Title

Pallidal spiking activity reflects learning dynamics and predicts performance

Journal

Proceedings of the National Academy of Sciences of the United States of America, 113(41)

ISSN

0027-8424

Authors

Schechtman, Eitan
Noblejas, Maria Imelda
Mizrahi, Aviv D
et al.

Publication Date

2016-10-11

DOI

10.1073/pnas.1612392113

Peer reviewed

Pallidal spiking activity reflects learning dynamics and predicts performance

Eitan Schechtman[a,1], Maria Imelda Noblejas[b], Aviv D. Mizrahi[b], Omer Dauber[b], and Hagai Bergman[a,b]

[a]Edmond and Lily Safra Centre for Brain Sciences, The Hebrew University of Jerusalem, Safra Campus, Jerusalem, Israel 9190401; and [b]Department of Neurobiology, Institute of Medical Research-Israel Canada, Hadassah Medical School, The Hebrew University of Jerusalem, Jerusalem, Israel 9112001

Edited by Peter L. Strick, University of Pittsburgh, Pittsburgh, PA, and approved August 15, 2016 (received for review July 28, 2016)

The basal ganglia (BG) network has been divided into interacting actor and critic components, modulating the probabilities of different state–action combinations through learning. Most models of learning and decision making in the BG focus on the roles of the striatum and its dopaminergic inputs, commonly overlooking the complexities and interactions of BG downstream nuclei. In this study, we aimed to reveal the learning-related activity of the external segment of the globus pallidus (GPe), a downstream structure whose computational role has remained relatively unexplored. Recording from monkeys engaged in a deterministic three-choice reversal learning task, we found that changes in GPe discharge rates predicted subsequent behavioral shifts on a trial-by-trial basis. Furthermore, the activity following the shift encoded whether it resulted in reward or not. The frequent changes in stimulus–outcome contingencies (i.e., reversals) allowed us to examine the learning-related neural activity and show that GPe discharge rates closely matched across-trial learning dynamics. Additionally, firing rates exhibited a linear decrease in sequences of correct responses, possibly reflecting a gradual shift from goal-directed execution to automaticity. Thus, modulations in GPe spiking activity are highest for attention-demanding aspects of behavior (i.e., switching choices) and decrease as attentional demands decline (i.e., as performance becomes automatic). These findings are contrasted with results from striatal tonically active neurons, which show none of these task-related modulations. Our results demonstrate that GPe, commonly studied in motor contexts, takes part in cognitive functions, in which movement plays a marginal role.

basal ganglia | learning | attention | globus pallidus | actor–critic model

The basal ganglia have long been implicated in learning new skills and associations (1). One of the most influential models of these structures distinguishes between two domains: the main axis, whose function is to execute and choose between different actions based on their expected outcome, and the neuromodulators, which act to shape the connectivity within said axis by incorporating information regarding the results of actions on the internal and external state. These components are aptly named the actor and critic, respectively (2).

The roles of the striatum, the largest structure in the main axis, and dopamine, the key neuromodulator serving as a critic in the basal ganglia framework, have been thoroughly studied. Dopamine modifies the efficacy of the corticostriatal synapses, as well as the striatal excitability directly, using signals that incorporate both expectation and external gains and losses (3–5). Decades of research have revealed the manner in which signals relaying information concerning expected and actual gains and costs are incorporated in the striatal dynamic system (6–9). However, these dopamine- and striato-centric views often fail to take into account our current understanding of the basal ganglia, which acknowledges that the actor (main axis) network has more than one input/output hub and employs multiple reciprocally and feed-forward connected computational components (10).

Considering downstream structures of the basal ganglia might enable us to expand these classic models and shed further light on the manner in which their function is implemented in the neural network. In this study, we focused on the role in learning of the external segment of the globus pallidus (GPe), a key player in the ganglia's computational physiology (10, 11). This nucleus receives inputs from both the striatum and subthalamic nucleus (STN) and is presumably the only to project strongly on all of the input and output structures of the basal ganglia main axis.

As is the case for the striatum, pallidal activity corresponds not only to movement planning and execution but also to high-level aspects of behavioral tasks such as predicted reward and context (10, 12–17). However, the role of the GPe in basic cognitive functions, in which the motor aspect plays a minor, marginal part and is not inherent to the task, has rarely been addressed directly (18), nor has its contribution in the context of actor–critic models. One of the few studies to focus on pallidal activity in a cognitive context in humans found that blood-oxygen-level-dependent (BOLD) activity correlated with working memory capacity and with the ability to disregard distractors (19). However, the BOLD signal does not fully reflect neural activity, and the fMRI resolution is not good enough to make the distinction between the two segments of the pallidum.

To reveal the extent of the GPe's involvement in the shaping of cognitive behavior, we recorded the extracellular spiking activity of high-frequency discharge GPe neurons [HFD, probably the prototypical cells (14, 20, 21)] from two monkeys (Fig. 1 *A* and *B*) while they were performing a deterministic three-choice reversal learning task (22, 23) (Fig. 1*C*). For each block, one of the three optional targets was randomly selected and associated with reward. In each trial (T_n) the monkeys had to choose a single target. Only one of the targets resulted in reward, and in next trial (T_{n+1}), the monkeys could either maintain their previous choice or switch to another target (Fig. 1*D*). The task required the monkeys to

Significance

The basal ganglia (BG) are a set of interconnected nuclei deeply buried within the brain that are involved in action selection and habit formation. Classically considered motor nuclei, their role in cognitive performance has become widely appreciated over time. Current models of learning in the BG focus on striatal neurons and the neurotransmitter dopamine, but these do not fully account for observed behaviors. In this paper, we considered the learning-related activity of the external globus pallidus (GPe), a downstream BG nucleus. We show that GPe spiking activity predicts future performance, corresponds with learning dynamics, and decreases as performance becomes more automatic. Taken together, our data reveal the role of GPe in learning and open new avenues for research.

Author contributions: E.S. and H.B. designed research; E.S., M.I.N., A.D.M., and O.D. performed research; E.S. analyzed data; and E.S. and H.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

[1]To whom correspondence should be addressed. Email: eitan.schechtman@mail.huji.ac.il.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1612392113/-/DCSupplemental.

