542. https://doi.org/10.1046/j.1469-7580.2000.19640527.x

Bouvier, J., Caggiano, V., Leiras, R., Caldeira, V., Bellardita, C., Balueva, K., Fuchs, A., Kiehn, O., 2015. Descending Command Neurons in the Brainstem that Halt Locomotion. Cell 163, 1191–1203. https://doi.org/10.1016/j.cell.2015.10.074

Caggiano, V., Leiras, R., Goñi-Erro, H., Masini, D., Bellardita, C., Bouvier, J., Caldeira, V., Fisone, G., Kiehn, O., 2018. Midbrain circuits that set locomotor speed and gait selection. Nature 553, 455–460. https://doi.org/10.1038/nature25448

Capelli, P., Pivetta, C., Soledad Esposito, M., Arber, S., 2017. Locomotor speed control circuits in the caudal brainstem. Nature 551, 373–377. https://doi.org/10.1038/nature24064

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., Shenoy, K.V., 2010. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? Neuron 68, 387–400. https://doi.org/10.1016/j.neuron.2010.09.015

Dahleh, M., Dahleh, M.A., Verghese, G., 2004. Lectures on Dynamic Systems and Control.

Drew, T., Prentice, S., Schepens, B., 2004. Cortical and brainstem control of locomotion. Prog. Brain Res. 143, 251–261. https://doi.org/10.1016/S0079-6123(03)43025-2

Eagle, D.M., Baunez, C., Hutcheson, D.M., Lehmann, O., Shah, A.P., Robbins, T.W., 2008. Stop-Signal Reaction-Time Task Performance: Role of Prefrontal Cortex and Subthalamic Nucleus. Cereb. Cortex 18, 178–188. https://doi.org/10.1093/cercor/bhm044

Economo, M.N., Viswanathan, S., Tasic, B., Bas, E., Winnubst, J., Menon, V., Graybuck, L.T., Nguyen, T.N., Smith, K.A., Yao, Z., Wang, L., Gerfen, C.R., Chandrashekar, J., Zeng, H., Looger, L.L., Svoboda, K., 2018. Distinct descending motor cortex pathways and their roles in movement. Nature 563, 79. https://doi.org/10.1038/s41586-018-0642-9

Ferreira-Pinto, M.J., Ruder, L., Capelli, P., Arber, S., 2018. Connecting Circuits for Supraspinal Control of Locomotion. Neuron 100, 361–374. https://doi.org/10.1016/j.neuron.2018.09.015

Fife, K.H., Gutierrez-Reed, N.A., Zell, V., Bailly, J., Lewis, C.M., Aron, A.R., Hnasko, T.S., 2017. Causal role for the subthalamic nucleus in interrupting behavior. eLife 6, e27689. https://doi.org/10.7554/eLife.27689

Freeze, B.S., Kravitz, A.V., Hammack, N., Berke, J.D., Kreitzer, A.C., 2013. Control of Basal Ganglia Output by Direct and Indirect Pathway Projection Neurons. J. Neurosci. 33, 18531–18539. https://doi.org/10.1523/JNEUROSCI.1278-13.2013

Giber, K., Diana, M.A., M Plattner, V., Dugué, G.P., Bokor, H., Rousseau, C.V., Maglóczky, Z., Havas, L., Hangya, B., Wildner, H., Zeilhofer, H.U., Dieudonné, S., Acsády, L., 2015. A subcortical inhibitory signal for behavioral arrest in the thalamus. Nat. Neurosci. 18, 562–568. https://doi.org/10.1038/nn.3951

Goulding, M., 2009. Circuits controlling vertebrate locomotion: moving in a new direction. Nat. Rev. Neurosci. 10, 507–518. https://doi.org/10.1038/nrn2608

Grätsch, S., Auclair, F., Demers, O., Auguste, E., Hanna, A., Büschges, A., Dubuc, R., 2019. A Brainstem Neural Substrate for Stopping Locomotion. J. Neurosci. 39, 1044–1057. https://doi.org/10.1523/JNEUROSCI.1992-18.2018

