



HHS Public Access

Author manuscript

Curr Biol. Author manuscript; available in PMC 2019 February 19.

Published in final edited form as:

Curr Biol. 2018 February 19; 28(4): 560–573.e5. doi:10.1016/j.cub.2018.01.031.

Inversely Active Striatal Projection Neurons and Interneurons Selectively Delimit Useful Behavioral Sequences

Nune Martiros^{1,2}, Alexandra A. Burgess^{1,2}, and Ann M. Graybiel^{1,2,3,*}

¹McGovern Institute for Brain Research, Massachusetts Institute of Technology, 43 Vassar St, Cambridge, MA 02139, USA

²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 43 Vassar St, Cambridge, MA 02139, USA

SUMMARY

Understanding neural representations of behavioral routines is critical for understanding complex behavior in health and disease. We demonstrate here that accentuated activity of striatal projection neurons (SPNs) at the beginning and end of such behavioral repertoires is a supraordinate representation specifically marking previously rewarded behavioral sequences independent of the individual movements making up the behavior. We recorded spike activity in the striatum and primary motor cortex as individual rats learned specific rewarded lever-press sequences, each one unique to a given rat. Motor cortical neurons mainly responded in relation to specific movements, regardless of their sequence of occurrence. By contrast, striatal SPN populations in each rat fired preferentially at the initiation and termination of its acquired sequence. Critically, the SPNs did not exhibit this bracketing signal when the same rats performed unreinforced sequences containing the same sub-movements that were present in their acquired sequence. Thus, the SPN activity was specifically related to a given repetitively reinforced movement sequence. This striatal beginning-and-end activity did not appear to be dependent on motor cortical inputs. However, strikingly, simultaneously recorded fast-spiking striatal interneurons (FSIs) showed equally selective but inverse firing patterns: they fired in-between the initiation and termination of the acquired sequences. These findings suggest that the striatum contains networks of neurons representing acquired sequences of behavior at a level of abstraction higher than that of the individual movements making up the sequence. We propose that such SPN-FSI networks of the striatum could underlie the acquisition of chunked behavioral units.

eTOC Blurbs

*Correspondence: graybiel@mit.edu.

³Lead Contact

AUTHOR CONTRIBUTIONS

N.M. and A.M.G. designed the experiments, N.M. and A.A.B. performed the experiments, N.M. analyzed the data with A.M.G. interactions, and N.M. and A.M.G. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Habitual behavioral routines acquired by repetition are ubiquitous in daily life. Martiros et al. demonstrate that such learned habits and skills become marked as behavioral units, regardless of their movement content, by intrastriatal neuronal networks. After learning, this neuronal signal is not strongly dependent on motor cortex.

Keywords

dorsolateral striatum; basal ganglia; habit; chunking; corticostriatal; sequence learning

INTRODUCTION

Libraries of learned behavioral repertoires serve as the building blocks of much of our daily activities, allowing allocation of attention to new or complex tasks while minimizing the effort toward accomplishing well-rehearsed, optimized tasks. The process of learning action sequences that are composed of multiple distinct actions results in these actions becoming “chunked” such that the individual actions making up the sequence are bound together into a single behavioral unit [1, 2], and the behavior becoming engrained with extensive repetition. Although habitual behavioral programs have powerful control over our daily lives, and their disruption is a key common feature in wide range of neuropsychiatric disorders [3–6], how the brain represents such chunked units of behavior is not well understood.

Most models suggest a prominent role for cortico-basal ganglia loops in the learning and performance of such behavioral routines [7–13]. In order to compare the neuronal representations of learned action sequences at different levels in the circuit, we recorded simultaneously from two nodes in this circuit — motor cortex and the dorsolateral striatum — and we performed simultaneous recordings of putative striatal projection neurons (SPNs) and fast-spiking interneurons (FSIs). The dorsolateral striatum has been found to be important in the transition of behavior from goal-directed to habitual, and in stereotypic, repetitive behaviors [8, 10, 14–20]. Recordings from the dorsolateral striatum, both during the running of a T-maze by rats and mice [18, 21–24] and during performance of a fixed-ratio lever-press task by mice [25, 26], have shown that neuronal firing peaks at the beginning and the end of the trials, potentially framing a chunked behavioral sequence. However, as in motor cortex, neurons in the dorsolateral striatum have also been shown to have patterns of activity that are correlated with movement [27–32]. This activity can change with learning, with the motor representations diminishing with experience [33–35].

Here we aimed to disambiguate movement-related activity from higher level “chunking” or habit-related activity in this corticostriatal circuit, and to test whether motor cortical representations of a learned action sequence are different from those in the dorsolateral striatum. To do so, we designed a task in which rats learned to perform a specific series of three lever-press actions in order to receive reward. The sequences were different for each rat, but for all of the sequences, lever-press movements were present at the beginning, middle, and end of the acquired sequence, as well as in trials during which rats performed unrewarded incorrect sequences. As a result of this experimental design, we could use the lever presses performed at different time points in reinforced and unreinforced sequences to

identify striatal spiking activity related to movements associated with specific lever presses as opposed to activity related to the supraordinate lever-press sequence. Crucially, because each rat was trained on a single correct lever-press sequence, we could explore whether there were patterns of neuronal activity that generalized across rats performing different learned action sequences.

The results of our simultaneous recordings in the striatum and motor cortex during the performance of this specially designed task demonstrate that the dorsolateral striatum contains complementary SPN and FSI representations of entire learned action sequences, whereas the motor cortex representations are more closely linked to the sub-components of the sequence regardless of reinforcement history. These results raise the possibility that striatal SPN-FSI networks are key to the performance of well-learned behavioral sequences and habits through complementary network activity patterns.

RESULTS

Rats were trained to perform single lever presses prior to implantation of the recording drives (see STAR Methods). After drive implantation, training on the three-lever-press sequence began in an operant chamber equipped with two levers (Figure 1A). One specific sequence of three lever presses requiring the use of both levers was assigned to each rat, and the sequences were different for different rats. The task was self-paced, and after every three lever presses, the rats were provided with a feedback cue. If the rat performed the correct sequence, it received an auditory click cue followed by a chocolate milk reward. If the rat performed any other sequence of three presses, it received a white noise auditory cue. The rats' performance began below chance level (12.5%, 1 of 8 possible sequences) and improved over the course of 5–6 weeks of training ($p < 0.001$, Wilcoxon rank-sum test; Figure 1B). Within single sessions, well-trained rats exhibited potential for high levels of performance during periods in which their performance was $> 80\%$ correct (Figure 1B), and they performed a correct trial every 11 s on average (Figure 1C), including an average of 5 s for reward consumption.

Rats Learn Individualized Stereotyped Movement Patterns to Execute Rewarded Lever-Press Sequence

While performing the learned sequence, the rats' head position moved in stereotyped patterns that were similar from trial to trial and from session to session (Figures 1D and 1E). These movement patterns varied greatly among rats, based on the sequence that they learned and the specific movements that they developed to successfully execute their assigned sequence. Day-to-day mean trajectories and trial-to-trial trajectories within rats performing their learned sequence were highly correlated (Figure 1F, Pearson's coefficients (r) of 0.92 and 0.83, respectively). These correlations were significantly higher than those found when rats were performing specific alternative incorrect sequences ($r = 0.83$ vs. $r = 0.67$, $p < 0.05$, Wilcoxon rank-sum test). We tested a separate group of four rats for devaluation resistance using a classic reward devaluation procedure as a test of habit learning [36] by providing them with unlimited chocolate milk in the home cage for 2 hr prior to placing them in the operant chamber. With training, the behavior of these rats became resistant to reward

devaluation (Figure 1G), resulting in their continuing to perform lever presses even when satiated on chocolate milk ($p < 0.05$, Wilcoxon rank-sum test). Together, the slow trial-and-error learning process, the development of stereotyped movement patterns, and the development of devaluation resistance suggested that the rats were undergoing habit formation throughout the course of the lever-press sequence training.

Motor Cortical Neurons Are More Likely than Striatal Neurons to Be Activated during Individual Lever-Press Actions Irrespective of the Motor Program within Which Those Actions Are Embedded

The question of how neuronal representations are transformed at each successive node in the corticobasal ganglia-thalamic circuit is a fundamental one for understanding the mechanisms underlying the function of these circuits. We recorded from two nodes within this circuit throughout the course of the lever-press sequence training: (1) the dorsolateral striatal region in which task-bracketing activity has been previously observed [21, 23], and (2) a motor cortical region selected using the results of anatomical tracing studies [37, 38] and our own test injections to identify an area of forelimb motor cortex, as confirmed by cortical microstimulation in two rats, with projections to the target striatal region.

We assessed lever-press-related spiking in putative projection neurons in both regions using a step-wise procedure (see STAR Methods). First, we constructed peri-event histograms for each unit for each of the three lever presses in the correct sequence and in any incorrect sequences the rat performed 10+ times within a training session. We then determined whether each unit was significantly activated (> 2 SDs above baseline) in 250-ms windows centered on six lever-press-related time-points: before (-200 ms), during (0 ms), and after ($+100$ ms) presses of lever 1 and lever 2, and during transitions from lever 1 to lever 2 or vice versa (in each case, 250 ms after the first press and 250 ms prior to the second press). If the unit was activated significantly and similarly in a press-related time-point regardless of whether it occurred as the first, second, or third press in the rewarded sequence, or in any unrewarded sequences, then we considered that the spiking response of the unit could be primarily accounted for by the occurrence of the given lever press (Figure 2A). If, however, the unit responded strongly during the lever press in one condition, but weakly when the same lever-press event occurred at other times in the sequence or in other sequences, we considered that we could not account for the response of this unit by the occurrence of this action only (Figure 2B).

Twenty percent of 616 putative pyramidal neurons in the motor cortex and 7% of 1303 putative SPNs in the dorsolateral striatum met the criteria for lever-press-related units ($p < 0.00001$, for difference between the two regions, Fisher's exact test). Further, in sessions in which we had recorded 10+ putative motor cortical pyramidal neurons and 10+ striatal SPNs simultaneously, the proportion of simple press-related units in motor cortex was significantly higher than that in the striatum ($n = 15$ sessions, $p < 0.05$, Wilcoxon rank-sum test; Figure 2C). This difference in proportions of lever-press-related units in the motor cortex and striatum occurred despite the fact that the overall distribution of z-scores of the responses of striatal and motor cortical neurons was similar (Figure 2D), with means of 0.97 and 0.72, respectively. Thus, motor cortical neurons were more likely to respond consistently during

individual motor actions than were striatal neurons, which were more likely to modulate their responses based on the motor program within which those actions occurred.

We constructed three general linear models for the binned raw firing rates of each putative striatal SPN and motor cortical pyramidal neuron (see STAR methods). In the first model, we included only behavioral variables related to lever presses. In the second model, we added contextual interaction factors between lever pressing, temporal position within the sequence, and whether the sequence was correct. In the third model, we added only variables indicating whether the rat was in the midst of performing the first press of the correct sequence or the last press of the correct sequence as related to task-bracketing previously observed in the striatum. To assess the explanatory power of these contextual interaction factors in the model, we compared the difference in the χ^2 statistic as a measure of model fit between the first and the second models, and between the first and the third models. The changes in the χ^2 statistic for SPNs were significantly greater than those for motor cortical pyramidal neurons after adding the contextual factors ($p < 0.000001$, Wilcoxon rank-sum test), and after adding the beginning and end factors ($p < 0.000001$, Wilcoxon rank-sum test; Figures 2E and 2F) further suggesting that striatal neurons were more strongly modulated by these factors than were the motor cortical neurons sampled.

