Neuron

Striatal Low-Threshold Spiking Interneurons **Regulate Goal-Directed Learning**

Graphical Abstract



Authors

Elizabeth N. Holly, M. Felicia Davatolhagh, Kyuhyun Choi, Opeyemi O. Alabi, Luigim Vargas Cifuentes, Marc V. Fuccillo

Correspondence

fuccillo@pennmedicine.upenn.edu

In Brief

The functional role of dorsomedial striatum low-threshold spiking interneurons (LTSIs) in behavior has been largely unexplored. Holly et al. show that LTSIs have reward outcome-associated activity that decreases with and causally modulates goal-directed learning.

Highlights

- Striatal low-threshold spiking interneurons (LTSIs) mediate goal-directed learning
- Reward-associated LTSI Ca²⁺ activity decreases with operant learning
- Decreasing LTSI activity accelerates and increasing LTSI activity slows learning
- Effects on learning are dependent on LTSI GABAergic signaling

July 3, 2019 © 2019 Elsevier Inc.







Striatal Low-Threshold Spiking Interneurons Regulate Goal-Directed Learning

Elizabeth N. Holly,¹ M. Felicia Davatolhagh,^{1,2} Kyuhyun Choi,¹ Opeyemi O. Alabi,^{1,2} Luigim Vargas Cifuentes,^{1,2} and Marc V. Fuccillo^{1,3,*}

¹Department of Neuroscience, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA ²Neuroscience Graduate Group, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA ³Lead Contact

*Correspondence: fuccillo@pennmedicine.upenn.edu

https://doi.org/10.1016/j.neuron.2019.04.016

SUMMARY

The dorsomedial striatum (DMS) is critically involved in motor control and reward processing, but the specific neural circuit mediators are poorly understood. Recent evidence highlights the extensive connectivity of low-threshold spiking interneurons (LTSIs) within local striatal circuitry; however, the in vivo function of LTSIs remains largely unexplored. We employed fiber photometry to assess LTSI calcium activity in a range of DMS-mediated behaviors, uncovering specific reward-related activity that is down-modulated during goal-directed learning. Using two mechanistically distinct manipulations, we demonstrated that this down-modulation of LTSI activity is critical for acquisition of novel contingencies, but not for their modification. In contrast, continued LTSI activation slowed instrumental learning. Similar manipulations of fast-spiking interneurons did not reproduce these effects, implying a specific function of LTSIs. Finally, we revealed a role for the γ -aminobutyric acid (GABA)ergic functions of LTSIs in learning. Together, our data provide new insights into this striatal interneuron subclass as important gatekeepers of goal-directed learning.

INTRODUCTION

The dorsal striatum is a central node for the integration of cortical, thalamic, and limbic inputs (Hunnicutt et al., 2016). Although historically associated with motor control (DeLong, 1990), recent work has also implicated this structure in predictive, appetitive, and reinforcement behaviors (Balleine et al., 2007; Redgrave et al., 2010). In particular, the caudate nucleus and its anatomical rodent equivalent, the dorsomedial striatum (DMS), have been functionally associated with goal-directed learning, action performance, and choice flexibility (Groman et al., 2011; Wang et al., 2013; Yin et al., 2005). Furthermore, clinical imaging studies have revealed caudate dysfunction in an array of neuropsychiatric disorders, including autism, schizophrenia, and major depression (Di Martino et al., 2011; Estes

et al., 2011; Kerestes et al., 2014). Nevertheless, how specific striatal circuits contribute to essential behavioral function and disease pathophysiology remains unclear.

The two main striatal cell types, dopamine D1 and D2 receptor-expressing spiny projection neurons (SPNs) are thought to encode initiation and suppression of movement, respectively (Kravitz et al., 2010). Recent work suggests that their activity may also encode the value of future actions and that they undergo divergent patterns of synaptic plasticity during goaldirected learning (Shan et al., 2014; Tai et al., 2012). Although SPNs comprise the vast majority of striatal neurons, their activity is strongly modulated by sparse, diverse populations of local circuit interneurons (Gittis et al., 2010; Kawaguchi et al., 1995; Tepper et al., 2010). Although striatal cholinergic interneurons (ChINs) have well-documented reward-related function (Zhang and Cragg, 2017), other striatal γ -aminobutyric acid (GABA)ergic interneurons have received less attention. Recent work implicates striatal parvalbumin (PV)-expressing fast-spiking interneurons (FSIs) in egocentric spatial learning (Owen et al., 2018), Pavlovian conditioned approach (Lee et al., 2017), and transitions to habitual responding (O'Hare et al., 2017), suggesting that further functional investigation of striatal interneuron subtypes is warranted.

Little is known about the in vivo function of DMS somatostatin (SST)-expressing low-threshold spiking interneurons (LTSIs). These interneurons are tonically active in slice (Bennett and Wilson, 1999) and in vivo (Sharott et al., 2012) and express neuropeptides (SST and neuropeptide Y [NPY]) and neuromodulators (nitric oxide) in addition to GABA (Kawaguchi, 1993). Within the local circuit, LTSIs make inhibitory synapses at distal dendrites of SPNs and partake in reciprocal connections with nearby ChINs (Elghaba et al., 2016; Straub et al., 2016). To further explore the functional role of this striatal subtype, we employed in vivo population Ca2+ imaging to characterize LTSI activity during a range of DMS-associated behaviors. We found that LTSIs exhibited specific reward-related activity that was robustly downregulated during operant learning. Using two mechanistically distinct in vivo manipulations (Kir2.1 overexpression and optogenetic inhibition), we demonstrated that down-modulation of LTSI activity during operant learning was critical for acquisition of goal-directed behavior but not for modification of previously learned contingencies. Furthermore, increasing LTSI activity was sufficient to slow instrumental responding. These behavioral effects could not be reproduced by similar manipulations of FSIs,



(legend on next page)

implying cell type functional specificity. Finally, we revealed a role of the GABAergic functions of LTSIs in control of learning. Together, our data provide novel evidence that LTSIs within the DMS serve as important gatekeepers of early goal-directed learning.

RESULTS

LTSI Activity Is Dynamically Altered during Goal-Directed Learning

Given the sparse striatal distribution of LTSIs, we employed fiber photometry to measure their population calcium activity (Figure 1A). We observed sustained spontaneous activity that was (1) not observed in control animals injected with Cre-dependent eGFP (Figures 1B and 1C) and (2) decreased when a Cre-sensitive inhibitory designer receptor exclusively activated by designer drugs (DREADD; hM4D) co-injected with GCaMP6f was activated by clozapine-N-oxide (Figure S1A). The frequency and amplitude of LTSI calcium transients was stable across 2-h sessions of continuous recording (Figures 1C and S1B).

Because the DMS plays an integral role in motor control (Kravitz et al., 2010; Tecuapetla et al., 2016), we first evaluated whether LTSI Ca²⁺ activity was modulated during locomotion. We did not observe significant changes in the frequency or amplitude of LTSI Ca²⁺ events when mice were immobile or engaging in small (forepaw) or large (walking and rearing) motor output (Figures 1D and S1C). To explore whether LTSI activity was specifically modulated during motor state transitions, we aligned our *in vivo* Ca²⁺ signals with these behavioral epochs. As with spontaneous activity, there were no detectable changes in LTSI Ca²⁺ activity as mice transitioned into either motor state category (Figure 1E). Additionally, there were no spontaneous activity changes when mice were introduced to novel social or environmental contexts (Figures 1F, 1G, S1D, and S1E).

The DMS is also responsive to appetitive actions, with both SPN subtypes inhibited during consumption (London et al., 2018). Given this, we recorded LTSI Ca2+ activity as mice obtained a liquid chocolate reward delivered once per minute in a lit reward port (Figure 1H). Although the overall frequency and amplitude of LTSI Ca²⁺ transients did not change across the session (Figures 1I and S1F), there was a reliable increase in the LTSI Ca²⁺ waveform as mice entered the magazine to retrieve the reward (Figure 1J), which did not change across the session (Figure S1G). In light of this activity and the significant role of DMS circuits in goal-directed behavior (Shan et al., 2014; Tai et al., 2012; Yin et al., 2005), we recorded in vivo LTSI activity as mice learned a self-initiated two-choice operant lever pressing task (Figures 1K and 1L). Interestingly, as mice learned this goal-directed task, the frequency of Ca²⁺ events significantly decreased from baseline, only returning to initial levels during a post-acquisition, non-operant epoch within the chamber (Figure 1 M). No change in average event amplitude was noted during acquisition (Figure S1H). In a subset of mice that did not acquire on the first testing day (<10 rewards in a 1-h session), we instead observed stable LTSI transient frequencies and amplitudes (Figures S1I and S1J). We additionally noted a significant negative correlation between the frequency of LTSI Ca²⁺ events and lever press rate (Figure 1N). To further characterize the relationship between Ca²⁺ event frequency, time, and other task parameters in individual subjects, we developed a multiple regression model with time (cumulative across 10-min bins) and measures of task performance (cumulative rewards obtained, incorrect presses, rates of initiation, omission, and responding) as regressors. This model ($R^2 = 0.486$) highlighted significant contributions of cumulative rewards obtained (t = -3.59, p < 0.0001), incorrect presses (t = 2.06, p = 0.044), and individual (t = 2.73, p = 0.008), but no contribution for time (t = 0.79, p = 0.434) or other performance measures, to changes in LTSI Ca2+ event frequency.

Figure 1. LTSI Activity Is Dynamically Modulated during Goal-Directed Learning

(A) Recording schematic and fiber placement (left), sample traces of the experimental GCaMP6f (black) or control eGFP (gray) signal, and order of imaging experiments (bottom).