Graybiel, A.M., 2000. The basal ganglia. Curr. Biol. 10, R509–R511. https://doi.org/10.1016/S0960-9822(00)00593-5

Grillner, S., 2003. The motor infrastructure: from ion channels to neuronal networks. Nat. Rev. Neurosci. 4, 573–586. https://doi.org/10.1038/nrn1137

Grillner, S., Wallén, P., Saitoh, K., Kozlov, A., Robertson, B., 2008. Neural bases of goal-directed locomotion in vertebrates—An overview. Brain Res. Rev., Networks in Motion 57, 2–12. https://doi.org/10.1016/j.brainresrev.2007.06.027

Hamani, C., Saint-Cyr, J.A., Fraser, J., Kaplitt, M., Lozano, A.M., 2004. The subthalamic nucleus in the context of movement disorders. Brain 127, 4–20. https://doi.org/10.1093/brain/awh029

Heindorf, M., Arber, S., Keller, G.B., 2018. Mouse Motor Cortex Coordinates the Behavioral Response to Unpredicted Sensory Feedback. Neuron 99, 1040–1054.e5. https://doi.org/10.1016/j.neuron.2018.07.046

Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., Yang, X.W., Shih, J.C., Toga, A.W., Dong, H.-W., 2016. The mouse cortico-striatal projectome. Nat. Neurosci. 19, 1100–1114. https://doi.org/10.1038/nn.4332

Hooks, B.M., Papale, A.E., Paletzki, R.F., Feroze, M.W., Eastwood, B.S., Couey, J.J., Winnubst, J., Chandrashekar, J., Gerfen, C.R., 2018. Topographic precision in sensory and motor corticostriatal projections varies across cell type and cortical area. Nat. Commun. 9, 3549. https://doi.org/10.1038/s41467-018-05780-7

Josset, N., Roussel, M., Lemieux, M., Lafrance-Zoubga, D., Rastqar, A., Bretzner, F., 2018. Distinct Contributions of Mesencephalic Locomotor Region Nuclei to Locomotor Control in the Freely Behaving Mouse. Curr. Biol. 28, 884–901.e3. https://doi.org/10.1016/j.cub.2018.02.007

Juvin, L., Grätsch, S., Trillaud-Doppia, E., Gariépy, J.-F., Büschges, A., Dubuc, R., 2016. A Specific Population of Reticulospinal Neurons Controls the Termination of Locomotion. Cell Rep. 15, 2377–2386. https://doi.org/10.1016/j.celrep.2016.05.029

Kiehn, O., 2016. Decoding the organization of spinal circuits that control locomotion. Nat. Rev. Neurosci. 17, 224–238. https://doi.org/10.1038/nrn.2016.9

Kiehn, O., 2006. Locomotor circuits in the mammalian spinal cord. Annu. Rev. Neurosci. 29, 279–306. https://doi.org/10.1146/annurev.neuro.29.051605.112910

Kim, L.H., Sharma, S., Sharples, S.A., Mayr, K.A., Kwok, C.H.T., Whelan, P.J., 2017. Integration of Descending Command Systems for the Generation of Context-Specific Locomotor Behaviors. Front. Neurosci. 11. https://doi.org/10.3389/fnins.2017.00581

Klemm, W.R., 2001. Behavioral arrest: in search of the neural control system. Prog. Neurobiol. 65, 453–471. https://doi.org/10.1016/S0301-0082(01)00016-8

Liang, F., Xiong, X.R., Zingg, B., Ji, X., Zhang, L.I., Tao, H.W., 2015. Sensory Cortical Control of a Visually Induced Arrest Behavior via Corticotectal Projections. Neuron 86, 755–767. https://doi.org/10.1016/j.neuron.2015.03.048

Liu, D., Ma, C., Zheng, W., Yao, Y., Dan, Y., 2018. Sleep and Motor Control by a Basal Ganglia Circuit. bioRxiv 405324. https://doi.org/10.1101/405324

Mandelbaum, G., Taranda, J., Haynes, T.M., Hochbaum, D.R., Huang, K.W., Hyun, M., Umadevi Venkataraju, K., Straub, C., Wang, W., Robertson, K., Osten, P., Sabatini, B.L., 2019. Distinct Cortical-Thalamic-Striatal Circuits through the Parafascicular Nucleus. Neuron 102, 636–652.e7. https://doi.org/10.1016/j.neuron.2019.02.035