SPNs in the Dorsolateral Striatum Are Most Active at the Beginning and End of Different Learned Lever-Press Sequences

Given the rare occurrence of simple lever-press-related SPNs in the dorsolateral striatum, we considered alternative accounts for the task-related spiking in this region. To do this, we constructed histograms of task-related activity for each unit by aligning consecutively the peri-event histograms for each of the task events in correct and incorrect trials. In each rat, we observed striatal SPNs that spiked preferentially at particular task times, most prominently those that were highly and selectively activated around the time of the first lever press and/or at the time of the last lever press (Figure 3A). When the same lever was pressed in the middle of the trial or repetitively in an incorrect trial, these units were less active (Figure 3B).

The mean firing rate of the entire SPN population across nine rats was highest around the time of the first and last lever presses in the correct sequence ($p < 0.01$ for first vs. second presses, and second vs. third presses, Wilcoxon rank-sum tests; Figures 3C and 3D, first column). This pattern occurred despite the fact that these nine rats were trained on different, unique sequences of lever presses. In addition to using fixed window-width peri-event histograms, we plotted the full time-course of the population activity across trial time by stretching or compressing the median time between each pair of successive lever presses to a 1-s duration (Figures S1A and S1B), and observed a similar beginning-and-end activation. This activation pattern was present early in training, as soon as there were sufficient numbers of correct trials for neural data analysis (Figure S1C), similar to the early development of the beginning-and-end pattern found in previous recording experiments with T-maze tasks [21, 23]. We further found that speed of movement was unlikely to account for the beginning-and-end activation due to a decrease in speed of the animals prior to each lever press in the sequence (Figures S1D and S1E). Importantly, it was the top 26% of the most responsive

SPNs, those for which the mean firing rates surpassed 5 Hz at some point in the course of the correct trial (644 of 2501 SPNs), that contributed to the increased spiking at the beginning and end of the learned sequences (Figures 3E and S1F). The mean firing rate of the other 1857 SPNs was 0.4 Hz, and was similar across the three presses (Figures 3D, first column dashed line, and S1F). These lesser responsive SPNs nevertheless had heterogeneous and significant modulations of firing rate which spanned task-time, but they were not the focus of our investigation here. Strikingly, the upper echelon of the highly task-responsive SPNs, those that exceeded 20-Hz spiking at some point in the correct trial ($n = 80$), did so almost exclusively during the time of the first (48%) or last (45%) lever presses (Figures 3F, S1F, and S1G).

Task-Boundary Activity Is Absent When Rats Perform Incorrect, Unrewarded Press Sequences

The same population of striatal SPNs that exhibited a preferential beginning-and-end response profile showed a remarkably different firing pattern when the rats were performing incorrect sequences within the same recording sessions. There were no significant differences between mean firing rates during the first, second, or third press in non-repeat incorrect sequences ($p > 0.6$ for first vs. second press, and second vs. third press, Wilcoxon rank-sum test; Figures 3C, 3D, S1A and S1B, second columns), or when they repeatedly pressed lever 1 or lever 2 (Figures 3C, 3D, S1A and S1B, third and fourth columns). The population firing rates of SPNs during the first and third lever press in incorrect trials were significantly lower than those in correct trials for the same SPNs, and their firing rates during the middle press were significantly higher than those in correct trials ($p < 0.0001$ for all three comparisons, Wilcoxon rank-sum test).

We performed several additional analyses (Figure S2) to assess the conditions that resulted in the expression of the striatal beginning-and-end activity. We found the beginning-and-end activation even in correct trials that occurred after two unrewarded trials; thus this activation was not dependent on recent rewards, nor was it due to the action of reward consumption prior to the start of the next trial (Figure S2, first column). However, in the trials during which the first two presses were performed correctly but the last press was incorrect, the SPN population spiking at the trial initiation was not significantly different from that in correct trials ($p = 0.35$, Wilcoxon rank-sum test), but the characteristic peak at the last lever press in the sequence was absent (Figure S2, second column, $p < 0.05$, Wilcoxon rank-sum test). Thus, if the rat failed to finish successfully the correct sequence, this putative completion signal was omitted. We then tested whether the task-boundary activation occurs in action sequences that were frequently performed but that lacked a history of reinforcement. In 89 sessions in which the most common incorrect sequence was performed frequently (a mean of 138 trials, as compared to a mean of 102 correct trials), there were no significant differences in the population spiking rates across the first, second, and third presses (Figure S2, third column), indicating that recent experience of frequent repetition is not sufficient for task-bracketing to occur. Differences in trial duration between correct and incorrect trials were also not responsible for the absence of the beginning-and-end activity during the performance of incorrect sequences (Figure S2, fourth column).

We also performed a second set of generalized linear regressions to fit the firing rate of striatal SPNs first with behavioral variables such as trial duration, number of trials, and others (see STAR methods), and subsequently added variables indicating whether the sequence was correct in a second model. After the addition of the variables indicating the correct/incorrect identity of the trial, we found significant increases in the fit of the model (mean χ^2 statistic increase of 164, $p = 0$, Wilcoxon signed rank test), and found that in 92% of SPNs (1572 of 1708) at least one of the coefficients for the reinforcement-related variables was significantly different from zero, indicating that the firing rate of the vast majority of the striatal SPNs was modulated depending on whether the rat was performing the correct sequence or not.

Finally, we took advantage of our task design, in which different rats were trained on different lever-press sequences, to compare striatal activity of rats trained on a particular sequence to that in other rats that were not trained on that sequence but performed it occasionally without receiving reward (Figure 4A). For each of the six possible sequences, the firing rates of SPNs in rats trained on the sequence were higher during the first and last presses, and lower during the middle lever press as compared to other rats performing the same sequence spontaneously (Figures 4B–4D).

Together, the absence of the task-boundary activation during incorrect trials within individual rats and the lack of this pattern in rats performing unreinforced matching sequences support the proposal that the beginning-and-end activation in the striatum is specifically related to the initiation and completion of learned reinforced behavioral programs.

The Dorsolateral Striatal Circuit Is Selectively Engaged during the Performance of Fixed Behavioral Movement Sequences

To assess the specificity of the engagement of this dorsolateral striatal habit-associated region in this task, we performed concurrent recordings from the dorsomedial striatum and dorsolateral striatum ($n = 3$ rats), and in prelimbic cortex ($n = 1$ rat), which projects to the dorsomedial striatum, but not to the dorsolateral striatum [37]. SPNs in the dorsomedial striatum were significantly less activated than those in the dorsolateral striatum ($p < 0.00001$, Wilcoxon rank-sum test; Figure S3A) during the performance of the learned sequences, despite the similar baseline firing rates in the two regions (0.47 Hz in dorsomedial striatum and 0.4 Hz in dorsolateral striatum). We did not observe clear task-boundary activation in the dorsomedial striatum (Figure S3A). Similarly, although baseline firing rates of putative pyramidal neurons in prelimbic cortex and motor cortex did not differ (1.9 Hz and 2.0 Hz, respectively), the activity in the motor cortex during the performance of the learned sequences was higher than that of prelimbic cortex ($p < 0.00001$, Wilcoxon rank-sum test; Figure S3B). These results add support for the hypothesis that the dorsolateral striatal circuit is highly and selectively engaged in marking the boundaries of learned behavioral repertoires, and that the behavioral task we developed here preferentially engaged this circuit.

Blanket Optogenetic Inhibition of Motor Cortical Cell Bodies and Terminals Had Weak Effects on Striatal Neuronal Firing

To assess the influence of motor cortex on striatal activity, we conducted optogenetic inhibition of motor cortex and its terminals in the striatum while we recorded in the dorsolateral striatum of the trained animals. We injected AAV5-CamKII-NpHR3-YFP virus into the forelimb motor cortex bilaterally, implanted recording drives after a four-week waiting period for halorhodopsin expression, and performed optogenetic manipulations during the 5–6 week training period. Six of the nine rats with striatal recordings also had these optogenetic manipulations and simultaneous striatal and motor cortical recordings. We confirmed that the motor cortical region that we targeted projected to the recording sites in the dorsolateral striatum via the overlap of the striatal tetrodes marked by post-experimental lesions and the terminal fields labeled by halorhodopsin-coupled yellow fluorescent protein (Figure 5A).

Pulsed yellow laser light (60 ms light on, 60 ms light off, 3 mW, in 10-s blocks) was delivered to a 200- μ m optical fiber placed at mid-depth in the motor cortex and surrounded by 10 tetrodes in freely moving rats (see STAR Methods). The mean firing rate of putative cortical pyramidal neurons was reduced from 3.3 Hz to 1.5 Hz at light onset (Figure 5B). Fifty-four percent of 383 putative cortical pyramidal neurons were significantly inhibited, and 8% were activated. The mean firing rate of 366 striatal SPNs recorded simultaneously was not changed by the cortical cell body inhibition (Figure 5B). The firing rates of the subset of SPNs with mean spiking rates of >1 Hz ($n = 106$, Figure 5B, inset) and the putative FSIs ($n = 109$, Figure 5B) were similarly unaffected by the cortical inhibition. However, 32% ($n = 34$) of these FSIs were slightly activated (Figure 5B, inset), possibly as a result of polysynaptic network effects. We repeated this manipulation using an optical fiber placed in the striatum and surrounded by 14 tetrodes to inhibit the corticostriatal terminals directly. The effect on striatal SPNs and FSIs was similarly weak, and a similar proportion of FSIs was subtly activated (Figure 5C). Thus, we found strong effects of the optogenetic silencing on the putative cortical pyramidal neurons, but only rare effects on SPNs (see also below).

Motor Cortical Axon-Terminal Silencing Did Not Detectably Affect Striatal Beginning-and-End Activity

We further assessed whether motor cortex might be driving the striatal beginning-and-end activity by using continuous laser light delivered to the intrastriatal optical fiber in-task to inhibit cortical terminals in a subset of trials. In some SPNs, the inhibition of the cortical terminals resulted in a significant task-time-specific firing rate modulation that recurred across 3+ recording sessions, suggesting a reliable effect (Figures 6A and S4A–S4D): 16 of 34 of the tetrodes from which we recorded 10+ SPNs had such recurring effects across 3+ sessions. However, for the majority of the SPNs, the task-related spiking in laser-off times was robustly replicated during laser-on times (Figure 6B). The proportion of SPNs that were significantly modulated by the inhibition of motor cortical terminals in the striatum at each time-point in the trial was $< 5\%$ (Figure 6C). As a result, the firing rates of the “start” and “end” SPNs (Figure 6D), and population SPN (Figures S4E and S4F) and FSI (Figure S4G) firing rates during laser-off and laser-on times were similar.

Neuronal Populations Recorded in the Motor Cortex Lacked Beginning-and-End Activity Comparable to That Observed in the Dorsolateral Striatum

We compared the normalized session-averaged spiking activity of motor cortical and striatal neurons by choosing sessions in which 5+ putative motor cortical pyramidal neurons and 5+ striatal SPNs were recorded simultaneously. The cortical population activity, recorded at mid to deep depths, was similar during each lever press during both correct and incorrect sequences (Figure 6E). Notably, the ensemble activity in the motor cortex did not peak selectively during the beginning and end of the learned sequences (Figures 6E, S3C and S3D), as did the simultaneously recorded ensemble activity in the striatum. These findings are consistent with the larger proportion of motor cortical neurons that were similarly activated during particular lever presses regardless of when in the trial the lever press occurred and regardless of which lever-press sequence was performed (Figure 2C). Overall, these results suggest that although motor cortex may be providing relevant information about ongoing movement to the striatum, movement itself and corresponding motor cortical activity are unlikely to be primarily responsible for the task-boundary activation in the striatal SPNs.

