

Robust Representation of Stable Object Values in the Oculomotor Basal Ganglia

Masaharu Yasuda,¹ Shinya Yamamoto,^{1,2} and Okihide Hikosaka¹

¹Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, Maryland 20892-4435, and ²System Neuroscience Group, Human Technology Research Institute, National Institute of Advanced Industrial Science and Technology, Tsukuba, 305-8568, Japan

Our gaze tends to be directed to objects previously associated with rewards. Such object values change flexibly or remain stable. Here we present evidence that the monkey substantia nigra pars reticulata (SNr) in the basal ganglia represents stable, rather than flexible, object values. After across-day learning of object–reward association, SNr neurons gradually showed a response bias to surprisingly many visual objects: inhibition to high-valued objects and excitation to low-valued objects. Many of these neurons were shown to project to the ipsilateral superior colliculus. This neuronal bias remained intact even after >100 d without further learning. In parallel with the neuronal bias, the monkeys tended to look at high-valued objects. The neuronal and behavioral biases were present even if no value was associated during testing. These results suggest that SNr neurons bias the gaze toward objects that were consistently associated with high values in one's history.

Introduction

Consider the experience of being in a grocery store. You may be bewildered by so many items around you. But somehow you choose several items from among many similar items. You may find, in retrospect, that you chose them because you previously chose them and liked them. Your everyday choices may thus depend on your emotional experiences repeated in your history. Cognitive processes relying on working memories may not be able to handle such choices because there are so many objects around you (Shiffrin and Schneider, 1977). Information processing for such everyday choices would instead require object-value memories with high-capacity and long-term retention. The process of choosing requires high capacity because so many objects must be assessed and compared. Choosing also requires long-term retention because many objects under consideration are encountered only rarely. Indeed, animal behavioral studies have reported a high-capacity long-term memory for choosing high-valued visual objects or associated motor responses (Vaughan and Greene, 1984; Fagot and Cook, 2006).

We hypothesized that the substantia nigra pars reticulata (SNr), a major output of the basal ganglia, is a key structure for the everyday choice mechanism. Anatomical studies suggest that

the SNr may receive visual object information from the infero-temporal cortical areas through the tail of the caudate nucleus (CDt) (Saint-Cyr et al., 1990). Neurons in the CDt indeed respond to visual objects selectively (Caan et al., 1984; Brown et al., 1995; Yamamoto et al., 2012). Furthermore, the SNr has efficient access to motor outputs. The SNr controls the initiation of saccadic eye movements through its inhibitory connections to the superior colliculus (SC) (Hikosaka and Wurtz, 1983). Electrical stimulation of the CDt induces saccadic eye movements, which are likely mediated by the CDt's connection to the SNr (Yamamoto et al., 2012). Therefore, the SNr appears equipped with the mechanisms necessary for choosing valuable objects.

Our hypothesis is in line with a general concept that the basal ganglia play an important role in the selection of behaviors (Redgrave et al., 1999), particularly behaviors that lead to better reward outcomes (Hikosaka et al., 2006). However, most studies examined flexible value coding in which subjects were provided with a small number of options whose reward values were changed frequently (Hikosaka et al., 2006; Lau and Glimcher, 2008; Hori et al., 2009). In contrast, little is known about the neuronal mechanisms underlying the stable value coding of many objects, and this is what we wanted to study. To this end, we devised a set of experimental procedures in which the monkey saw many computer-generated visual objects, each associated with either a large or small reward consistently across many daily sessions, and then the neuronal activity as well as saccadic behavior was tested. We found that SNr neurons indeed encode the values of a remarkable number of visual objects and then retain the object value information for a long time with no further learning.

Materials and Methods

Animals. We used two male rhesus monkeys (*Macaca mulatta*), D and G. All animal care and experimental procedures were approved by the National Eye Institute Animal Care and Use Committee and complied with

Received July 18, 2012; revised Sept. 5, 2012; accepted Sept. 26, 2012.

Author contributions: M.Y., S.Y., and O.H. designed research; M.Y. performed research; M.Y. analyzed data; M.Y., S.Y., and O.H. wrote the paper.

This work was supported by the Intramural Research Program at the National Institutes of Health, National Eye Institute. We thank E. Bromberg-Martin, H.F. Kim, I. Monosov, S. Hong, and Y. Tachibana for discussions; and A.M. Nichols, A.V. Hays, B. Nagy, D. Parker, G. Tansey, J.W. McClurkin, L.P. Jensen, M.K. Smith, and T.W. Ruffner for technical assistance.

This article is freely available online through the *JNeurosci* Open Choice option.

Correspondence should be addressed to Masaharu Yasuda, Laboratory of Sensorimotor Research, National Eye Institute, National Institute of Health, 49 Convent Drive, Bldg. 49, Rm. 2C28, Bethesda, MD 20892-4435. E-mail: yasudama@mail.nih.gov.

DOI:10.1523/JNEUROSCI.3438-12.2012

Copyright © 2012 the authors 0270-6474/12/3216917-16\$15.00/0

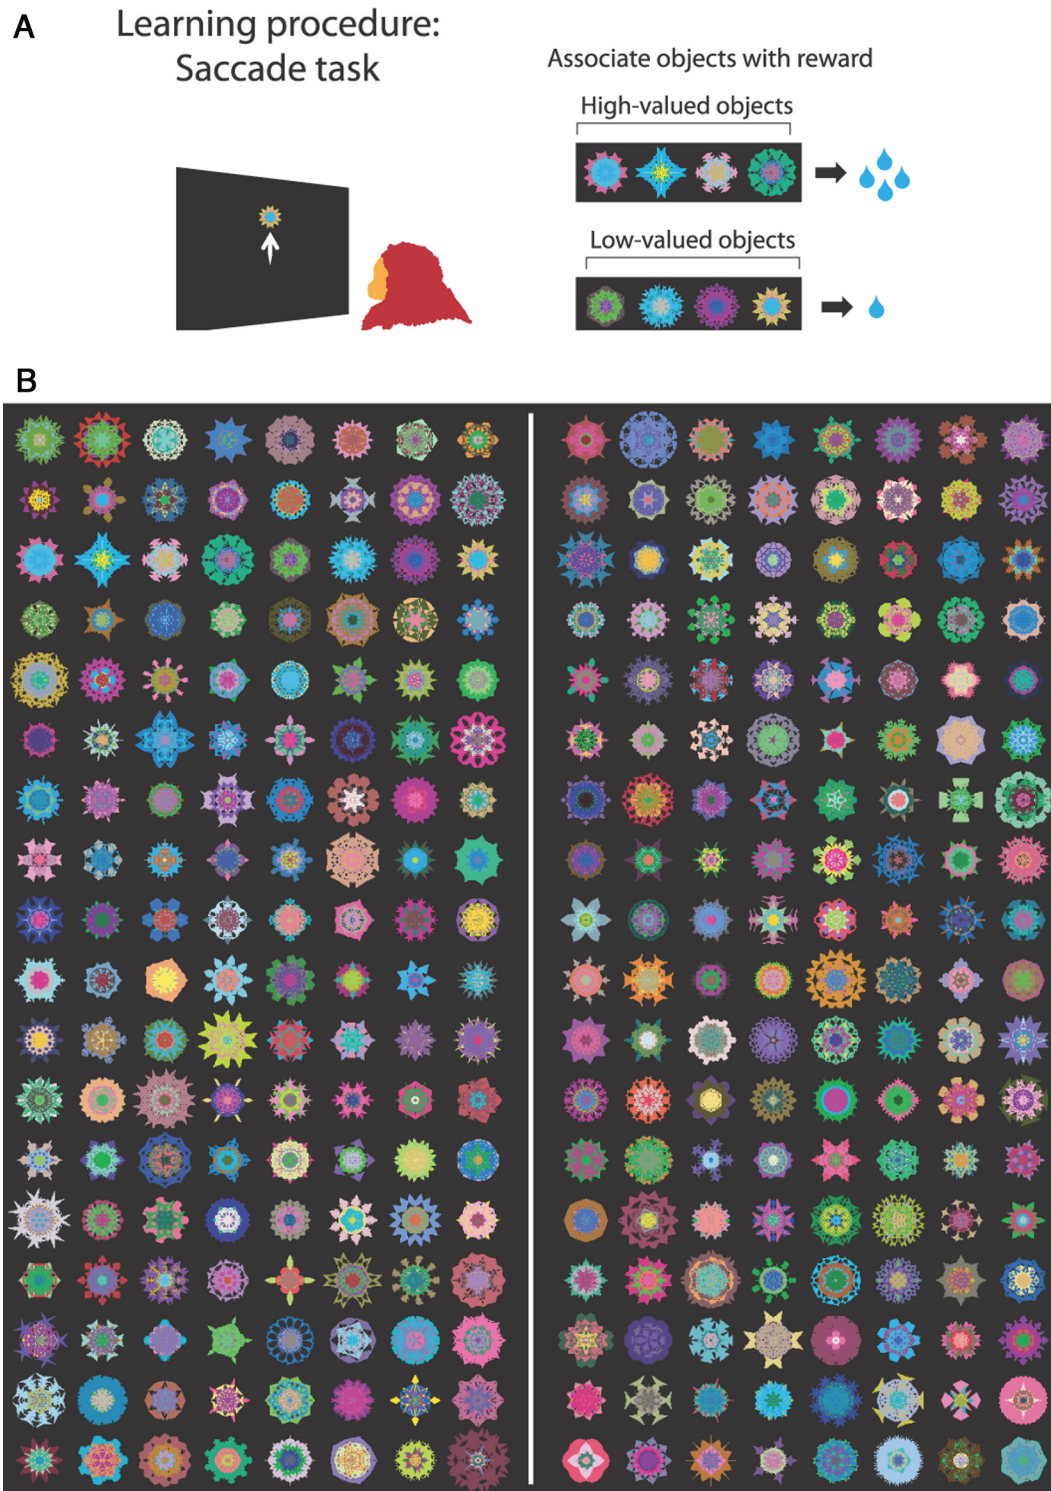


Figure 1. Procedure for learning stable object values. **A**, Stable object–reward association task. Among a set of eight fractal objects, four were assigned as “high-valued” objects and the other four were assigned as “low-valued” objects. On each trial, one of the eight objects was presented at one of four positions, and the monkey made a saccade to it. If the object was a high-valued object, a large reward was delivered. If the object was a low-valued object, a small reward was delivered. For each set of visual objects, the learning procedure was done repeatedly across many daily sessions, during which each fractal remained to be either a high-valued object or a low-valued object. **B**, Well learned fractals (≥ 5 learning sessions) used for monkey G. Each row of eight fractals (shown on each of left and right sides) were used as a set of objects. Among them, the left four fractals were high-valued objects (associated with a large reward) and the right four fractals were low-valued objects (associated with a small reward). Monkey G learned 648 fractals, among which 288 (shown here) were well learned. Monkey D learned 696 fractals, among which 280 were well learned. Many of the fractals were shared by both monkeys.