Marple-Horvat, D.E., Amos, A.J., Armstrong, D.M., Criado, J.M., 1993. Changes in the discharge patterns of cat motor cortex neurones during unexpected perturbations of on-going locomotion. J. Physiol. 462, 87–113. https://doi.org/10.1113/jphysiol.1993.sp019545

Meyer, H.C., Bucci, D.J., 2016. Neural and behavioral mechanisms of proactive and reactive inhibition. Learn. Mem. Cold Spring Harb. N 23, 504–514. https://doi.org/10.1101/lm.040501.115

Nachev, P., Kennard, C., Husain, M., 2008. Functional role of the supplementary and pre-supplementary motor areas. Nat. Rev. Neurosci. 9, 856–869. https://doi.org/10.1038/nrn2478

Nachev, P., Wydell, H., O'Neill, K., Husain, M., Kennard, C., 2007. The role of the pre-supplementary motor area in the control of action. Neuroimage 36, T155–T163. https://doi.org/10.1016/j.neuroimage.2007.03.034

Nambu, A., Tokuno, H., Hamada, I., Kita, H., Imanishi, M., Akazawa, T., Ikeuchi, Y., Hasegawa, N., 2000. Excitatory Cortical Inputs to Pallidal Neurons Via the Subthalamic Nucleus in the Monkey. J. Neurophysiol. 84, 289–300. https://doi.org/10.1152/jn.2000.84.1.289

Nambu, A., Tokuno, H., Takada, M., 2002. Functional significance of the cortico–subthalamo–pallidal 'hyperdirect' pathway. Neurosci. Res. 43, 111–117. https://doi.org/10.1016/S0168-0102(02)00027-5

Oppenheim, A.V., Willsky, A.S., Nawab, S.H., 1996. Signals & Systems (2Nd Ed.). Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

Roseberry, T., Kreitzer, A., 2017. Neural circuitry for behavioural arrest. Philos. Trans. R. Soc. Lond. B. Biol. Sci. 372. https://doi.org/10.1098/rstb.2016.0197

Roseberry, T.K., Lee, A.M., Lalive, A.L., Wilbrecht, L., Bonci, A., Kreitzer, A.C., 2016. Cell-Type-Specific Control of Brainstem Locomotor Circuits by Basal Ganglia. Cell 164, 526–537. https://doi.org/10.1016/j.cell.2015.12.037

Ryczko, D., Dubuc, R., 2013. The multifunctional mesencephalic locomotor region. Curr. Pharm. Des. 19, 4448–4470. https://doi.org/10.2174/1381612811319240011

Schmidt, R., Leventhal, D.K., Mallet, N., Chen, F., Berke, J.D., 2013. Canceling actions involves a race between basal ganglia pathways. Nat. Neurosci. 16, 1118–1124. https://doi.org/10.1038/nn.3456

Schweizer, N., Pupe, S., Arvidsson, E., Nordenankar, K., Smith-Anttila, C.J.A., Mahmoudi, S., Andrén, A., Dumas, S., Rajagopalan, A., Lévesque, D., Leão, R.N., Wallén-Mackenzie, Å., 2014. Limiting glutamate transmission in a Vglut2-expressing subpopulation of the subthalamic nucleus is sufficient to cause hyperlocomotion. Proc. Natl. Acad. Sci. 111, 7837–7842. https://doi.org/10.1073/pnas.1323499111

Shenoy, K.V., Sahani, M., Churchland, M.M., 2013. Cortical Control of Arm Movements: A Dynamical Systems Perspective. Annu. Rev. Neurosci. 36, 337–359. https://doi.org/10.1146/annurev-neuro-062111-150509

Svoboda, K., Li, N., 2018. Neural mechanisms of movement planning: motor cortex and beyond. Curr. Opin. Neurobiol., Neurobiology of Behavior 49, 33–41. https://doi.org/10.1016/j.conb.2017.10.023