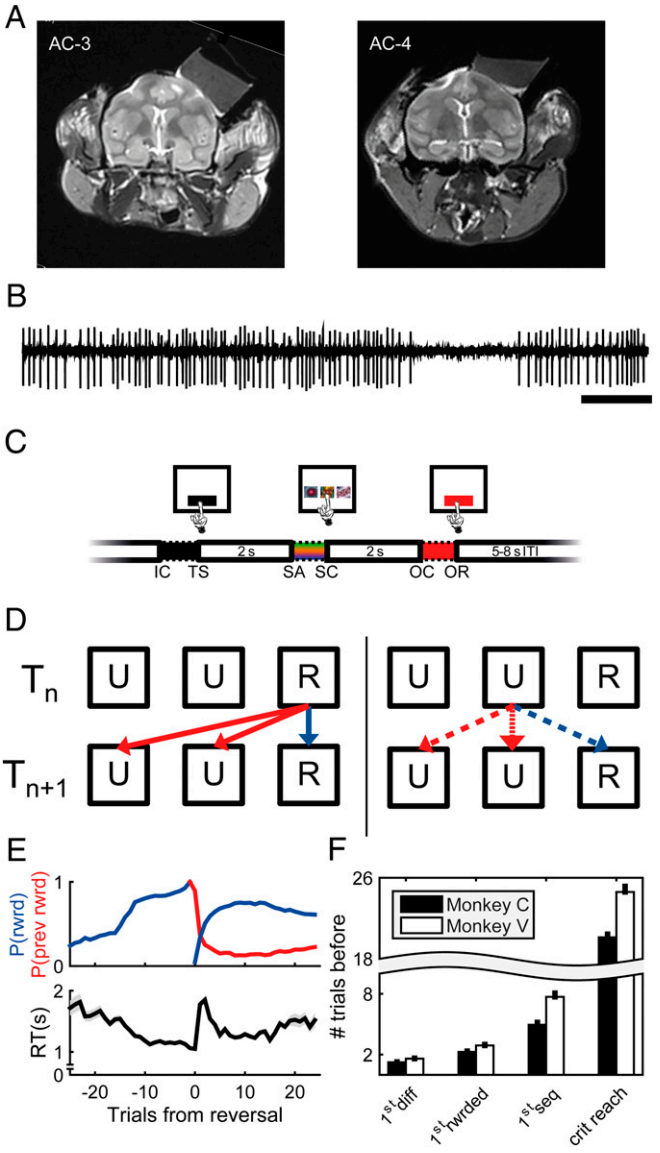

Fig. 1. Experimental paradigm and behavior. (*A*) Coronal MRI scans showing the locations of the recording chambers. (*Left*) Monkey C. (*Right*) Monkey V. (*B*) The 300- to 6,000-Hz filtered traces from a single electrode showing spontaneous GPe spiking. Note the "pause" in high frequency firing of the recorded unit, implicating it as one belonging to GPe (14, 54). (Scale bar, 100 ms.) (*C*) Behavioral paradigm. Trials were preceded by a presentation of an initiation cue (black rectangle). Two seconds after the monkeys touched the cue, three square fractal stimuli appeared, one of which was deterministically associated with reward. Two seconds after choosing a stimulus, an outcome cue (red rectangle) appeared. When the cue was pressed, the monkeys received either liquid reward or no reward, depending on the selected stimulus. Once a predefined learning criterion was reached, a different stimulus was associated with reward (reversal). There was no explicit cue for the reversal, and the monkeys had to learn by trial and error. IC, initiation cue; OC, outcome cue; OR, outcome revealed; SA, stimuli appear; SC, stimulus chosen; TS, trial start. (*D*) Behavioral patterns in sets of two trials. If the first trial (T_n) was rewarded (*Left*), it can be followed by either another rewarded trial (R⇒R; solid blue) or a different unrewarded one (R⇒U, solid red). If T_n was unrewarded (*Right*), it could be followed in T_{n+1} by an identical choice (U⇒U; dotted red), a different unrewarded choice (U⇒U; dashed red), or a different rewarded choice (R⇒R, dashed blue). (*E*) Success rates and response times per trial before and after reversal. (*Upper*) The blue (red) line signifies the probability of choosing the currently (previously) rewarded stimulus. (*Lower*) Response times. Values represent the mean of 314–1,411 repetitions per trial position. Shaded areas indicate SEM. (*F*) Mean performance metrics after reversals for both monkeys. From left to right, number of trials until first choosing a stimulus other than the previously rewarded one, first choosing the newly rewarded stimulus, first choosing the newly rewarded stimulus for at least three sequential trials, and reaching the criterion for reversal. Values represent the mean results of 662/750 blocks for monkeys V/C, respectively. Error bars indicate SEMs.

reacquire new stimulus–outcome (S-O) associations once a predefined learning criterion was reached (Fig. 1 *E* and *F*; correct choice for 12–15 trials out of the last 25, details in *Materials and Methods*). Our results revealed that switches in the chosen stimulus involved predictive changes in an index of GPe discharge rate modulation. Considering the neural activity involved in the monkeys' reaction to changes in S-O contingencies, we show that the temporal dynamics of neuronal discharge rates across trials matches the slope of the learning curves and that spiking rates linearly decline as rewarded responses are repeated.

Results

A total of 306 GPe HFD units met our inclusion criteria (118 and 188 units for monkeys C and V, respectively). The average (± SEM) isolation quality for these units was 0.91 ± 0.003 (24) and their average firing rate (FR) was 66.3 ± 1.61 spikes per s. The mean time period for which units were stably held was 32.2 ± 1.77 min (97.09 ± 5.57 trials and 4.27 ± 0.28 changes in S-O associations per unit).

GPe Spiking Activity Predicts Subsequent Performance. Following reversals of the rewarded stimulus (i.e., changes in S-O associations), the monkeys rapidly adapted their behavior toward the newly rewarded choice, influencing both their response times and error rates (Fig. 1 *E* and *F*). However, as shown in Fig. 1*E*, *Upper*, erroneous choices were not uncommon, not only immediately following the reversal but also later. GPe discharge rates were modulated not only throughout the course of each trial but also across trials. We focused our analysis on pairs of consecutive trials (T_n and T_{n+1}; Fig. 1*D*) and considered the dynamic spiking activity in four different 1-s epochs: directly after the outcome was revealed (OR) toward the end of T_n, during the last second of the intertrial interval (ITI) between T_n and T_{n+1} just before the initiation cue (IC) of T_{n+1}, just before the monkeys made their choices (stimulus choice; SC) for T_{n+1}, and directly after the OR of T_{n+1} (Fig. 2*A*).

GPe neurons encode information using both decreases and increases in their discharge rates (e.g., refs. 18, 25, and 26). Moreover, the range of the spontaneous discharge rate of different pallidal neurons is very broad (in this study 20.2–148.7 spikes per s; see, e.g., refs. 27 and 28). Therefore, we calculated the absolute Z-scores (baseline mean and SD calculated over last second of the ITI preceding T_n) of the FRs during these four periods. We then considered whether the means of these scores differed according to the behavioral pattern the monkeys displayed (Fig. 2 *B* and *C*). The stimulus chosen in T_n could either have been rewarded or not, and the choice was either maintained or switched in T_{n+1} (Fig. 1*D*). Switching from an unrewarded choice may have resulted in either a different unrewarded choice or a rewarded one (Fig. 1*D*, *Right*, dashed red and blue arrows, respectively); switching from a rewarded choice always results in an unrewarded choice (*Materials and Methods* and Fig. 1*D*, *Left*, red arrows). Distinguishing between the average absolute Z-scores obtained for these five behavioral patterns enabled us to examine the manner in which GPe spiking data encoded the recently received outcome (i.e., whether a liquid reward was administered or not) and the subsequent switch of the chosen stimulus (i.e., whether a switch would be made or not). Switching related changes can be observed in all fours panel of Fig. 2*C*. In both the right set of bars (representing previously rewarded trials) and the