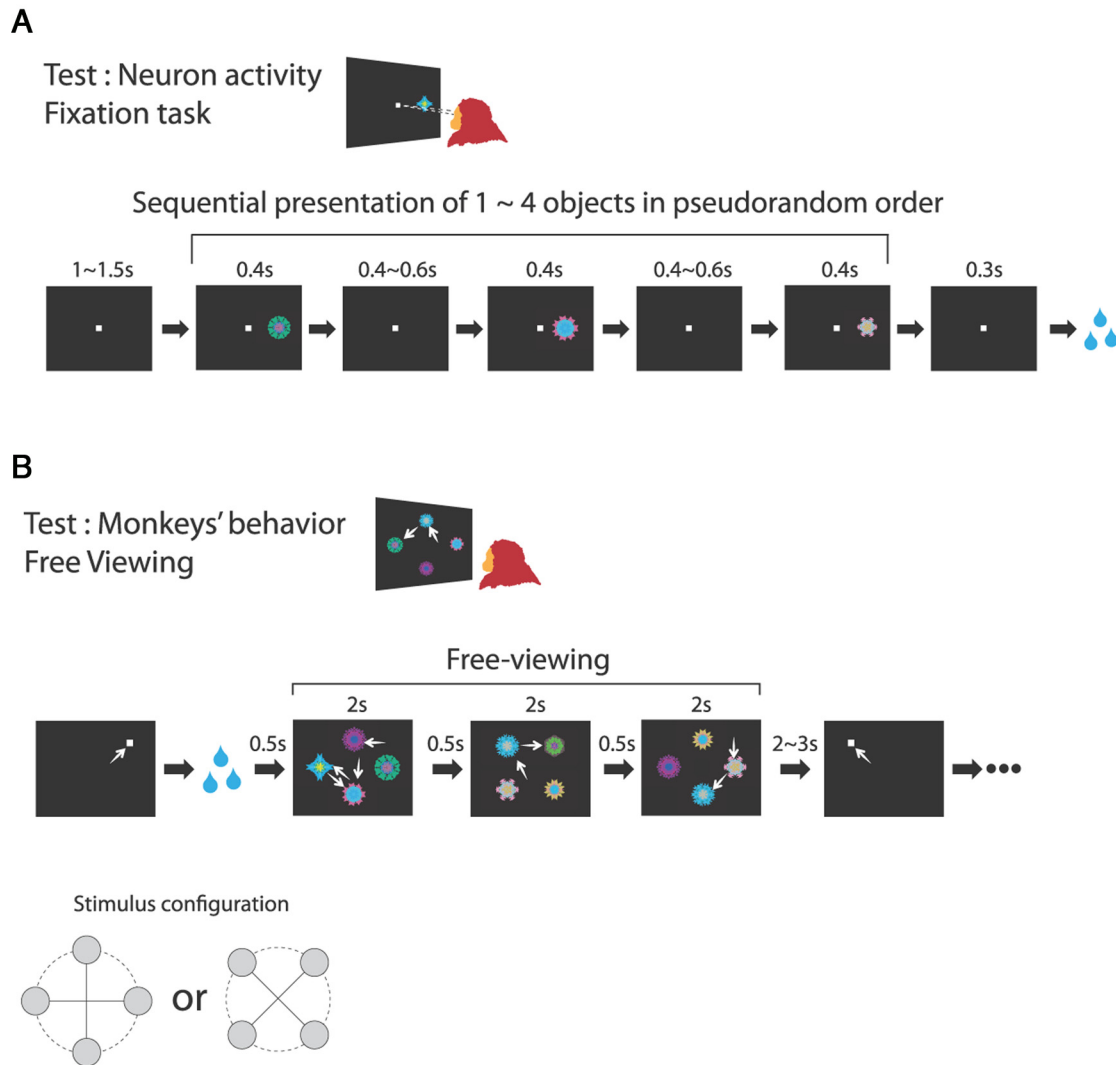


Figure 2. Procedures for testing the stable object value coding. **A**, Procedure for testing SNr neurons’ responses to high-valued objects and to low-valued objects. While the monkey was fixating a central spot of light, 1–4 fractal objects (pseudorandomly chosen from a set of 8 objects) were presented sequentially in the neuron’s preferred location. The monkey was rewarded 300 ms after the final object disappeared. The reward was thus not associated with particular objects. **B**, Free-viewing task for testing saccade behavior. In one session, a set of eight fractal objects were presented and the monkey was free to look at these objects (or somewhere else) by making saccades between them. For one presentation, four objects were chosen randomly from the eight objects and for the next presentation the other four objects were chosen, and this cycle was repeated. No reward was delivered during the free viewing. Occasionally, a white small dot was presented at one of four positions. If the monkey made a saccade to it, a reward was delivered to the monkey. The testing procedures were done on separate days from the learning procedure (Fig. 1) to avoid possible influences of working memory or short-term memory.

the Public Health Service Policy on the humane care and use of laboratory animals.

General behavioral procedure. Behavioral tasks were controlled by a QNX-based real-time experimentation data acquisition system [REX; Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health (LSR/NEI/NIH)]. The monkey sat in a primate chair, facing a frontoparallel screen 33 cm from the monkey’s eyes in a sound-attenuated and electrically shielded room. Stimuli generated by an active-matrix liquid crystal display projector (PJ550, ViewSonic) were rear-projected on the screen. We created visual stimuli using fractal geometry (Miyashita et al., 1991). One fractal was composed of four point-symmetrical polygons, which were overlaid around a common center such that smaller polygons were positioned more toward the front. The size of each fractal was $\sim 8^\circ \times 8^\circ$.

Our main aim was to study the responses of SNr neurons to visual objects when they are associated with different reward values. Specifically, we hypothesized that SNr neurons encode objects’ reward values differently depending on whether the values remain stable or change

flexibly. To examine the stable and flexible value coding, we devised two behavioral procedures, as shown below.

Stable object–value association procedure. To examine behavioral and neuronal encoding of stable object values, we conducted the learning procedure (see Fig. 1) and the testing procedure (see Fig. 2) separately on different days. In the learning procedure, the monkey experienced visual objects repeatedly in association with consistent reward values and thus learned their stable values (see Fig. 1). In the testing procedure, neuronal activity and monkey’s saccade behavior were examined using different tasks (see Fig. 2). The separation of the learning procedure and the testing procedure precluded possible influences of working memory or short-term memory. Also importantly, the testing procedure was done in a neutral condition: the monkey obtained no reward when learned objects were presented (in case of behavioral testing) or the monkey did obtain reward but not in association with particular objects (in case of neuronal testing). The “neutral” condition during testing was critical to examine the effect of long-term reward experiences (stable values) because, oth-

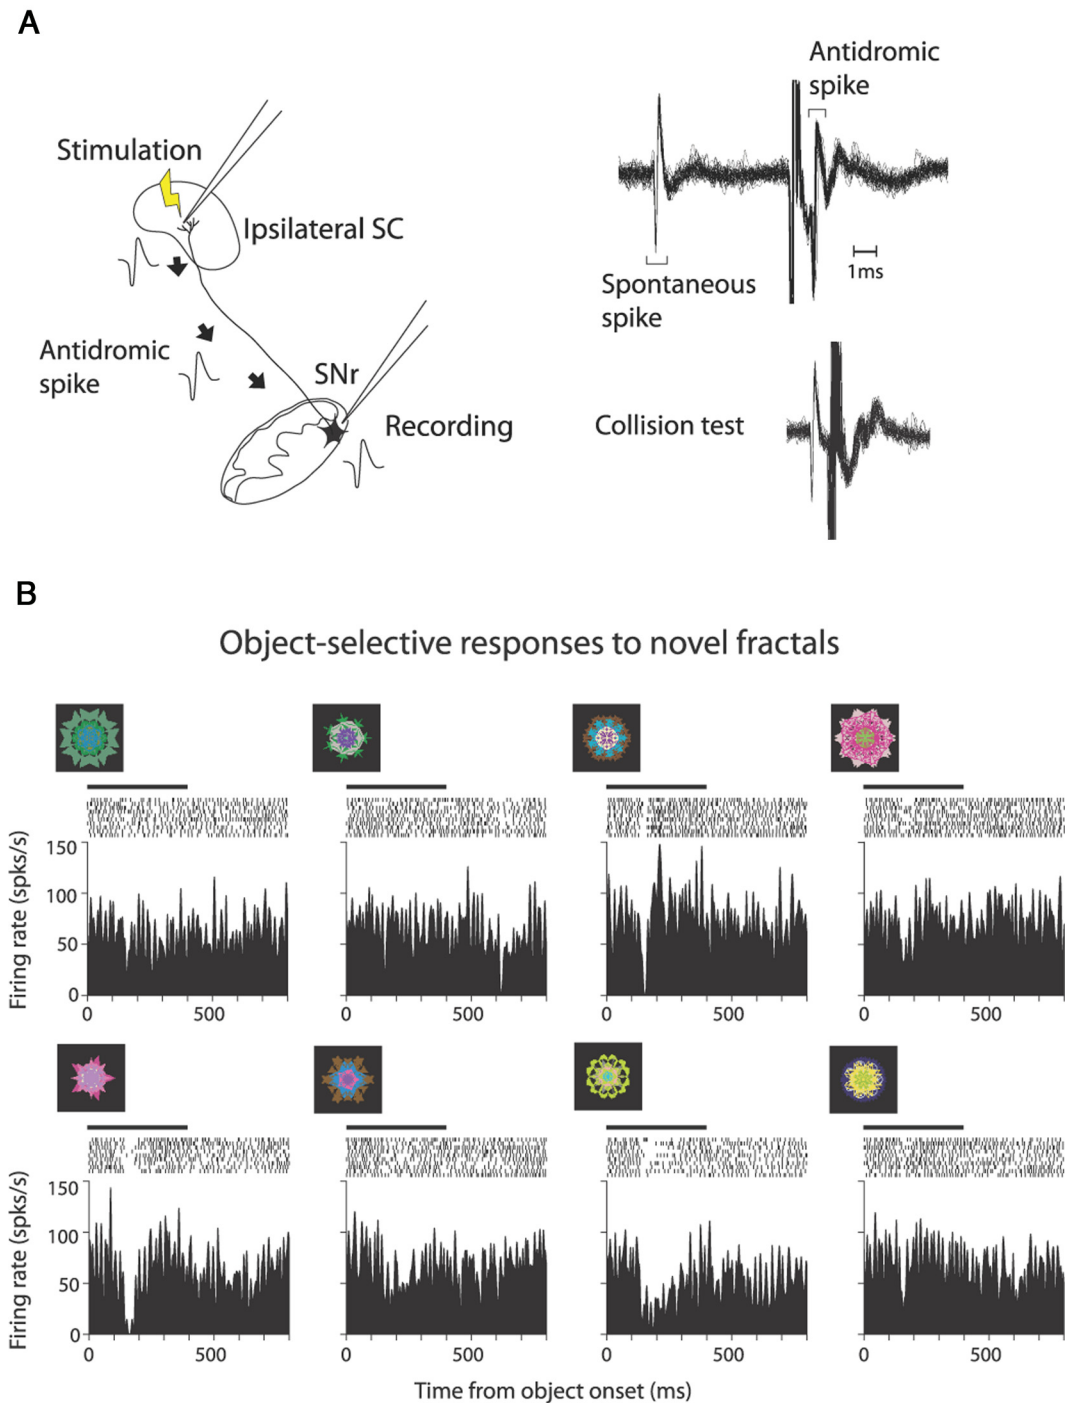


Figure 3. Object-selective responses of an SNr neuron to novel fractal objects. **A**, Antidromic activation of an SNr neuron to electrical stimulation of the ipsilateral SC. The antidromic nature of the spikes was confirmed by a collision test (right, bottom). **B**, Responses of the SNr neuron to novel fractal objects, shown by rasters of dots (top) and spike density functions (bottom). The objects were presented at the neuron’s preferred location while the monkey was fixating at the central fixation point (see Fig. 2A). The same neuron as shown in Figure 10. Top and bottom four objects were assigned as high-valued and low-valued objects in the later learning.

erwise, the effect of short-term reward experiences (flexible values) could change the monkey’s behavior or neuronal activity.

In the following we will explain in detail (1) the learning procedure, (2) the procedure for testing neuronal activity, and (3) the procedure for testing saccade behavior.

1. Procedure for learning stable object values (see Fig. 1). The goal of this procedure was to create a fixed bias among fractal objects in their reward values (i.e., high-valued objects and low-valued objects). For this purpose we used an object-directed saccade task (see

Fig. 1A). In each session, a set of eight fractal objects was used as the target. On each trial, one of the fractal objects was chosen pseudo-randomly as the target and was presented at one of four positions (right, up, left, and bottom). The monkey was required to make a saccade to the target to obtain a liquid reward. Importantly, half of the fractal objects were associated with a large reward (0.11 ml) (i.e., high-valued objects), whereas the other half were associated with a small reward (0.02 ml) (i.e., low-valued objects). One learning session consisted of 64 trials (8 trials for each object) for monkey D and