(C) Frequency of detected Ca²⁺ peaks (events per minute) during 2-h baseline recording.

(D) Frequency of detected peaks in different locomotor states (immobile, small forepaw movements, or large walking and rearing movements) in 12-min recording in the home cage.

(E) Peri-event temporal histograms (PETHs) as mice transition from immobile to small (left) or large (center) motor states, quantified by area under the curve (AUC) analysis (right).

(F and G) Event frequency upon exposure to a novel social (F) or environmental context (G).

(H) Schematic of reward magazine training (RMT) sessions.

(I) Event frequency during the third RMT session.

(J) PETHs of signal in response to light illumination (left) and reward retrieval (entering the reward magazine; center), summarized by AUC (right); ***p < 0.0001 versus light-on control.

(K) Self-initiated operant task.

(L) Cumulative rewards versus minutes for the first 50 rewards obtained in the session where acquisition occurred.

(M) Event frequency as mice learn the operant task. POST refers to the period after termination of the operant task, not necessarily after the 50th reward for all mice, and was therefore not included in the statistical analysis. **p < 0.01 versus BL within GCaMP6f group.

(N) Correlation between event frequency and lever press rate. Circles represent the mean of 7 mice for each 10-reward bin, and the dotted line represents linear regression for fit.

(O) PETHs for the first (red) and last (blue) 10 rewards of acquisition during mid-ITI, initiation, correct press, and reward retrieval. AUC analysis reveals a significantly elevated signal during reward retrieval in the first 10 rewards (far right). **p < 0.01, ****p < 0.00001 versus reward retrieval in first 10 rewards. Dotted lines represent the mean of baseline data points.

All data are represented as mean ± SEM. See also Table S1 for detailed statistics and Figure S1.

⁽B) Mean amplitude of Ca²⁺ peaks in (A); ***p < 0.0001 versus eGFP control.

To uncover whether this decrease in LTSI activity was tied to a specific behavioral epoch, we aligned Ca2+ waveforms to key task events (initiation, correct lever press, and reward retrieval) or task non-engagement (mid-inter-trial interval [ITI]), averaging trials from beginning (rewards 1-10, red) and end (rewards 41-50, blue) of acquisition. We found a robust LTSI activity signal upon reward retrieval in early acquisition, which was significantly suppressed by the end of action-outcome learning (Figure 10). We did not uncover reliable event-related peaks or modulation for other components of the task. We suspect that this modulation was operant learning-dependent because LTSI Ca²⁺ activity during non-operant reward presentation did not decrease as a function of reward number (Figure S1G). Finally, we did not observe a statistically significant modulation of LTSI Ca²⁺ events or a reward-associated peak on 2 subsequent operant days following acquisition, where animals received unlimited rewards under the same contingency for 1 h (Figures S1K–S1N).

Bidirectional LTSI Manipulation Modulates Goal-Directed Learning

Our calcium imaging showed that both global and rewardassociated LTSI activity decreased as mice learned a goaldirected task, suggesting a potential role of these cells in regulating operant learning. To explore a causal relationship, we experimentally reduced DMS LTSI activity in SST-Cre mice through viral overexpression of Kir2.1, an inwardly rectifying potassium channel that decreases cellular excitability (Figure 2A; Lin et al., 2010). This strategy permitted highly specific and penetrant access to DMS LTSIs, which, surprisingly, extend projections throughout the compartment's anteriorposterior (A-P) extent (Figures S2A–S2E). Kir2.1 overexpression reliably reduced spontaneous firing of LTSIs, their response to depolarizing current injection, as well as other proxies of cellular excitability (Figures 2A and S2C).

We bilaterally injected a Cre-dependent DIO-Kir2.1 or control DIO-eGFP virus into the DMS of SST-Cre mice, incubated for 2 weeks, and trained animals as in the prior Ca²⁺ imaging experiment (Figure 2B). Kir2.1-mediated LTSI inhibition significantly reduced the time to acquire the task, operationally defined as 50 rewards (Figure 2C). These results suggest that the activity of striatal LTSIs may act as a brake on operant learning, with enhanced acquisition possibly arising from (1) fewer trial omissions (failure to press a lever within 10 s of trial initiation), (2) fewer incorrect responses, and/or (3) enhanced motor efficiency. To examine acquisition behavior, we divided all trials into 10-reward bins. Mice injected with Kir2.1 had significantly fewer omissions in the first reward bin compared with eGFP-injected controls (Figure 2D), whereas there was no difference in incorrect presses (Figure 2E). This was paired with shorter latencies to initiate during the first 10-reward bin (Figure 2F), with no significant differences in latencies to lever press or retrieve the reward (Figures S2F and S2G). This suggests that the faster acquisition rate observed in LTSI-Kir2.1 mice was due to a combination of increased completed trials and a shorter time to initiate individual trials. To investigate a role of LTSIs in generating light cue-mediated Pavlovian interference of operant responding, we repeated the experiment without the light cue during the reward magazine training sessions and obtained identical results (Figures S2J-S2Q).

Because Kir2.1-mediated inhibition of LTSIs was persistent and irreversible, our behavioral effects may be the result of local circuit adaptation. To test whether within-task inhibition of LTSI activity could modulate goal-directed learning, we employed the lightactivated chloride channel halorhodopsin (Gradinaru et al., 2008) to optogenetically inhibit LTSIs (Figures 2G and 2H). We tested the efficacy of this manipulation in acute striatal slices, demonstrating that halorhodopsin completely prevented tonic LTSI firing during optical activation without altering subsequent firing rates (Figure 2G). Given the dynamic modulation of LTSI Ca²⁺ activity during reward retrieval, we tested whether inhibition during this epoch was sufficient to modulate learning. Unilateral optogenetic inhibition of LTSIs during the first 4 s of the reward period recapitulated the faster acquisition observed with constitutive Kir2.1-mediated inhibition (Figure 2I). As with Kir2.1, optogenetic inhibition during reward retrieval significantly decreased omissions in the first reward bin (Figure 2J) without affecting incorrect presses (Figure 2K). In contrast to Kir2.1 manipulation, optogenetic inhibition did not significantly affect latency to initiate (Figure 2L), instead significantly decreasing latencies to lever press (Figure S2R) and retrieve the reward (Figure S2S). We do not believe that the learning changes result from global changes in activity levels or reward processing because (1) LTSI manipulation did not significantly alter locomotor activity (Figures S2H and S2V), (2) Kir2.1-mediated LTSI inhibition did not affect progressive ratio performance (Figure S2I) or free reward consumption (data not shown), (3) optogenetic inhibition of LTSIs was not more rewarding than the chocolate reward (Figure S2W), and (4) mice did not self-stimulate for optogenetic inhibition of LTSIs in a spatial task (Figures S2T and S2U; Carta et al., 2019).

Because reward-related LTSI activity decreases with an increased learning rate, we next probed whether increasing LTSI activity could prolong the time to acquire goal-directed behaviors. To do this, we employed channelrhodopsin to artificially increase LTSI activity during learning. In acute slice, we demonstrated that LTSIs can follow 4 s of 10-Hz stimulation without producing prolonged suppression of activity following laser termination (Figure 3A). Mimicking our halorhodopsin experimental design (Figure 3B), we demonstrate that optogenetic excitation specifically during the reward period significantly increased the time to acquire the operant task (Figure 3C). In contrast with LTSI inhibition, reward-associated stimulation did not affect the number of omissions or incorrect choices across learning (Figures 3D and 3E), instead modulating the latencies to initiate (Figure 3F) and retrieve the reward (Figure S3B) without affecting latency to press (Figure S3A). Together, these data show that LTSIs can bidirectionally affect goal-directed instrumental learning.

LTSIs Modulate Novel Contingency Acquisition in a Cell-Type-Specific Manner

Striatal LTSIs represent one subtype of a heterogenous group of local circuit interneurons (Tepper et al., 2010). Recent work has suggested that PV+ FSIs, a similarly sparse local striatal inhibitory component, can strongly modulate aspects of reward processing and action selection (Lee et al., 2017; O'Hare et al., 2017; Owen et al., 2018). To assess whether acceleration of goal-directed learning was specific to LTSIs, we inhibited FSIs



Figure 2. Suppression of LTSI Activity Accelerates Goal-Directed Learning

(A) Viral injection followed by whole-cell acute slice recordings of LTSIs in response to increasing current injection (left). Spontaneous firing frequency (hertz) in slice with Kir2.1-mediated LTSI inhibition (center) and current step-action potential (AP) plot with Kir2.1-mediated LTSI inhibition (right). *p < 0.05, ****p < 0.0001 versus eGFP control. eGFP control, n = 12–14 cells/3 mice; Kir2.1, n = 13 cells/3 mice.

(B) Histological confirmation of Kir2.1 targeting across experimental animals (left) and experimental timeline (right).

(C) Time-to-acquire instrumental task (left) and cumulative rewards versus minutes for eGFP-expressing (gray, n = 9) and Kir2.1-expressing (green, n = 9) mice (right). **p < 0.01 versus eGFP control.

(D–F) Omissions (D), incorrect presses (E), and initiation latencies (F) across acquisition in 10-reward bins. ***p < 0.001 versus eGFP control within the reward bin. (G) Representative trace of cell-attached recording during 530-nm halorhodopsin illumination (left) and LTSI spontaneous firing frequency (hertz) before, during, and after 530-nm illumination (right). ****p < 0.0001 versus baseline (before); n = 15 cells/3 mice.

(H) Histological confirmation of the fiberoptic tract above the eGFP+ (gray, n = 8) and NpHR3.0+ (green, n = 8) viral penumbra.