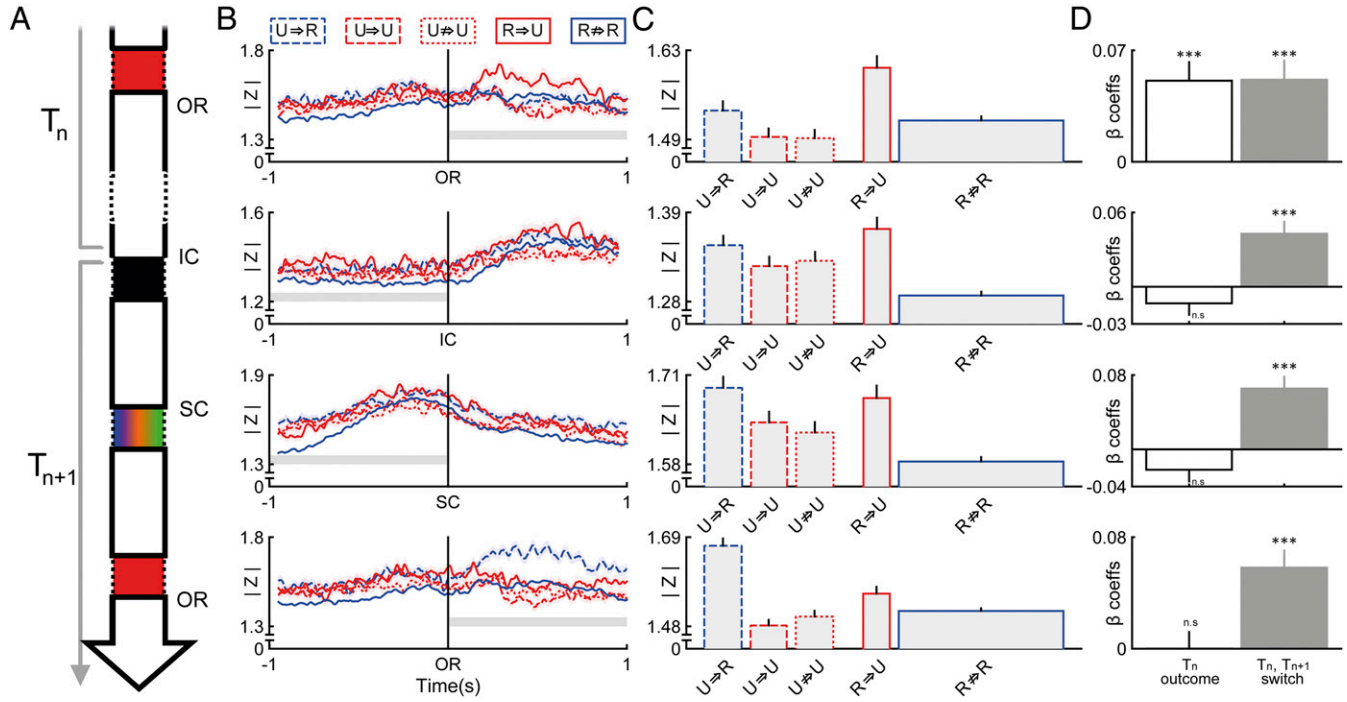

Fig. 2. Subsequent switching of the chosen stimuli is encoded by absolute modulation of GPe FRs. (*A*) To examine the predictive nature of modulations in GPe discharge rates, we considered different epochs throughout the course of pairs of consecutive trials (i.e., T_n and T_{n+1}). Top to bottom: the first second after the outcome of T_n was revealed (OR); the last second of the ITI, before the initiation cue (IC); the last second before a stimulus was chosen in T_{n+1} (SC); and the first second after the OR of T_{n+1}. (*B*) Per epoch, we examined the mean absolute Z-scores of the FR for different response patterns: no reward in T_n, switch to rewarded choice in T_{n+1} (U⇒R; dashed blue; $n = 3{,}161$–$3{,}384$ trials); no reward in T_n, switch to different unrewarded choice in T_{n+1} (U⇒U; dashed red; $n = 3{,}090$–$3{,}227$ trials); no reward in T_n, identical unrewarded choice in T_{n+1} (U⇏U; dotted red; $n = 3{,}231$–$3{,}349$ trials); reward in T_n, switch to unrewarded choice in T_{n+1} (R⇒U; solid red; $n = 2{,}261$–$2{,}367$ trials); and reward in T_n, identical rewarded choice in T_{n+1} (R⇏R; solid blue; $n = 13{,}979$–$14{,}666$ trials). Gray areas signify the examined epoch for *C* and *D*. (*C*) Mean values for each epoch and condition. Bar widths reflect the portion of trials that corresponded to each response pattern (e.g., most trials belonged to the R⇏R group). (*D*) β coefficient estimates for a generalized linear model predicting the absolute Z-scores based on the previous outcome and on whether the choice in T_{n+1} differed from that in T_n. n.s, $P > 0.05$, ***$P < 0.001$. Shaded areas/error bars in all panels represent SEMs across trials. Data for this figure were extracted from all 306 GPe units that passed the inclusion criteria.

left ones (representing previously unrewarded trials) a clear increase in absolute Z-score is associated with subsequent switching.

We used a generalized linear model (*Materials and Methods*) to examine the contribution of T_n outcome and T_n–T_{n+1} switching to the obtained values (i.e., the β coefficients of each effect). Our results revealed unique temporal dynamics throughout the different epochs (Fig. 2*D*). In the epoch following the outcome or lack thereof (OR) for T_n, both previous outcome and subsequent switching significantly contribute to the observed values ($P < 0.001$). The effect of the previous outcome is somewhat obvious, because this epoch directly follows reward delivery or withholding, yet the predictive nature of this time period regarding the choice in the next trial, which is seconds away, is less trivial. As T_n ends and T_{n+1} starts, the previous outcome is no longer encoded in the GPe discharge rates, yet the encoding of the subsequent switch remains intact throughout the course of the trial ($P < 0.001$). Higher absolute Z-score values consistently predicted that the monkeys would choose a stimulus different from the one chosen in the last trial.

The lower panels of Fig. 2, which display the absolute Z-scores following OR in T_{n+1}, reveal another interesting feature of GPe neural activity. Following a switch from an unrewarded choice in T_n to a rewarded one in T_{n+1}, discharge rates were strongly modulated (dashed blue line). This activity is probably not associated with the received reward itself, because its temporal profile differs greatly from that observed for trials in which the reward was repeatedly delivered (solid blue line). The pattern of activity therefore resembles a reinforcement cue that could serve as a learning signal shaping subsequent performance. Taken together with the elevated scores during the first outcome-related epoch for rewarded but unrepeated choices (solid red line in the top panels of Fig. 2 *B* and *C*), it seems that GPe spiking activity is highly modulated following rewarded outcome that is either a result of, or will be followed by, a behavioral shift.

It should be noted that some of the inspected epochs directly follow external stimuli such as the presentation of visual (SC, stimulus choice following the stimulus appearance, SA, Fig. 1*C*) or visual and auditory cues (OR). However, these stimuli are maintained constant irrespective of the monkeys' behavior and any experimental condition. Therefore, the external stimuli should not have had any major effect on our results, which focus on differences between conditions and consider any constant effect of external stimuli as noise. It should, however, be acknowledged that the epoch preceding choice (SC) comprises both decision making and execution (i.e., hand movement toward the chosen stimulus), which were not teased apart in our experimental design. Nevertheless, because our results were collapsed across stimulus positions, the reported results do not merely manifest motion-related neural activity.

GPe Spiking Activity Correlates to Learning Slopes. In the previous paragraph, we emphasized the correlation between the modulation of GPe spiking activity and subsequent behavior; however, we neglected the temporal structure of our behavioral task that included repeated learning of new stimulus–action association. To further explore the role of GPe in learning, we considered the relation between the changes in the task-related behavior of the monkeys and the corresponding discharge rates. For this analysis, we did not use absolute Z-score, because we were interested in the

slow temporal dynamics of the baseline GPe discharge rates across long time spans, including many trials, and not the phasic modulation within a trial. Using Z-scores requires the definition of an appropriate baseline timeframe that is not available for the kind of analysis we conducted here. We concentrated on the discharge rate dynamics during the epoch just before the choice of the monkeys (SC, stimulus chosen, Fig. 3).

Across sets of rewarded choices, discharge rate dynamics depended on the number of trials that have occurred since the last reversal (Fig. 3*A*). To explore this effect, we first calculated the change in probability to choose the newly rewarded stimulus following reversal (i.e., the learning curve) and its derivative [i.e., the learning slope (29)]. We then correlated these two measures to the discharge rates of the GPe units just before the monkeys made a choice in each of the first 10 trials following reversal. We found that both measures correlated strongly to the discharge rates ($r_{pearson}$ to the curve, slope: −0.938 and 0.934, respectively, $P < 0.001$ for both; Fig. 3*B*). A similar pattern emerged when we restricted our analysis to the first five trials after reversal, before a learning plateau was reached ($r_{pearson}$ to the curve, slope: −0.9 and 0.91, respectively, $P < 0.05$ for both).

Due to the highly anticorrelated nature of the learning curve and slope ($r_{pearson} = -0.99$, $P < 0.001$ for all data points and for the first five trials), it is not surprising that a strong (positive) correlation to one would bring about a strong (negative) correlation to the other as well. This correlation between learning curve and slope arises from the relatively fast exponential rise in the learning curve over the first few trials. Therefore, we analyzed a subset of postreversal trials that were characterized by slower learning, nonexponential learning curves, and lower correlation between the curve and slope. Only blocks in which the first sequence of three (or more) correct choices started relatively late (i.e., at least four trials) after reversal were considered. Consequently, the initial segments of the learning curve and slope (i.e., the first five trials after reversal) were no longer significantly correlated ($r_{pearson} = -0.68$, $P > 0.05$, Fig. 3*C*). The relation between the discharge rates in this reduced dataset and the learning curve and slope allowed us to show that neural activity during the initial learning segment correlates to the learning slope and not to the learning curve (Fig. 3*C*, $r_{pearson}$ to the curve, slope: −0.73 and 0.99, $P > 0.05$ and $P < 0.001$, respectively). The correlation coefficient between the slope and the discharge rate was significantly stronger than that between the curve and the discharge rate ($P < 0.05$, one-way correlation comparison using Fisher r-to-z transformation).