Swann, N.C., Cai, W., Conner, C.R., Pieters, T.A., Claffey, M.P., George, J.S., Aron, A.R., Tandon, N., 2012. "Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: electrophysiological responses and functional and structural connectivity." NeuroImage 59, 2860–2870. https://doi.org/10.1016/j.neuroimage.2011.09.049

Uslaner, J.M., Robinson, T.E., 2006. Subthalamic nucleus lesions increase impulsive action and decrease impulsive choice − mediation by enhanced incentive motivation? Eur. J. Neurosci. 24, 2345–2354. https://doi.org/10.1111/j.1460-9568.2006.05117.x

Wessel, J.R., Aron, A.R., 2017. On the Globality of Motor Suppression: Unexpected Events and Their Influence on Behavior and Cognition. Neuron 93, 259–280. https://doi.org/10.1016/j.neuron.2016.12.013

Wong, A.L., Haith, A.M., Krakauer, J.W., 2015. Motor Planning. Neurosci. Rev. J. Bringing Neurobiol. Neurol. Psychiatry 21, 385–398. https://doi.org/10.1177/1073858414541484

Yamawaki, N., Radulovic, J., Shepherd, G.M.G., 2016. A Corticocortical Circuit Directly Links Retrosplenial Cortex to M2 in the Mouse. J. Neurosci. 36, 9365–9374. https://doi.org/10.1523/JNEUROSCI.1099-16.2016

Zhang, S., Xu, M., Chang, W.-C., Ma, C., Hoang Do, J.P., Jeong, D., Lei, T., Fan, J.L., Dan, Y., 2016. Organization of long-range inputs and outputs of frontal cortex for top-down control. Nat. Neurosci. 19, 1733–1742. https://doi.org/10.1038/nn.4417

Zhang, S., Xu, M., Kamigaki, T., Do, J.P.H., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L., Dan, Y., 2014. Long-range and local circuits for top-down modulation of visual cortex processing. Science 345, 660–665. https://doi.org/10.1126/science.1254126

Methods

Animals

All procedures were approved by the Massachusetts Institute of Technology's Animal Care and Use Committee and conformed to the National Institutes of Health guidelines. Adult mice (>2 months old) from a C57BL/6J background were used in this study. Male or female mice were randomly selected for each experiment. We also used the Ai148D (Ai148(TIT2L-GC6f-ICL-tTA2)-D, Jackson Laboratory) mouse line for the two-photon imaging experiment.

Surgery

The preparation of the animal is common to all surgical procedures. The procedures were performed under isoflurane anesthesia while maintaining body temperature at 37.5 °C using an animal temperature controller (ATC2000, World Precision Instruments). After deep anesthesia was confirmed, mice were placed in a stereotaxic frame (51725D, Stoelting), scalp hairs were removed with hair-remover cream, the underlying skin was cleaned with 70% alcohol and betadine, and an incision was made in the scalp. The conjunctive tissue was removed by rubbing hydrogen peroxide on the skull. The skull was positioned such that the lambda and bregma marks were aligned on the anteroposterior and dorsoventral axes. The skull was further leveled along the mediolateral axis. Animals were given analgesia (slow release buprenex, 0.1mg/kg and Meloxicam 0.1mg/kg) before surgery and their recovery was monitored daily for 72 hours.

For viral injections, a small hole was drilled through the skull at the location of interest. Viruses were delivered with a thin glass pipette at a rate of 75 nL/min by an infuser system (QSI 53311, Stoelting). The following viruses (titer: ~10e-12 virus genomes per ml) were injected for optogenetics and imaging experiments, respectively: AAV5-CamKII-ChR2-mCherry (pAAV-CamKIIa-hChR2-(H134R)-mCherry-WPRE-pA, UNC Vector Core) AAVrg-EF1a-Cre-mCherry (pAAV-Ef1a-mCherry-IRES-Cre, Addgene). For optogenetic experiments, after injection, the skin was sutured and we let mice recover for 4–6 weeks for optimal opsin expression and then performed another surgical procedure to implant an optic fiber. For imaging experiments, a cranial window was placed during the same procedure as the injection.