(I–L) Same as in (C)–(F) but for halorhodopsin (eNpHR3.0)-mediated inhibition.

All data represented as mean \pm SEM. See also Table S1 for detailed statistics and Figure S2.

within the DMS. Bilateral injection of DIO-Kir2.1 into the DMS of PV-Cre mice significantly suppressed evoked firing (Figure 4A) and other measures of cellular excitability (Figure S4A) but failed to alter the acquisition rate of operant learning (Figures 4B and S4B–S4G).

The down-modulation of LTSI activity observed during operant acquisition may place constraints on the relevance of this population after learning. To explore this, we tested whether LTSI inhibition similarly accelerates the modification of previously established action-outcome contingencies via reversal of



Figure 3. LTSI Activation Slows Goal-Directed Learning

(A) Representative trace of cell-attached recording during 470-nm channelrhodopsin illumination (left) and LTSI spontaneous firing frequency (hertz) before, during, and after 470-nm illumination (right). *p < 0.05, ****p < 0.0001 versus baseline (before); n = 17 cells/3 mice.

(B) Histological confirmation of the fiberoptic tract above the eGFP+ (gray, n = 10) and ChR2+ (blue, n = 9) viral penumbra (left) and experimental timeline (right). (C) Time-to-acquire instrumental task (left) and cumulative rewards versus minutes (right). *p < 0.05 versus eGFP control.

(D-F) Omissions (D), incorrect presses (E), and latency to initiate (F) across bins of 10 rewards. ***p < 0.001 versus eGFP control within the reward bin. All data are represented as mean ± SEM. See also Table S1 for detailed statistics.

the rewarded lever. The time to choice reversal was not affected by LTSI-specific Kir2.1 (Figure 4C) or optogenetic inhibition during reward retrieval (Figure 4D). These results suggest that LTSI modulation may exert significant behavioral control only during initial learning of novel contingencies. To test this, we ran a cohort of mice that had previously undergone optogenetic LTSI inhibition during lever press reward association on a novel contingency requiring a nosepoke on the opposite chamber wall. Mice were subdivided into two groups: one receiving halorhodopsin activation during reward as in prior tests and one receiving no illumination. Optogenetic inhibition of LTSIs during reward again accelerated the acquisition of the novel nosepoke-reward contingency compared with eGFP controls, whereas the subgroup receiving no laser learned at a similar rate as the control (Figure 4E). From this experiment we conclude that (1) LTSI-mediated enhancement of acquisition does not carry over to alternative action-outcome contingencies and that (2) LTSI modulation exerts the largest behavioral effects on goal-directed learning of novel associations.

LTSI GABAergic Signaling Is Critical for the Regulation of Goal-Directed Learning

In addition to providing dendritic GABAergic inhibition, LTSIs express multiple neuropeptides associated with learning. To test the functional relevance of LTSI GABAergic inhibition in the modulation of operant acquisition, we employed a genetic-viral strategy to specifically delete VGAT, a gene essential for packaging of GABA into synaptic vesicles, from DMS LTSIs. SST-FIp mice were crossed with *Slc32a1* (VGAT) conditional knockout mice,

and a Flp-sensitive CreGFP virus was injected into the DMS to permit cell-type-specific deletion. To validate VGAT deletion, we performed *in situ* hybridization combined with immunohistochemistry on SST-Flp^{+/-};VGAT^{+/+} or SST-Flp^{+/-};VGAT^{C/C} mice injected with fDIO-CreGFP and observed a significant reduction of VGAT mRNA in GFP+/SST-Flp+ cells of VGAT-conditional animals (Figure 5A). To test the physiological effectiveness of this manipulation, we co-injected fDIO-CreGFP and Cre-sensitive ChR2(H134R)-eYFP into the DMS of SST-Flp^{+/-};VGAT^{+/+} or SST-Flp^{+/-};VGAT^{C/C} mice and performed whole-cell recordings on neighboring SPNs while optogenetically recruiting LTSIs (Figure 5B, left). Conditional deletion of VGAT from LTSIs significantly decreased optically evoked inhibitory postsynaptic current (oIPSC) amplitude (confirmed by sensitivity to picrotoxin; Figure 5B, right) across light intensities.

To probe the effects of DMS LTSI-specific VGAT deletion on behavior, we next injected SST-FIp^{-/-}; VGAT^{C/C} or SST-FIp^{+/-}; VGAT^{C/C} mice with fDIO-CreGFP and trained them in our operant task (Figure 5C). Consistent with prior circuit manipulations of LTSI activity, conditional deletion of VGAT in LTSIs significantly decreased the acquisition time (Figure 5D). This was accompanied by significant decreases in early-stage omissions (Figure 5E) and latencies to initiate, lever press, and retrieve the reward (Figures 5G–5I) but no change in incorrect lever presses (Figure 5F). LTSI-specific VGAT deletion mice also did not exhibit enhanced flexibility in a single reversal session (Figure S5). These data suggest that GABAergic transmission from LTSIs is central in modulating acquisition of instrumental learning but not modification of existing contingencies.



Figure 4. LTSIs Modulate Novel Contingency Acquisition in a Cell-Type-Specific Manner

(A) Viral injection followed by whole-cell acute slice recordings of FSIs in response to increasing current injection (left). Current-AP plot for eGFP- and Kir2.1expressing FSIs (right). *p < 0.05 versus eGFP control. eGFP control, n = 15 cells/3 mice; Kir2.1, n = 14 cells/3 mice.

(B) Time to acquire (left) and cumulative rewards versus minutes for eGFP-expressing (gray, n = 9) and Kir2.1-expressing (purple, n = 10) mice (right).

(C and D) Time to complete reversal with Kir-mediated (C) and halorhodopsin-mediated (D) inhibition.

(E) Schematic (left) and time to acquire (right) novel nosepoke contingency in eGFP-expressing (n = 8) or halorhodopsin-expressing mice receiving laser (n = 4) or no laser (n = 4) during reward. *p < 0.05 versus the halorhodopsin + laser group.

All data are represented as mean ± SEM. See also Table S1 for detailed statistics and Figure S3.

DISCUSSION

Here we explore the *in vivo* activity and function of SST+ LTSIs, a major striatal inhibitory subtype. We find that these cells exhibit circumscribed reward-related activity that is robustly modulated during operant learning. Reductions in reward-related LTSI activity are functionally relevant because manipulations suppressing this subpopulation accelerate the rate of operant acquisition, whereas activation of this subpopulation conversely slows learning. Furthermore, we reveal functional specificity for LTSIs because manipulation of striatal FSIs does not similarly alter behavior. Finally, we demonstrate the importance of LTSI GABAergic signaling in this modulatory control of reward-based learning.

In vivo population Ca²⁺ imaging of LTSIs revealed relative specificity of DMS LTSIs for reward-associated behaviors. We found no evidence of correlation of LTSI activity with locomotor state, motor transition, or exposure to novel environmental or social contexts. In contrast, we found robust responses to both delivered and operantly obtained appetitive rewards. This reward-related LTSI activity could result from increased phasic excitatory drive, reduced inhibition, or local neuromodulation. Cell-type-specific retrograde tracing of LTSIs showed a particularly high convergence of medial orbital frontal cortical neurons (Choi et al., 2019), an excitatory population exhibiting significant outcome value encoding (Schoenbaum et al., 2011). Although striatal FSIs and PV+ projections from the globus pallidus exhibit sparse GABAergic connectivity onto LTSIs (Saunders et al., 2016; Szydlowski et al., 2013), most LTSIs are inhibited by striatal tyrosine hydroxylase interneurons (Assous et al., 2017). Other major GABAergic sources of innervation are presently unknown. Regarding modulation of reward-related activity with learning, it is interesting to consider that serotonin, a striatal neuromodulator whose signal integrates positive outcomes over many trials (Cohen et al., 2015), strongly decreases LTSI activity (Cains et al., 2012).

Multiple manipulations reducing striatal LTSI activity dramatically enhanced operant acquisition, whereas increasing LTSI activity slowed acquisition. The early stages of instrumental learning are marked by a transition from non-focused, highly variable motor output to behavioral stability and efficient motor expression when desirable outcomes are uncovered (Dezfouli and Balleine, 2012). We suggest that LTSI inhibition may regulate this transition by facilitating enhancement of specific motor sequences, enhancing task-relevant attention, or enhancing overall motor efficiency. Future work exploring how LTSI GABAergic modulation of striatal circuits mediates these behavioral mechanisms must consider the extensive axonal projections along the entire anterior-posterior DMS axis. Dendritic striatal LTSI activity is ideally situated to control the flow of cortical input (Fino et al., 2018), suggesting a role in regulating the plasticity of corticostriatal excitatory synapses that mediate reward-driven motor learning



Figure 5. LTSI GABAergic Signaling Is Critical for the Regulation of Goal-Directed Learning

(A) Viral-genetic approach for LTSI VGAT deletion (top left), showing selective reduction of the VGAT *in situ* hybridization (ISH) signal in CreGFP+ neurons of VGAT^{C/C} animals. ****p < 0.001 versus VGAT^{+/+}; VGAT^{+/+}, n = 4 mice; SSF^{+/-}; VGAT^{C/C}, n = 4 mice.

(B) Left: physiological demonstration of VGAT deletion. Shown are representative optically evoked LTSI GABAergic currents (oIPSCs) recorded in SPNs and plot of oIPSC amplitude across light emitting diode (LED) intensities (center). **p < 0.01, ***p < 0.001 versus VGAT^{+/+} control. Right: sensitivity of optical responses to picrotoxin. SSF^{+/-};VGAT^{+/+}, n = 23 cells/4 mice; SSF^{+/-};VGAT^{C/C}, n = 25 cells/4 mice).