The basal ganglia play an important role in habit formation and automaticity (30). Following reversal and learning of the new S-O contingencies, the monkeys commonly engaged in long sequences of repeatedly correct trials. These can hardly be deemed "habits" in their strict sense, but the repetitive nature of these series of trials may serve to expose the GPe activity underlying the development of automaticity. We revealed that GPe FRs decreased over sequences of at least 10 correct trials ($r_{pearson} = -0.93$, $P < 0.001$, Fig. 4 *A* and *B*). We ruled out the possibility that this was due to adaptation to the repeated motion toward a specific stimulus by examining the change in activity for incorrect sequences (Fig. 4*C*). Indeed, the decrease in activity was specific only to repeatedly rewarded sequences (for sequences of four, five, and six unrewarded choices, $r_{pearson} = -0.81$, −0.15, and 0.28; $P = 0.19$, 0.81, and 0.58, respectively).

The interpretation of these results as evidence for a decrease in GPe activity as behavior becomes automated can be challenged by another explanation. It could be that the decline manifests the increase in predictability (and decrease in surprise) of the subsequent outcome. Although we cannot rule this possibility out, there is some evidence against it. In Pavlovian conditioning paradigms in which the presented cue deterministically predicts the outcome, learning involves a gradual increase in outcome predictability around the cue epoch. The dopamine response, for example, first follows the rewarded outcome but gradually shifts to the cue epoch after learning the contingency (3, 31). Similar results have also been shown for striatal tonically active neurons (TANs; refs. 32 and 33),

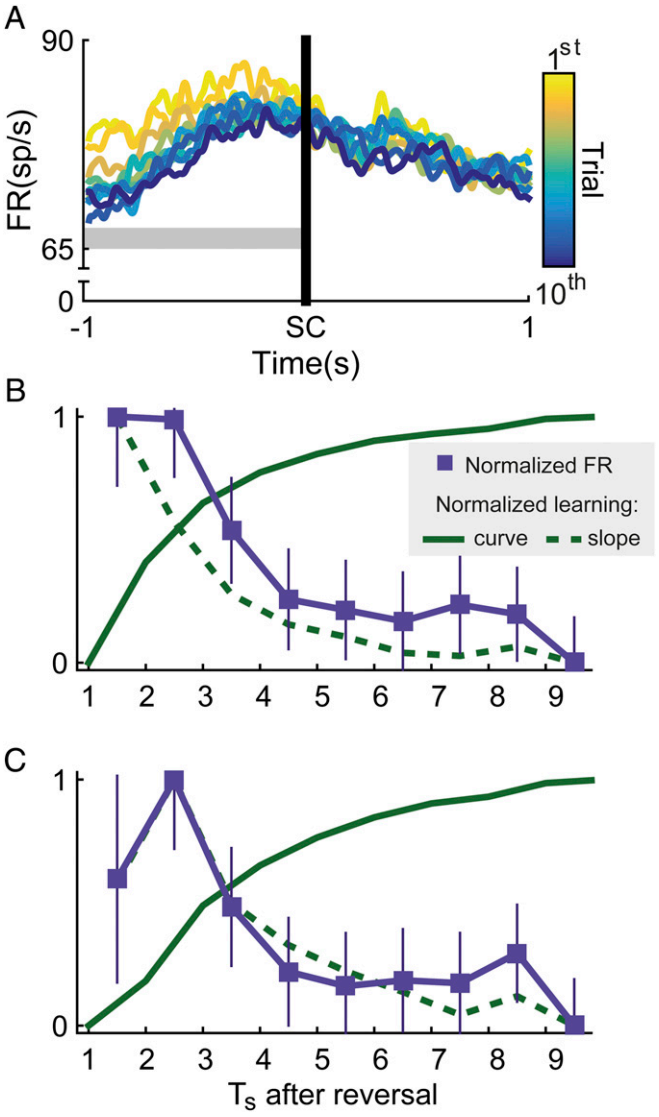

Fig. 3. GPe FR correlates with learning curve slope. (*A*) Population mean of unit discharge rates for the time segment around stimulus choice (SC) for rewarded trials following reversal. Color codes from yellow to blue indicate the indices of the trial positions following reversal (e.g., yellow signifies the first trial after reversal). The number of trials averaged ranged between 475 and 1,019. The gray area signifies the examined epoch for *B* and *C*. (*B*) Purple squares denote the normalized mean GPe discharge rates during the last second before SC for all rewarded trials that follow reversal (e.g., the leftmost value signifies the normalized mean FR of all rewarded trials around the first SC following reversal). The solid green line depicts the normalized probability of success following reversal (i.e., the learning curve). The dashed green line depicts the normalized derivative of the learning curve (i.e., the learning slope). For clarity, all three plots were normalized, so the *y* axis ranges from 0 to 1. (*C*) Same as *B*, except only blocks in which the first sequence of at least three correct trials began more than three trials after reversal were considered (i.e., blocks with nonoptimal reversal). This selection dissociates between the hypotheses that FR is correlated to the learning curve and that it is correlated to the learning rate. The number of trials for which the FR was averaged ranged between 159 and 749. Error bars in *B* and *C* represent SEMs across trials. Although some bars are truncated, they are all symmetric around the means so no data are obscured. Data for this figure were extracted from all 306 GPe units that passed the inclusion criteria.

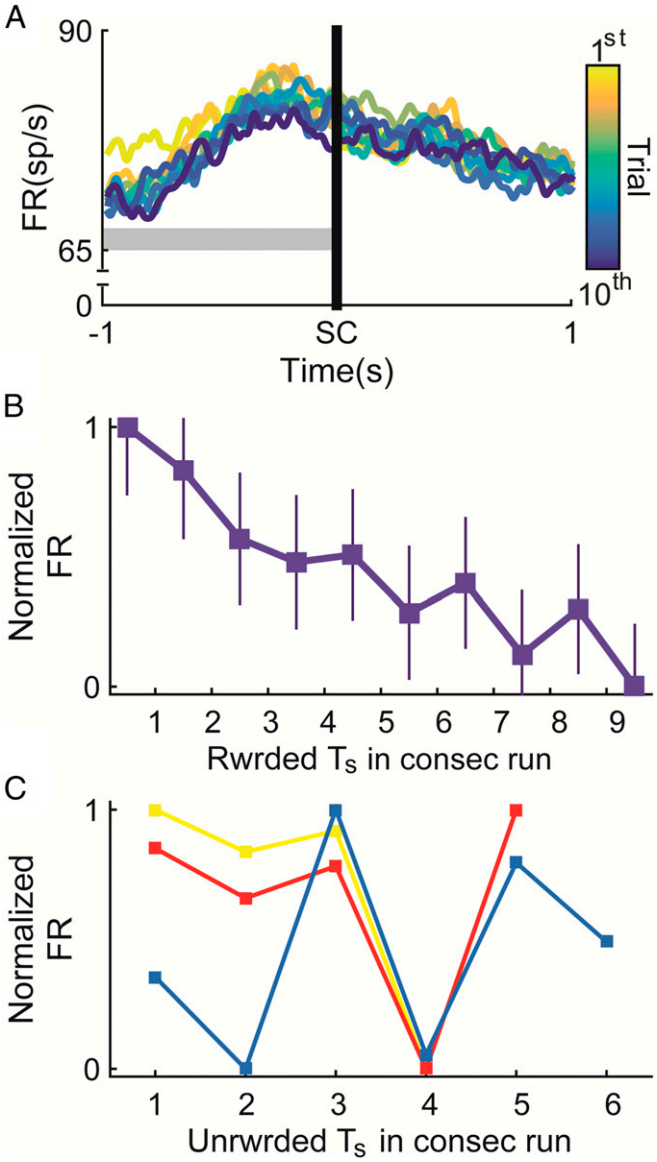

Fig. 4. GPe FR decreased linearly during consecutive correct sequences of trials. (*A*) Population mean of unit discharge rates for the time segment around stimulus choice (SC) for the first sequence of 10 consecutive correct choices of each block. Color codes from yellow to blue indicate the indices of the trial positions along the consecutive run (e.g., yellow signifies the first trial in the sequence). The number of trials averaged ranged between 862 and 883. The gray area signifies the examined epoch for *B* and *C*. (*B*) Normalized mean GPe discharge rates during the last second before SC for all trials in sequences of at least 10 rewarded trials($r_{pearson}$= −0.93; $P < 0.001$). Error bars represent SEMs across trials. (*C*) Normalized mean GPe FR during the last second before SC for all trials in sequences of four, five, or six (yellow, red, and blue lines, respectively) unrewarded trials in which the same stimulus was repeatedly chosen ($r_{pearson}$ = −0.81, −0.15, and 0.28; P = 0.19, 0.81, and 0.58, n = 240–242, 87–89, and 36–40, respectively). Unlike in *B*, no decrease was observed throughout the sequence. Although some bars are truncated, they are all symmetric around the means so no data are obscured. Data for this figure were extracted from all 306 GPe units that passed the inclusion criteria.

a group of presumably cholinergic interneurons (34, 35) that are considered part of the BG critic system (36). Basically, this implies that signals reflecting the predictive nature of a cue (or outcome-related choice) should show learning dynamics that are in the opposite direction relative to the postoutcome dynamics. In the context of our task, we would expect the gradual decrease in GPe activity around the SC epoch to be mirrored by an increase in the outcome-related activity across trials as the association is learned. However, our data show the opposite pattern. Just like the pre-choice discharge rates, the postoutcome activity decreases in a linear fashion across consecutively correct trials ($r_{pearson}$ = −0.89; $P < 0.001$, Fig. S1), thus lending support to the alternative hypothesis that GPe activity encodes automaticity and not predictability.