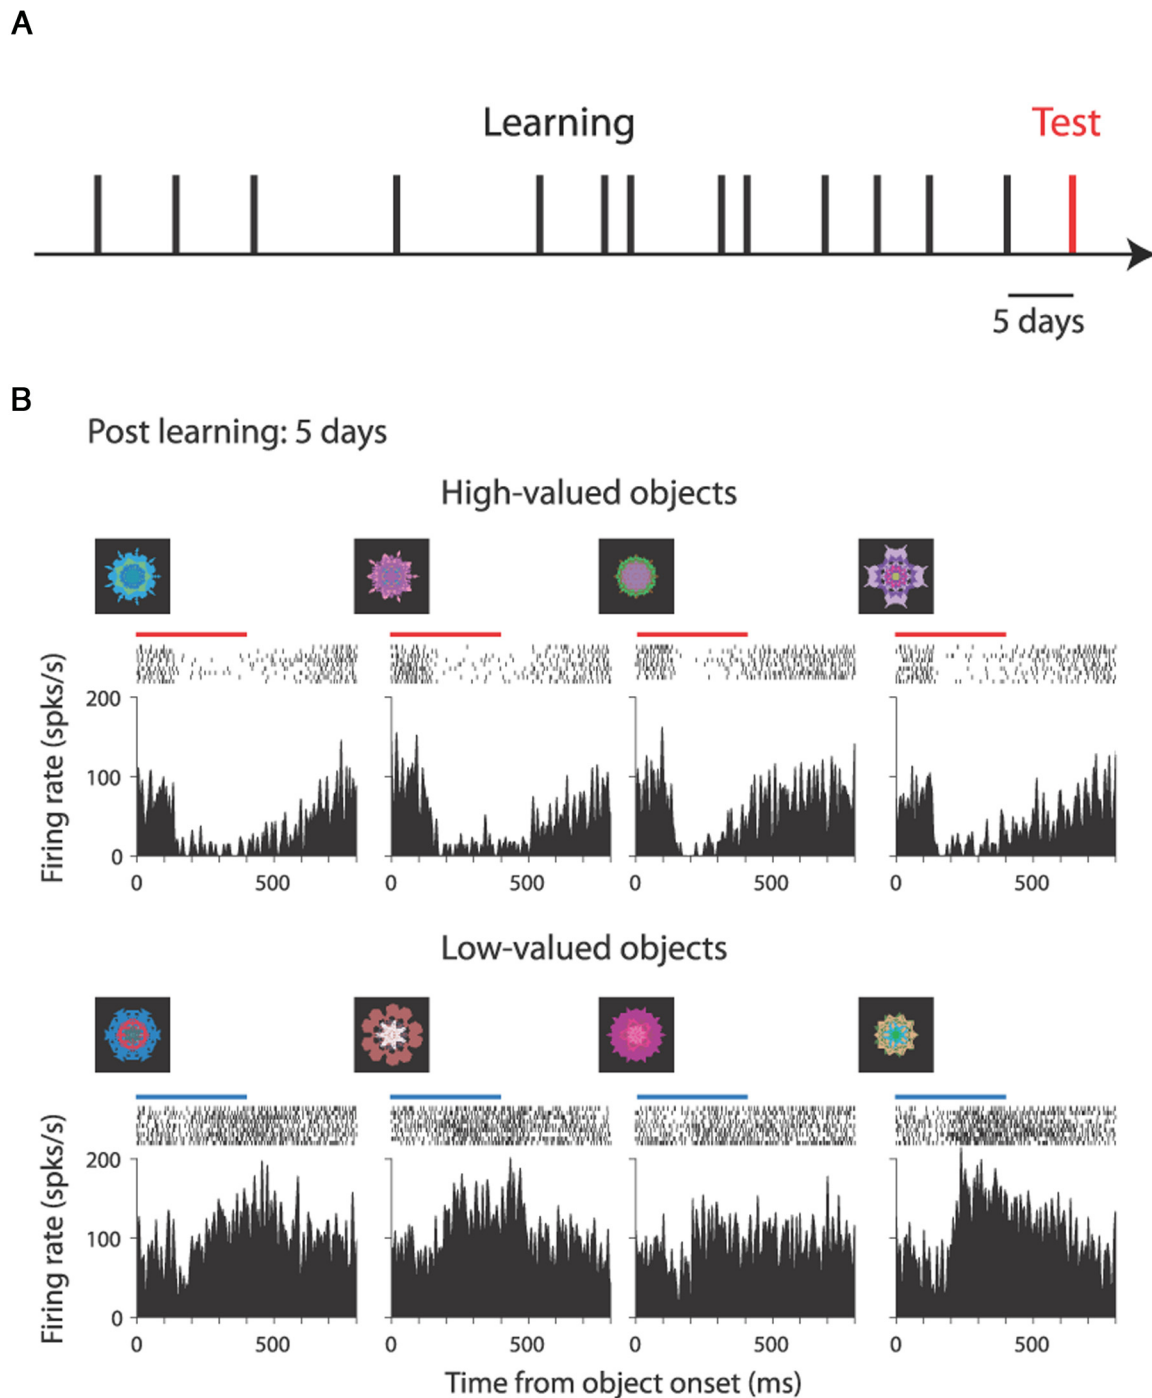


Figure 4. Stable object value coding of an SC-projecting SNr neuron. *A*, Learning schedule for a set of fractal objects used to test the responses of the SNr neuron (as shown in *B*). Monkey D learned the set of objects for 13 sessions across days (or months). Five days after the last learning session, the SNr neuron was recorded (red bar). One learning session was done in 1 d. *B*, The responses of the SNr neuron to the learned object set: high-valued objects (top) and low-valued objects (bottom). For each fractal, the activity is aligned on the onset of the fractal (time 0). The horizontal bar shown above each raster plots indicate the duration of object presentation (0–400 ms).

128 trials (16 trials for each object) for monkey G. On each learning day, monkey D and G learned 1–22 and 1–25 sets of objects, respectively (see Fig. 6 for monkey G). Each set was learned in one learning session in 1 d. The same sets of fractals were used repeatedly for learning across days (or months), throughout which each object remained to be either a high-valued object or a low-valued object.

2. Procedure for testing neuronal activity (see Fig. 2*A*). To examine the neuronal coding of stable object values, while excluding any effect of flexible value coding, we used a passive viewing task. In each session a set of eight fractal objects was used as the visual

stimuli. While the monkey was fixating on a central spot of light, the fractal objects were presented sequentially in the neuron’s preferred location in a pseudorandom order (presentation time: 400 ms, interobject time interval: 400–600 ms). After every 1–4 object presentation, a reward was delivered 300 ms later. The reward was thus not associated with particular objects. Each object was presented at least six times in one session.

3. Procedure for testing saccade behavior (see Fig. 2*B*). To examine the monkey’s memory of stable object values, while excluding any effect of flexible value coding, we used a free-viewing task. In fact,

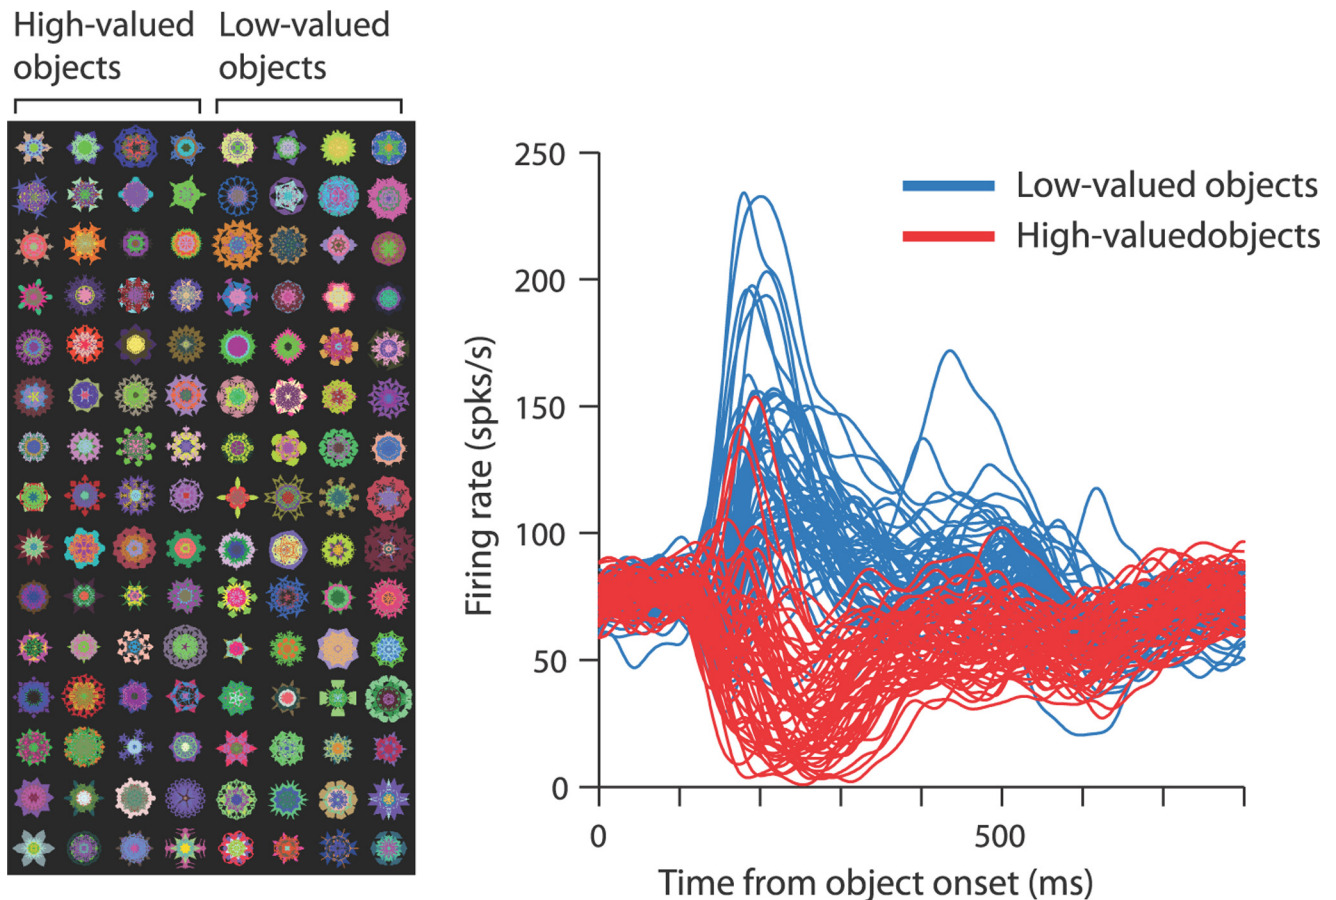


Figure 5. An SC-projecting SNr neuron encoding stable values of many objects. We tested the neuron's responses to 120 well learned fractal objects (≥ 5 learning sessions) (left). Shown superimposed on right are the neuron's responses to 60 high-valued objects (red) and 60 low-valued objects (blue). Firing rates (shown by spike density functions) are aligned on the onset of the object (time 0). The object disappeared at 400 ms. The number of learning sessions for these objects ranged from 5 to 22. The last learning of object–reward association had been done 3 d before the neuronal recording.

there was no rule for the monkey to follow while fractal objects were presented. On each trial, four of a set of eight fractal objects were chosen pseudorandomly and were presented simultaneously for 2 s. The monkey was free to look at these objects (or something else) by making saccade between them, but no reward was given. After a blank period (0.5 s) another four objects were presented. Occasionally, a white small dot instead was presented at one of four positions. If the monkey made a saccade to it and held the gaze on it for 300–600 ms, a reward was delivered. Each object was presented at least 25 times in each session. We also used a modified version, which was the same as the original task except that, before each fractal presentation, the white dot appeared at the center and the monkey was required to fixate it for 1000–1500 ms.

We arranged the learning procedure and the testing procedure in various ways to examine (1) the neuronal-behavioral changes after long-term learning, (2) the retention of the object–reward association information, and (3) the time course of the learning. To examine the neuronal-behavioral changes, we trained the monkey to learn multiple object sets for multiple days, and then tested the activity of SNr neurons and the monkey's free viewing. To examine the retention of the object–reward association information, we stopped the learning of particular object sets that had been learned, and then, after a no-learning period of days or months, tested SNr neurons' responses to the previously learned objects (see Fig. 8). To examine the time course of the learning, we added new sets of objects on different learning days such that, days or months later, the monkey had a repertoire of multiple sets of fractal objects with different degrees of learning (see Fig. 10A). This allowed us to test the

response of a single SNr neuron to objects with different degrees of object–reward association as well as the monkey's free viewing.

The actual schedule in stable object-value learning for monkey G is shown in Figure 6. The learning and testing sessions were often alternated, but in other cases each was performed consecutively across many days. As evident in this scheme, this experimental procedure required the monkeys to learn many fractal objects: 87 sets (696 objects) for monkey D and 81 sets (648 objects) for monkey G. The numbers of well learned object sets (≥ 5 learning sessions) were 35 sets (280 objects) for monkey D and 36 sets (288 objects) for monkey G. The 288 well learned objects for monkey G are shown in Figure 1B.

Flexible object–value association procedure. This procedure allowed us to examine behavioral and neuronal encoding of flexible object values as they were being updated in blocks of trials. Therefore, learning (of object values) and testing (of the monkey's behavior and of SNr neuron's activity) were done in one task procedure (object-directed saccade task), as illustrated in Figure 13A. For each monkey, a fixed set of two fractal objects (say, A and B) was used as the saccade target. Each trial started with the appearance of a central white spot, which the monkey had to fixate on. After 1000–1500 ms while the monkey was fixating on the central spot, one of the two fractal objects was chosen pseudorandomly and was presented at the neuron's preferred location. The fixation spot disappeared 400 ms later, and then the monkey was required to make a saccade to the object within 700 ms. If the gaze was held on the object for 400 ms, a liquid reward was delivered. The monkey received a large reward (0.20 ml) after making a saccade to one object (e.g., A) and received a small reward (0.016

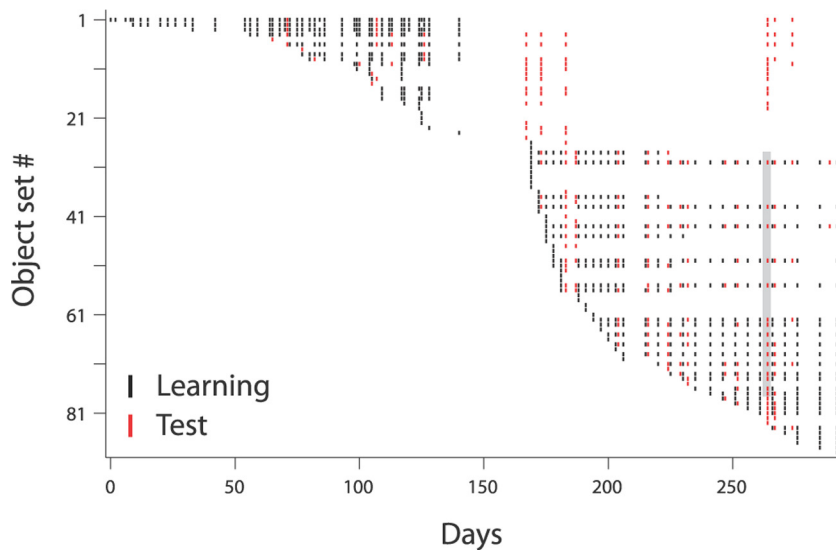


Figure 6. Actual experimental schedule for monkey G. The monkey experienced 712 fractal objects grouped into 89 sets (8 fractals each). Each row indicates the schedule for each object set (black, learning; red, testing of neuronal activity) across the days of the monkey's career in this study. Among the 89 object sets, 81 sets were used for object–reward association learning; eight sets were used only for testing. As the monkey learned many fractal objects, new object sets were added for object–reward association. Thus, at any point during the career, the monkey had learned many object sets in different degrees (i.e., different number of learning sessions). This means that the neuronal learning curve can be inferred from weak to strong response biases of one SNr neuron (recorded on a particular day) to the object sets with small to large amount of learning, as illustrated in Figure 10A. Many learned object sets were omitted from learning sessions and, after many days (up to 165 d) of no further learning, were used for testing SNr neurons' responses and the monkey's saccade behavior. The shaded area indicates the experiments shown in Figure 5.

ml) to the other object (e.g., B). During a block of 30 trials, the object–reward contingency was fixed, but it was reversed in the following block (e.g., B-high-valued, A-low-valued) without any external cue. While an SNr neuron was being recorded, these two blocks (A-high-valued/B-low-valued and B-high-valued/A-low-valued) were alternated in blocks (their order counterbalanced across neurons). Most trials (24 of 30 trials) were forced trials: one of the two objects was presented and the monkey had to make a saccade to it. The rest of trials (6 of 30 trials) were choice trials: two objects were presented at the same time. Their locations were chosen randomly from four possible locations (right, up, left, and bottom). The monkey had to choose one of the objects by making a saccade to it to obtain the reward associated with the chosen object. If the monkey failed to make a saccade correctly on either forced or choice trials, the same trial was repeated. In each recording session, these two types of block were repeated at least twice.