For optogenetics experiments, the optic fiber cannula (Thorlabs) was held by a stereotaxic manipulator and inserted at the needed position. The cannula was attached to the skull with dental cement (C&B Metabond, Parkell). To avoid light reflection and absorption, dental cement was mixed with black ink pigment (Black Iron Oxide 18727, Schmincke). A custom-designed headplate (eMachineShop) was implanted at the end of the surgery for head fixation.

Two-photon calcium imaging was performed through a cranial window. We drilled a 3-mm circular window centered over the midline (AP: 1mm) to expose M2 on each hemisphere. We stacked two 3-mm coverslips centered on a 5-mm coverslip (CS-5R and CS-3R, Warner Instruments) and glued the three together with ultraviolet adhesive (NOA 61 UV adhesive, Norland Products). The fabricated window was positioned over the craniotomy and attached to the skull using dental cement (C&B Metabond, Parkell). The dental cement was mixed with black ink pigment (Black Iron Oxide 18727, Schmincke) to block light leakage during imaging. We then attached a custom-designed headplate (eMachineShop) to the skull for head fixation.

Behavioral equipment and training.

Mice were head-fixed using optical hardware (Thorlabs) and positioned on a custom-fabricated rubber treadmill (LEGO Technic + M-D Building Products) coupled with a rotary encoder (Signswise). The virtual runway was constructed using two parallel rails (Thorlabs) each equipped with high-density LED PCB bars (DotStar, Adafruit). Reward consisted of 5-10ul of water and was delivered through a lick spout using a solenoid valve (Parker). Licks were collected, when needed, using a capacitive touch sensor (Adafruit). The behavioral apparatus was controlled through custom-written code deployed on a microcontroller (Arduino). The microcontroller interfaced, through serial connection, with custom-written scripts (Python) running on a custom-designed multi-purpose computer (operating under GNU/Linux) to execute the task.

The landmark consisted a contiguous set of LEDs that are lit blue, on both rails. The length of the landmark amounted to 6cm, and the position of the animal in track units (t.u.) was referenced to the mouse’s nose. Specifically, the landmark interval 200-300 corresponded to the mouse’s noise being within the lit LED range. Track units (t.u.) were calibrated to have 200t.u. correspond to 12cm. The rotation of the treadmill was coupled to the movement of the landmark so that the linear velocity of the outer-edge of the treadmill equaled the linear velocity of the landmark.

Mice underwent water-regulation and obtained water reward during behavior. They were then habituated to the treadmill, and underwent a shaping procedure that rewards a stop after crossing a certain distance. The landmark was then introduced and the required waiting time to collect reward was gradually increased across days from 0.6s to 1.5s until they reached their full performance. Experiments began once the number of successful stops at the landmark was consistently above 100 within a 30 minute session.

Behavioral Model

The behavior of the animal in a single trial was modeled as a minimum-time optimal-control problem, of the following form:

$$\begin{aligned}
\arg\min_{u_t} \quad & \mathrm{T} \\
\text{subject to:} \quad & \frac{dv_t}{dt} = -\frac{1}{\tau} v_t + u_t && (1) \\
& u_{\min} \leq u_t \leq u_{\max} && (2) \\
& v_T = 0 \\
& d_T = 0 \\
& d_0 = \text{initial position, with the landmark as origin}
\end{aligned}$$

For all fixed values of $u_{min} < 0$ and $u_{max} > 0$, this problem admits an optimal solution for any starting point. The solution can be arrived to by construction the Hamilton-Jacobi-Bellman (HJB) equation and solving it. The optimal solution will be a bang-bang policy of control. For instance, if $d_0 = -250$, namely the mouse is starting at position 250 away from the landmark, then the optimal solution consists of forcing u_t equals u_{max} till the animal reaches a switching time point, then u_t becomes u_{min} so that the animal stops at position 0, taking the middle of the landmark as 0. For a linear system, as is the case here, the solution is generally known to be bang-bang (Bertsekas, 1995). As the linear system is first-order, the solution is known to admit at most one switching point (Bertsekas, 1995):Example 3.4.3.