Neither the correlation between the postreversal learning curve and discharge rates nor the linear decrease in spiking in conjunction with sequences of consecutively correct choices is limited to the population level. Analysis of correlation on a single-unit level revealed that the R^2 distributions for both effects are skewed to the right relative to shuffled data (Fig. S2). We therefore conclude that the observed results on the population level do not arise from a small group of units with high correlation coefficients, but rather from a small yet consistent effect across the population. We found no evidence for any region-specific clustering of highly correlated units in either monkey (Fig. S3).

Striatal TAN Activity Neither Predicts Subsequent Behavior nor Correlates to Learning Slopes. To consider the striatal learning-related activity we recorded TANs from monkey V, 79 of which met our inclusion criteria. The average (± SEM) isolation quality (24) for these units was 0.88 ± 0.006 and their average FR was 5.23 ± 0.17 spikes per s. The mean time period for which units were stably held was 37.5 ± 3.87 min.

We conducted the exact same analysis for the TANs. To assess the effects of switching behavior and of the previous outcome (i.e., as depicted for GPe units in Fig. 2), we used unaltered Z-scores and not absolute ones (but results were similar for absolute Z-scores). Throughout most of the considered epochs (with the exception of the ITI-related epoch) TANs did not significantly encode subsequent switches (Fig. S4*A*). Comparing the effect sizes of switching and of the previous outcome (i.e., the portion of variance explained by changes in discharge rates; η^2) between TANs and GPe units reveals a contrast between the two: Whereas TAN activity accounts mostly for the variance explained by the outcome of T_n, GPe activity mostly accounts for the variance explained by the T_{n+1}–T_n switch (Fig. 5*A*).

Next, we examined whether TAN discharge rates corresponded to learning curves in a manner similar to that of GPe cells. The TAN FR did not correlate to the learning curve or slope (Fig. 5*B* and Fig. S4*B*, *Left*; $r_{pearson}$ to the curve, slope: −0.46 and 0.51, respectively, $P > 0.05$ for both). The same pattern of nonsignificant correlations was also observed for the first five trial segment following reversal ($r_{pearson}$ to the curve, slope: −0.57 and 0.61, respectively, $P > 0.05$ for both). Similarly, the gradual decrease observed for sequences of correct choices explained only a relatively small part of the variance of TAN FR (Fig. 5*C* and Fig. S4*B*, *Right*, $r_{pearson}$ = 0.56, $P > 0.05$).

The response pattern of TANs is often shorter than that for GPe cells (e.g., see Fig. S4*B*). To confirm that our results are not affected by these differences, we ran all TAN-related analyses using shorter, 500-ms-long time windows. The results were identical to those received for the 1-s-wide time epochs considered for the previous analyses.

Discussion

We examined the correspondence between GPe neural activity and the cognitive aspects of behavior in monkeys performing a deterministic reversal task. Our results show that GPe spiking activity predicted whether the monkeys would switch or maintain their previous stimulus choice. This foreshadowing activity was observed throughout the time period between the consecutive trials, starting immediately after the outcome of the first trial and lasting up until the foretold choice was made. Additionally, GPe

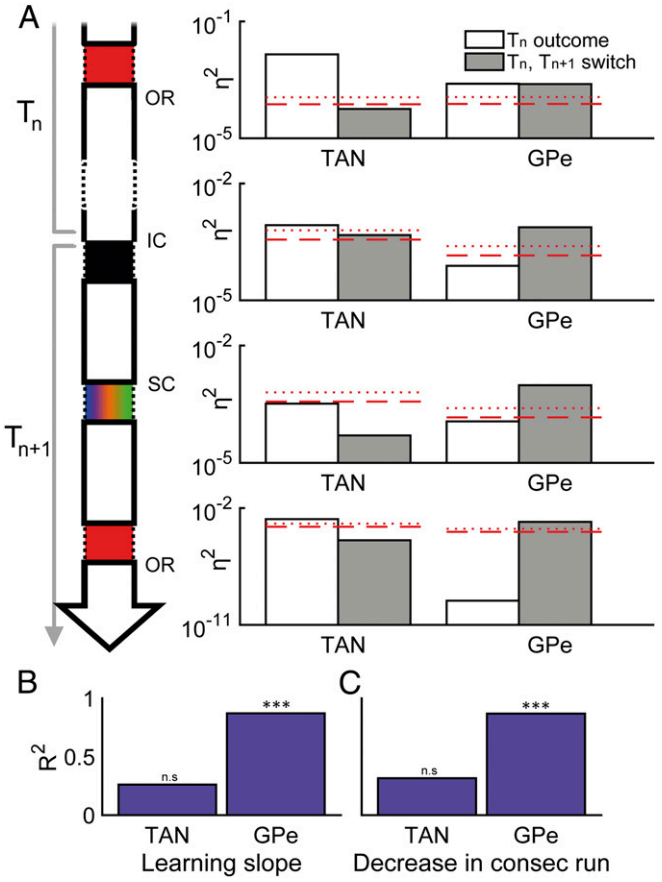

Fig. 5. Unlike GPe units, TANs do not predict switching throughout most of the trial, do not correlate to learning, and do not linearly decrease in repeatedly correct sequences. (*A*) We conducted an ANOVA to compare the effect size between both cell types. We considered the portions of variance (η^2; log scale) explained by changes in TAN FRs (n = 10,506–10,561 trials) and absolute changes in GPe FRs (n = 25,722–26,993 trials) for encoding of previous outcome (white) and future switching (gray). Dashed/dotted lines represent significance threshold for P = 0.05 and P = 0.01, respectively. (*B*) Comparison between the variance explained using linear regression (R^2) between observed learning slopes and either TAN or GPe discharge rates (FRs) for rewarded trials during the last second before the screen was touched (SC, stimulus chosen). (*C*) Comparison between the variance explained (R^2) by fitting a linear function to either TAN or GPe FR during the last second before SC in sequences of 10 repeatedly correct trials. n.s, P > 0.05, ***P < 0.001. Data for this figure were extracted from all 306 GPe units and all 79 TANs that passed the inclusion criteria.

spiking activity following outcome delivery encoded whether such switches were successful.

As performance became more automated and less demanding—either following reversal or in consecutively correct trial runs—GPe discharge rates in the context of choice selection decreased. These patterns of correlation were evident at the population level but also at the level of single units. They do not seem to be limited to a specific topographical region of GPe. Such correlations could be interpreted in two ways: It could either be that changes in GPe discharge rates reflect an increase in reward predictability, which is not action-dependent, or that they reflect a shift from goal directed behavior to automaticity. Although we cannot rule either out, we believe the latter explanation to be more plausible. An increase in reward predictability is mirrored by a decrease in the surprise following outcome delivery. We should therefore expect opposite patterns of activity to emerge around stimulus choice and reward administration. However, our data show that both decline linearly, in a manner that befits gradual habituation.