(C) Histological confirmation of fDIO-CreGFP targeting across experimental animals (left) and experimental timeline (right).

(D) Left: time to acquire with conditional VGAT deletion in LTSIs. *p < 0.05 versus SSF^{-/-} control. Right: cumulative rewards versus minutes for control (SSF^{-/-}, gray) and experimental (SSF^{+/-}, green) mice.

(E–I) Omissions (E), incorrect presses (F), and latencies to initiation (G), press (H), and reward (I) across acquisition in bins of 10 rewards. **p < 0.01, ***p < 0.001 versus SSF^{-/-} control.

All data are represented as mean ± SEM. See also Table S1 for detailed statistics and Figure S4.

(Bar-Ilan et al., 2013; Straub et al., 2016). Consistent with this, cue-associated punishment and stimulus selectivity learning in the primary visual cortex are accompanied by reduced SST+ interneuron activity and their decorrelation from pyramidal networks, respectively (Khan et al., 2018; Makino and Komiyama, 2015). In addition to targeting SPN dendrites, LTSIs make strong GABAergic projections to ChINs (Holley et al., 2015). Although the behavioral implications of these interactions are unclear, the location of these two inhibitory subclasses in zones surrounding striatal patches suggests compartment-specific control of striatal output (Brimblecombe and Cragg, 2017). Furthermore, although we observed that LTSI GABAergic function is essential for behavioral control, we do not rule out the possibility that LTSI-derived neuromodulators also have effects. Nitric oxide and somatostatin

receptor signaling, both implicated in the plasticity of local striatal circuitry (Lopez-Huerta et al., 2008; Rafalovich et al., 2015), may also contribute to goal-directed control.

We show that the effects of LTSI suppression on operant acquisition cannot be achieved by similar manipulations in DMS PV+ FSIs. Together with previous work (Lee et al., 2017; O'Hare et al., 2017; Owen et al., 2018), our data suggest that striatal interneurons may contribute to distinct aspects of goaldirected behavior. In the dorsal striatum, FSIs appear to be pleiotropic regulators of goal-directed behavior with effects on egocentric action selection, value-insensitive habitual responding, and acquisition of Pavlovian reward-conditioned responses. We find that, in contrast to FSIs, dorsal striatal LTSIs control the initial stages of goal-directed learning. Furthermore, although FSI manipulations exert behavioral effects through biasing choice selection, the enhanced acquisition we observed did not result from improved choice accuracy, instead occurring via increased progression from initiation to choice states as well as via increased efficiency of individual motor components. Further work will be required to parse the potential role of LTSIs in action-value encoding or value-based comparisons.

In summary, we provide novel insight into the *in vivo* role of striatal SST+ LTSIs. Our data reveal a reward-responsive neuronal population essential for modulating goal-directed learning. We believe that these findings contribute to our understanding of how animals engage in reward-yielding instrumental actions.

STAR*METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - General Methods
 - Specific Experimental Methods
 - Kir2.1 Overexpression-Mediated Inhibition
 - eNpHR3.0-Mediated Inhibition
 - ChR2-Mediated Excitation
 - Conditional VGAT Deletion
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j. neuron.2019.04.016.

ACKNOWLEDGMENTS

This work was supported by grants from the NIMH (F32-MH114506 to E.N.H., F31-MH114528 to O.O.A., and R00-MH099243 and R01-MH115030 to M.V.F.), Howard Hughes Medical Institute (Gilliam Fellowship to M.F.D.), and the Whitehall Foundation (to M.V.F.). We thank Boris Heifets, Elizabeth Steinberg, Rosemary Bagot, and Jessie Muir for assistance with the MATLAB code for photometry analysis, Bo Li for providing the fDIO-CreGFP plasmid, Byung-kook Lim for providing the DIO-Kir2.1 plasmid, and Amelia Eisch for contributing transgenic mice. We also thank Patrick Rothwell, Emily Newman, and Hiruy Meharena for insightful feedback on the manuscript draft.

AUTHOR CONTRIBUTIONS

Conceptualization, E.N.H. and M.V.F.; Methodology, E.N.H., M.F.D., K.C., and M.V.F.; Formal Analysis, E.N.H., M.F.D., and K.C.; Investigation, E.N.H., M.F.D., K.C., O.O.A., and L.V.C.; Writing – Original Draft, E.N.H. and M.V.F.; Writing – Review and Editing, E.N.H., M.V.F., M.F.D., and L.V.C.; Funding Acquisition, E.N.H., M.F.D., O.O.A., and M.V.F.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 18, 2018 Revised: March 4, 2019 Accepted: April 9, 2019 Published: May 13, 2019

REFERENCES

Assous, M., Kaminer, J., Shah, F., Garg, A., Koós, T., and Tepper, J.M. (2017). Differential processing of thalamic information via distinct striatal interneuron circuits. Nat. Commun. 8, 15860.

Balleine, B.W., Delgado, M.R., and Hikosaka, O. (2007). The role of the dorsal striatum in reward and decision-making. J. Neurosci. 27, 8161–8165.

Bar-Ilan, L., Gidon, A., and Segev, I. (2013). The role of dendritic inhibition in shaping the plasticity of excitatory synapses. Front. Neural Circuits 6, 118.

Bennett, B.D., and Wilson, C.J. (1999). Spontaneous activity of neostriatal cholinergic interneurons in vitro. J. Neurosci. 19, 5586–5596.

Brimblecombe, K.R., and Cragg, S.J. (2017). The striosome and matrix compartments of the striatum: A path through the labyrinth from neurochemistry toward function. ACS Chem. Neurosci. 8, 235–242.

Cains, S., Blomeley, C.P., and Bracci, E. (2012). Serotonin inhibits low-threshold spike interneurons in the striatum. J. Physiol. 590, 2241–2252.

Carta, I., Chen, C.H., Schott, A.L., Dorizan, S., and Khodakhah, K. (2019). Cerebellar modulation of the reward circuitry and social behavior. Science *363*, eaav0581.

Choi, K., Holly, E.N., Davatolhagh, M.F., Beier, K.T., and Fuccillo, M.V. (2019). Integrated anatomical and physiological mapping of striatal afferent projections. Eur. J. Neurosci. *49*, 623–636.

Cohen, J.Y., Amoroso, M.W., and Uchida, N. (2015). Serotonergic neurons signal reward and punishment on multiple timescales. eLife 4.

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.

DeLong, M.R. (1990). Primate models of movement disorders of basal ganglia origin. Trends Neurosci. *13*, 281–285.

Dezfouli, A., and Balleine, B.W. (2012). Habits, action sequences and reinforcement learning. Eur. J. Neurosci. 35, 1036–1051.

Di Martino, A., Kelly, C., Grzadzinski, R., Zuo, X.N., Mennes, M., Mairena, M.A., Lord, C., Castellanos, F.X., and Milham, M.P. (2011). Aberrant striatal functional connectivity in children with autism. Biol. Psychiatry *69*, 847–856.

Elghaba, R., Vautrelle, N., and Bracci, E. (2016). Mutual control of cholinergic and low-threshold spike interneurons in the striatum. Front. Cell. Neurosci. *10*, 111.

Estes, A., Shaw, D.W., Sparks, B.F., Friedman, S., Giedd, J.N., Dawson, G., Bryan, M., and Dager, S.R. (2011). Basal ganglia morphometry and repetitive behavior in young children with autism spectrum disorder. Autism Res. *4*, 212–220.

Fino, E., Vandecasteele, M., Perez, S., Saudou, F., and Venance, L. (2018). Region-specific and state-dependent action of striatal GABAergic interneurons. Nat. Commun. *9*, 3339.

Friard, O., and Gamba, M. (2016). BORIS: a free, versatile open-source event-logging software for video/audio coding and live observations. Methods Ecol. Evol. 7, 1325–1330.

Gittis, A.H., Nelson, A.B., Thwin, M.T., Palop, J.J., and Kreitzer, A.C. (2010). Distinct roles of GABAergic interneurons in the regulation of striatal output pathways. J. Neurosci. *30*, 2223–2234.

Gradinaru, V., Thompson, K.R., and Deisseroth, K. (2008). eNpHR: a Natronomonas halorhodopsin enhanced for optogenetic applications. Brain Cell Biol. *36*, 129–139.

Groman, S.M., Lee, B., London, E.D., Mandelkern, M.A., James, A.S., Feiler, K., Rivera, R., Dahlbom, M., Sossi, V., Vandervoort, E., and Jentsch, J.D. (2011). Dorsal striatal D2-like receptor availability covaries with sensitivity to positive reinforcement during discrimination learning. J. Neurosci. *31*, 7291–7299.

Holley, S.M., Joshi, P.R., Parievsky, A., Galvan, L., Chen, J.Y., Fisher, Y.E., Huynh, M.N., Cepeda, C., and Levine, M.S. (2015). Enhanced GABAergic inputs contribute to functional alterations of cholinergic interneurons in the R6/2 mouse model of Huntington's Disease. eNeuro 2, e0008. Hunnicutt, B.J., Jongbloets, B.C., Birdsong, W.T., Gertz, K.J., Zhong, H., and Mao, T. (2016). A comprehensive excitatory input map of the striatum reveals novel functional organization. eLife 5, 5.

Kawaguchi, Y. (1993). Physiological, morphological, and histochemical characterization of three classes of interneurons in rat neostriatum. J. Neurosci. *13*, 4908–4923.

Kawaguchi, Y., Wilson, C.J., Augood, S.J., and Emson, P.C. (1995). Striatal interneurones: chemical, physiological and morphological characterization. Trends Neurosci. *18*, 527–535.