From each behavioral session we defined landmark-stop windows (rewarded, stops between positions 200-300). Each time window corresponds to 2.5s, aligned to a switching point. The switching point was defined as the last major velocity peak (above 25% of the maximum velocity within 200ms window before sustained zero velocity). The time window then consisted of 1s before the switching point, and 1.5s after. We pooled together all landmarks-stop windows, across sessions and animals, and averaged them to get the average traces, presented in (**Figure 2-C**). We then fixed $u_{min} = 0$ to approximate the problem, and then derived

τ and u_{max} from the peak velocity of the animal before stopping and the slope of the logarithm of the velocity trace during decay.

Optogenetics manipulations

Wild-type mice were unilaterally (left hemisphere) injected with 400nL of AAV5-CaMKII-ChR2-mCherry in M2 (centered at AP: +1mm, ML: +0.5mm, DV: +0.5mm). A fiber-optic cannula (Thorlabs) with 200um diameter (0.39NA) was implanted above STN (AP: -1.9mm, ML: +1.5mm, DV: +4.3mm).

Blue light (473nm) was delivered using a diode-pumped solid-state laser (Optoengine). Laser stimulation was triggered using a custom-designed source-follower circuit driven by a microcontroller (Arduino) dictating the stimulation pattern. A fiber-optic patch cable with a ferrule end (200um, 0.39NA) was coupled to the implanted fiber optic cannula with a ferrule mating sleeve (Thorlabs). A piece of black electrical tape was wrapped around the connection between the patch cable and the implanted ferrule, to block any light emitted from that interface.

Each trial had a 30% chance of being selected for photo-stimulation. We further imposed a condition that no two consecutive trials could be selected for laser stimulation. On a trial selected for photo-stimulation, light was delivered once the animal reached position 100, for 500ms at 20Hz, 20% duty-cycle (PW:10ms, and T:50ms) with a peak power of 10-15mw (average power of 2-2.5mw). Each animal underwent three 30min behavioral sessions of photo-stimulation. We kept the session where the behavioral performance was deemed adequate (> 80 hits) with some exception sessions of ~50 hits.

Calcium imaging and neuronal response analysis

Ai148D mice were bilaterally injected with 300nL of AAVrg-EF1a-Cre-mCherry in STN (AP: -1.9mm, ML: +1.5mm, DV: +4.6mm) and a cranial window centered at (AP: 1mm, ML: 0mm) was implanted above M2 in both hemispheres.

GCaMP6f fluorescence was imaged through a 25x/1.05NA objective (Olympus) using a custom-configured two-photon microscope (MOM, Sutter Systems). Excitation light at 910nm was delivered with a Ti:Sapphire laser (Mai-Tai eHP, Spectra-Physics) equipped with dispersion compensation (DeepSee, Spectra-Physics). Emitted light was bandpass filtered and collected with a GaAsP photomultiplier tube (Hamamatsu). STN projecting neurons were imaged between 400-500um below the surface at 10Hz using galvo scanning, and images were acquired by

ScanImage (Vidrio) to generate a TIFF stack. Power at the objective ranged from 15 to 30 mW depending on GCaMP6f expression level and depth.

Neuronal ROI selection and calcium signal extraction was performed using CaImAn (Flatiron Institute) (Giovannucci et al., 2019) implementing a constrained nonnegative matrix factorization approach (Pnevmatikakis et al., 2016). The obtained ROIs and signals were additionally hand-curated to leave out any false-positives.

The animals performed 30min behavioral sessions, and fluorescence was acquired for 1600s (~26.6 mins) during that period. The behavioral data included a reference signal derived from the microscope acquisition trigger signal, and that signal was used to align to the behavioral signals to the neural signals. The neural signals were then upsampled using piecewise-constant interpolation to match the temporal resolution of 200Hz of the behavior.

From each behavioral session we defined landmark-stop windows (rewarded, stops between positions 200-300) and spontaneous-stop windows (non-rewarded, stops before position 150). Each time window corresponds to 2.5s, aligned to a switching point. The switching point was defined as the last major velocity peak (above 25% of the maximum velocity within 200ms window before sustained zero velocity). The time window then consisted of 1s before the switching point, and 1.5s after.