Striatal Narrow Spike-Waveform FSIs Have Activity That Is Inverse to That of SPN Task-Boundary Activation Patterns

Finally, we examined an alternative possible contributor to the striatal task-boundary activity within the striatal microcircuitry. Among the units recorded in the striatum, we found a clear bimodal distribution in spike waveform width, with a small cluster of units with high firing rates and narrow spike-width and a large cluster of units with low firing rates and wide spike-width (Figures 7A, S5A, and S5B). These differences in spike width were not driven by unit distance from the tetrode as reflected in spike amplitude (Figure S5C), and are likely a distinction in the physiological properties of the striatal SPNs and FSIs [39–41].

In a sharp contrast to the task-boundary activation in the SPNs (Figure 3C), the firing of the narrow spike-width FSIs in the same nine rats was concentrated in the mid-task period during correct trials (Figure 7B). The inverse firing patterns of the FSIs to those of the SPNs were apparent across multiple analyses. There was no significant mid-task activation of these FSIs in incorrect trials (Figure 7C, left), corresponding to the absence of the task-boundary activation in the SPNs during incorrect sequence performance. Further, in the partially correct trials in which the rats had started by performing the correct first two lever presses but pressed the incorrect last lever, the SPN “end” activity peak was missing and corresponded to an extra activity peak in the FSIs (Figure 7C, right). In sessions in which 2+ FSIs and 5+ SPNs were recorded simultaneously, the session-averaged activation patterns of the FSIs were similarly clearly antagonistic to those of the SPNs (Figures 7D and 7E).

Although the baseline firing rates of these distinct narrow-waveform units ranged widely (Figure S5B, red cluster) across the total range of firing rates of this neuronal class, these neurons shared the property of preferential mid-task activation (Figure S5E). By contrast, neurons with high firing rates but with wide spike-width similar to those of the SPN units (Figure S5B, green cluster) were most highly activated during the first and last lever presses, in patterns similar to those of the SPN population (Figure S5F). These results suggest that

narrow spike-width and not high firing rate was the relevant electrophysiological property related to the opposite task-related firing properties of these striatal units.

SPNs with task-related activity that was negatively correlated with that of a simultaneously recorded FSI were highly activated at the start and end of correct trials, but those that were positively correlated with an FSI did not exhibit task boundary activity patterns (Figure 7F). Conversely, the distribution of correlation coefficients of SPN-FSI pairs in the specific “start” and “end” SPNs was skewed toward negative correlations (Figure 7G), significantly different from those of the “middle” responsive SPNs ($p < 0.00001$, two sample t-test). Additionally, there was higher incidence of a significant interaction in FSI-SPN spike cross-correlograms (Figures S5G and S5H) for the mid-task responsive FSIs (131 of 933 of FSI-SPN pairs, 14%) than for other FSIs (418 of 4937 FSI-SPN pairs, 8.4%, Fisher’s exact test $p < 0.0001$).

Based on the opposite task-related responses of the FSI and SPN populations across multiple conditions, the specific increase in the interaction of the mid-task responsive FSIs with SPNs, and the inhibitory influence of FSIs on SPN firing documented in prior experiments [39, 41, 42], it is likely that these narrow spike-width FSIs participate in shaping the task-boundary activity of the SPNs.

DISCUSSION

Here we explicitly attempted to determine whether high-level encoding of entire action sequence entities occurs in the dorsolateral striatum, as opposed to responsivity to individual movements performed, movement contingencies, or the beginning or end of any behavior, regardless of action sequence chunking due to learning. We recorded from two nodes in the dorsolateral corticostriatal circuit during the performance of different reinforced and unreinforced action sequences and compared the neuronal representations of the behavior in these two nodes. Our findings provide key characterizations of the striatal bracketing activation, and contrast this level of neuronal representation with that of motor cortex. (1) We rarely found SPNs in the dorsolateral striatum that fired similarly in response to individual lever presses regardless of the sequence within which the presses were embedded, suggesting that SPN activity was highly modulated by parameters other than movement itself. (2) The SPN population spiking peaked at the start and end of the learned lever-press sequence across different rats, although these individual rats had learned different sequences of lever presses and had developed very different movement patterns in order to execute them. (3) It was a specific subset of “expert neurons” in each rat, making up 12% of the total SPN population, that were the critical contributors to the beginning-and-end activity. These SPNs were more strongly activated than other task-responsive SPNs spiking at other times in the trial, leading to their dominance of the population activity. (4) The population of narrow-width FSIs, presumptive striatal interneurons, fired in an inverse, mid-task pattern, with equal selectivity for the acquired sequence. (5) The start- and end-related SPNs recorded on the same day during trials in which the rats pressed the same levers in alternative incorrect sequences did not spike selectively at the beginning and end of the trials for the incorrect sequences, nor did the FSIs fire at mid-sequence. Although we are unable to control for all potential behavioral differences between correct and incorrect trials, we provide several sets

of evidence that support the hypothesis that the selective expression of the task-boundary activity during correct trials, and not alternative sequences, is due to extended history of reinforcement for the performance of those specific sequences. These findings indicate that striatal network neurons selectively represent practiced and reinforced action sequences, not, for example, a certain number of presses. We further note that the end-related spiking occurred around the time of the completion of sequence execution, and not in response to the reward notification cue (or conditioned reinforcer) or reward delivery itself, as demonstrated by imposition of temporal delays between the last press in the trial, the auditory feedback cue, and the reward delivery. Taken together, these findings argue for an abstract encoding of successfully practiced behavioral sequences in the striatum — a level of neuronal representation that has been proposed for highly honed skills and habits, but that had not yet been directly tested.

To assess possible sources of the task-boundary activity in the striatum, we first addressed the potential role of primary motor cortex. The population of putative pyramidal neurons that we recorded in the motor cortex was not preferentially active at the beginning or end of the correct sequences, but became activated prior to each of the three lever presses similarly in correct and incorrect sequences. Optogenetic inhibition of motor cortical cell bodies and their terminals in the striatum, a strategy that has been previously demonstrated to effectively inhibit excitatory inputs in other corticostriatal pathways [42, 43], also failed to affect large numbers of SPNs. The striatal bracketing activity was robustly preserved during laser-on times. Our findings are in accord with recent studies demonstrating that the primary motor cortex may not be required for the execution of motor sequences that do not require dexterous movement of the digits [44], as has been recently directly demonstrated via motor cortical lesions in rats performing stereotyped movement sequences [45]. Additionally, it has been reported that the fraction of quiescence-active motor cortical neurons increased after the first two days of training on a lever-press task, a time frame that coincided with an increase in movement stereotypy [46], and that calcium signals in motor cortical terminals in the striatum diminished upon motor training [47]. In our recordings, bracketing activity was present in the striatum early, indicating that the striatum may already have undergone plasticity resulting in the reduced influence of motor cortex due to initial behavioral shaping in which we trained rats to press single levers.

We also made a focused attempt to record from presumptive striatal interneurons that are known to be capable of inhibiting striatal SPNs in context-specific modes [42, 48, 49], and have recently been shown to be important for the expression of habitual behavior [50]. We found that a distinct population of striatal narrow-spike-width units resembling FSIs [40, 41] fired strongly at mid-sequence of the acquired behavioral sequences, during the time-window when SPN spiking was reduced relative to their firing at the start and end of the sequences. This selective mid-task activation was observed in a distinct group of narrow spike-width, high-firing units, and not in high-firing units with wide spike-width, two populations not previously segregated in studies demonstrating task-bracketing [22, 51, 52]. Given that the bracketing pattern we observed is so dominant in a sub-population of SPNs, and that we could observe inverse SPN-FSI firing in simultaneously recorded pairs of bracketing SPNs and mid-sequence active FSIs, we raise the hypothesis that the remarkable inverse correlation of the FSI and SPN firing patterns across multiple conditions, and the

known inhibitory influence of FSIs onto SPNs [39, 41, 42, 50], could mean that these narrow spike-width FSIs interact with the SPNs to shape their task-boundary activation pattern.

The critical finding of this study is the fundamental strengthening of the hypothesis that the dorsolateral striatum may fulfill its function of promoting habitual behavior in part by representing repeatedly rewarded, thus potentially useful, action sequences as single units [10, 53]. This task-chunking activation in the striatum may be related to the acquisition of skills or habits, which have both overlapping and distinct features, but which have both been linked to basal ganglia function [10]. We propose that the presence of this robust task-boundary activation, together with the low numbers of SPNs whose spiking pattern could be accounted for by single motor events, favors the idea that the functions of the dorsolateral striatum in behavior should be considered to include context-specific representations of motor programs as higher-order units. We introduce the possibility that the striatum could operationalize such representations through the action of intrastriatal networks involving interneurons and projection neurons. The existence of such a general-purpose signal for delimiting chunked action sequences raises numerous questions about how such activity could facilitate the acquisition and expression of behavioral routines favored by virtue of their value. In one possible model (Figure S6), the bracketing activity could serve to initiate learned action programs in cortical or subcortical motor regions that have been found to have long-term memory of learned action sequences [54]. This possibility would be in accord with a role of the dorsolateral striatum in the direct control of learned action programs proposed in action-selection models of the striatum [7, 24, 55, 56]. An interesting possibility raised by our findings is that interneurons could serve to inhibit further activation of the striatal SPNs after a behavioral program is initiated until the ongoing behavioral program has finished being executed. The dorsolateral striatum might also, or instead, provide a teaching signal that facilitates the chunking of action repertoires.

Given that packaged behavioral routines make up much of normal behavior in humans and other animals, and that disorders of sequential behavior and repetitiveness are features of a large range of neurological and neuropsychiatric disorders, the identification of this signal as a general biomarker of chunked behavioral programs is fundamental for characterizing the circuit mechanisms underlying habitual motor programs.

STAR METHODS

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for reagents may be directed to and will be fulfilled by the corresponding author Dr. Ann Graybiel (graybiel@mit.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All procedures conformed to NIH Guidelines for the Care and Use of Laboratory Animals and were approved by the Committee on Animal Care at the Massachusetts Institute of Technology. Twelve adult male Long Evans rats (8 wildtype and 4 ChAT-Cre rats) were studied in recording experiments, and 4 additional wildtype Long Evans rats were tested in behavior experiments without recording. Three of the 12 rats from which recordings were

made were not included in the data set — one due to the loss of the implanted drive, one due to poor striatal recordings, and one due to very few correct trials performed. They were single-housed under a 12-hr reverse light/dark cycle and were trained and tested during the dark cycle. After the start of behavioral training, rats were kept on mild food restriction with 15 g of food/day, and were allowed to be maintained at 85% of their free-feeding weight.

METHOD DETAILS

Recording Drive—Custom-built recording drives with 28 independently moveable microdrives were assembled to hold 24 tetrodes, each house-made with four twisted tungsten wires, and two 200- μ m optical fibers with a zirconia ferrule (Doric Lenses). Tubes running from the microdrives to drive tip were arranged so that each one of them contained an optical fiber at the center of the bundle, surrounded by 10–14 tetrode tubes all positioned ~150 μ m away from the edge of the optical fiber.