Kerestes, R., Harrison, B.J., Dandash, O., Stephanou, K., Whittle, S., Pujol, J., and Davey, C.G. (2014). Specific functional connectivity alterations of the dorsal striatum in young people with depression. Neuroimage Clin. *7*, 266–272.

Khan, A.G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T.D., and Hofer, S.B. (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. Nat. Neurosci. *21*, 851–859.

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.

Lee, K., Holley, S.M., Shobe, J.L., Chong, N.C., Cepeda, C., Levine, M.S., and Masmanidis, S.C. (2017). Parvalbumin interneurons modulate striatal output and enhance performance during associative learning. Neuron *93*, 1451–1463.e4.

Lin, C.-W., Sim, S., Ainsworth, A., Okada, M., Kelsch, W., and Lois, C. (2010). Genetically increased cell-intrinsic excitability enhances neuronal integration into adult brain circuits. Neuron *65*, 32–39.

London, T.D., Licholai, J.A., Szczot, I., Ali, M.A., LeBlanc, K.H., Fobbs, W.C., and Kravitz, A.V. (2018). Coordinated ramping of dorsal striatal pathways preceding food approach and consumption. J. Neurosci. *38*, 3547–3558.

Lopez-Huerta, V.G., Tecuapetla, F., Guzman, J.N., Bargas, J., and Galarraga, E. (2008). Presynaptic modulation by somatostatin in the neostriatum. Neurochem. Res. *33*, 1452–1458.

Makino, H., and Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. Nat. Neurosci. *18*, 1116–1122.

Muir, J., Lorsch, Z.S., Ramakrishnan, C., Deisseroth, K., Nestler, E.J., Calipari, E.S., and Bagot, R.C. (2018). In vivo fiber photometry reveals signature of future stress susceptibility in nucleus accumbens. Neuropsychopharmacology *43*, 255–263.

O'Hare, J.K., Li, H., Kim, N., Gaidis, E., Ade, K., Beck, J., Yin, H., and Calakos, N. (2017). Striatal fast-spiking interneurons selectively modulate circuit output and are required for habitual behavior. eLife *6*, e26231.

Owen, S.F., Berke, J.D., and Kreitzer, A.C. (2018). Fast-spiking interneurons supply feedforward control of bursting, calcium, and plasticity for efficient learning. Cell *172*, 683–695.e15.

Rafalovich, I.V., Melendez, A.E., Plotkin, J.L., Tanimura, A., Zhai, S., and Surmeier, D.J. (2015). Interneuronal nitric oxide signaling mediates post-

synaptic long-term depression of striatal glutamatergic synapses. Cell Rep. 13, 1336–1342.

Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M.C., Lehericy, S., Bergman, H., Agid, Y., DeLong, M.R., and Obeso, J.A. (2010). Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease. Nat. Rev. Neurosci. *11*, 760–772.

Saunders, A., Huang, K.W., and Sabatini, B.L. (2016). Globus pallidus externus neurons expressing parvalbumin interconnect the subthalamic nucleus and striatal interneurons. PLoS ONE *11*, e0149798.

Schoenbaum, G., Takahashi, Y., Liu, T.L., and McDannald, M.A. (2011). Does the orbitofrontal cortex signal value? Ann. N Y Acad. Sci. *1239*, 87–99.

Shan, Q., Ge, M., Christie, M.J., and Balleine, B.W. (2014). The acquisition of goal-directed actions generates opposing plasticity in direct and indirect pathways in dorsomedial striatum. J. Neurosci. 34, 9196–9201.

Sharott, A., Doig, N.M., Mallet, N., and Magill, P.J. (2012). Relationships between the firing of identified striatal interneurons and spontaneous and driven cortical activities in vivo. J. Neurosci. *32*, 13221–13236.

Straub, C., Saulnier, J.L., Bègue, A., Feng, D.D., Huang, K.W., and Sabatini, B.L. (2016). Principles of synaptic organization of GABAergic interneurons in the striatum. Neuron *92*, 84–92.

Szydlowski, S.N., Pollak Dorocic, I., Planert, H., Carlén, M., Meletis, K., and Silberberg, G. (2013). Target selectivity of feedforward inhibition by striatal fast-spiking interneurons. J. Neurosci. *33*, 1678–1683.

Tai, L.H., Lee, A.M., Benavidez, N., Bonci, A., and Wilbrecht, L. (2012). Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. Nat. Neurosci. *15*, 1281–1289.

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.

Tepper, J.M., Tecuapetla, F., Koós, T., and Ibáñez-Sandoval, O. (2010). Heterogeneity and diversity of striatal GABAergic interneurons. Front. Neuroanat. *4*, 150.

Tischfield, D.J., Saraswat, D.K., Furash, A., Fowler, S.C., Fuccillo, M.V., and Anderson, S.A. (2017). Loss of the neurodevelopmental gene Zswim6 alters striatal morphology and motor regulation. Neurobiol. Dis. *103*, 174–183.

Wang, A.Y., Miura, K., and Uchida, N. (2013). The dorsomedial striatum encodes net expected return, critical for energizing performance vigor. Nat. Neurosci. *16*, 639–647.

Yin, H.H., Ostlund, S.B., Knowlton, B.J., and Balleine, B.W. (2005). The role of the dorsomedial striatum in instrumental conditioning. Eur. J. Neurosci. *22*, 513–523.

Zhang, Y.F., and Cragg, S.J. (2017). Pauses in striatal cholinergic interneurons: What is revealed by their common themes and variations? Front. Syst. Neurosci. *11*, 80.

STAR***METHODS**

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Antibodies		
Rat monoclonal anti-somatostatin, clone YC7	EMD Millipore	MAB354; RRID: AB_2255365
Chicken polyclonal anti-GFP	abcam	ab13970; RRID: AB_300798
Goat anti-rat, Alexa Fluor 555 conjugated	Invitrogen	A-21434; RRID: AB_141733
Goat anti-chicken, Alexa Fluor 488 conjugated	Jackson ImmunoResearch	#103-545-555; RRID: AB_2337390
Bacterial and Virus Strains		
AAV5.Syn.Flex.GCaMP6f.WPRE.SV40	Penn Vector Core	N/A
AAV1.CAG.Flex.eGFP.WPRE.bGH	Penn Vector Core	AllenInstitute854
AAVDJ.EF1a.DIO.zsGreen.p2A.ES.Kir2.1.WPRE.hGH	Fuccillo Lab	N/A
AAV5.hSyn.eNpHR3.0-eYFP.WPRE.hGH	Penn Vector Core	Addgene26972P
AAV5.EF1a.DIO.hChR2(H134R)-eYFP.WPRE.hGH	Penn Vector Core	Addgene20298P
AAVDJ.EF1a.fDIO.CreGFP	Fuccillo Lab	N/A
Chemicals, Peptides, and Recombinant Proteins		
Picrotoxin	Sigma Aldrich	P1675
NBQX Disodium salt	Abcam	ab120046
Critical Commercial Assays:		
RNAScope Multiplex Fluorescent V2 Kit	Advanced Cell Diagnostics, Inc	323110
RNAscope probe – Mm-Slc32a1-C2	Advanced Cell Diagnostics, Inc	319191-C2
RNAscope Negative Control probe – DapB-C2	Advanced Cell Diagnostics, Inc	310043-C2
RNAscope Positive Control probe – Mm-Ubc-C2	Advanced Cell Diagnostics, Inc	310871-C2
RNAscope Positive Control probe – Mm-Polr2a-C2	Advanced Cell Diagnostics, Inc	312471-C2
Experimental Models: Organisms/Strains		
Mouse: Sst ^{tm2.1(cre)Zjh} /J (SST-ires-Cre)	Jackson Laboratory	JAX:013044; RRID: IMSR_JAX:013044
Mouse: Sst ^{tm3.1(flop)Zjh} /J (SST-ires-Flp)	Jackson Laboratory	JAX:028579; RRID: IMSR_JAX:028579
Mouse: Slc32a1 ^{tm1Lowl} /J (VGAT ^{C/C})	Jackson Laboratory	JAX:012897; RRID: IMSR_JAX:012897
Mouse: B6.Cg-Pvalb ^{tm1.1(cre)Aibs} /J (PV-2a-Cre)	Jackson Laboratory	JAX:012358; RRID: IMSR_JAX:012358
Recombinant DNA		
pAAV.DJ.EF1a.DIO.eGFP.p2A.ES.Kir2.1	B. Lim (UCSD)	N/A
pAAV.EF1a.fDIO.CreGFP	Bo Li Lab (CSHL)	N/A
pAAV.DJ.hSyn.DIO.mEGFP-2a-Synaptophysin-mRuby (SynaptoTAG2.0)	Kevin Beier (Stanford)	N/A

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Marc Fuccillo (fuccillo@ pennmedicine.upenn.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Prior to experimental manipulation, all mice (SST-IRES-Cre, Jackson stock number 013044, RRID:IMSR_JAX:013044; PV-2A-Cre, Jackson stock number 012358, RRID:IMSR_JAX:012358; SST-IRES-Flp, Jackson stock number 028579, RRID:IMSR_JAX:028579; VGAT (SIc32a1) conditional allele, Jackson stock number 012897, RRID:IMSR_JAX:012897; PV-2A-Cre, Jackson stock number 012358; SST-IRES-Flp, Jackson stock number 028579; VGAT (SIc32a1) conditional allele, Jackson stock number 012897; Were group housed with littermates on a 12:12 light-dark cycle and provided *ad libitum* food and water. Unless otherwise noted, all experiments were conducted on naive adult male mice, which were randomly assigned to experimental groups. After surgical implantation

of optical cannulas, mice in photometry and optogenetic experiments were individually housed. All experiments were conducted in accordance with the *National Institutes of Health Guidelines for the Use of Animals*, and all procedures approved by the Institutional Animal Care and Use Committee of the University of Pennsylvania (protocol 805643). Sample sizes are detailed in figure legends and Table S1 and were determined by a priori power analyses.