Electrophysiology. Based on a stereotaxic atlas (Saleem and Logothetis, 2007), we placed two rectangular chambers, one targeting the SNr and the other targeting the SC. The SNr chamber was placed over the frontoparietal cortex, tilted laterally by 35°. The SC chamber was placed over the midline of the parietal cortex, tilted posteriorly by 40°. MR images (4.7 T, Bruker) were then obtained along the direction of the recording chamber, which was visualized with gadolinium that filled grid holes and the space outside the grid and inside chamber. Single-unit recordings and electrical stimulations were performed using tungsten electrodes (Frederick Haer) advanced by an oil-driven micromanipulator (MO-97A, Narishige). The recording and stimulation sites were determined by using a grid system, which allowed recordings at every 1 mm between penetrations. The electrode was introduced into the brain through a stainless steel guide tube, which was inserted into one of the grid holes and then to the brain via the dura. The electrical signal from the electrode was amplified with a bandpass filter (200–5 kHz; BAK) and collected at 1 kHz. Spike potentials of single neurons were isolated online using a custom voltage-time window discrimination software (MEX, LSR/NEI/NIH).

Identification of SNr–SC neurons. We identified SC-projecting SNr neurons by antidromic activation (Hikosaka and Wurtz, 1983). In each experiment, we inserted a tungsten microelectrode in the SC. After pass-

ing through the surface of the SC, we lowered the electrode while testing neuron's visual and saccade-related response. When we found SC neurons that were predominantly activated in the presaccadic period, we fixed the electrode in that position for stimulation. We then lowered another electrode in the SNr while stimulating the SC until spikes with fixed latencies were detected. The antidromic nature of the spikes was confirmed using a collision test (Hikosaka and Wurtz, 1983). A biphasic pulse with cathodal and anodal components was used for the stimulation. The currents for the cathodal pulse ranged from 20 to 180 μ A; the currents for the anodal pulse were made lower than those for the cathodal pulse.

Data analysis. We analyzed the effects of object-value association learning on neuronal and behavioral discriminations of high-valued and low-valued objects. To assess the neuronal discrimination, we first measured the magnitude of the SNr neuron's response to each fractal object by counting the numbers of spikes within a test window in individual trials. For stable object-value learning, the test window was 100–400 ms after the onset of the object in the passive-viewing task. For flexible object-value learning, the test window was 100–300 ms after the onset of the object in the object-directed saccade task. The neuronal discrimination was defined as the area under the receiver operating characteristic (ROC) based on the response magnitudes of the SNr neurons to

high-valued objects versus low-valued objects. It was computed for each object set for stable object-value learning and for each SNr neuron for flexible object-value learning. The statistical significance of the neuronal discrimination was tested using Wilcoxon rank-sum test.

To assess the behavioral discrimination, we computed the choice rate for large-reward-associated objects, which was defined as follows: $nSACg/(nSACg + nSACb)$, where $nSACg$ and $nSACb$ are the numbers of saccades toward high-valued and low-valued objects, respectively. For stable object-value learning, the choice rate was computed by the monkey's saccade choice in the free-viewing task between four high-valued objects and the other four low-valued objects. For flexible object-value learning, the choice rate was computed by the monkey's saccade choice in the choice trials of the object-directed saccade task between the high-valued object (object that was recently associated with a large reward) and the low-valued object (object that was recently associated with a small reward).

To examine possible contributions of SNr neuronal activity to object-value learning, we first focused on the periods in which the monkey showed clear saccade biases (well learned periods): ≥ 5 daily learning sessions in stable object-value learning; fifth or later trials after the reversal of the object–reward contingency in flexible object-value learning. We then examined changes in the neuronal and behavioral discriminations during learning (learning curves). To this end, the neuronal and behavioral discriminations were computed individually after n th learning session in stable object-value learning (neuronal, see Fig. 10; behavioral, see Fig. 12) and for n th trials in flexible object-value learning (see Fig. 14G,H). For the neuronal learning in flexible object-value learning, we quantified the magnitude of SNr neuron's responses in each trial as the difference of firing rate between a test window (100–300 ms after object onset) and a baseline window (500–0 ms before object onset).

Results

SNr neurons robustly encode stable object values

To test the hypothesis that SNr neurons encode stable object values, we used the stable object–value association procedure. This procedure consisted of (1) object–reward association learn-

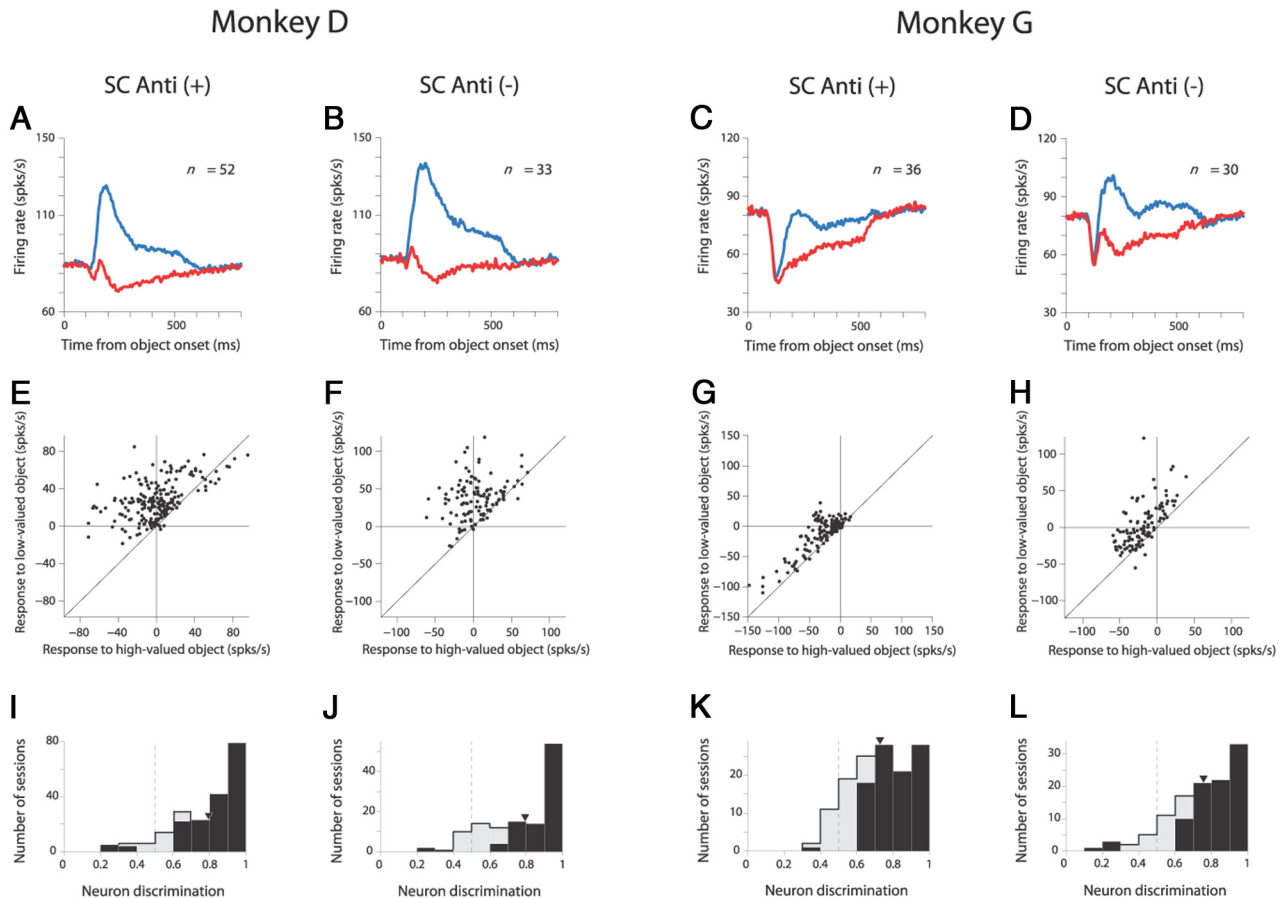


Figure 7. Stable object value coding of SNr neurons in monkey D (left) and G (right). They are shown separately for neurons antidromically activated from the SC (left) ($D: n = 52$; $G: n = 36$) and neurons not antidromically activated from the SC (right) ($D: n = 33$; $G: n = 30$). **A–D**, Averaged responses to high-valued objects (red) and low-valued objects (blue). The data were obtained for well learned sets (≥ 5 learning sessions). **E–H**, Comparison between the responses to high-valued objects and the responses to low-valued objects. Plotted for each recording session (each dot) are the average response to high-valued objects (abscissa) and the average response to low-valued objects (ordinate). The response was measured as the difference in firing rate between the test window (100–400 ms after object onset) and the baseline window (500–1000 ms before fixation point onset). **I–L**, Discrimination between high-valued and low-valued objects in individual recording sessions measured by ROC area based on spike counts in a test window. ROC 0.5 would mean no differentiation; ROC 1.0 would mean that the neuronal activity was always lower (or more inhibited) by high-valued objects than low-valued objects; ROC 0.0 would mean the opposite. Black bars indicate neurons with statistically significant discrimination assessed by Wilcoxon rank-sum test ($p < 0.05$). A triangle indicates the mean of the ROC areas: 0.79 for **I** and **J**, 0.73 for **K**, and 0.76 for **L**.

ing (Fig. 1A) and (2) neuronal and behavioral testing (Fig. 2). These subprocedures were done on separate days so that any effect of recently updated object values (i.e., flexible values) was excluded (see Materials and Methods).

For object–reward association learning, the monkeys were trained to learn many sets of fractal objects with different amounts of reward by using a saccade task (Fig. 1A). Among a set of eight fractal objects, four were always associated with a large reward (“high-valued objects”); the other four were always associated with a small reward (“low-valued objects”). Each object set was learned in one session in 1 d. Other object sets were added for learning on subsequent days. Thus, the monkey learned multiple object sets on most experimental days, each set being learned on consecutive days (see Fig. 6). Some of the objects monkey G learned are shown in Figure 1B.

To test neuronal responses, we recorded the spike activity of single SNr neurons using a passive viewing task (Fig. 2A). To evaluate only the effects of stably maintained object values while avoiding any effect of recently updated (i.e., flexible) object values, we used two specific procedures: (1) the neuronal testing was done at least 1 d after the last learning; (2) the reward given to the monkey during the neuronal (Fig. 2A) and behavioral (Fig. 2B)

testing was not associated with either high-valued or low-valued objects.

We recorded the spike activity of single SNr neurons, many of which projected to the SC (Fig. 3A). First we tested the responses of SNr neurons to novel fractal objects. The responses of an SC-projecting SNr neuron are shown in Figure 3B. The SNr neuron responded to the eight novel objects differentially, mostly with inhibitions in varying magnitudes and time courses.

We found that the object–reward association across many daily learning sessions induced differential responses in SNr neurons. We tested their responses to the learned fractal objects while the monkey was passively viewing them. A typical example is shown in Figure 4B. This SC-projecting SNr neuron was strongly inhibited by all of the high-valued objects (Fig. 4B, top); in contrast, it was mostly excited by the low-valued objects (Fig. 4B, bottom). At the time of neuronal testing, the monkey had completed 13 learning sessions for this set of fractals, but had not seen the fractals for 5 d before the recording of this SNr neuron (Fig. 4A).