Surgical Procedures—Rats were initially anesthetized with 3% isoflurane gas and 0.2 ml of ketamine (I.P.), and were thereafter kept under 1–2% isoflurane gas anesthesia. Rats were secured in a stereotaxic frame, and small drill holes (1.5 mm wide) were made over left caudoputamen (AP 0.5 mm, ML 3.5 mm, tetrodes lowered to DV –4 mm) and le[^] primary motor cortex (AP 0 mm, ML 1.5 mm, tetrodes lowered to DV –1 mm). Some of the rats had additional drill holes and tetrode implants in left prelimbic cortex (AP 3 mm, ML 0.4 mm, tetrodes lowered to DV –3 mm) and the le[^] dorsomedial striatum (AP 1.7 mm, ML 1.7 mm, tetrodes lowered to DV –4 mm). Several small burr holes were also made in other sites on the skull for anchor jeweler’s screws before the recording drive was attached to the skull using dental cement. On the day of the implantation, all tetrodes were lowered to within ~200 μ m of the target depths. Optical fibers were lowered 1 mm on the day of the implantation and then were further lowered gradually on the days following surgery to reach their target depth. Tetrodes and optical fibers were left in place at the target depth except for rare small adjustments of < 125 μ m made to maintain a high yield of single unit recordings.

Four weeks prior to the recording drive implantation, 0.5 μ l of halorhodopsin-containing AAV5-CaMKII-eNpHR3.0-EYFP virus (University of North Carolina Vector Core) was injected in 6 rats bilaterally into primary motor cortex (AP 0 mm, ML 1.5 mm, DV –1 mm) at a rate of 0.05 μ l/min.

Behavioral Training—The rats were trained in a square custom-built operant chamber (floor size of 40 cm X 40 cm), with walls that were slanted near the floor to provide extra head space to prevent bumping of the recording drive on the chamber walls. The outer dimensions of the chamber walls were 70 cm X 70 cm. The chamber was painted with black anti-static paint, was enclosed inside of a cabinet, and was lit only with infrared light during behavioral training. Rats were habituated to being handled, to drinking chocolate milk, and to being in the operant chamber prior to the start of lever-press training. They were then taught to associate an auditory click tone with chocolate milk delivery; and then learned to press the levers in two training days in which they were rewarded for pressing either of the two levers. Rats were then trained for two days on which they were rewarded after any three lever presses. Each rat was then assigned a single correct three-press sequence. Upon the

initiation of the training session, both levers were made available to the rat and stayed extended throughout the duration of the 90-min training session. After every three presses, 200 ms after the third lever press, an auditory feedback cue was played. If the sequence performed was the correct sequence, a click tone was played, and with another 200-ms delay a pump containing a syringe of chocolate milk was activated delivering 0.2 ml of chocolate milk to the reward well positioned between the two levers. In the rare cases that the rat performed a fourth lever press in the intervals prior to the reward delivery, the reward was omitted. If any other press sequence was performed, a white noise was played 200 ms after the third lever press. In a random 20% of incorrect trials, rats received a small amount (0.05 ml) of chocolate milk reward, and the frequency of this small reward was gradually decreased to a random 5% of trials as their performance improved. These random small reinforcements were critical for continued performance of the rats due to the low probability of performing the correct sequence by chance (1/8). No timing requirements were imposed on the rats for lever pressing; however, if the rat began a trial by performing one or two lever presses, but more than 20 s passed prior to the completion of three lever presses, the trial was automatically terminated. Post-hoc, any trials with a duration of longer than 8 s between the first lever press and the third lever press were excluded from data analysis. Rats' behavior was monitored remotely using video from the operant chamber, but no experimenters were present in the experimental room other than at the start and end of the training sessions, and occasionally to correct unplugged preamplifiers or other issues. Rats were left in the operant chamber for an entire 90 min, although in some cases they did not continue to press the levers during the entire period, or often took rest breaks before returning to lever pressing.

Devaluation Procedure—Four rats were given a reward-satiation devaluation test. Each of the rats received 3–4 probe sessions across training in which unlimited chocolate milk was made accessible to them for 2 hr prior to the probe session. The devaluation probe sessions were 30 min in duration, during which time rats were placed in the normal operant chamber with levers extended and normal auditory feedback cues were provided in response to lever pressing. However, although there was chocolate milk in the reward tube, the chocolate milk syringe was not placed in the pump, and therefore no chocolate milk was delivered to the rat in response to correct presses.

Optogenetic Inhibition of Cortical Terminals during Task Performance—In six rats, we used the 200- μ m optical fiber placed at the center of the array of 14 tetrodes in the dorsolateral striatum to inhibit halorhodopsin-expressing cortical terminals. A laser patch cord, running from the 589-nm yellow laser located in an adjacent room, was routed to the recording room and through the commutator on the ceiling of the operant chamber targeting 3 mW of laser light to the dorsolateral striatum. The laser shutter was controlled by transistor-transistor logic (TTL) pulses initiated by custom MATLAB behavior software. TTL pulses were copied to two channels in the Neuralynx Data Acquisition System and saved as timestamps for laser on/off times. The laser shutter was opened in one of four periods: (1) from the end of the previous trial to the first press of the next trial, (2) from the first press to the second press in the trial, (3) from the second press to the third press in the trial, or (4) for 3 s from the third press. These laser-trial types were randomly intermixed

within the training session, and no laser light was used in 20% of the trials. In all cases, when the laser light was illuminated for 3 s, it was automatically turned off to prevent heating. For the analysis of the effect of the laser manipulation on neural activity in the striatum, all laser-off periods were combined and shown together, and all laser-on periods were combined and shown together. In post-hoc analysis, any periods during which the laser had been on for longer than 2 s were excluded from further analysis.

Optogenetic Inhibition of Cortical Cell Bodies and Cortical Terminals Out-of-Task—After some behavioral training and recording sessions, we tested the responsiveness of the recorded units to cortical cell body and terminal inhibition. We used 3 mW of yellow light targeted to the 200- μ m optical fiber located in dorsolateral striatum and then on the optical fiber located in motor cortex. In each test session, 500 trials of 70-ms pulses were used. During this time, rats were free to move around the operant chamber, but the levers were retracted and not available to them.

Data Acquisition—Lever-press events were recorded with Med Associates hardware and Med Associates MATLAB toolboxes along with custom MATLAB behavioral control software that received timestamps from the Neuralynx Data Acquisition System. Two video cameras connected to the Neuralynx system were mounted above the operant chamber. One was located at the top of the chamber and was used to collect head-position data with the use of red and green LEDs mounted onto a preamplifier positioned directly above the recording drive. A second infrared camera was used to record video images in the dark environment and was attached to the wall closest to the levers and the reward well. Unit activity (gain: 200–10000, filter: 600–6000 Hz) was recorded with the Neuralynx Data Acquisition System. Spikes exceeding a preset voltage threshold on any of the four tetrode-channels triggered the waveform to be sampled at 32 kHz and to be stored on all four channels.

Histology—At the end of training, rats were deeply anesthetized (sodium pentobarbital, 50–100 mg/kg, I.P.), and lesions were made to mark the final recording sites (25 μ A, 10 s). Rats were then perfused with 4% paraformaldehyde in 0.1 M sodium-potassium phosphate buffer, and 30- μ m thick transverse frozen sections were stained for CD11 to identify lesion sites and by immunohistochemistry for green fluorescent protein to detect viral expression in cell bodies and terminals.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analysis of Head Movement Trajectories—The x and y positions in each trial from the first lever-press event to the last event of the trial (reward delivery or white noise error tone) were downsampled (time normalized) into fixed-length vectors with 70 points each, forming the matrix T_i for trial i .

$$T_i = [x_0 \dots x_{70} \ y_0 \dots y_{70}]$$

The head tracking data from two of the nine rats were not used in this analysis due to unacceptable noise in head tracking data resulting from the frequent blocking of the head-

mounted LEDs from the view of the over-head camera by the tethers of the recording system and from frequent glitches in the tracked location. The similarity between these trial-to-trial and mean session-to-session trajectories was then assessed for each pair of trajectories using the following equation for the correlation coefficient, with the MATLAB function `corrcoef`.

$$R(i, j) = \text{covariance}(T_i, T_j) \div (\sigma T_i \times \sigma T_j)$$

Spike Sorting and Quality Assessment—Spike data were manually sorted into single units using Plexon Offline Sorter. Additionally, data from 1/4 of the sessions were sorted using an automated clustering procedure [57]. All sorted clusters were graded according to a custom algorithm. Each cluster received a grade on waveform quality based on similarity in the waveforms and presence of a valley, and a cluster quality grade based on L-ratio, the distance from other sorted units and from the noise cluster, percent of the short interspike intervals, percent of the cluster below threshold, and the continuity of the spikes throughout the recording session. Based on these measures, each cluster was assigned a grade of 1–5, and only clusters with grades ≥ 3 were accepted for analysis (72% of the manually sorted units yielding 3545 striatal units and 926 motor cortical units, and 32% of the automatically sorted units yielding 925 striatal units and 194 motor cortical units). In both striatum and motor cortex, the electrophysiological properties and task-responses of the putative cell types that we later classified were very similar in the manually and automatically spike sorted units; thus, the manually and automatically sorted units were combined in further analyses.

Classifying Putative Cell Types—We used an in-house MATLAB toolbox to depict the distribution of various electrophysiological properties of the single units in each recorded brain region. We found that waveform width was one primary factor along which distinct clusters appeared in the set of recorded units. In the striatum, we classified 3595 units in the larger cluster with half-peak waveform widths of 70–125 μs and baseline firing rates of < 3.5 Hz as putative SPNs. We classified 414 units in the smaller cluster with waveform widths of < 70 μs , baseline firing rates of > 3.5 Hz, and less than 10% of ISIs > 1 s as putative FSIs (Figure 3A). As a precaution against the potential of distant spikes to appear narrower, low-amplitude narrow-spike units were not included in the analysis. Based on comparison with data previously collected in our laboratory, we found that very few ($n = 23$) of the units that we recorded satisfied the criterion of putative tonically active neurons (putative cholinergic interneurons), and these neurons were not included in the analysis.

In the motor cortex, we also found a division of units across spike width, and we assigned 1060 units with spike widths at half peak of 70–125 μs as putative pyramidal neurons. The small number ($n = 36$) of narrow waveform putative interneurons that we identified were not included in the analysis.

Analysis of Neural Spike Activity—Peri-event spike histograms were created with custom MATLAB code using a sliding window with a step size of 50 ms and a bin size of 250 ms for each event (lever presses, click or white noise auditory feedback, and reward

delivery) in each trial type. In individual examples, the peri-event histograms were pasted together using window sizes corresponding to the median time between each set of events. In the population activity plots, peri-event histograms were pasted together using 0.5-s windows centered at each lever-press event due to the variability across rats and sessions in the time between successive lever presses. In a complimentary analysis (Figure S1), we plotted the task-related population activity by stretching or compressing the times between lever presses to a standard 1-s period to avoid excising portions of the trials excluded in the pasted window histograms. This was done by first generating session-by-session peri-event spike histograms for each single unit with window sizes based on the median time between each set of consecutive events in the session. To construct population plots, these session-by-session windows were stretched or compressed to a standard 1-s window before calculating averages. In color plots displaying the SPN and FSI population activity of individual rats in rows, the average neural activity in each rat was normalized by subtracting the mean firing rate of the population activity in the given trial type.

To compare SPN and FSI task-related activity, we identified sessions in which at least 5 SPNs and 2 FSIs were recorded simultaneously. Normalized session averages were used in the comparisons of their task-related activity in order to reduce the possible impact of differences in numbers of units and of differences in firing rate between cell types. To do this, we computed the mean population firing rate of the two cell types in the trial in each session and normalized these session averages by subtracting the mean population firing rate and dividing by the maximum population firing rate. The same method was used for the comparison of motor cortical and striatal task-related activity in sessions in which at least 5 putative motor cortical pyramidal neurons and 5 striatal SPNs were recorded simultaneously.