METHOD DETAILS

General Methods

Viral Injection and Fiberoptic Cannula Implantation

Viral injections and fiber implantations were performed on a stereotaxic frame (Kopf Instruments, Model 1900) under isoflurane anesthesia (1.5%-2% + oxygen at 1 L/min). Mouse body temperature was maintained at 30°C throughout surgery (Harvard Apparatus, #50722F). Briefly, fur over the skull was removed with depilatory cream, and the skin cleaned with 70% isopropyl alcohol and betadine, after which a small anterior/posterior incision was made to expose the skull. Small (0.5 mm) holes were drilled above the target coordinates and a pulled glass needle was lowered into the injection site (AP: 0.75 mm, ML: 1.25 mm, DV: -2.85 mm, unilaterally (fiber photometry and optogenetic experiments) or bilaterally (all other experiments)). 500 nL of specific adeno-associated virus (see individual experimental details below) was infused at 125 nl/min using a microinfusion pump (Harvard Apparatus, #70-3007), and the injection needle was removed 10 min after termination of viral infusion. For fiberoptic cannula implantation, 2 small screws were secured into the skull and a 200 μ m (optogenetic) or 400 μ m (photometry) fiberoptic cannula was lowered into the injection site and held with dental cement (Den-Mat, Geristore A and B). Mice were given a minimum of 7d (Kir2.1 overexpression) or 3 weeks (all other experiments) to recover from surgery prior to subsequent experimental testing.

General Operant Behavior Methods

Here we outline the general operant behavior methods which are shared by all experimental manipulations. Specific details regarding each experimental manipulation are provided afterward. Unless specifically noted, all operant experiments employed the same behavioral methods for initial goal-directed acquisition and single reversal.

Mice were first food deprived to ~85% of free feeding weight prior to beginning of operant training. Operant experiments were conducted in a modular chamber (21.59 × 18.08 × 12.7 cm; Med Associates Inc, Model ENV307W). Each chamber was equipped with a modified reward magazine, where a pump (Med Associates Inc, Model PHM-100) delivered 10 μ L chocolate liquid reward (Nestlé Boost) into custom-made receptacles. Retractable levers were located on either side of the magazine port, and nosepokes were located on the opposite wall (see Figure 4E).

Reward Magazine Training (RMT) Sessions

In order to familiarize the mice with the operant chambers and liquid reward, mice were initially delivered $10 \,\mu$ L reward once per minute, paired with 10 s magazine light illumination (Figure 1H). Sessions were 40 min in length, and were conducted for a minimum of 2 days until mice had fewer than 10 omissions (trials in which mice did not retrieve the reward within 10 s of magazine light illumination).

In order to rule out any effects of possible Pavlovian conditioning of this task on subsequent operant learning, the light was omitted during the RMT sessions in a separate cohort of mice (n = 10 eGFP controls, 9 Kir2.1 overexpression). Other than the omission of the light pairing, this cohort of mice underwent identical procedures to the initial Kir2.1 overexpression experiment (see below).

Fixed Ratio FR1 Self-Initiated Two-Choice Operant Task

After completion of one session of RMT session with less than 10 omissions (2 or 3 days), mice were then trained in an FR1 self-initiated two-choice operant task (Figure 1K). The task structure was identical for all experiments and consisted of four phases: (1) Intertrial Interval (ITI): All stimulus lights were off for 5 s between each trial. (2) Initiation: After the 5 s ITI, the magazine light was illuminated, and nosepoke in the magazine was considered an initiation. (3) Choice: Initiation prompted the extension of both retractable levers. Mice had 10 s to press a lever (choice). If the mouse did not press a lever within 10 s (omission), levers were retracted and the ITI phase began. (4) Outcome: Mice were randomly assigned a correct lever (left or right). Pressing the correct lever resulted in 10 μ L chocolate reward, paired with 5 s magazine port illumination, followed by ITI. Pressing the incorrect lever ended the trial, leading to the ITI.

FR sessions were 60 min in duration, and acquisition was operationally defined as the 50th reward. Following acquisition, mice were given 2 more days to perform the initially acquired action-outcome contingency task. The following session, mice underwent a single reversal, whereby the initially correct lever was now incorrect and unreinforced, and the previously incorrect lever was then reinforced. Successful reversal was operationally defined as 8/10 choices to the new correct side.

Following experimental testing, mice were transcardially perfused with 4% formalin/PBS and 50 µm slices were cut on a vibratome (Vibratome, model 1000plus) and mounted for verification of targeting. Post hoc re-assignment of animals based on histology was performed in a blind fashion by an independent investigator.

General Electrophysiology Methods

Specific electrophysiological methods will be detailed in sections below for each individual experiment. Our general procedures have been described previously (Choi et al., 2019). Briefly, mice were deeply anesthetized and transcardially perfused with ice-cold aCSF (124 mM NaCl, 1.2 mM NaH₂PO₄, 2.5 mM NaHCO₃, 5 mM HEPES, 13 mM glucose, 1.3 mM MgSO₄, 2.5 mM CaCl₂). The brain was then quickly removed and coronally sectioned (250-300 µm) on a vibratome (Leica, Model VT1200s). Slices were incubated for

12-15 min at 32°C in an NMDG-based recovery solution (92 mM NMDG, 2.5 mM KCl, 1.2 mM NaH₂PO₄, 30 mM NaHCO₃, 20 mM HEPES, 25 mM glucose, 5 mM sodium ascorbate, 2 mM thiourea, 3 mM sodium pyruvate, 10 mM MgSO₄, 0.5 mM CaCl₂), then transferred to room temperature (20-22°C) aCSF for at least 1h before recording. For recording, slices were placed in a recording chamber, fully submerged in oxygenated (95% O₂, 5% CO₂) aCSF at a flow rate of 1.4-1.6 mL/min, and maintained at 29-30°C. Voltage clamp (VC) recordings were conducted with pulled borosilicate glass (World Precision Instruments, TW150-3) recording pipettes with tip resistance of 3-5 MΩ when filled with internal solution (135 mM CsCl, 10 mM HEPES, 2.5 mM MgCl, 0.6 mM EGTA, 1 mM QX-314, 10 mM Na-phosphocreatine, 4 mM NaATP, 0.3 mM NaGTP, 0.1 mM spermine, pH adjusted to 7.3-7.4 with CsOH). Current clamp (CC) recordings were made with electrodes filled with 140 mM K-gluconate, 5 mM KCl, 2 mM MgCl₂, 0.2 mM EGTA, 10 mM HEPES, 10 mM Na-phosphocreatine, 4 mM Mg-ATP, 0.3 mM NaGTP, pH adjusted to 7.3 with KOH. Cell-attached recordings were made with electrodes filled with 20 mM K-gluconate, 5 mM KCl, 2 mM MgCl₂, 0.2 mM EGTA, 10 mM HEPES, 10 mM Na-phosphocreatine, 4 mM Mg-ATP, 0.3 mM NaGTP, pH adjusted to 7.3 with KOH. Cell-attached recordings were made with electrodes filled with 140 mM K-gluconate, 5 mM KCl, 2 mM MgCl₂, 0.2 mM EGTA, 10 mM HEPES, 10 mM Na-phosphocreatine, 4 mM Mg-ATP, 0.3 mM NaGTP, pH adjusted to 7.3 with KOH. Cell-attached recordings were made with electrodes filled with aCSF. All electrophysiology recordings were sampled at 20kHz, filtered at 2.8kHz and analyzed offline with Recording Artist (Rick Gerkin) running on Igor 6.37 (Wavemetrics).

Immunohistochemistry

Somatostatin immunohistochemistry methods were described previously (Choi et al., 2019). Briefly, mice were trans-cardially perfused with 4% formalin/PBS and 30 µm slices were cut on a vibratome (Vibratome, Model 1000plus). Free floating slices were permeabilized in 0.6% Triton X-100 and blocked in 3% normal goat serum in PBS for 1h. Primary antibody (rat monoclonal anti-somatostatin, 1:500, Millipore, #MAB354, RRID:AB_2255365) was incubated overnight in 1% normal goat serum and 0.2% Triton X-100 in PBS. Slices were washed then incubated in secondary antibody (goat anti-rat IgG (H+L), Alexa Fluor 555 conjugate, RRID:AB_141733, 1:500) for 2h, then mounted and scanned on a standard epifluorescent microscope (Olympus, BX63) under 10x (Olympus, 0.4NA) and 20x (Olympus, 0.75NA) objectives. Colocalization of SST immunoreactivity and virally-expressed eGFP was performed on representative images at 20x and quantified manually.

Specific Experimental Methods

Fiber Photometry

Mice were unilaterally injected with AAV5.Syn.Flex.GCaMP6f.WPRE.SV40 (Penn Vector Core, n = 7) or AAV1.CAG.Flex.eGFP. WPRE.bGH (Penn Vector Core, n = 4) and implanted with a fiberoptic cannula (400 μ m, 0.48NA, 4.1mm in length; Doric Lenses). Viruses were allowed to express for three weeks. Mice initially underwent baseline imaging sessions in their home cage to assess LTSI Ca²⁺ activity in response to locomotor movement, novel environment, and social interaction (see 'Baseline Sessions' below). Next, mice were food deprived and underwent reward magazine training (RMT) sessions and operant training (see 'Appetitive Reward' and 'Operant Task Acquisition' below).