Single SNr neurons differentiated many fractal objects into high-valued and low-valued objects, as exemplified in Figure 5. We tested the responses of this SC-projecting SNr neuron to 120

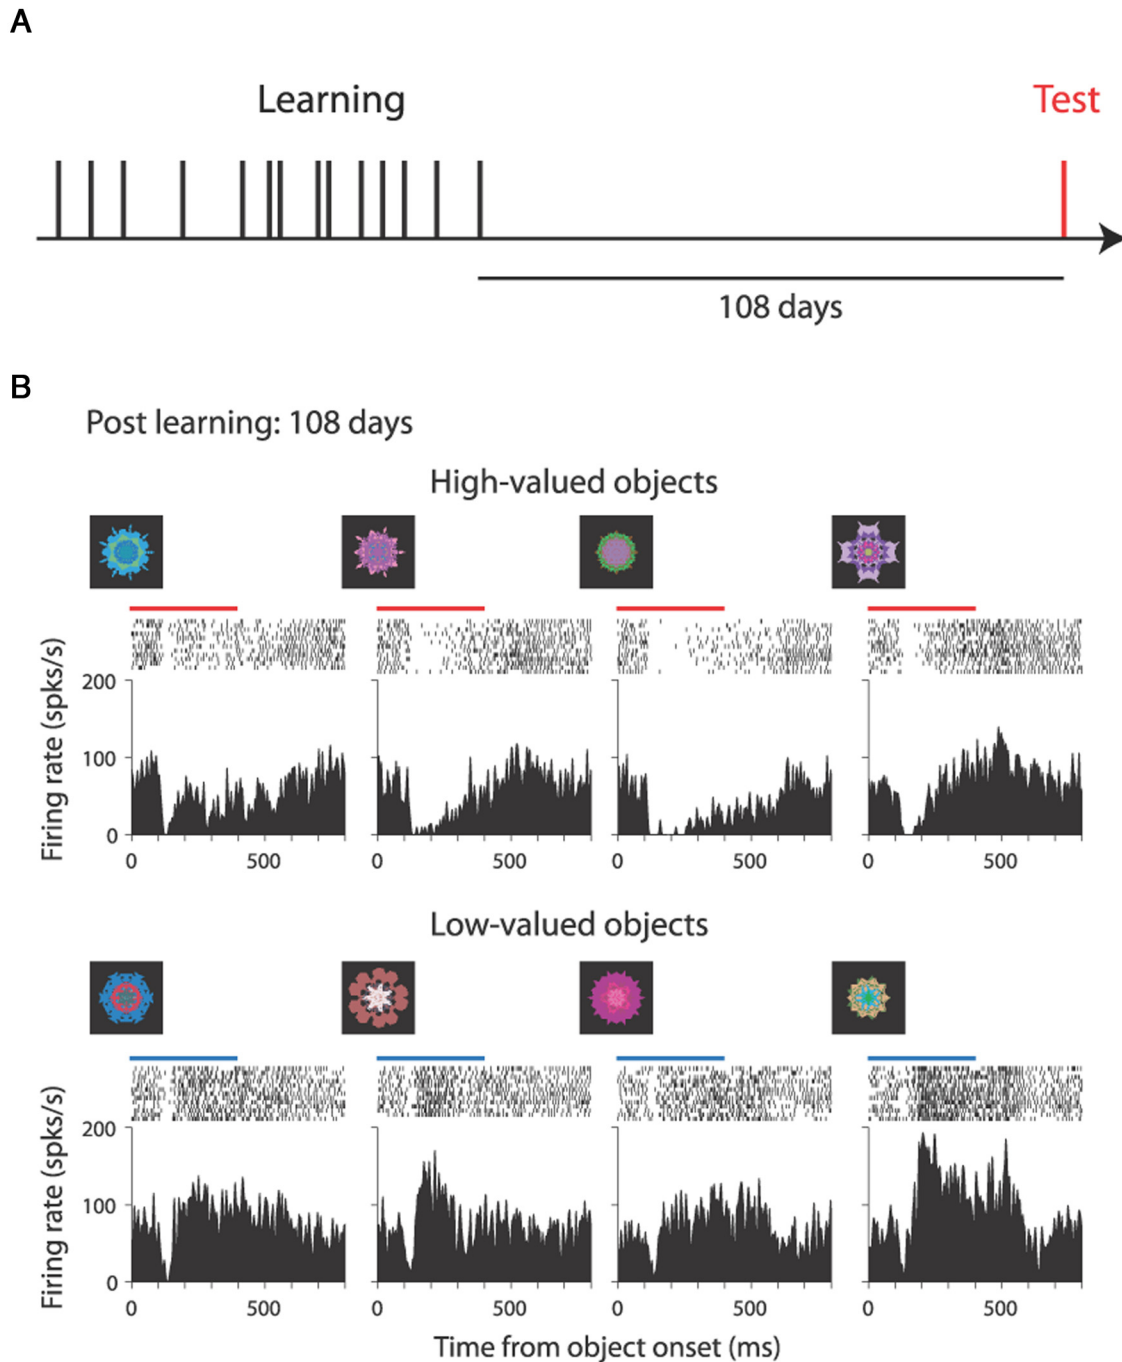


Figure 8. An SC-projecting SNr neuron retaining object values after a long period of no learning. *A*, Learning schedule for a set of fractal objects. Monkey D learned the set of objects for 14 sessions across days (black bars). Then the learning for this object set was stopped for 108 d, until the SNr neuron was tested (red bar). The same format as Figure 4 *A*. *B*, After the long retention period, the SNr neuron clearly discriminated between high-valued and low-valued objects. Note that during the retention period, the monkey learned many other fractals.

well learned fractal objects (≥ 5 learning sessions) (Fig. 5, left). The SNr neuron was mostly inhibited by the 60 high-valued objects (red lines) and mostly excited by the 60 low-valued objects (blue lines) (Fig. 5, right). The last learning sessions for these objects were done 3 d before the recording of this SNr neuron. Thus, the SNr neuron robustly encoded reward values of many objects after they were associated with the same reward values stably across days. In other words, the responses of the SNr neuron represented the stable memories of object values. For each SNr neuron, we tested 8–184 well learned objects (mean, 32 objects), but did not see a clear memory limit.

Throughout the whole experimental period (10 months for both monkeys), the monkey learned many fractal objects: 696 objects (87 sets) for monkey D and 648 objects (81 sets) for monkey G. Among them, we tested neural responses to 216 and 240 well learned objects (≥ 5 learning sessions) for monkey D and G (Fig. 1 *B*), respectively. The actual learning and recording schedule for monkey G is shown in Figure 6. The experiments shown in Figure 5 is indicated by a shaded area in which 15 sets of fractal objects (120 objects) were presented while the SNr neuron was recorded.

We examined the responses of 151 SNr neurons in two monkeys to the well learned objects. Among them, 93 neurons were

shown to project to the ipsilateral SC. As a population, they were clearly inhibited by high-valued objects and excited by low-valued objects (Fig. 7*A–D*). The responses to high-valued and low-valued objects are plotted for individual sessions in Figure 7*E–H*. Most data points are located above the unity line, indicating that the firing rates were lower in response to high-valued objects than to low-valued objects. To quantify neuron's discriminability between high-valued and low-valued objects, we calculated the ROC area for each session. In a large majority of the recording sessions, the ROC area was larger than 0.5 and its distribution was heavily skewed toward 1.0 (Fig. 7*I–L*). In 130 of 151 neurons (86.1%), the discrimination was statistically significant (Wilcoxon rank-sum test, $p < 0.05$), indicating that a majority of SNr neurons reliably categorized visual objects based on long-term reward experiences. These results were consistent across monkeys and across the two groups of SNr neurons with and without antidromic activation from the SC. Thus, the information on stable object values is conveyed to the SC by SNr neurons that were antidromically activated from the SC, and possibly to other brain areas by SNr neurons that were not antidromically activated from the SC.

SNr neurons retain object-value information for a long time

Our results have suggested that high-capacity memories of object values are represented in SNr. However, for the brain to have high-capacity memories, the values of individual objects need to be retained for a long time; otherwise, memories would not accumulate. To test this hypothesis, we stopped the learning of the monkey on some of the learned sets, and after a long delay (108–165 d) tested the object-value information in SNr neurons (Fig. 6, upper one-third). An example is shown in Figure 8. Before recording from this neuron, the monkey had learned the object set for 14 sessions and then learning was stopped for 108 d (Fig. 8*A*). During this period, the monkey never saw the objects but continued to learn many other objects. This SC-projecting SNr neuron was clearly inhibited by the high-valued objects (previously associated with a large reward) (Fig. 8*B* top) and excited by the low-valued objects (previously associated with a small reward) (Fig. 8*B*, bottom).

A majority of SNr neurons tested after the long retention period (>100 d) showed similarly clear reward biases (Fig. 9, right), which were comparable with the biases obtained after short retention periods (<20 d) (Fig. 9, left). In particular, the mean ROC area after long-retention periods (0.74) was not significantly

< 20 Days after
the last learning session

>100 Days after
the last learning session

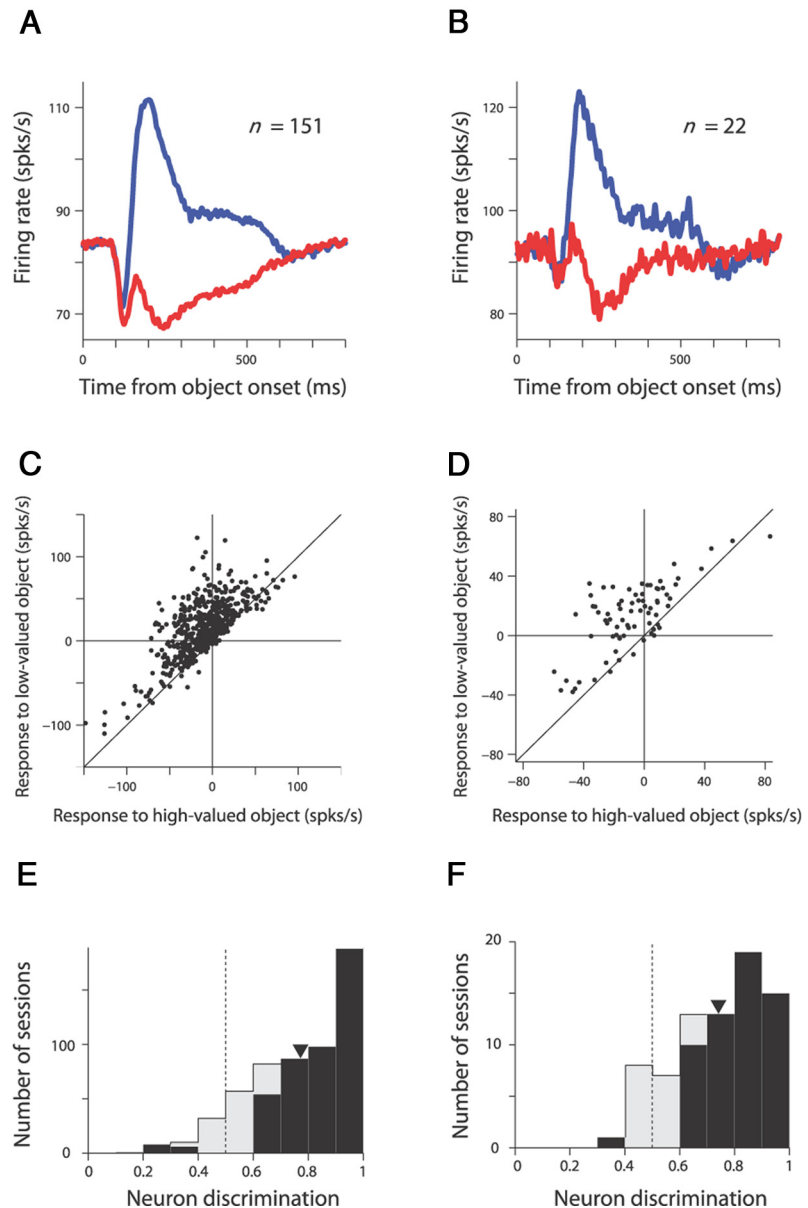


Figure 9. Long-term retention of object values by SNr neurons. Comparison between short retention periods (<20 d, left) and long retention periods (>100 d, right). The data were obtained from 151 and 22 SNr neurons for short and long retention periods, respectively, in two monkeys for well learned sets (≥ 5 learning sessions). The same format as in Figure 7. *A, B*, Averaged responses to high-valued objects (red) and low-valued objects (blue). *C, D*, Comparison between the responses to high-valued objects and the responses to low-valued objects. *E, F*, Discrimination between high-valued and low-valued objects in individual recording sessions measured by ROC area. Black bars indicate neurons with statistically significant discrimination assessed by Wilcoxon rank-sum test ($p < 0.05$). A triangle indicates the mean of the ROC areas: 0.77 for short retention periods and 0.74 for long retention periods. There was no statistical difference in the mean ROC area between the two groups of data (t test, $p = 0.18$). Data in *C–F* are based on each recording session (short retention period, $n = 563$; long retention period, $n = 76$).

cantly different from that after short retention periods (0.77) (t test, $p = 0.2$). These results indicate that SNr neurons, as a population as well as individually, retained the information of object values >100 d with little degradation. Notably, the object-value information in SNr neurons was not disrupted even though, during the retention period, the monkeys learned many other objects and neuronal responses were tested repeatedly without reward feedback (Fig. 6).