Comparison of Task Representations in Motor Cortex and Striatum—To assess whether the task-related firing of single units could be accounted for with the occurrence of single lever-press events, we compared the task responses of each neuron to lever 1 press and lever 2 press events in 250-ms bins centered on three time-points around the lever press (−200 ms, 0 ms, and +100 ms). We also assessed the task responses of each neuron during transitions from lever 1 to lever 2 or vice versa in 250 ms bins centered on two time-points in the transition (250 ms after the press of the first lever and 250 ms prior to the press of the next lever). To assess whether the units responded similarly in every instance of these events, we identified all the different contexts in which the lever press occurred — such as the first, second, or third press in a correct sequence, or the first, second, or third press in each of the incorrect sequences. If the unit fired 2SDs above baseline in each different context within which the given lever press or transition occurred and the firing rates were similar in each case (the maximum response – minimum response must be smaller than half of the median response), we included this unit in the list of units whose task-related firing could be well accounted for by single lever-press event occurrences. Baseline activity was defined as the firing rate in the 2-s interval 6–8 s after reward delivery (after reward consumption had occurred) and ending at least 1.5 s before the next lever press. To compare the incidence of such lever-press-related units in motor cortex and striatum, we identified sessions in which we recorded at least 10 units simultaneously in both motor cortex and striatum and compared the proportion of such motor units for both regions.

Generalized Linear Models—The time series of the raw firing rate in consecutive 500-ms bins were calculated for each SPN recorded in the dorsolateral striatum and pyramidal neuron recorded in the motor cortex with a mean firing rate > 0.3 Hz, and they were fitted with the following three generalized linear models. The firing rate was best approximated with a Poisson distribution, so a logarithmic link function was used in the generalized linear model analysis.

Model 1:

Firing rate = intercept + c1×time from last lever 1 press + c2×time to next lever 1 press + c3×time from last lever 2 press + c4×time to next lever 2 press

Model 2:

Firing rate = variables from Model 1 + c5 × temporal position in trial × if trial is correct + c6 × time from last lever 1 press × temporal position in trial × if trial is correct + c7 × time to next lever 1 press × temporal position in trial × if trial is correct + c8 × time from last lever 2 press × temporal position in trial × if trial is correct + c9 × time to next lever 2 press × temporal position in trial × if trial is correct

Model 3:

Firing rate = variables from Model 1 + c5×beginning + c6×end

where the categorical “beginning” variable is 1 if the rat is in the midst of performing the first press of a correct sequence, and 0 otherwise. Similarly, the “end” variable is 1 if the rat is in the midst of performing the last press of the correct sequence, and 0 otherwise.

The differences between the χ^2 statistic in Model 1 and Model 2 were considered as a measure of the dependence of the firing rate of the neurons on contextual variables such as temporal position in the trial and whether the trial was correct. The differences between the χ^2 statistic in Model 1 and Model 3 were considered as a measure of the degree to which neurons were modulated during the more specific epochs of the beginning and end of the correct sequence.

In the generalized linear model analysis performed to assess firing rate modulations that occurred as a result of the correct/incorrect identity of the sequence as compared to other potentially co-varying behavioral variables, we used the following two models.

Model 1:

Firing rate = intercept + c1×time from last lever 1 press + c2×time to next lever 1 press + c3×time from last lever 2 press + c4×time to next lever 2 press + c5×temporal position in trial + c6×number of trials of current sequence in the session + c7×trial duration + c8×if previous trial rewarded + c9×if previous trial rewarded×temporal position in trial.

Model 2:

Firing rate = variables from Model 1 + c10×if trial is correct + c11×beginning + c12×end

Statistical Tests—To test the statistical significance of the improvement in the rats' behavioral performance across training, we used the Mann-Whitney Wilcoxon rank-sum non-parametric test to compare the percentage of correct trials between the first 10 days of training and on the days 31–40 of training. For the comparison of population activity in the striatum between different lever presses within a trial type and between lever presses across trial types, we used the Mann-Whitney Wilcoxon rank-sum nonparametric test on the raw firing rates of each unit in the 250-ms time bins around each of the lever presses. For the comparison of task-related responses of striatal SPNs and motor cortical neurons, and between SPNs and FSIs, we used the Mann-Whitney Wilcoxon rank-sum non-parametric test to compare the session-normalized activity in the 250-ms time bins around each of the three lever presses in the trial. For the comparison of the proportions of simple lever-press-related (“motor-type”) units across all of the putative projection neurons recorded in motor cortex and striatum, we used Fisher's exact test. For the comparison of the proportions of the “motor-type” units in the 15 sessions with 10+ units recorded simultaneously in motor cortex and striatum, we used the Mann-Whitney Wilcoxon rank-sum non-parametric test for each of the lever-press-related time-points and for the proportions of all “motor-type” units. To determine the effect of optogenetic inhibition on striatal activity during the laser pulse tests, we used the Mann-Whitney Wilcoxon rank-sum non-parametric test to compare the per-trial firing rate in 40-ms bins prior to laser onset and 40-ms bins after laser onset. To determine the laser effect during in-task laser experiments, we used the Mann-Whitney Wilcoxon rank-sum non-parametric test to compare the laser-on and laser-off per-trial firing rates for each single unit in each 250-ms time bin in the peri-event trial histograms. For this analysis, we calculated the statistical power required to observe a significant effect of the laser manipulation, and thus included in the analysis only time bins in which 20+ laser-off trials existed, 10+ laser-on trials existed, and the sum of the spikes during the laser-off trials was greater than 20.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Leif Gibb and Henry Hall for designing the operant chamber behavioral room, Jessica Pourian and Devin Ahern for conducting some of the recording experiments and spike sorting, Alexander Friedman for developing the automated spike clustering toolbox used for analyzing part of the single unit data, Michael Riad and Jannifer Lee for performing the histology, and Yasuo Kubota for his extensive help in manuscript preparation. This work was funded by NIH/NIMH (R01 MH060379 to A.M.G. and F31 MH099782 to N.M.), Office of Naval Research (N00014-04-1-0208 to A.M.G.), Nancy Lurie Marks Family Foundation (to A.M.G.), support from R. Rae Pourian and Julia Madadi (to A.M.G.), and the McGovern Institute Mark Gorenberg Fellowship (to N.M.).

References

1. Lashley KS. The problem of serial order in behavior. *Cerebral mechanisms in behavior*. 1951:112–131.
2. Rosenbaum DA, Cohen RG, Jax SA, Weiss DJ, van der Wel R. The problem of serial order in behavior: Lashley's legacy. *Hum Mov Sci*. 2007; 26:525–554. [PubMed: 17698232]
3. Berridge KC, Aldridge JW, Houchard KR, Zhuang X. Sequential super-stereotypy of an instinctive fixed action pattern in hyper-dopaminergic mutant mice: a model of obsessive compulsive disorder and Tourette's. *BMC Biol*. 2005; 3:4. [PubMed: 15710042]
4. Hyman SE, Malenka RC, Nestler EJ. Neural mechanisms of addiction: the role of reward-related learning and memory. *Annual review of neuroscience*. 2006; 29:565–598.
5. Everitt BJ, Robbins TW. Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature neuroscience*. 2005; 8:1481–1489. [PubMed: 16251991]
6. Voon V, Derbyshire K, Ruck C, Irvine MA, Worbe Y, Enander J, Schreiber LR, Gillan C, Fineberg NA, Sahakian BJ, et al. Disorders of compulsivity: a common bias towards learning habits. *Molecular psychiatry*. 2015; 20:345–352. [PubMed: 24840709]
7. Friend DM, Kravitz AV. Working together: basal ganglia pathways in action selection. *Trends in neurosciences*. 2014; 37:301–303. [PubMed: 24816402]
8. Graybiel AM. Habits, rituals, and the evaluative brain. *Annual review of neuroscience*. 2008; 31:359–387.
9. Stephenson-Jones M, Samuelsson E, Ericsson J, Robertson B, Grillner S. Evolutionary conservation of the basal ganglia as a common vertebrate mechanism for action selection. *Current biology : CB*. 2011; 21:1081–1091. [PubMed: 21700460]
10. Graybiel AM, Grafton ST. The striatum: where skills and habits meet. *Cold Spring Harb Perspect Biol*. 2015; 7:a021691. [PubMed: 26238359]
11. Smith KS, Graybiel AM. Investigating habits: strategies, technologies and models. *Front Behav Neurosci*. 2014; 8:39. [PubMed: 24574988]
12. Dolan RJ, Dayan P. Goals and habits in the brain. *Neuron*. 2013; 80:312–325. [PubMed: 24139036]
13. Daw ND, Niv Y, Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*. 2005; 8:1704–1711. [PubMed: 16286932]
14. Yin HH, Ostlund SB, Knowlton BJ, Balleine BW. The role of the dorsomedial striatum in instrumental conditioning. *The European journal of neuroscience*. 2005; 22:513–523. [PubMed: 16045504]
15. Yin HH, Knowlton BJ, Balleine BW. Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. *Behavioural brain research*. 2006; 166:189–196. [PubMed: 16153716]
16. Yin HH, Knowlton BJ, Balleine BW. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *The European journal of neuroscience*. 2004; 19:181–189. [PubMed: 14750976]
17. Yin HH, Knowlton BJ, Balleine BW. Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. *The European journal of neuroscience*. 2005; 22:505–512. [PubMed: 16045503]

18. Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM. Building neural representations of habits. *Science*. 1999; 286:1745–1749. [PubMed: 10576743]
19. Aldridge JW, Berridge KC, Rosen AR. Basal ganglia neural mechanisms of natural movement sequences. *Canadian journal of physiology and pharmacology*. 2004; 82:732–739. [PubMed: 15523530]
20. Aldridge JW, Berridge KC. Coding of serial order by neostriatal neurons: a “natural action” approach to movement sequence. *J Neurosci*. 1998; 18:2777–2787. [PubMed: 9502834]
21. Thorn CA, Atallah H, Howe M, Graybiel AM. Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. *Neuron*. 2010; 66:781–795. [PubMed: 20547134]
22. Kubota Y, Liu J, Hu D, DeCoteau WE, Eden UT, Smith AC, Graybiel AM. Stable encoding of task structure coexists with flexible coding of task events in sensorimotor striatum. *Journal of neurophysiology*. 2009; 102:2142–2160. [PubMed: 19625536]
23. Barnes TD, Kubota Y, Hu D, Jin DZ, Graybiel AM. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature*. 2005; 437:1158–1161. [PubMed: 16237445]
24. Smith KS, Graybiel AM. A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*. 2013; 79:361–374. [PubMed: 23810540]
25. Jin X, Costa RM. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*. 2010; 466:457–462. [PubMed: 20651684]
26. Jin X, Tecuapetla F, Costa RM. Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nature neuroscience*. 2014; 17:423–430. [PubMed: 24464039]
27. Carelli RM, West MO. Representation of the body by single neurons in the dorsolateral striatum of the awake, unrestrained rat. *J Comp Neurol*. 1991; 309:231–249. [PubMed: 1885787]
28. DeLong MR. Putamen: activity of single units during slow and rapid arm movements. *Science*. 1973; 179:1240–1242. [PubMed: 4631890]
29. Kim N, Barter JW, Sukharnikova T, Yin HH. Striatal firing rate reflects head movement velocity. *Eur J Neurosci*. 2014; 40:3481–3490. [PubMed: 25209171]
30. Rueda-Orozco PE, Robbe D. The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nature neuroscience*. 2015; 18:453–460. [PubMed: 25622144]
31. Desmurget M, Turner RS. Motor sequences and the basal ganglia: kinematics, not habits. *J Neurosci*. 2010; 30:7685–7690. [PubMed: 20519543]
32. Desmurget M, Grafton ST, Vindras P, Grea H, Turner RS. The basal ganglia network mediates the planning of movement amplitude. *Eur J Neurosci*. 2004; 19:2871–2880. [PubMed: 15147320]
33. Carelli RM, Wolske M, West MO. Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 1997; 17:1804–1814. [PubMed: 9030639]
34. Tang CC, Root DH, Duke DC, Zhu Y, Teixeira K, Ma S, Barker DJ, West MO. Decreased firing of striatal neurons related to licking during acquisition and over-training of a licking task. *J Neurosci*. 2009; 29:13952–13961. [PubMed: 19890005]
35. Tang C, Pawlak AP, Prokopenko V, West MO. Changes in activity of the striatum during formation of a motor habit. *The European journal of neuroscience*. 2007; 25:1212–1227. [PubMed: 17331217]
36. Balleine B, Dickinson A. Signalling and incentive processes in instrumental reinforcer devaluation. *Q J Exp Psychol B*. 1992; 45:285–301. [PubMed: 1475401]
37. McGeorge AJ, Faull RL. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience*. 1989; 29:503–537. [PubMed: 2472578]
38. Ebrahimi A, Pochet R, Roger M. Topographical organization of the projections from physiologically identified areas of the motor cortex to the striatum in the rat. *Neuroscience research*. 1992; 14:39–60. [PubMed: 1380687]
39. Mallet N, Le Moine C, Charpier S, Gonon F. Feedforward inhibition of projection neurons by fast-spiking GABA interneurons in the rat striatum in vivo. *J Neurosci*. 2005; 25:3857–3869. [PubMed: 15829638]