The function of GPe in behavior has been studied for decades (12–14), but historically the focus has been on its role in motor or motor-related contexts. The shift toward models exploring the cognitive role of the basal ganglia (i.e., its role in learning and decision making) has focused much of the attention on the striatum and dopamine, often overlooking underlying downstream structures (3, 37, 38). The fact that GPe activity encodes more than just motor execution has been well established (15, 17), but this study explores the GPe function in a purely cognitive context (18). Although we acknowledge the motor role of GPe, we intentionally neglected it in our analyses. This was done, for example, by focusing on the information carried by GPe during time epochs that were seconds ahead of motor execution. Additionally, we averaged upon different motor actions with identical cognitive characteristics in all our analyses.

In this paper, we show the neural activity of GPe units during learning. Our monkeys were overtrained in the task but still had to continuously acquire which stimulus was rewarded on a block-by-block basis. Our results are somewhat surprising, because they differ from results obtained for striatal spiking activity in similar tasks. Along with the STN, the striatum is the main input structure to the GPe. Studies examining striatal activity in cognitively demanding tasks did not show any modulation in the activity during the time epoch preceding the choice of the animal as choices became more and more habitual and automated (7, 29, 39). However, there is consensus that the value associated with an outcome, its cost, and their interaction are all represented in striatal neuronal activity (7). This pattern is reflected in our own data, obtained from striatal TANs. Although these neurons are not considered part of the basal ganglia main axis (36), their response is still indicative of the basal ganglia critic system (40) and striatal information processing. We show that TAN spiking activity does not encode subsequent behavioral shifts but does encode the outcome itself throughout most of the trial. Additionally, TAN discharge rates neither correlate to learning curves nor decrease across trials that are sequentially rewarded.

Unlike the striatum, STN activity was shown to convey information on executive function and attentional constraints (41–43). The only study to date examining the dynamics of activity during learning in the STN of healthy humans showed a decrease in prechoice activity similar to the one we have shown in GPe spiking activity (44) and is in line with our results and with the robust reciprocal connections between the GPe and the STN (45).

The established connection between the STN and executive function and the resemblance between its activity and that of GPe units suggests that perhaps our results in the pallidum reflect some aspects of higher cognitive function. Specifically, the observed decline in discharge rates as new associations are acquired suggests that GPe activity is negatively associated with the level of attention dedicated to the task (or to concentration in general). This could also explain the higher discharge modulation observed in the context of choice switching. To shift from one choice to the other, more attention needs to be dedicated to the task. We acknowledge that attention is a multilayered, highly complex, and barely definable construct. To better understand the involvement of GPe in task-related attention, a more nuanced experimental paradigm should be designed. The GPe's integration of attention-related STN inputs with valence-related striatal activity, along with our findings that hint toward attention-related representations, raise exciting new avenues of research regarding the possible cognitive role of the GPe in the basal ganglia framework.

In conclusion, unlike STN and the internal segment of the globus pallidus, relatively little attention has been dedicated to GPe activity in dysfunctional brains. In recent years, there has been a growing focus on the role of GPe in the cognitive aspects of behavior, including its putative involvement in cognitive deficits identified in psychiatric disorders. Several studies have shown that negative and cognitive symptoms of schizophrenia are correlated

with altered GPe activity. A recent study also identified lower interhemispheric GPe functional connectivity in schizophrenia, associated with the cognitive symptoms of the disorder (46–48). Other psychiatric disorders, such as depression and attention deficit hyperactivity disorder, have been associated with smaller GPe volumes (49–51). Our results reveal yet another facet of the role of the GPe in cognition and behavioral control. We believe that future studies should focus on the causal role of the normal and dysfunctional GPe in cognitive behavior, hopefully paving the way for medical interventions (such as deep brain stimulation) that could improve the prognosis for the cognitive aspects of psychiatric disorders.

Materials and Methods

Animal Training and Behavioral Tasks. Data were obtained from two female vervet monkeys (*Cercopithecus aethiops aethiops*, monkeys C and V), weighing 3.5–4 kg. All data were pooled for both monkeys for all analyses. Results obtained for individual monkeys are shown in Fig. S5. Care and surgical procedures were in accordance with the National Research Council *Guide for the Care and Use of Laboratory Animals* (52) and the Hebrew University guidelines for the use and care of laboratory animals in research, supervised by the institutional animal care and use committee of the Hebrew University and Hadassah Medical Center. The Hebrew University is an Association for Assessment and Accreditation of Laboratory Animal Care internationally accredited institute.

The behavioral paradigm used was a multiblock three-choice reversal learning task (Fig. 1 *C* and *D*). The monkeys used their right (contralateral to the recording side, discussed below) hands to touch stimuli presented on a screen that was located ∼16 cm from their heads (Elo 1939L 19-inch open-frame touchmonitor; Elo Touch Solutions Limited). In each daily session, three square fractal images were randomly selected as stimuli out of a set of 10 possible stimuli. The stimuli were presented in fixed positions throughout the session, either on the left or right or in the center of the white screen. Each trial began with a presentation of a black horizontal box on the bottom of the screen. After the monkeys touched the box, it disappeared. Two seconds later, the three fractal stimuli appeared. After touching (choosing) one of the three, all three stimuli disappeared. Two seconds later, a red horizontal box appeared on the bottom of the screen. Touching this box was followed by three simultaneous results: The box would disappear, a banana-flavored liquid reward was either delivered or not (according to the monkey's choice, discussed in the next paragraph), and an ∼80-ms auditory stimulus was played, obscuring any acoustic artifacts of the food pump. The sound was independent of the trial's outcome. Trials were aborted if no choice was made within 30 s or if the red box was not touched within 30 s. All trials (correct, incorrect, and aborted) were followed by a variable intertrial interval (ITI) lasting 5–8 s.

For each block, only one of the three stimuli was deterministically rewarded. No reward was delivered for the choice of other stimuli. The monkeys had no guiding information pointing them toward the correct stimulus, but eventually learning was established and the probability of choosing the correct stimulus reached a plateau (Fig. 1*E*). The criterion for learning was reached once the monkey chose the rewarded stimulus for 12–15 trials out of the last 25 (the criterion was randomly selected per block). Once this happened, an uncued switch in the identity of the rewarded stimulus occurred (i.e., reversal) and a new block started. Each daily session ended once the monkeys no longer initiated trials. After long periods in which the monkeys did not work, the experimenter occasionally delivered a free reward to remotivate them. The behavioral paradigm was designed and run using the Psychophysics toolbox (53) for MATLAB 2012a (The MathWorks, Inc.).

Monkeys were trained for 5–6 d per week and were allowed free access to water in their home cages. Supplementary food was delivered when the monkeys did not reach the predefined daily calorie minimum. Monkeys were given free access to food on the weekends.

Surgery and MRI. The monkeys were fully trained on the task (4–9 mo) before the recording chamber was implanted (Fig. 1*A*). After the training period, they were operated on under full anesthesia and in sterile conditions. In the surgery, an MRI-compatible Cilux head holder (Crist Instrument) and a square Cilux recording chamber (AlphaOmega) with a 27-mm (inner) side, located above a burr hole in the skull, were attached to the heads of the monkeys. The recording chamber was attached to the skull tilted ∼45° laterally in the coronal plane with its center targeted at the stereotaxic coordinates of the left GPe. All surgical procedures were performed under aseptic conditions and general isoflurane and N_2O deep anesthesia. Analgesia and antibiotics were administered during surgery and continued postoperatively. Recording began after a postoperative recovery period of several days. We estimated the stereotaxic coordinates of the recording target using MRI scans. The MRI scan (General Electric 3 tesla system, T2 sequence) was performed under i.m. Domitor and ketamine light anesthesia.

Recording and Data Acquisition. During recording sessions, the heads of the monkeys were immobilized, and up to eight glass-coated tungsten microelectrodes (impedance 0.15–1 MΩ at 1,000 Hz), confined within a cylindrical guide (1.65-mm inner diameter), were advanced separately (EPS; Alpha-Omega Engineering) into the GPe and the striatum. The electrical activity was amplified with a gain of 20 then filtered using hardware Butterworth filters (high-passed at 0.075 Hz, two poles; low-passed at 10,000 Hz, three poles) and finally sampled at 44.6 kHz (SnR; Alpha-Omega Engineering). TANs and GPe units were identified by the stereotaxic and MRI coordinates, by the electrophysiological hallmarks of the encountered structures along the penetration, and by their own unique characteristics such as GPe neurons' high FRs and pausing behavior (14). Neuronal activity was sorted and classified online using a template-matching algorithm (SnR; Alpha-Omega Engineering).