For each neuron, we computed the average DFF response over a landmark-stop window, normalized each to a maximum DFF of 1, and formed an N-by-T (N: number of neurons and T: time) matrix M where each row corresponded to the normalized average DFF of a neuron. To capture most of the energy in the responses in a low-dimensional space, we reduced neuronal dimensions by performing a low-rank approximation. Specifically, we performed a singular value decomposition of M as $M=USV^T$ where the matrix S is a rectangular diagonal matrix of singular values, and the matrices U and V are orthonormal matrices (Dahleh et al., 2004; Horn and Johnson, 2012; Strang, 2016). We then kept only the highest 4 singular values in S, and set the remaining ones to zero, to get a matrix S_{approx}. The matrix obtained by the multiplication $M_{approx} = US_{approx}V^T$ is then a low-rank approximation of M. Each of the four non-zero singular values in S corresponded to a temporal neuronal response in V^T, which together span a 4-dimensional subspace. The subspace contained more than 80% of the calcium signal energy for the whole population (squared Frobenius norm of M). Specifically, the sum of squares of all the entries (energy) in $M - M_{approx}$ is less than 20% of the sum of squares of all the entries (energy) in M, indicating that the low-dimensional subspace indeed captured 80% of the energy in the neuronal responses (Horn and Johnson, 2012; Strang, 2016).

The four temporal neuronal responses obtained in V^T corresponding the non-zero singular values in S_{approx} span a 4-dimensional space, and form a basis to that space. Each DFF response can then be written as a weighted combination of these four basis responses and an additional response outside (orthogonal) to that space that is considered as noise. As a space can admit multiple bases, we decided to find a basis that represents stopping epochs and decompose our neuronal responses onto it. We then sought to define ideal neuronal templates of pre-stop, stop and post-stop neurons. We began with three square signals with value 1 between [-1, 0.5]s , [-0.25, 0.25] and [0.5, 1.5]s respectively, and 0 elsewhere. We then projected each of the three square signals to the 4-dimensional subspace defined by the low-rank approximation, and ensured orthogonality using gram-schmidt orthonormalization process (Strang, 2016). We were then left with three templates of ideal pre-stop, stop and post-stop neurons, spanning a three-dimensional subspace. We then kept the imaged neurons whose 85% of their calcium signal energy (area of the squared signal) belonged to the subspace, and discarded the rest from the analysis. This kept about 64% of all neurons we started with, and we performed clustering on them as described below.

Each neuronal response can then be written as a weighted combination of the three templates, plus some additional component orthogonal to the subspace (noise). We recovered the weights for each response by taking the inner-product (dot-product) with each template. We multiplied the weight by the maximum value in the template to account for the template width, to get a corrected weight. We then attributed a neuron to one of three classes (pre-stop, stop and post-stop) whose corresponding corrected-weight is highest.

Most importantly, the calcium responses in each class are not reduced in dimensions. They are the original averages of the raw DFF traces taken over landmark stop windows. The dimensionality reduction is only used to cluster the neuronal responses.

For the class of stop-neurons, we computed the weight with respect to the stop-neuron template for the average response in landmark-stop windows and spontaneous-stop windows, by taking the inner-product (dot-product) with the stop-neuron template.

Methods References

Bertsekas, D.P., 1995. Dynamic programming and optimal control. Athena scientific Belmont, MA.

Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxidis, J., Najafi, F., Gauthier, J.L., Zhou, P., Khakh, B.S., Tank, D.W., Chklovskii, D.B., Pnevmatikakis, E.A., 2019. CaImAn an open source tool for scalable calcium imaging data analysis. eLife 8, e38173. https://doi.org/10.7554/eLife.38173

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., Ahrens, M., Bruno, R., Jessell, T.M., Peterka, D.S., Yuste, R., Paninski, L., 2016. Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron 89, 285–299. https://doi.org/10.1016/j.neuron.2015.11.037

Strang, G., 2016. Introduction to linear algebra, 6th edition. Wellesley-Cambridge Press, Wellesley ; Massachusetts.

Horn, R.A., Johnson, C.R., 2012. Matrix analysis, 2nd ed. ed. Cambridge University Press, Cambridge ; New York.