Signal collection

Mice were attached via an optical fiber (400 μ m core, 0.48 NA; Doric Lenses) to a Doric 4-port minicube (FMC4, Doric Lenses). Blue (470 nm wavelength for GCaMP6f stimulation, ThorLabs #MF470F3) and violet (405 nm wavelength for artifact control fluorescence, ThorLabs #MF405FP1) LED light was delivered to the brain at 10-100 μ W (LED driver: Thor Labs, Model DC4104). Emissions passed through a dichroic mirror, a 500-550 nm cut filter and were then detected by a femtowatt silicon photoreceiver (Newport, Model 2151). Analog signals were demodulated and recorded (Tucker Davis Technologies, RZ5 processor and Synapse Software). To reduce the autofluorescence of patchcord fibers, 470nm light was passed through for a minimum of 4 hours prior to recording. Mice were always hooked up to the optical fiber with the LEDs on for at least 10min prior to recordings began. Once signal collection recordings were started, there was a baseline period of at least 6 min prior to the introduction of any stimuli or the initiation of any task in the operant chamber.

Baseline Sessions

Mice initially underwent baseline imaging in their homecage. After three days of habituation to handling and optical fiber tethering, photometry signal and behavioral video was collected in a 12 min session where mice were allowed to freely move around the homecage. Videos were subsequently analyzed for immobility, small (head or forepaw), or large (walking and rearing) movements with Behavioral Observation Research Interactive Software (BORIS; Friard and Gamba, 2016). Next, three separate 12 min sessions were recorded with experimental mice exposed to an unknown male, unknown female, or novel bedding (Shepherd Specialty Papers, ALPHA-dri bedding) introduced at 6 min. Finally, on a separate day, a subset of 4/7 mice were recorded for a 2h session in their homecage while freely moving to assess the stability of the signal over prolonged periods of recording.

Appetitive Reward

After the aforementioned baseline testing, mice were food deprived to 85% of free feeding weight over the course of 1 week. To habituate mice to the operant chambers, they were tethered to the optical fiber and allowed to explore the chamber for 60 m. The following day, they were attached to the optical fiber and placed in the operant chamber and allowed to habituate for at least 10min with the LEDs on prior to the recording session began. After the recording began, there was a 10 min baseline period, followed by the 40min reward magazine training (RMT) session described above. This was repeated for 3d.

Operant Task Acquisition

After 3d of the appetitive RMT sessions, mice began the self-initiated two-choice FR1 operant task, as described above. Again, mice were allowed to habituate to the chambers for at least 10 min with the LEDs on prior to beginning the photometric recordings, and the first 6 min of the recordings was a baseline period before the operant task was started. The first session in which a mouse obtained 50 rewards was used for analysis; some mice took multiple sessions to acquire the task. Sessions were allowed to extend beyond

60 min (but not longer than 130 min) if the mouse had started to acquire the task but had not obtained 50 rewards by the end of 60 min. 5/7 mice were recorded an additional two days, but did not complete reversal training.

Signal analysis

Demodulated 470nm and 405nm signals were processed and analyzed using custom scripts written in MATLAB (MathWorks, Version R2017b). Data were down-sampled to 40 samples/s and digitally filtered. To account for a steady decrease in baseline auto-fluorescence over prolonged recording sessions, the signals (405 and 470nm channels) at the end of the recording were baselined to zero and the data fit with a cubic polynomial curve, which was then subtracted from both respective raw signals. Afterward, the control signal (405nm) was subtracted from the GCaMP6f signal (470nm) to output the Δ F/F.

In order to assess general, non-event-locked Ca^{2+} events and how they changed over time, we used custom peak detection scripts, which were adapted from Muir et al. (2018) to use a 10 s moving window for thresholding. High amplitude events (local maxima greater than two median average deviations (MADs) above the median of a 10 s moving window) were removed and a baseline moving median was calculated. Peaks were considered to be events with local maxima greater than 3 MADs above the baseline moving median. Peak amplitude was calculated as the difference between the maxima and the local median. Three 5 min recordings were conducted of optical fibers in darkness without any mouse attached, and detected peak amplitude was 0.536 ± 0.007 (mean ± SEM). Therefore, in subsequent analyses for event frequency and amplitude, only peaks greater than 0.536 were included.

In order to assess the Ca²⁺ activity tied to discrete behavioral events, we used peri-event temporal histogram (PETH) analysis as reported previously (Cui et al., 2013). The Δ F/F signal was aligned to time 0 for each behavioral timestamp, and a 10 s window (5 s before and after the event) was extracted. Positive areas under the curve (AUC) were calculated for each 10 s PETH trace.

Kir2.1 Overexpression-Mediated Inhibition

Mice were bilaterally injected with AAVDJ.EF1a.DIO.zsGreen.p2A.ESKir2.1.WPRE.hGH (plasmid generously provided by B. Lim, virally packaged in house; n = 9, behavior, n = 3 electrophysiology) or AAV1.CAG.Flex.eGFP.WPRE.bGH (Penn Vector Core; n = 9 behavior, n = 3 electrophysiology). After 7d viral expression, mice either underwent electrophysiological recordings or operant training.

Electrophysiology

To determine whether SST-Cre+ cells were LTSIs, we patched eGFP+ neurons in SST-IRES-Cre mice and searched for the following previously described characteristics (Tepper et al., 2010): (1) depolarized resting membrane potential at break-in, (2) high input resistance, measured as the slope of the I-V plot at hyperpolarizing current injections; (3) presence of Ih-mediated voltage sag, revealed by measuring the membrane potential difference between the beginning and end of a 500ms –70pA current step; (4) presence of plateau depolarizing envelope or spiking activity following hyperpolarizing current step, measured at current injections of –50, –60 and –70pA.

To test the effects of Kir2.1 on the excitability of striatal interneurons, zsGreen+ or eGFP+ neurons were patched in whole cell VC configuration in the absence of synaptic blockers and (1) resting membrane potential, input resistance and spontaneous activity shortly after break-in were noted; (2) 500ms current steps were injected and number of resulting APs was noted. All comparisons were made between Kir2.1 and eGFP-expressing neurons.

Operant Behaviors

Mice were food deprived to 85% free feeding weight and underwent 2-3 sessions of reward magazine training (RMT) followed by FR1 self-initiated two-choice operant acquisition, as described in 'General Operant Behavior Methods' above. Briefly, mice first underwent 2-3 reward magazine training sessions, followed by the FR1 self-initiated two-choice operant task. After acquisition, mice were allowed two additional days at FR1 at their acquired lever pressing contingency, after which the contingency was reversed in a single reversal session.

Following the single reversal session, mice were tested on a progressive ratio PR4 (after each lever press, the required number of presses to obtain reinforcement increased by 4) contingency in three 60 min sessions to determine whether LTSI inhibition altered motivational responding, and the average number of rewards across the three sessions was averaged (Figure S2I).

To assess whether the RMT sessions generated a potentially competing Pavlovian association, a separate cohort of mice (n = 10 eGFP, n = 9 Kir2.1) underwent 3d of modified RMT sessions, where the reward was delivered every 60 s without associated magazine light illumination (Figure S1J). After 3d RMT, mice were trained in the same operant task as above.

Force Plate Actometer for Locomotor Behavior

Force-plate actometer assays were performed using previously methods previously described in detail (Tischfield et al., 2017). Mice were placed on an open field (42cm x 42cm) with 4 transducers (sampled at 200 scans/s) and allowed to freely explore for two 20 min habituation sessions prior to the test day. On the test day, mice were allowed to acclimate to the open field for 10 min, followed by a 20 min recorded session.

eNpHR3.0-Mediated Inhibition

Mice were unilaterally injected with AAV5.hSyn.eNpHR3.0-eYFP.WPRE.hGH (Penn Vector Core; n = 8 behavior, n = 3 electrophysiology) AAV1.CAG.Flex.eGFP.WPRE.bGH (Penn Vector Core; n = 8 behavior, n = 3 electrophysiology) and implanted with a 200 μ m optic fiber in the DMS (except for electrophysiological validation). Virus was allowed to express for 3 weeks prior to testing.

Electrophysiology

To test the efficacy of halorhodopsin-mediated optogenetic inhibition, cell attached recordings of EYFP-expressing LTSIs were made in the absence of synaptic blockers. As we were particularly interested in the temporal precision and potential after-effects of a brief optical manipulation, we employed cell-attached recording configuration, as those permitted the most stable assessment of LTSI firing over the entire optogenetic and post-optogenetic period. We compared two 10 s windows of no illumination that surrounded a 4 s window of full-field 530nm illumination through the 40x objective (Olympus, 0.8NA water immersion). The total duration of each sweep was set to match the average trial duration during our halorhodopsin behavioral sessions. Neuronal spiking was detected by Neuromatic (v.3.0, Jason Rothman) for Igor 6.3 and firing frequency for each bin was calculated by counting the total number of spikes in 4-5 traces and dividing by the total recording time. These individual bin rates for each neuron were used to create averages and error terms across recorded cells.

Operant Behaviors

Mice were food deprived to 85% free feeding weight and habituated to handling. Mice were attached to an optical fiber, which was attached to a 1x1 fiberoptic rotary joint (Doric Lenses FRJ_1x1_FC-FC), which was connected to a 532nm laser (Shanghai Dream Lasers, SDL-532-100T). The lasers were controlled by Arduino Uno microcontroller boards and programmed to deliver 4 s constant illumination upon TTL input from the MedAssociates breakout box. Lasers were turned on 30 min prior to use to enhance output stability and adjusted to deliver \sim 5mW light output at the mouth of the optical fiber.