Analysis of Behavior. To evaluate animal behavior in the task, we measured response times and examined their cognitive performance. Performance was assessed both on a trial-by-trial basis (Fig. 1*D*) and on a block-by-block basis (Fig. 1*E*). Trials in which the monkey made no choice and trials in which the experimenter delivered free rewards were omitted from analysis. Per trial, we measured the probability of making a correct choice before and after reversal. Additionally, we assessed the probability of choosing the stimulus that was rewarded during the previous block. Per block, we used several measures of behavior: (*i*) the number of trials until first choosing a stimulus other than the one rewarded in the previous block (perseverative errors), (*ii*) the number of trials until first choosing the newly rewarded stimulus, (*iii*) the number of trials until first choosing the newly rewarded stimulus for at least three consecutive trials, and (*iv*) and the number of trials until the criterion for reversal is reached. Response times were defined as the length of time it took for the monkeys to make a choice upon presentation of the three stimuli (i.e., reaction plus movement time). By the time the recordings started, the monkeys were overtrained and no across-session differences in the monitored measures of performance were observed (Fig. S6). All data analysis was performed using MATLAB 2014b (The MathWorks, Inc.).

Analysis of Single Unit Firing Activity. Data analysis was conducted only for units that were stably held, well-isolated [isolation score (24) ≥0.8], and unquestionably identified as either TANs or GPe units. Only GPe unit that fired ≥20 spikes per s on average were included. Continuous traces showing GPe spiking activity (Fig. 1*B*) were obtained using a two-pole Butterworth filter (1,000–3,000 Hz).

For the plots included in Fig. 2*B*, showing the dynamic temporal patterns of GPe activity (i.e., post stimulus time histograms, PSTH), we defined a baseline period of one second for calculating Z-scores. We divided that second to 10 nonoverlapping bins and calculated the spike count per bin. This data, per trial, were used as a baseline for the PSTHs. For each time point of the PSTH, the baseline discharge rate for that trial was subtracted from the average count per bin and the difference was divided by the baseline SD. Finally, the absolute value of the obtained Z-score was calculated per time point. We used the last second of the ITI before T_n as the baseline for all epochs. Trial pairs in which a reversal occurred between T_n and T_{n+1} were not considered in this analysis. In all cases, the obtained values were filtered using a 50-ms-wide Gaussian Kernel (MATLAB's filtfilt function).

The values for each time epoch shown in Fig. 2*B* (gray bars) were averaged over time per trial. The resultant values were grouped according to the conditions detailed in the text and averaged (Fig. 2*C*). For Fig. 2*D*, these values were fitted using a generalized linear model (GLM), using the following equation:

$$|z| = \beta_0 + \beta_1 O(T_n) + \beta_2 S(T_n, T_{n+1}),$$

where $O(T_n)$ is the outcome of trial T_n (0/1: un/rewarded) and $S(T_n, T_{n+1})$ indicates whether a switch took place (0/1: un/switched). Estimates for the β coefficients for both effects were obtained and statistically compared with the null hypothesis that $\beta = 0$. To assess the effect size (i.e., the portion of variance explained by switching and the outcome of T_n), we used an ANOVA (which is equivalent to the GLM) to calculate the η^2 of each effect (Fig. 5).

The exact same approach was used for TANs (Fig. 5*A* and Fig. S4*A*), with one exception: The sign of the obtained Z-scores was not altered. Unlike GPe, there is no empirical reason to assume that TANs encode information identically by increases and decreases in discharge rates. The same analysis was also run using absolute Z-scores with similar results.

To consider the relationship between performance and prechoice discharge rates, we calculated the learning curve (the probability of choosing the correct stimulus as a function of trials after reversal) and learning slope (the change in probability of choosing the correct stimulus between trials) for the first 10 trials after reversal. We focused on prechoice activity for these 10 trials and therefore omitted all data from the first postrewarded trial, because the choice leading to it was made before any information regarding the change in stimulus–reward contingencies was delivered to the monkeys. We therefore considered the first nine prechoice periods following the disclosure for most analyses and also focused on the first five trials in cases where we aimed to limit our analyses to the primary acquisition phase. In all these analyses, we calculated the linear correlation between learning curves, learning slopes, and the mean discharge rates directly preceding rewarded choices, calculated for the last second before the monkeys touched the preferred stimulus. This calculation was conducted for the population average.

To tease apart the correlation of discharge rates to the learning curve and slope, we conducted a new population-wide correlation analysis, introducing a new omission criterion. For this analysis, we excluded all blocks in which the first sequence of at least three correct trials was initiated in one of the first three trials after reversal (thus considering only blocks characterized by a slow learning rate).

We also conducted analysis examining the change in discharge rate over repeated correct trials. For this, we fitted a linear function to the prechoice discharge rates of sequences of at least 10 repeatedly correct trials and considered the P value for the beta coefficient and the explained variance (R^2). The same analysis was conducted for repeatedly incorrect sequences of four, five, and six trials. The exact same approaches were used for the analysis of TANs.

In addition to analyzing the correlations detailed in the last three paragraphs on the population level, we considered the same correlations for single GPe units (Fig. S2). We only considered units for which sufficient data were available (i.e., data for at least a single trial were recorded for each of the 10 trial positions following reversal or following the initiation of a rewarded sequence). The discharge rate for each trial position was averaged over repetitions per neuron and the results were correlated to the relevant measure (either the learning curve or a linear slope). A histogram was created for the resultant R^2 values. These results were compared with results obtained by random shuffling of the data. For each unit, the labels for all associated trials, indicating the trial position, were randomly shuffled and the correlation coefficient for the newly arranged data was calculated. The labels for each unit were shuffled 100 times and the results were averaged.

Next, we considered the spatial distribution of correlation coefficient (Fig. S3). Our methods do not provide us with exact localization of the units, yet we do know the 2D stereotactic position in which our electrodes were injected. Roughly speaking, this translates to an anterior–posterior, medial–lateral map of the GPe. However, it should be noted that the flexibility of our electrodes, as well as possible shifts in the brain of the animals over recording sessions, make this estimation highly prone to error.

All data analyses were performed using MATLAB 2014b (The MathWorks, Inc.). The highest/lowest 1% of all data for all tests was omitted to avoid the effects of outliers.

Data Availability. All relevant data are available from the authors upon request.

ACKNOWLEDGMENTS. The authors thank the anonymous referees for extremely helpful remarks and acknowledge the contributions of Sharon Freeman, Zvi Israel, Anatoly Shapochnikov, and Hila Gabbay to this project. This work was supported by research grants from the European Research Council, the Israel Science Foundation, the United States–Israel Binational Science Foundation, the German–Israeli Foundation for Scientific Research and Development, the Adelis Foundation, the Rosetrees Trust, and the Vorst Foundation. E.S. is a fellow of the Harry and Silvia Hoffman leadership program and the Adams Fellowship Program of the Israel Academy of Sciences.