40. Berke JD, Okatan M, Skurski J, Eichenbaum HB. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*. 2004; 43:883–896. [PubMed: 15363398]
41. Gittis AH, Leventhal DK, Fensterheim BA, Pettibone JR, Berke JD, Kreitzer AC. Selective inhibition of striatal fast-spiking interneurons causes dyskinesias. *J Neurosci*. 2011; 31:15727–15731. [PubMed: 22049415]
42. Friedman A, Homma D, Gibb LG, Amemori K, Rubin SJ, Hood AS, Riad MH, Graybiel AM. A Corticostriatal Path Targeting Striosomes Controls Decision-Making under Conflict. *Cell*. 2015; 161:1320–1333. [PubMed: 26027737]
43. Tye KM, Prakash R, Kim SY, Fenno LE, Grosenick L, Zarabi H, Thompson KR, Gradinaru V, Ramakrishnan C, Deisseroth K. Amygdala circuitry mediating reversible and bidirectional control of anxiety. *Nature*. 2011; 471:358–362. [PubMed: 21389985]
44. Peters AJ, Liu H, Komiyama T. Learning in the Rodent Motor Cortex. *Annual review of neuroscience*. 2017; 40:77–97.
45. Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Olveczky BP. Motor cortex is required for learning but not for executing a motor skill. *Neuron*. 2015; 86:800–812. [PubMed: 25892304]
46. Peters AJ, Lee J, Hedrick NG, O’Neil K, Komiyama T. Reorganization of corticospinal output during motor learning. *Nature neuroscience*. 2017; 20:1133–1141. [PubMed: 28671694]
47. Kupferschmidt DA, Juczewski K, Cui G, Johnson KA, Lovinger DM. Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron*. 2017; 96:476–489. e475. [PubMed: 29024667]
48. Berke JD. Functional properties of striatal fast-spiking interneurons. *Frontiers in systems neuroscience*. 2011; 5:45. [PubMed: 21743805]
49. Gage GJ, Stoetzer CR, Wiltschko AB, Berke JD. Selective activation of striatal fast-spiking interneurons during choice execution. *Neuron*. 2010; 67:466–479. [PubMed: 20696383]
50. O’Hare JK, Li H, Kim N, Gaidis E, Ade K, Beck J, Yin H, Calakos N. Striatal fast-spiking interneurons selectively modulate circuit output and are required for habitual behavior. *eLife*. 2017; 6
51. Hernandez LF, Kubota Y, Hu D, Howe MW, Lemaire N, Graybiel AM. Selective Effects of Dopamine Depletion and L-DOPA Therapy on Learning-Related Firing Dynamics of Striatal Neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2013; 33:4782–4795. [PubMed: 23486949]
52. Thorn CA, Graybiel AM. Differential entrainment and learning-related dynamics of spike and local field potential activity in the sensorimotor and associative striatum. *J Neurosci*. 2014; 34:2845–2859. [PubMed: 24553926]
53. Graybiel AM. The basal ganglia and chunking of action repertoires. *Neurobiol Learn Mem*. 1998; 70:119–136. [PubMed: 9753592]
54. Matsuzaka Y, Picard N, Strick PL. Skill representation in the primary motor cortex after long-term practice. *Journal of neurophysiology*. 2007; 97:1819–1832. [PubMed: 17182912]
55. Da Cunha C, Gomez AA, Blaha CD. The role of the basal ganglia in motivated behavior. *Rev Neurosci*. 2012; 23:747–767. [PubMed: 23079510]
56. Mink JW, Thach WT. Basal ganglia intrinsic circuits and their role in behavior. *Current opinion in neurobiology*. 1993; 3:950–957. [PubMed: 8124079]
57. Friedman A, Keselman MD, Gibb LG, Graybiel AM. A multistage mathematical approach to automated clustering of high-dimensional noisy data. *Proc Natl Acad Sci U S A*. 2015; 112:4477–4482. [PubMed: 25831512]

HIGHLIGHTS

- Striatal projection neurons selectively mark bounds of acquired movement sequences
- Such sequence-boundary activity does not occur for random, unreinforced sequences
- Fast-spiking interneurons fire inversely to this sequence-bracketing pattern
- These activity patterns were rare in simultaneously recorded motor cortical units

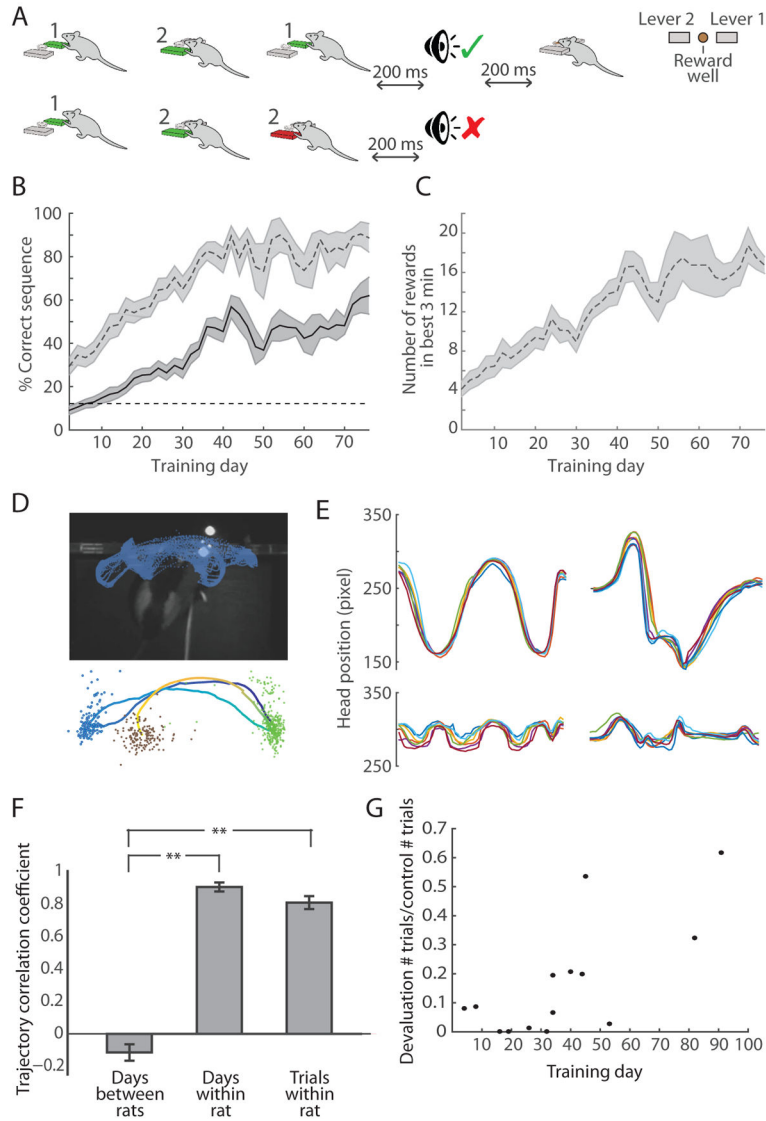


Figure 1. Rats Learn to Perform Specific Three-Step Lever-Press Sequences

(A) Task procedure with examples of a correct sequence (1–2–1, top) and an incorrect sequence (1–2–2, bottom), and top view of operant chamber configuration (right).

(B) Percent correct in total session (solid) and in the best 3-min performance period (dashed) across training days (n = 13 rats). Only 4 unimplanted rats were trained past day 40. Shading indicates SEM.

(C) Number of correct trials performed during best 3-min performance period across training days (n = 13 rats).

(D) Head tracking data from 30 consecutive correct trials in one training session (top). Average session trajectory (dark blue to yellow) of rat performing the 1–2–1 sequence (bottom) with head locations for lever 1 presses (green), lever 2 presses (blue), and chocolate milk delivery (brown).

(E) Side-to-side (top) and forward-and-back (bottom) session average head trajectories from 7 consecutive training days for a rat performing the 1-2-1 sequence (left) and another rat performing the 1-2-2 sequence (right).

(F) Correlation coefficients of session average head trajectories between pairs of rats that learned different sequences (left, $n = 42$ pairs), session average trajectories within single rats (middle, $n = 7$ rats), and trial average trajectories within rats (right, $n = 7$ rats). Only trajectories from days 10+ of training with successful video tracking were used. Error bars indicate SEM. $**p < 0.001$ (Wilcoxon rank-sum test).

(G) Ratio of trials performed in 30-min devaluation probe session and number of trials performed in the first 30 min of prior day's training session as a function of number of days trained. Each data point corresponds to a single session.