Mice were placed in the chamber for 40min for habituation on the day before training began. Next, mice underwent 2-3 reward magazine training (RMT) sessions as described in 'General Operant Behaviors'. Lasers were not used during these sessions. During the FR1 self-initiated two-choice operant task (described in 'General Operant Behavior Methods'), response on the correct lever triggered 4 s constant illumination coinciding with reward delivery. As with the Kir2.1 overexpression experiments, after mice acquired they were given an additional two days at the initial contingency, followed by a single reversal day.

After reversal mice were given one additional 60min FR1 session where one lever resulted in chocolate reward and the other resulted in laser illumination, in order to test preference for reward over laser stimulation. After this session, mice were trained on a new contingency, where they then had to nosepoke on the opposite side of the box to receive reward in the magazine port (Figure 3E). eNpHR3.0 mice were randomly assigned to receive laser illumination or no laser during reward; all eGFP mice received laser during reward.

Force Plate Actometer

As described in the 'Force Plate Actometer' section of the 'Kir2.1-Mediated Inhibition' experimental details, mice were habituated to the force plate for two 20 min sessions prior to the test day. Mice were attached to a 1x1 fiberoptic rotary joint (Doric Lenses FRJ_1x1_FC-FC) above the center of the force plate, which was connected to a 532nm laser (Shanghai Dream Lasers, SDL-532-100T). Lasers were controlled by Arduino Uno microcontroller boards and programmed to deliver 4 s constant illumination every 15 s during minutes 5-15 of the 20 min session. This laser pattern was selected to model the patterns of laser illumination in the operant task.

Spatial Self-Stimulation

In order to assess whether halorhodopsin-mediated LTSI inhibition was itself rewarding or reinforcing, a separate cohort of female mice (n = 8) were injected with AAV5.hSyn.eNpHR3.0-eYFP.WPRE.hGH (Penn Vector Core) and implanted with a 200 μ m optic fiber in the DMS. After 3 weeks of viral expression, mice were allowed to freely explore the open field (42cm x 42cm) while attached to the optical fiber in two 20 min baseline habituation sessions. The optical fiber was attached to a 1x1 fiberoptic rotary joint (Doric Lenses FRJ_1x1_FC-FC) above the center of the open field, which was connected to a 532nm laser (Shanghai Dream Lasers, SDL-532-100T). The open field was divided into four equal quadrants (see Figure S2T), and mice were tracked using Biobserve Viewer3 software.

Methods for the spatial self-stimulation experiment were adapted from Carta et al. (2019). On test day, one of the four quadrants was randomly assigned for each mouse as the stimulation quadrant prior to the beginning of the experiment. Stimulation quadrants were balanced across the 8 mice. Upon placement in the open field, mice were first allowed to habituate for 10 min without recording. Then, a 10 min baseline period was recorded, where mice were tracked during free exploration of the open field. After the baseline, a 10 min stimulation period was recorded, where crossing into the pre-assigned stimulation quadrant triggered a loop of 4 s constant illuminations, separated by 10 s. Crossing out of the stimulation quadrant immediately terminated the loop. This later pattern was selected because (1) it models the patterns of laser illumination period, an additional 10 min post-test period was recorded to measure any long-lasting preference for the stimulation quadrant. Carta et al. (2019) demonstrated that mice robustly demonstrate a strong place preference for a stimulation quadrant within one day when stimulation was delivered for either the entire duration or just upon crossing into the quadrant. Data were analyzed as proportion of time spent in the stimulation quadrant, compared to the average of the proportions of time spent in the remaining three quadrants.

ChR2-Mediated Excitation

In order to determine whether excitation of LTSIs altered goal-directed learning, we injected SST-Cre mice with AAV5.EF1a.DIOhChR2(H134R)-eYFP.WPRE.hGH (Penn Vector Core; n = 9 behavior, n = 3 electrophysiology) or AAV1.CAG.Flex-eGFP.WPRE.bGH (n = 10 behavior, n = 3 electrophysiology) and implanted them with a 200 μ m optic fiber in the DMS (except for electrophysiological validation). Virus was allowed to express for at least 3 weeks prior to testing.

Electrophysiology

To delineate the appropriate stimulation protocol for maintaining LTSI activity during our behavior, we recorded from ChR+ neurons in acute striatal slices from SST-IRES-Cre mice, looking for patterns that (1) could be followed during stimulation and (2) did not lead to prolonged post-stimulation periods of spiking reduction. Cell attached recordings of ChR-expressing neurons were made in the absence of synaptic blockers (see rationale for cell-attached above), with a 10 s pre-stimulus window followed by 4 s long full-field 470nm illumination through the 40x objective (Olympus, 0.8NA water immersion) at either 10Hz or 20Hz (1ms pulse width). Analysis of firing rates was conducted as for NpHR experiment.

Operant behaviors

Procedures were identical to those detailed in 'eNpHR3.0-Mediated Inhibition', except that 473nm lasers (LaserGlow Technologies, LRS-0473-GFM-00100-0) were triggered to deliver 10Hz stimulation for 4 s upon reward delivery (1ms pulse width), and experiments did not continue past the single reversal.

Conditional VGAT Deletion

For electrophysiological and *in situ* hybridization validation that conditional VGAT deletion from LTSIs reduced both VGAT mRNA expression and GABAergic release, SST^{Flp/-};VGAT^{+/+} (n = 4 *in situ* hybridization, n = 4 electrophysiology) and SST^{Flp/-};VGAT^{c/c} (n = 4 *in situ* hybridization, n = 4 electrophysiology) mice were bilaterally injected with AAVDJ.EF1a.fDIO.GFP-Cre (plasmid kindly provided by Bo Li, virally packaged in house) and AAV5.EF1a.DIO-hChR2(H134R)-eYFP.WPRE.hGH (Penn Vector Core). For behavioral experiments, SST^{-/-};VGAT^{c/c} (n = 12) and SST^{Flp/-};VGAT^{c/c} (n = 12) mice were injected with AAVDJ.EF1a.fDIO.GFP-Cre to determine if the effects of LTSI inhibition on goal-directed learning were mediated at least in part by GABAergic release. Viruses were allowed to express at least three weeks prior to experimentation.

Electrophysiology

To test the efficacy of VGAT deletion on removing LTSI GABAergic transmission, EYFP-negative neurons (putative SPNs) were patched in whole cell VC configuration (VC hold = -70 mV) without synaptic blockers. A 1ms pulse of 470nm light through the 40x objective was used to recruit LTSIs while recording inward currents (VC internal Cl⁻ reversal ~0mV) in neighboring patched cells. Currents were recorded at 3 increasing LED intensities. Sequential application of NBQX (50 μ M) and picrotoxin (100 μ M) was used during optogenetic stimulation to confirm the identity of inward synaptic currents. For all measures, liquid junction potentials were not corrected but measured to be ~-10mV.

In Situ Hybridization and Immunohistochemistry

To further validate that conditional VGAT deletion from SST cells eliminated VGAT mRNA expression, we performed *in situ* hybridization (ISH) using RNAscope Multiplex Fluorescent Reagent Kit v2 (Advanced Cell Diagnostics, #323100) combined with immunohistochemistry. SST^{Flp/+};VGAT^{-/-} and SST^{Flp/+};VGAT^{+/+} mice were injected with 300nl AAVDJ.EF1a.fDIO.GFP-Cre into the DMS. As the fDIO.GFP-Cre expression is nuclear, it was mixed with AAV1.CAG.Flex.GFP.WPRE.bGH (1:1) to generate a somatic mask for subsequent analysis of SIc32a1 (VGAT) mRNA expression. Virus was allowed to express for at least three weeks, after which mice were perfused with 10% formalin and brains extracted, stored in formalin overnight, followed by 30% sucrose. Brains were frozen and 25 μ m cryosections were collected directly onto charged slides and stored at -80° C until use. The RNAscope Multiplex Fluorescent v2 assay was performed according to manufacturer's instruction for fresh-fixed tissue, and all incubation steps were conducted in a hybridization oven (HybEZ, ACD). Target probes for *Slc32a1 (Vgat)* were used, as well as two highly characterized housekeeping genes for positive control (*UBC*, ubiquitin C; *POLR2A*, DNA-directed RNA polymerase II subunit RPB1), and one negative control (*DAPB*, dihydrodipicolinate reductase, a gene from the soil bacterium *Bacillus subtilis*, strain SBY, that is not present in mouse tissue).

In order to recover the GFP signal that was largely quenched by ISH protocol, we next performed immunohistochemistry for GFP. Tissue was first blocked with 10% normal goat serum+0.05% triton-x. Primary antibody (chicken anti-GFP, abcam ab13970, RRID:AB_300798) was diluted 1:500 in 1% normal goat serum + 0.05% triton-x and incubated overnight at room temperature. Secondary antibody (goat anti-chicken, conjugated to Alexa Fluor 488, Jackson ImmunoResearch #103-545-555, RRID:AB_2337390) was diluted 1:250 in 1% normal goat serum + 0.05% triton-x and incubated for two hours at room temperature.

Slides were scanned on a standard epifluorescent microscope (Olympus, BX63) under a 20x (Olympus, 0.75NA) objective. 2-3 representative images were captured from each section, and quantification of GFP+;VGAT+ and GFP+;VGAT- cells was conducted manually.

QUANTIFICATION AND STATISTICAL ANALYSIS

All data were analyzed with Graphpad Prism v7.0. Appropriate t tests (paired and unpaired) and ANOVAs (one-way, two-way, and repeated-measures) were performed as indicated in the results and supplementary table. ANOVAs with significant main effects/interactions were followed up with Bonferroni multiple comparisons post hoc analyses. Significant effects and p values are indicated in the figures and legends, and specific statistical methods, sample sizes, and effect sizes are located in Table S1.