1. Smith KS, Graybiel AM (2016) Habit formation. *Dialogues Clin Neurosci* 18(1):33–43.
2. Houk JC, Adams JL, Barto AG (1995) A model of how the basal ganglia generate and use neural signals that predict reinforcement. *Models of Information Processing in the Basal Ganglia*, Computational Neuroscience, eds Houk JC, Davis JL, Beiser DG (MIT Press, Cambridge, MA), pp 249–270.
3. Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275(5306):1593–1599.
4. Morris G, Nevet A, Arkadir D, Vaadia E, Bergman H (2006) Midbrain dopamine neurons encode decisions for future action. *Nat Neurosci* 9(8):1057–1063.
5. Parker NF, et al. (2016) Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nat Neurosci* 19(6):845–854.
6. Schouppe N, Demanet J, Boehler CN, Ridderinkhof KR, Notebaert W (2014) The role of the striatum in effort-based decision-making in the absence of reward. *J Neurosci* 34(6):2148–2154.
7. Desrochers TM, Amemori K, Graybiel AM (2015) Habit learning by naive macaques is marked by response sharpening of striatal neurons representing the cost and outcome of acquired action sequences. *Neuron* 87(4):853–868.
8. Yamada H, et al. (2013) Coding of the long-term value of multiple future rewards in the primate striatum. *J Neurophysiol* 109(4):1140–1151.
9. Yanike M, Ferrera VP (2014) Representation of outcome risk and action in the anterior caudate nucleus. *J Neurosci* 34(9):3279–3290.
10. Goldberg JA, Bergman H (2011) Computational physiology of the neural networks of the primate globus pallidus: Function and dysfunction. *Neuroscience* 198(0):171–192.
11. Kita H (2007) Globus pallidus external segment. *Prog Brain Res* 160:111–133.
12. Turner RS, Anderson ME (1997) Pallidal discharge related to the kinematics of reaching movements in two dimensions. *J Neurophysiol* 77(3):1051–1074.
13. Georgopoulos AP, DeLong MR, Crutcher MD (1983) Relations between parameters of step-tracking movements and single cell discharge in the globus pallidus and subthalamic nucleus of the behaving monkey. *J Neurosci* 3(8):1586–1598.
14. DeLong MR (1971) Activity of pallidal neurons during movement. *J Neurophysiol* 34(3):414–427.
15. Arkadir D, Morris G, Vaadia E, Bergman H (2004) Independent coding of movement direction and reward prediction by single pallidal neurons. *J Neurosci* 24(45):10047–10056.
16. Morris G, Hershkovitz Y, Raz A, Nevet A, Bergman H (2005) Physiological studies of information processing in the normal and Parkinsonian basal ganglia: Pallidal activity in Go/No-Go task and following MPTP treatment. *Prog Brain Res* 147:285–293.
17. Brotchie P, Iansek R, Horne MK (1991) Motor function of the monkey globus pallidus. 2. Cognitive aspects of movement and phasic neuronal activity. *Brain* 114(Pt 4):1685–1702.
18. Joshua M, Adler A, Rosin B, Vaadia E, Bergman H (2009) Encoding of probabilistic rewarding and aversive events by pallidal and nigral neurons. *J Neurophysiol* 101(2):758–772.
19. McNab F, Klingberg T (2008) Prefrontal cortex and basal ganglia control access to working memory. *Nat Neurosci* 11(1):103–107.
20. Abdi A, et al. (2015) Prototypic and arkypallidal neurons in the dopamine-intact external globus pallidus. *J Neurosci* 35(17):6667–6688.
21. Mallet N, et al. (2012) Dichotomous organization of the external globus pallidus. *Neuron* 74(6):1075–1086.
22. Izquierdo A, Jentsch JD (2012) Reversal learning as a measure of impulsive and compulsive behavior in addictions. *Psychopharmacology (Berl)* 219(2):607–620.
23. Izquierdo A, Brigman JL, Radke AK, Rudebeck PH, Holmes A (March 12, 2016) The neural basis of reversal learning: An updated perspective. *Neuroscience* S0306-4522 (16)00244-X (abstr).
24. Joshua M, Elias S, Levine O, Bergman H (2007) Quantifying the isolation quality of extracellularly recorded action potentials. *J Neurosci Methods* 163(2):267–282.
25. Yoshida A, Tanaka M (2016) Two types of neurons in the primate globus pallidus external segment play distinct roles in antisaccade generation. *Cereb Cortex* 26(3):1187–1199.
26. Adler A, et al. (2012) Temporal convergence of dynamic cell assemblies in the striatopallidal network. *J Neurosci* 32(7):2473–2484.
27. Miller WC, DeLong MR (1987) Altered tonic activity of neurons in the globus pallidus and subthalamic nucleus in the primate MPTP model of Parkinsonism. *The Basal Ganglia II: Structure and Function—Current Concepts*, eds Carpenter MB, Jayaraman A (Springer, Boston), pp 415–427.
28. Starr PA, et al. (2005) Spontaneous pallidal neuronal activity in human dystonia: Comparison with Parkinson's disease and normal macaque. *J Neurophysiol* 93(6):3165–3176.
29. Williams ZM, Eskandar EN (2006) Selective enhancement of associative learning by microstimulation of the anterior caudate. *Nat Neurosci* 9(4):562–568.
30. Ashby FG, Turner BO, Horvitz JC (2010) Cortical and basal ganglia contributions to habit learning and automaticity. *Trends Cogn Sci* 14(5):208–215.
31. Rescorla RA, Wagner AW (1972) A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory*, eds Black AH, Prokasy WF (Appleton-Century-Crofts, New York), pp 64–99.
32. Aosaki T, et al. (1994) Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. *J Neurosci* 14(6):3969–3984.
33. Apicella P, Ravel S, Sardo P, Legallet E (1998) Influence of predictive information on responses of tonically active neurons in the monkey striatum. *J Neurophysiol* 80(6): 3341–3344.
34. Inokawa H, Yamada H, Matsumoto N, Muranishi M, Kimura M (2010) Juxtacellular labeling of tonically active neurons and phasically active neurons in the rat striatum. *Neuroscience* 168(2):395–404.
35. Reynolds JN, Hyland BI, Wickens JR (2004) Modulation of an afterhyperpolarization by the substantia nigra induces pauses in the tonic firing of striatal cholinergic interneurons. *J Neurosci* 24(44):9870–9877.
36. Cragg SJ (2006) Meaningful silences: How dopamine listens to the ACh pause. *Trends Neurosci* 29(3):125–131.
37. Alexander GE, Crutcher MD, DeLong MR (1990) Basal ganglia-thalamocortical circuits: Parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. *Prog Brain Res* 85:119–146.
38. Graybiel AM (1997) The basal ganglia and cognitive pattern generators. *Schizophr Bull* 23(3):459–469.
39. Barnes TD, Kubota Y, Hu D, Jin DZ, Graybiel AM (2005) Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* 437(7062):1158–1161.

40. Morris G, Arkadir D, Nevet A, Vaadia E, Bergman H (2004) Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron* 43(1):133–143.
41. Jahanshahi M, et al. (2000) The impact of deep brain stimulation on executive function in Parkinson's disease. *Brain* 123(Pt 6):1142–1154.
42. Smeding HM, et al. (2006) Neuropsychological effects of bilateral STN stimulation in Parkinson disease: A controlled study. *Neurology* 66(12):1830–1836.
43. Aron AR, Poldrack RA (2006) Cortical and subcortical contributions to Stop signal response inhibition: Role of the subthalamic nucleus. *J Neurosci* 26(9):2424–2433.
44. Lehéricy S, et al. (2005) Distinct basal ganglia territories are engaged in early and advanced motor sequence learning. *Proc Natl Acad Sci USA* 102(35):12566–12571.
45. Carpenter MB, Carleton SC, Keller JT, Conte P (1981) Connections of the subthalamic nucleus in the monkey. *Brain Res* 224(1):1–29.
46. Mwansisya TE, et al. (2013) The diminished interhemispheric connectivity correlates with negative symptoms and cognitive impairment in first-episode schizophrenia. *Schizophr Res* 150(1):144–150.
47. Early TS, Reiman EM, Raichle ME, Spitznagel EL (1987) Left globus pallidus abnormality in never-medicated patients with schizophrenia. *Proc Natl Acad Sci USA* 84(2):561–563.
48. Galeno R, Molina M, Guirao M, Isoardi R (2004) Severity of negative symptoms in schizophrenia correlated to hyperactivity of the left globus pallidus and the right claustrum. A PET study. *World J Biol Psychiatry* 5(1):20–25.
49. Aylward EH, et al. (1996) Basal ganglia volumes in children with attention-deficit hyperactivity disorder. *J Child Neurol* 11(2):112–115.
50. Griffiths KR, Lagopoulos J, Hermens DF, Hickie IB, Balleine BW (2015) Right external globus pallidus changes are associated with altered causal awareness in youth with depression. *Transl Psychiatry* 5:e653.
51. Baumann B, et al. (1999) Reduced volume of limbic system-affiliated basal ganglia in mood disorders: preliminary data from a postmortem study. *J Neuropsychiatry Clin Neurosci* 11(1):71–78.
52. National Research Council (2011) *Guide for the Care and Use of Laboratory Animals* (National Academies, Washington, DC), 8th Ed.
53. Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10(4):433–436.
54. Elias S, et al. (2007) Statistical properties of pauses of the high-frequency discharge neurons in the external segment of the globus pallidus. *J Neurosci* 27(10):2525–2538.