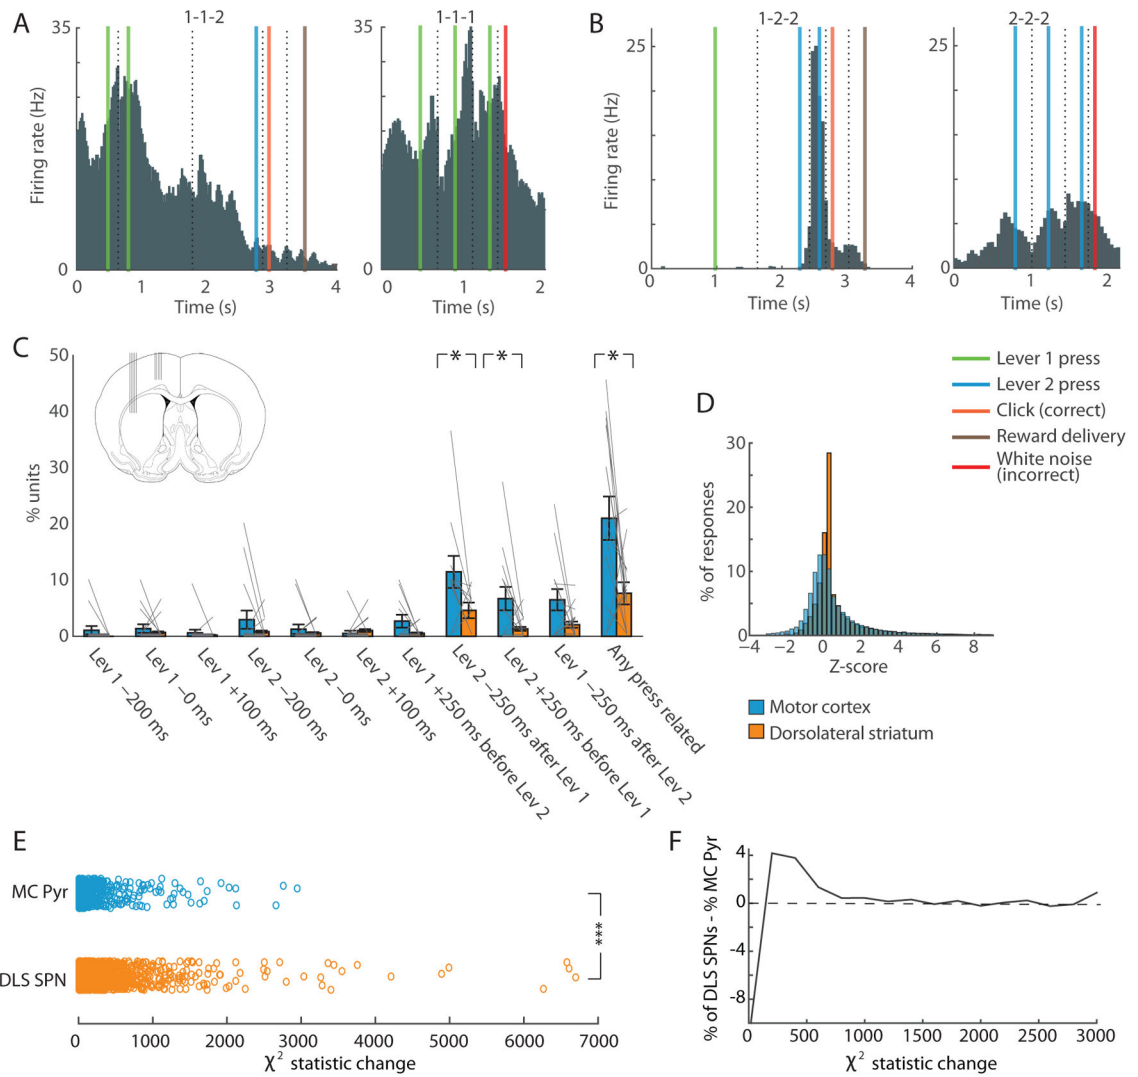


Figure 2. Simple Lever-Press-Responsive Neurons Are More Prevalent in Forelimb Motor Cortex than in Dorsolateral Striatum

(A) Activity of a lever-press-related unit recorded in motor cortex during 1-1-2 correct (left) and 1-1-1 incorrect (right) sequences. Peri-event histograms for individual trial events were pasted together with window sizes determined by the session median time between each pair of successive events. Colored lines indicate events, and dashed lines indicate the time-window borders. This unit fired during lever 1 presses in both sequences.

(B) Activity of a ‘non-motor-type’ unit in dorsolateral striatum during 1-2-2 correct (left) and 2-2-2 incorrect (right) sequences, showing selective spiking during the last press of the correct sequence, but not during repeated presses of the same lever.

(C) Proportions of motor-type units in ten press-related time-points in 15 sessions with 10+ units recorded simultaneously in motor cortex and striatum (inset). Gray lines represent individual sessions. Error bars indicate SEM. * $p < 0.05$.

(D) Distribution of overall z-scores for motor cortical pyramidal neurons (blue) and dorsolateral striatal SPNs (orange) in response to the events indicated in C.

(E) Changes in the χ^2 statistic for motor cortical pyramidal neurons (MC Pyr; n = 670) and SPNs in dorsolateral striatum (DLS SPNs; n = 1708) when the beginning and end variables were added to the generalized linear model. ***p < 0.000001.

(F) The difference in the proportions of SPNs and pyramidal neurons with given χ^2 statistic changes (x-axis) after adding the beginning and end variables to the model.

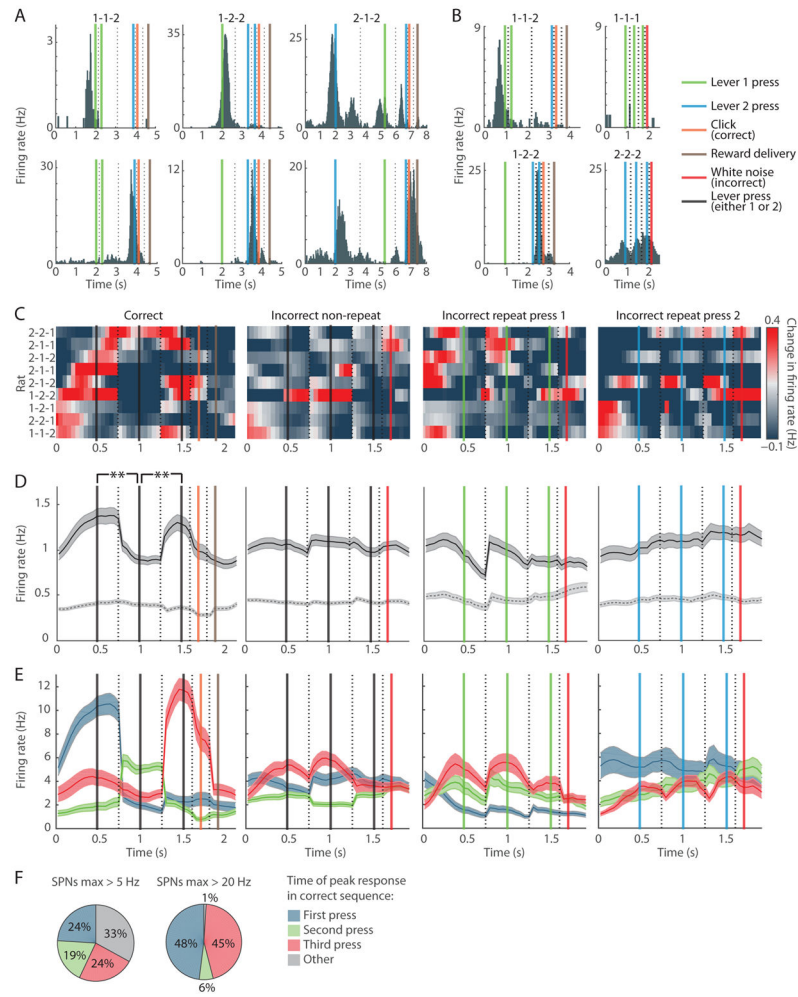


Figure 3. Striatal SPN Population Spiking Is Concentrated around the Initiation and Termination of the Learned Lever-Press Sequence but Not during Incorrect Lever-Press Sequences, Regardless of the Lever-Press Sequence Learned

(A) Activity of start (top) and end (bottom) SPNs in rats that learned 1-1-2 (left), 1-2-2 (middle) or 2-1-2 (right) lever-press sequences, plotted as in Figure 2A.

(B) Activity of start SPN (top) that spiked during the first lever 1 press in the correct sequence, but not during the middle lever 1 press (left). This SPN did not respond during incorrect trials with repetitive lever 1 pressing (right). Activity of end SPN (bottom) that fired during the last lever 2 press in the correct sequence but not during the middle lever 2 press during the correct sequence (left). This SPN responded weakly during repetitive lever 2 pressing (right).

(C) Mean peri-event SPN spiking during correct (first column), non-repeat incorrect (second column), lever 1 repeat (third column), and lever 2 repeat (fourth column) trials in 9 rats that learned different lever-press sequences indicated on the left.

(D) Mean peri-event activity for all SPNs for the same trial types as in C (solid, $n = 2501$, 2501, 1143 and 1338, respectively) and for SPNs that did not surpass 5-Hz firing rate during correct trials (dashed, $n = 1857$, 1857, 842 and 1002, respectively). Shading indicates SEM.

** $p < 0.01$.

(E) Activity of sub-groups of SPNs responsive to first (blue), second (green) and third (red) lever press in the same trial types (n = 154, 125 and 157, respectively).

(F) Proportion of highly responsive SPNs that fired maximally in the 0.5-s window around the first (blue), second (green) or third (red) lever press, or at other times in the correct trial (gray).

See also Figures S1–3.

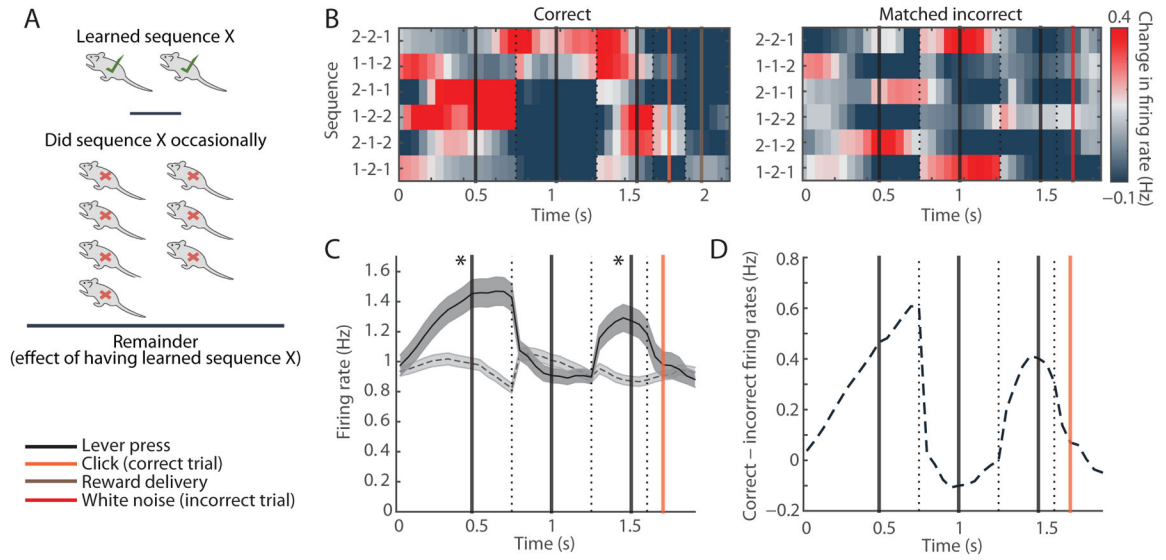


Figure 4. Beginning and End of Specific Sequences Are Emphasized Selectively in Rats That Were Trained on That Specific Sequence

(A) We calculated the difference between the mean population SPN activity in the one or two rats that were trained on a particular sequence and the population SPN activity in the rest of the 9 rats that were not trained to perform it, but did so occasionally as an incorrect trial.

(B) Population spiking of SPNs recorded in dorsolateral striatum in all six matched correct (left) and incorrect (right) sequences.

(C) SPN mean firing rates in correct sequences (solid) and the matched incorrect sequences (dashed). Shading indicates SEM. * $p < 0.001$ (Wilcoxon rank-sum test).

(D) Difference in the SPN firing rates in correct sequences and matched incorrect sequences.