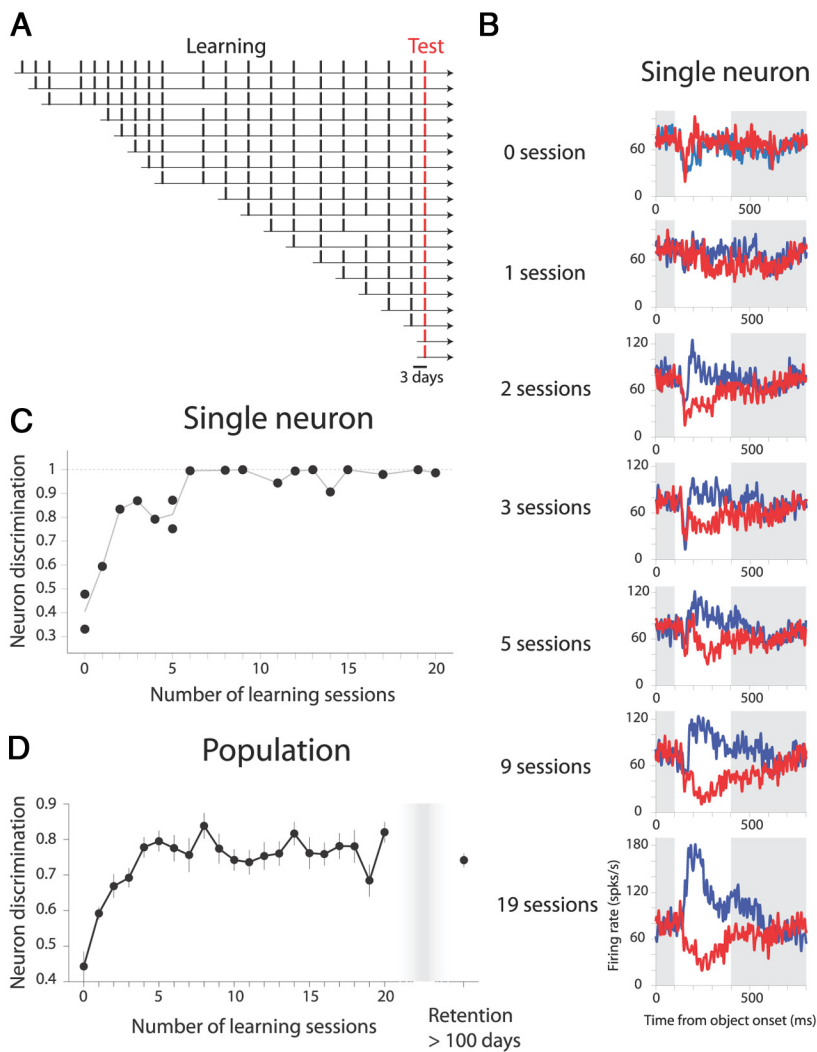


Figure 10. SNr neurons learn stable object values gradually. **A**, Learning schedule for 19 sets of fractal objects used for testing the responses of an SNr neuron (as shown in **B** and **C**). Note that multiple sets of fractals were added for learning at different points in the monkey's career, so that when an SNr neuron was recorded, the amount of learning varied across the fractal objects. **B**, The responses of an SNr neuron to fractal object sets with seven different degrees of learning (0–19 sessions). The responses are shown for each set, separately for high-valued objects (red) and low-valued objects (blue). **C**, Learning curve of the neuron shown in **B**. The neuron's discrimination between high-valued and low-valued objects for each object set (expressed as an ROC area) is plotted against the number of learning sessions for the object set. Data points are connected by lines to illustrate the neuron's learning curve. For sessions 0 and 5, we used two object sets, and the mean ROC area is used for the learning curve. **D**, Averaged learning curve across 118 SNr neurons in two monkeys. Data point on right indicates the average of neuronal biases that were obtained >100 d after the last learning. Error bars, ± 1 SEM.

SNr neurons learn object values slowly

Our results so far have shown that SNr neurons encode object values after long-term learning (≥ 5 learning sessions). We then asked how long it takes for SNr neurons to acquire the object value. To answer this question, we introduced many sets of novel fractals on different learning days, while the monkey continued to learn other object–reward associations (Fig. 10A). On the day of testing, we recorded from single SNr neurons and examined their responses to the multiple object sets with different degrees of learning.

The responses of a single SNr neuron to the multiple object sets (Fig. 10A) are shown in Figure 10B. In responses to a new object set (Fig. 10B, 0 session), the neuron showed no differential responses (as there had been no learning). The neuron showed a hint of differential responses to another object set that had been learned once (Fig. 4B, 1 session). The difference became larger

when the tested set had been learned in more sessions (Fig. 10B, 2–19 sessions). We used 19 object sets with different degrees of learning, and the neuron's learning curve is illustrated as changes in the ROC area (Fig. 10C). We examined 118 SNr neurons in this manner (Fig. 11). For most of the neurons, the ROC area tended to be larger for experienced in more learning sessions object sets. The change in the averaged ROC areas (Fig. 10D) suggests that the object-value information in SNr neurons developed gradually in the first several daily sessions before reaching a plateau level.

Since many of the SNr neurons projected to the SC, the differential responses of SNr neurons are likely to lead to differential responses of SC neurons. Specifically, because the SNr–SC connection acts as an inhibitory gate (Hikosaka et al., 2000), the inhibitory responses of SNr neurons to high-valued objects should open the gate to facilitate saccades to high-valued objects. To test this hypothesis, we examined the monkey's saccadic eye movements in a free-viewing condition while learned objects were presented (Fig. 2B). We found that, after the object–reward association learning, the monkey made saccades more frequently to high-valued objects than to low-valued objects (Fig. 12A). Notably, the preferential looking toward high-valued objects developed gradually, and its average time course was very similar to the development of the average response bias of SNr neurons (compare Figs. 10D and 12B, $r = 0.89$).

Note that both SNr neuronal responses in passive-viewing task and saccades in the free-viewing condition are likely automatic rather than goal-directed because rewards were not associated with particular objects during the two tasks. These results suggest that SNr is involved in preferential looking toward valuable objects based on long-term learning and

that the SNr–SC inhibition may underlie the reward-based, but nongoal-directed, gaze bias. Interestingly, after the long retention period, the object-value information in the SNr remained intact (Fig. 10D, data point on right), but the saccade bias became weaker (Fig. 12B, data point on right). This suggests that, even though the neuronal memories were retained, the monkey's gaze was influenced by unknown factors that developed during the retention period.

SNr neurons are weakly influenced by flexible object values

The results so far have indicated that, as SNr neurons slowly acquired discriminative responses to many visual objects depending on their stable values, the monkey also acquired discriminative gaze responses to these objects. The remarkable neuronal–behavioral correlation might be explained if the activity of SNr neurons simply reflected saccadic motor outputs. To

test this possibility, we examined whether the SNr neurons also changed their responses flexibly as the monkey changed its gaze preference frequently based on working memory of object values.

Figure 13A shows the flexible object–value association procedure in which two fractal objects alternated their values in blocks of trials. On each trial, one of the two objects was presented at the neuron’s preferred location, and the monkey made a saccade to it. One object was associated with a large reward and the other object was associated with a small reward. The reward amounts for these objects were fixed within a block of 30 trials, but were reversed in the following block. These two blocks were repeated at least twice in each recording session.

To test the monkey’s preference, we included a choice trial in which the two fractal objects were presented simultaneously and the monkey was required to choose one by making a saccade to it. The reward amount was determined by the chosen object. As shown in Figure 13B, the monkey typically chose the object that was recently associated with a large reward. The choice tendency developed quickly in several trials after the reversal of the object–reward contingency. These results indicate that the monkeys flexibly changed their preference as the objects changed their values.

The responses of a single SC-projecting SNr neuron are shown in Figure 13C. The neuron was inhibited by both fractal objects, but with different magnitudes. The responses showed no robust changes across blocks of trials regardless of whether the object was high-valued (i.e., recently associated with a large reward) (indicated by red bars) or low-valued (i.e., recently associated with a small reward) (blue bars). The average firing rates (shown at bottom) showed that the neuron was slightly more inhibited by the high-valued object (red) than by the low-valued object (blue) in the later part of object presentation (i.e., after 200 ms). Notably, these changes were much smaller than the differential responses of the same SNr neuron to the objects that had been assigned stable values (Fig. 5).

The averaged activity of all SNr neurons tested also showed slightly stronger inhibitory responses on large-rewarded than small-rewarded trials (Fig. 14A,B). To quantify individual neuron’s discriminability, we calculated ROC area based on the responses between large-rewarded and small-rewarded trials (Fig. 14C,E). The average ROC areas were significantly larger than 0.5 (monkey D: mean, 0.58; t test, $p = 1.43 \times 10^{-5}$; monkey G: mean, 0.54; t test, $p = 0.55 \times 10^{-2}$). A small number of SNr neurons (monkey D, 11 of 39; monkey G, 6 of 37) showed value differential responses (shown in black). Overall, the differential responses of SNr neurons in this condition (i.e., flexibly changing values) were much smaller than in the stable value condition in both monkeys. This was true even though the monkeys showed strong reaction time bias (Fig. 14D,F, paired t test, $p \approx 0$ in both monkeys). The discrepancy between SNr neuron activity and saccade behavior is also indicated by monkey D’s gaze preferences to the high-valued object. Every time the values of the objects were

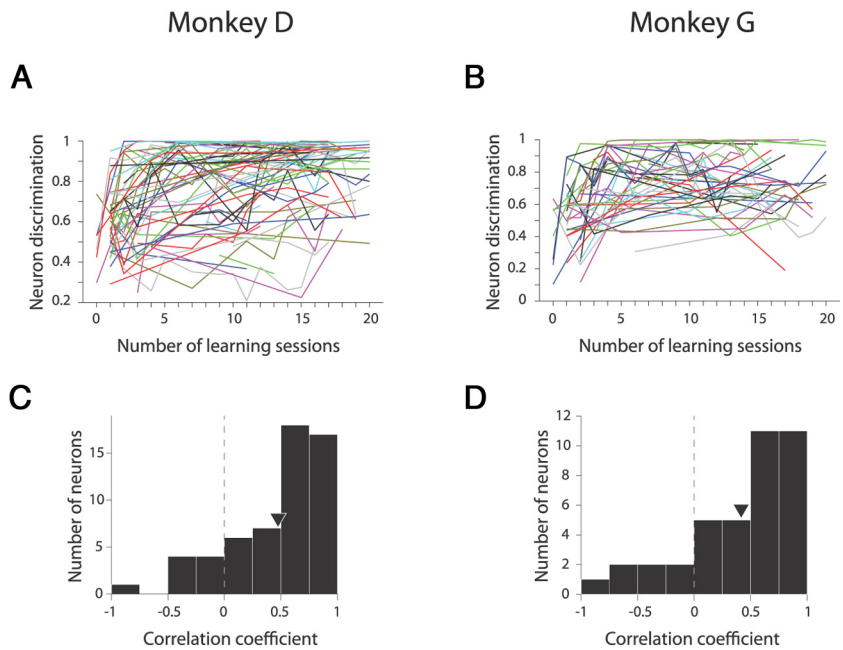


Figure 11. Learning curves for individual SNr neurons in monkey D (left) and monkey G (right). *A, B*, For each neuron, its discrimination between high-valued and low-valued objects (expressed as an ROC area) is plotted against the number of learning sessions for all sets of fractals tested, and the data points thus obtained are connected with lines in a unique color (which reflects the neuron’s learning curve). An example is shown in Figure 10C. *C, D*, Distribution of the linear correlation coefficients between the ROC area and the number of learning sessions obtained for individual neurons. The distribution was shifted toward positive values, indicating that each neuron’s discrimination between high-valued and low-valued objects became stronger with object–reward association learning. The mean correlation coefficients were 0.48 for monkey D and 0.41 for monkey G (triangles). The neuronal data obtained by <3 sets were excluded.

reversed, the monkey quickly learned to choose the high-valued object across trials (Fig. 14G, black curve). SNr neurons also showed a sign of learning (Fig. 14G, red curve), but much more slowly than the monkey’s choice rate. For example, between the third and seventh trials, the monkey’s choice was already biased toward the high-valued object (t test, $p \approx 7.99 \times 10^{-5}$), yet the SNr neurons’ activity still showed a weak bias toward the other object that had been the high-valued object in the preceding block of trials (t test, $p = 0.03$). Monkey G also showed a weak but consistent choice bias toward the high-valued object after 15th trial, but the neuronal bias was unclear. These results indicate that SNr neurons learned object values slowly and therefore were unable to keep track of flexibly changing values.

In summary, it is unlikely that the activity of SNr neurons simply reflect saccadic motor outputs. The monkey prefers to look at high-valued objects whether their values have been stably maintained or flexibly changed, but SNr neurons would contribute to the preferential looking based on stable object values, but not based on flexible object values.