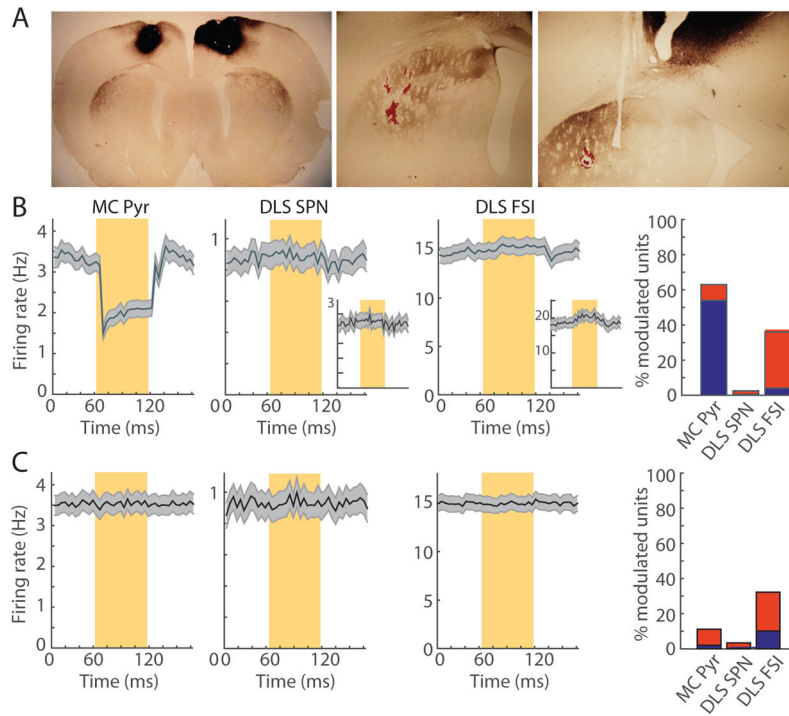


Figure 5. Inhibition of Motor Cortical Cell Bodies and Terminals Has Little Effect on Spiking in Dorsolateral Striatum

(A) Motor cortical injections result in selective halorhodopsin-expressing cortical terminals in dorsolateral striatum (left). Lesion marks (red) indicate tetrode tips within this termination zone (center and right).

(B) Mean firing rate of 383 putative cortical pyramidal neurons (MC Pyr, left), 366 striatal SPNs (DLS SPN, second) and 106 striatal FSIs (DLS FSI, third) during pulses of laser light (yellow shading) delivered to cortical cell bodies in the freely moving rat, and proportions of significantly inhibited (blue) and activated (red) units (right). Insets in the second and third columns show, respectively, the subsets of 96 SPNs with firing rates of >1 Hz and 33 significantly activated FSIs. Shading indicates SEM.

(C) Mean firing rate and proportions of modulated units for the same three cell groups (n = 383, 311 and 108, respectively) when the laser pulses were delivered to the corticostriatal terminals in the striatum, plotted as in B.

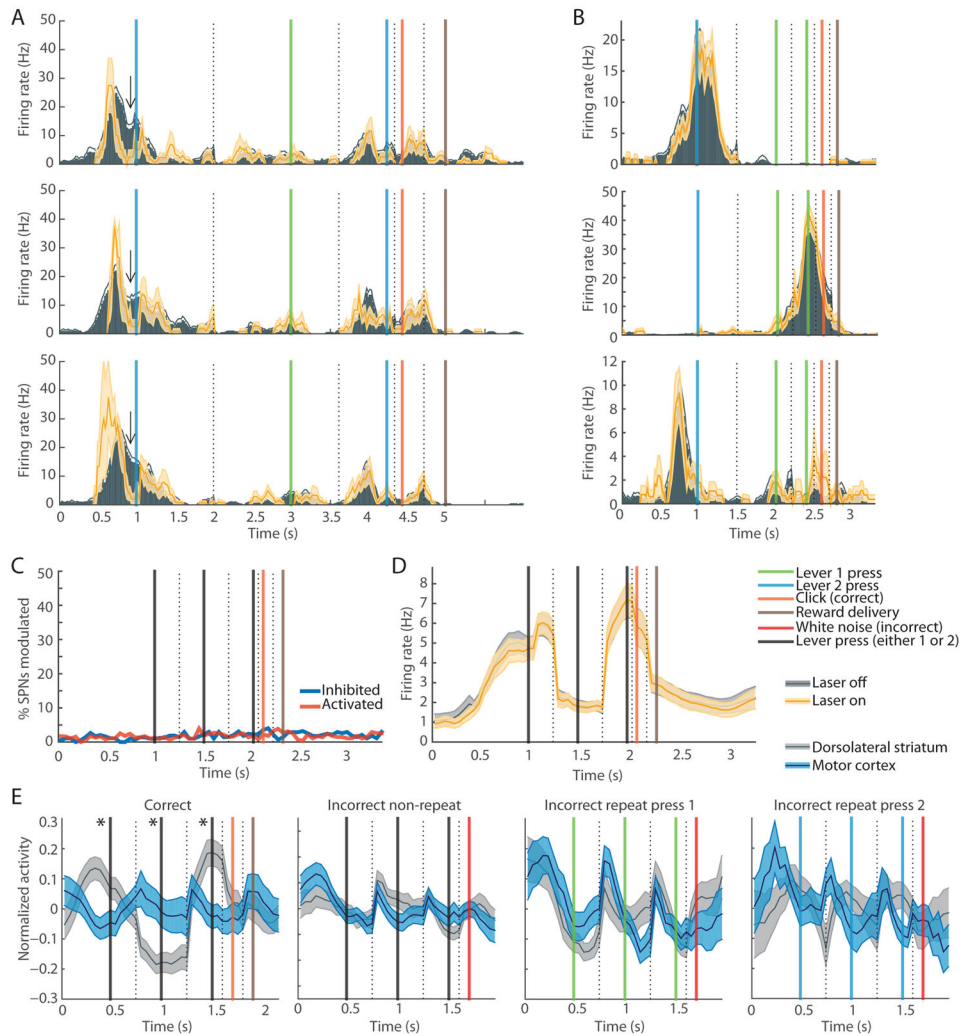


Figure 6. Motor Cortex Is Unlikely to Be Driver of Task-Boundary Activation in Dorsolateral Striatum

(A) Activity of an SPN during laser-off (gray) and laser-on (yellow) times across three training days (rows), showing task-time-selective modulation (arrow) by cortical terminal inhibition. Shading indicates SEM.

(B) Activity of three SPNs recorded in one rat, with no differences in task-related spiking in laser-off (gray) and laser-on (yellow) times.

(C) Proportions of all SPNs inhibited (blue) and activated (red) by cortical terminal inhibition in striatum during correct trial performance.

(D) Beginning-and-end SPN firing during correct trials with (yellow) and without (gray) cortical terminal inhibition ($n = 163$ SPNs).

(E) Session average activity of simultaneously recorded motor cortical pyramidal (blue) and striatal SPN (gray) units in correct trials ($n = 51$ sessions), non-repeat incorrect trials ($n = 57$), lever 1 repeat press trials ($n = 34$), and lever 2 repeat press trials ($n = 20$).

See also Figure S4.

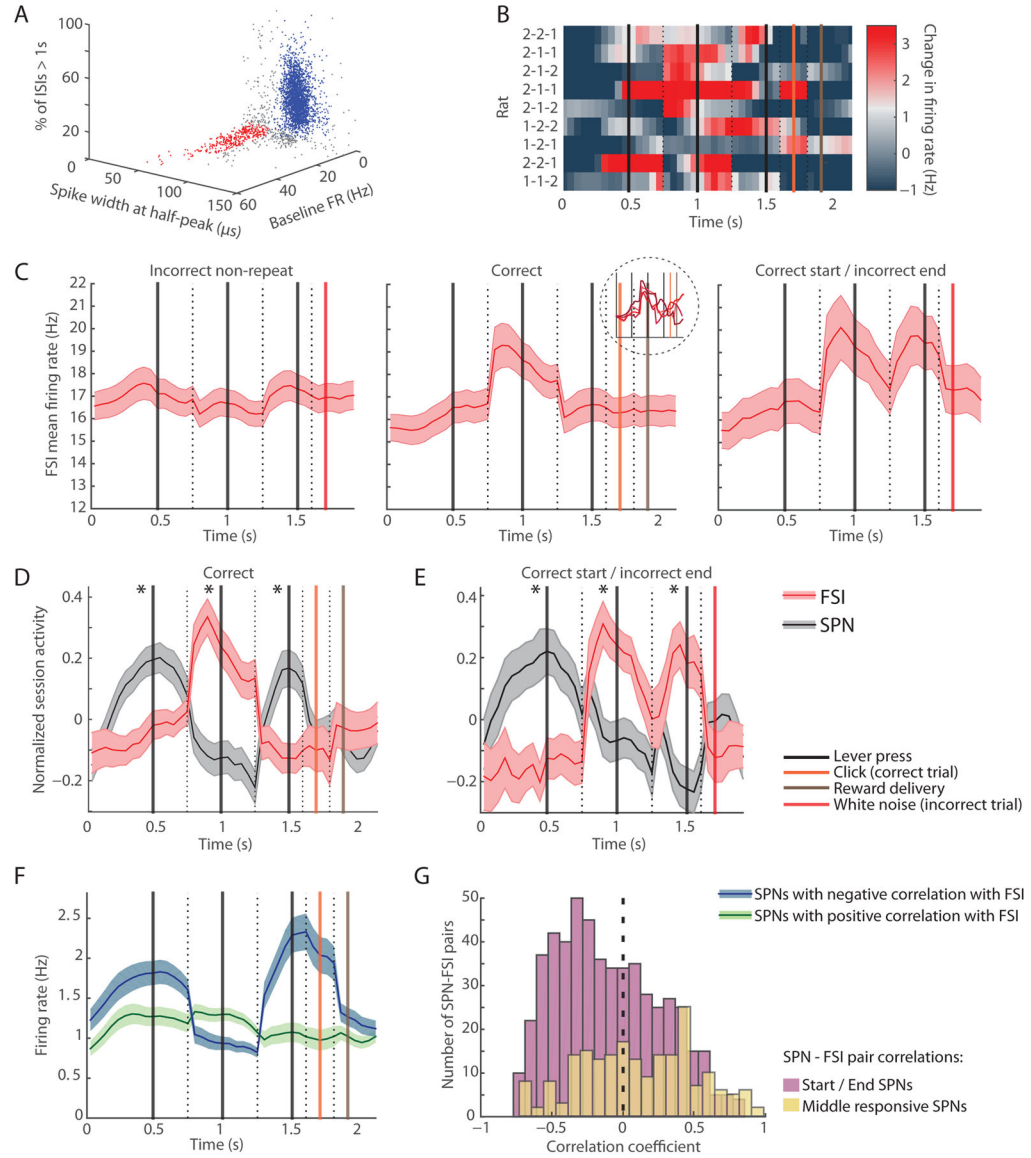


Figure 7. Narrow Spike-Width FSIs Fire in an Opposing Pattern to SPNs during Correct Sequence Performance

(A) Single units recorded in dorsolateral striatum putatively classified as narrow spike-width FSIs (red) and SPNs as used in previous analyses (blue) along features of spike waveform width, firing rate (FR), and proportion of interspike intervals (ISIs) greater than 1 s.

(B) Normalized FSI activity recorded in nine rats performing different learned sequences indicated at left, plotted as in Figure 3C.

(C) Mean FSI firing rates during incorrect trials (left, n = 311), correct trials (middle, n = 284), and trials in which the first two presses were correct but the last press was incorrect (right, n = 153). Inset shows activation of sub-groups of FSIs with different baseline firing rates in correct trials. Shading indicates SEM.

(D and E) Normalized session average activity of simultaneously recorded SPNs (gray) and FSIs (red) in correct trials (D, n = 69 sessions) and correct start trials in which the last lever press was incorrect (E, n = 42 sessions). Shading indicates SEM. * $p < 0.001$.

(F) Activity of SPNs with task-related spiking that was significantly correlated positively (green, n = 977) or negatively (blue, n = 1007) with that of simultaneously recorded FSIs.

(G) Distribution of SPN-FSI pair correlation coefficients for start and end SPNs (purple) and middle SPNs (yellow).

See also Figure S5.