Discussion

We found that SNr neurons robustly encoded the stable values of many visual objects. Most SNr neurons, including many SC-projecting neurons, showed opposite responses (i.e., inhibitory and excitatory) to high-valued objects (i.e., previously associated with large rewards) and low-valued objects (i.e., previously associated with small rewards). Since the SNr–SC connection is inhibitory (Hikosaka and Wurtz, 1983), the SNr neurons’ inhibitory responses to high-valued objects would lead to disinhibitions of saccadic neurons in the SC, thus leading to more saccades to high-valued objects. In contrast, the SNr neurons’ excitatory responses to low-valued objects would lead to enhanced inhibitions

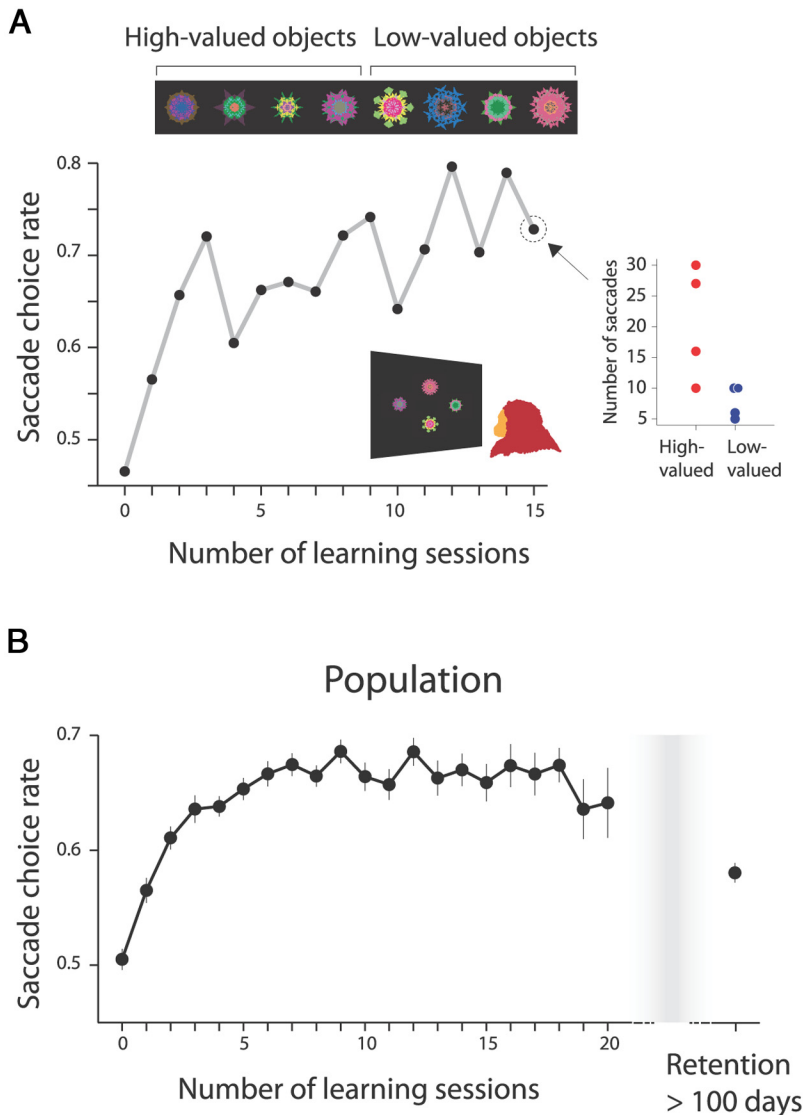


Figure 12. Monkeys learned stable object values gradually. **A**, Example behavioral learning curve. The monkey’s preference among a set of eight objects (shown above) is plotted against the number of learning sessions. The monkey’s preference was tested in a free-viewing task (Fig. 2B). The testing was done at least 1 d after the last learning. Inset graph shows, as an example, the total numbers of saccades to four high-valued objects (red) and four low-valued objects (blue) in the free-viewing task. **B**, The averaged behavioral learning curve across 128 sets of fractals in two monkeys. Data point on right indicates the average of behavior biases obtained > 100 d after the last learning. Error bars, ± 1 SEM.

of saccadic neurons in the SC, thus leading to fewer saccades to low-valued objects. Indeed, the monkeys developed a strong bias in gaze toward high-valued objects versus low-valued objects in the free-viewing condition. Moreover, the monkeys’ gaze bias toward high-valued objects developed slowly across daily learning sessions and its time course was very similar to the development of the SNr neurons’ discriminative responses to high-valued versus low-valued objects. In contrast, when objects changed their values frequently, the responses of SNr neurons to the objects showed only minor changes, although the monkeys changed their gaze bias flexibly. Our findings thus suggest that SNr neurons may be predominantly involved in the choice of objects based on their stable values, rather than flexible values.

How might the choices based on stable values be useful? When an animal forages in the forest, it may encounter many objects, but would choose appetizing objects and avoid bitter, nauseating, or inedible objects. Such choices need to be prompt and efficient.

The animal may not thrive well if it makes trial-and-error tasting or ponders too much. Rather, efficient choices are likely based on the long-term memories of these objects, as suggested in behavioral studies in humans (Standing, 1973; Brady et al., 2008) and animals (Vaughan and Greene, 1984; Fagot and Cook, 2006). The object-value memories represented by SNr neurons would be ideal for such prompt and efficient choices. SNr neurons encoded the values of a remarkable number of objects. This would be essential because the animal experiences so many objects throughout its life. SNr neurons categorized objects into high-valued and low-valued ones even when no action or no reward feedback was expected, as if the information were processed automatically. This would be suitable because the animal may encounter valuable objects unexpectedly. Furthermore, the SNr neurons’ discriminative responses were retained for a long time. This would be beneficial because the values of many objects may remain unchanged.

We predicted that SNr neurons would influence the choice of gaze, since many of the SNr neurons projected to the SC. Indeed, gaze (as well as attention) is attracted by reward-associated objects in monkeys (Hikosaka et al., 2006), rodents (Gallagher et al., 1990), and humans (Theeuwes and Belopolsky, 2012). Moreover, the association with rewards induces lasting effects on attentional processing in monkeys (Bichot and Schall, 1999; Peck et al., 2009) and humans (Hickey et al., 2010; Kristjánsson et al., 2010; Anderson et al., 2011; Awh et al., 2012) even several days later (Della Libera and Chelazzi, 2009; Anderson et al., 2011). It also promotes visual perceptual learning in humans even when the stimuli are processed unconsciously (Seitz et al., 2009). We speculate that the SNr may be a key mechanism to

guide gaze and possibly attention based on learned object values. The SNr may also contribute to manual choice (e.g., reach out and manipulate objects), because gaze choice precedes or guides manual choice in humans (Johansson et al., 2001; Hayhoe and Ballard, 2005) and monkeys (Miyashita et al., 1996).

The memories represented by SNr neurons are akin to skill memories. A skill develops gradually with repeated practice in a consistent context, has high capacity compared with working memory (Shiffrin and Schneider, 1977), and is retained for a long time with little gradation (Hikosaka et al., 2002), as observed in our monkeys’ gaze behavior and SNr neurons’ activity. Another feature of a skill is its speed: a motor response is evoked quickly by a particular sensory input, often beyond subjective control (Hikosaka et al., 2002). Since SC-projecting SNr neurons differentiate between high-valued and low-valued objects < 150 ms after they appear (Fig. 7A), saccades could be biased to rewarded objects within 200 ms at most after they appear. These consider-

ations suggest that the SNr could be a mediator of visuomotor skills. This hypothesis is consistent with a general theory that the basal ganglia are essential for skill or habit memories (Mishkin et al., 1984; Yin and Knowlton, 2006; Graybiel, 2008; Ashby et al., 2010; Seger and Spiering, 2011).

Perhaps most relevant to our study is a series of seminal studies by Mishkin and colleagues suggesting that the basal ganglia are responsible for the choice of objects based on stable values (Mishkin et al., 1984). In their experiment, the monkey was presented with pairs of objects in succession, and was asked to choose one object from each pair. One object in each pair was consistently associated with a reward. By repeating choices across daily learning sessions, the monkey gradually became able to choose reward-associated objects (Malamut et al., 1984). Notably, this learning was not impaired by lesions of hippocampus-amygdala regions (Malamut et al., 1984), but was impaired by lesions of the caudate tail (CDt) (Fernandez-Ruiz et al., 2001), which is known to receive inputs from the inferotemporal cortex (Saint-Cyr et al., 1990; Yeterian and Pandya, 1995) and send outputs to the SNr (Szabo, 1972; Saint-Cyr et al., 1990; François et al., 1994). Our results are consistent with the basal ganglia hypothesis offered by Mishkin and colleagues (1984) and may further reveal underlying single cellular mechanisms.

What is striking about the responses of SNr neurons was that almost invariably they were more inhibited by stably high-valued objects than stably low-valued objects (Fig. 7). Such a homogeneous bias among neurons is present in brain areas involved in motivational control for flexibly valued objects (Schultz, 1998; Matsumoto and Hikosaka, 2009), but has rarely been found in sensorimotor areas involved in decision-making (Schultz et al., 2003), even among SNr neurons (Sato and Hikosaka, 2002). This is particularly remarkable because the same SNr neurons responded to visual objects variably if the monkey had not experienced the objects with long-term reward biases (Fig. 3B). It is as if visual information were replaced with value information. Importantly, however, this cannot occur as a single global process because the stable value information is attached to the objects individually and separately. We thus propose the following hypothesis: the visual information is gradually but robustly modified through long-term reward experiences and this occurs separately for individual objects depending on their associated values.

The above considerations lead to an important question: do SNr neurons store the memories of stable object values? SNr neurons are likely to receive object-selective signals from the inferotemporal cortical areas through the CDt (Saint-Cyr et al.,

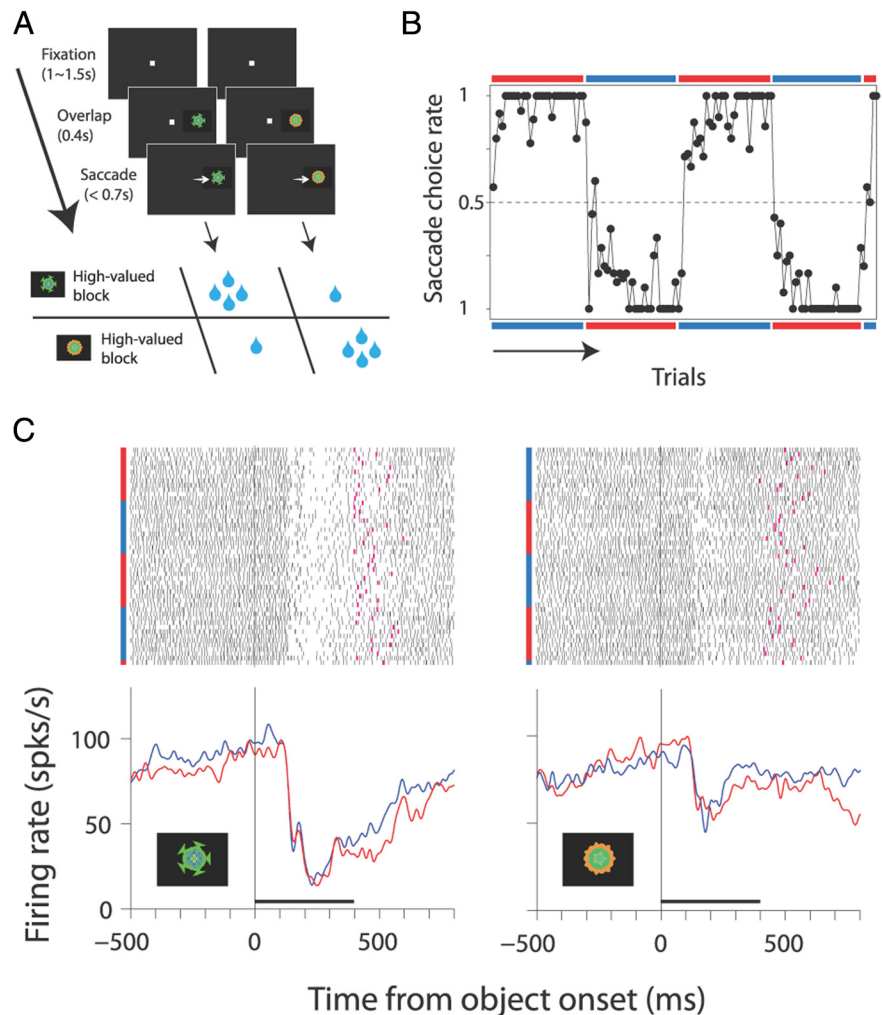


Figure 13. Weak influences of flexibly changing object values on activity of an SNr neuron. **A**, Flexible object–reward association task. The monkey was required to hold gaze on the fixation spot until it was turned off (overlap period, 400 ms) and then make a saccade to the object. Two fractal objects were associated with large and small rewards in a reversible manner in blocks of trials (1 block, 30 trials). On most trials (4 of 5 trials) one of the two objects was presented at the neuron’s preferred position (forced trials). Occasionally (1 of 5 trials) two objects were presented and the monkey had to choose one of the objects (choice trials, not shown). **B**, Changes in the choice of monkey D on choice trials across four blocks. The plots close to the horizontal red bars indicate choices preferring the high-valued object (i.e., whichever object was recently associated with the large reward). **C**, Spike activity of a single SC-projecting SNr neuron during the flexible object–reward association task aligned on the onset of the object (time 0). Vertical red and blue bars on the left of the rasters indicate large-reward and small-reward trials. The purple tick in each raster line indicates the onset of the saccade toward the presented object. The averaged responses are shown separately for the high-valued object (red) and the low-valued object (blue). Black horizontal bars indicate the overlap period. The data were obtained from the same neuron as shown in Figure 5.

1990) and the SNr neuronal signals may be fed back to the inferotemporal cortex (Middleton and Strick, 1996). Neurons in the inferotemporal cortex are known to change their responses such that, through natural experience, they can respond to particular visual objects regardless of their appearances (Li and DiCarlo, 2008, 2012), but this “tolerance” learning occurs in an unsupervised manner and does not require rewarding experiences (Li and DiCarlo, 2012). Some neurons in the inferotemporal cortex change their responses to visual objects depending on their stable reward values, but the changes are minor (Jagadeesh et al., 2001; Mogami and Tanaka, 2006). In contrast, the object-selective responses of CDt neurons are modulated by stable values, similarly to SNr neurons but less robustly (S. Yamamoto, H. F. Hyoung, and O. Hikosaka, unpublished observation). These results raise the possibility that synaptic plasticities occur within the SNr, per-

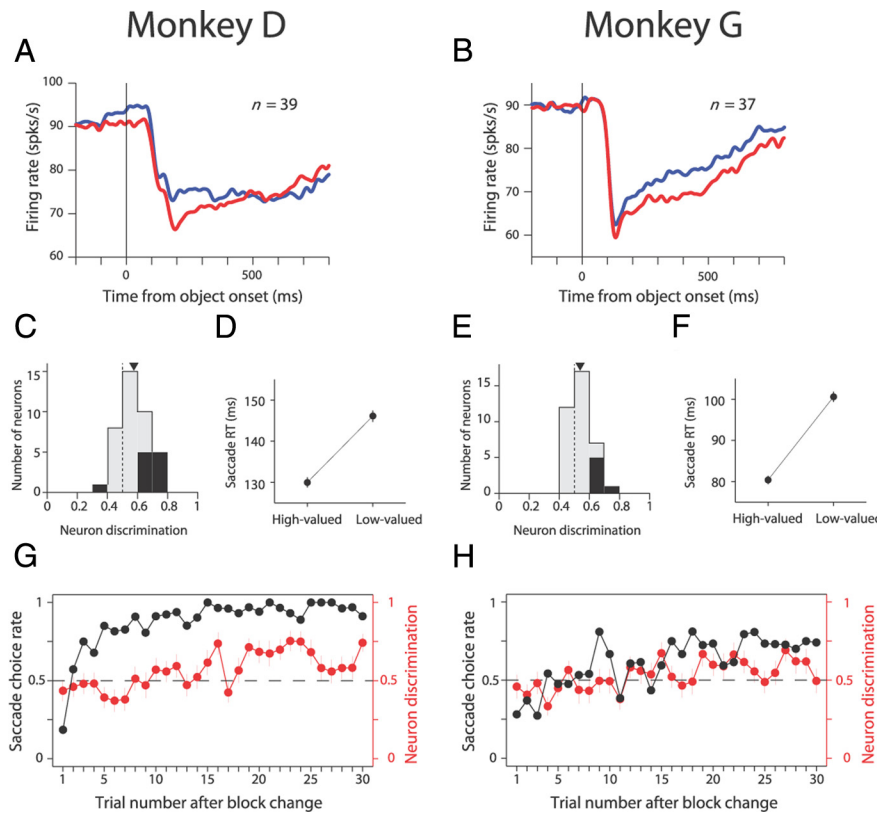


Figure 14. Flexible object value coding of SNr neurons shown separately for two monkeys. **A, B**, Average responses of SNr neurons to the high-valued object (i.e., recently associated with a large reward) (red) and the low-valued object (i.e., recently associated with a small reward) (blue). The data were obtained from 39 neurons in monkey D (left) and 37 neurons in monkey G (right). **C, E**, SNr neurons' discriminations between high-valued and low-valued objects measured by ROC area based on spike counts in a test window. Black bars indicate neurons with statistically significant discrimination assessed by Wilcoxon rank-sum test ($p < 0.05$). A triangle indicates the mean of the ROC areas (monkey D: 0.58; monkey G: 0.54). **D, F**, The mean saccade reaction times (Saccade RT) for large-reward trials were significantly shorter for high-valued than for low-valued objects (paired t test, $p \approx 0$ in both monkeys). Note that the reaction time was measured from the offset of the fixation point. For (**A–F**), the neuron's responses in the initial four trials in each block have been excluded. **G, H**, Behavioral (black) and neuronal (red) learning curves associated with flexible object values. The behavioral learning curve is the choice rate for the high-valued object. The neuronal learning curve is the averaged ROC areas across trials after the reversal of the object–reward contingency.

haps in addition to the CDt or the inferotemporal cortex. Such synaptic plasticities might be achieved by dopaminergic inputs to the SNr (Ruffieux and Schultz, 1980; Waszczak and Walters, 1983; Radnikow and Misgeld, 1998), as dopamine neurons encode reward prediction errors (Schultz, 1998) and can induce synaptic plasticities (Wickens et al., 2003).

References

Anderson BA, Laurent PA, Yantis S (2011) Value-driven attentional capture. *Proc Natl Acad Sci U S A* 108:10367–10371. [CrossRef Medline](#)

Ashby FG, Turner BO, Horvitz JC (2010) Cortical and basal ganglia contributions to habit learning and automaticity. *Trends Cogn Sci* 14:208–215. [CrossRef Medline](#)

Awh E, Belopolsky AV, Theeuwes J (2012) Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn Sci* 16:437–443. [CrossRef Medline](#)

Bichot NP, Schall JD (1999) Effects of similarity and history on neural mechanisms of visual selection. *Nat Neurosci* 2:549–554. [CrossRef Medline](#)

Brady TF, Konkle T, Alvarez GA, Oliva A (2008) Visual long-term memory has a massive storage capacity for object details. *Proc Natl Acad Sci U S A* 105:14325–14329. [CrossRef Medline](#)

Brown VJ, Desimone R, Mishkin M (1995) Responses of cells in the tail of the caudate nucleus during visual discrimination learning. *J Neurophysiol* 74:1083–1094. [Medline](#)

Caan W, Perrett DI, Rolls ET (1984) Responses of striatal neurons in the behaving monkey. 2. Visual processing in the caudal neostriatum. *Brain Res* 290:53–65. [CrossRef Medline](#)

Della Libera C, Chelazzi L (2009) Learning to attend and to ignore is a matter of gains and losses. *Psychol Sci* 20:778–784. [CrossRef Medline](#)

Fagot J, Cook RG (2006) Evidence for large long-term memory capacities in baboons and pigeons and its implications for learning and the evolution of cognition. *Proc Natl Acad Sci U S A* 103:17564–17567. [CrossRef Medline](#)

Fernandez-Ruiz J, Wang J, Aigner TG, Mishkin M (2001) Visual habit formation in monkeys with neurotoxic lesions of the ventrocaudal neostriatum. *Proc Natl Acad Sci U S A* 98:4196–4201. [CrossRef Medline](#)

François C, Yelnik J, Percheron G, Fénelon G (1994) Topographic distribution of the axonal endings from the sensorimotor and associative striatum in the macaque pallidum and substantia nigra. *Exp Brain Res* 102:305–318. [Medline](#)

Gallagher M, Graham PW, Holland PC (1990) The amygdala central nucleus and appetitive Pavlovian conditioning: lesions impair one class of conditioned behavior. *J Neurosci* 10:1906–1911. [Medline](#)

Graybiel AM (2008) Habits, rituals, and the evaluative brain. *Annual review of neuroscience* 31:359–387. [CrossRef Medline](#)

Hayhoe M, Ballard D (2005) Eye movements in natural behavior. *Trends Cogn Sci* 9:188–194. [CrossRef Medline](#)

Hickey C, Chelazzi L, Theeuwes J (2010) Reward guides vision when it's your thing: trait reward-seeking in reward-mediated visual priming. *PLoS One* 5:e14087. [CrossRef Medline](#)

Hikosaka O, Wurtz RH (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *J Neurophysiol* 49:1285–1301. [Medline](#)

Hikosaka O, Takikawa Y, Kawagoe R (2000) Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev* 80:953–978. [Medline](#)

Hikosaka O, Rand MK, Nakamura K, Miyachi S, Kitaguchi K, Sakai K, Lu X, Shimo Y (2002) Long-term retention of motor skill in macaque monkeys and humans. *Exp Brain Res* 147:494–504. [CrossRef Medline](#)

Hikosaka O, Nakamura K, Nakahara H (2006) Basal ganglia orient eyes to reward. *J Neurophysiol* 95:567–584. [Medline](#)

Hori Y, Minamimoto T, Kimura M (2009) Neuronal encoding of reward value and direction of actions in the primate putamen. *J Neurophysiol* 102:3530–3543. [CrossRef Medline](#)

Jagadeesh B, Chelazzi L, Mishkin M, Desimone R (2001) Learning increases stimulus salience in anterior inferior temporal cortex of the macaque. *J Neurophysiol* 86:290–303. [Medline](#)

Johansson RS, Westling G, Bäckström A, Flanagan JR (2001) Eye-hand coordination in object manipulation. *J Neurosci* 21:6917–6932. [Medline](#)

Kristjánsson A, Sigurjónsdóttir O, Driver J (2010) Fortune and reversals of fortune in visual search: reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Atten Percept Psychophys* 72:1229–1236. [CrossRef Medline](#)

Lau B, Glimcher PW (2008) Value representations in the primate striatum during matching behavior. *Neuron* 58:451–463. [CrossRef Medline](#)

Li N, DiCarlo JJ (2008) Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science* 321:1502–1507. [CrossRef Medline](#)

Li N, DiCarlo JJ (2012) Neuronal learning of invariant object representation in the ventral visual stream is not dependent on reward. *J Neurosci* 32:6611–6620. [CrossRef Medline](#)

- Malamut BL, Saunders RC, Mishkin M (1984) Monkeys with combined amygdalo-hippocampal lesions succeed in object discrimination learning despite 24-hour intertrial intervals. *Behav Neurosci* 98:759–769. [CrossRef Medline](#)
- Matsumoto M, Hikosaka O (2009) Representation of negative motivational value in the primate lateral habenula. *Nat Neurosci* 12:77–84. [CrossRef Medline](#)
- Middleton FA, Strick PL (1996) The temporal lobe is a target of output from the basal ganglia. *Proc Natl Acad Sci U S A* 93:8683–8687. [CrossRef Medline](#)
- Mishkin M, Malamut B, Bachevalier J (1984) Memories and habits: two neural systems. In: *Neurobiology of Human Learning and Memory* (Lynch G, McGaugh JL, Weinberger NM, eds), pp 65–77. New York: The Guilford.
- Miyashita K, Rand MK, Miyachi S, Hikosaka O (1996) Anticipatory saccades in sequential procedural learning in monkeys. *J Neurophysiol* 76:1361–1366. [Medline](#)
- Miyashita Y, Higuchi S, Sakai K, Masui N (1991) Generation of fractal patterns for probing the visual memory. *Neurosci Res* 12:307–311. [CrossRef Medline](#)
- Mogami T, Tanaka K (2006) Reward association affects neuronal responses to visual stimuli in macaque TE and perirhinal cortices. *J Neurosci* 26:6761–6770. [CrossRef Medline](#)
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J (2009) Reward modulates attention independently of action value in posterior parietal cortex. *J Neurosci* 29:11182–11191. [CrossRef Medline](#)
- Radnikow G, Misgeld U (1998) Dopamine D₁ receptors facilitate GABA_A synaptic currents in the rat substantia nigra pars reticulata. *J Neurosci* 18:2009–2016. [Medline](#)
- Redgrave P, Prescott TJ, Gurney K (1999) The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89:1009–1023. [CrossRef Medline](#)
- Ruffieux A, Schultz W (1980) Dopaminergic activation of reticulata neurons in the substantia nigra. *Nature* 285:240–241. [CrossRef Medline](#)
- Saint-Cyr JA, Ungerleider LG, Desimone R (1990) Organization of visual cortical inputs to the striatum and subsequent outputs to the pallidum in the monkey. *J Comp Neurol* 298:129–156. [CrossRef Medline](#)
- Saleem KS, Logothetis NK (2007) A combined MRI and histology atlas of the rhesus monkey brain in stereotaxic coordinates. London: Academic.
- Sato M, Hikosaka O (2002) Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22:2363–2373. [Medline](#)
- Schultz W (1998) Predictive reward signal of dopamine neurons. *J Neurophysiol* 80:1–27. [Medline](#)
- Schultz W, Tremblay L, Hollerman JR (2003) Changes in behavior-related neuronal activity in the striatum during learning. *Trends Neurosci* 26:321–328. [CrossRef Medline](#)
- Seger CA, Spiering BJ (2011) A critical review of habit learning and the basal ganglia. *Front Syst Neurosci* 5:66. [Medline](#)
- Seitz AR, Kim D, Watanabe T (2009) Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron* 61:700–707. [CrossRef Medline](#)
- Shiffrin RM, Schneider W (1977) Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychol Rev* 84:127–190. [CrossRef](#)
- Standing L (1973) Learning 10,000 pictures. *Q J Exp Psychol* 25:207–222. [CrossRef Medline](#)
- Szabo J (1972) The course and distribution of efferents from the tail of the caudate nucleus in the monkey. *Exp Neurol* 37:562–572. [CrossRef Medline](#)
- Theeuwes J, Belopolsky AV (2012) Reward grabs the eye: oculomotor capture by rewarding stimuli. *Vision Res*. Advance online publication. Retrieved August 8, 2012. doi:10.1016/j.visres.2012.07.024. [CrossRef Medline](#)
- Vaughan W, Greene SL (1984) Pigeon visual memory capacity. *J Exp Psychol Anim Behav Process* 10:256–271. [CrossRef](#)
- Waszczak BL, Walters JR (1983) Dopamine modulation of the effects of gamma-aminobutyric acid on substantia nigra pars reticulata neurons. *Science* 220:218–221. [CrossRef Medline](#)
- Wickens JR, Reynolds JN, Hyland BI (2003) Neural mechanisms of reward-related motor learning. *Curr Opin Neurobiol* 13:685–690. [CrossRef Medline](#)
- Yamamoto S, Monosov IE, Yasuda M, Hikosaka O (2012) What and where information in the caudate tail guides saccades to visual objects. *J Neurosci* 32:11005–11016. [CrossRef Medline](#)
- Yeterian EH, Pandya DN (1995) Corticostriatal connections of extrastriate visual areas in rhesus monkeys. *J Comp Neurol* 352:436–457. [CrossRef Medline](#)
- Yin HH, Knowlton BJ (2006) The role of the basal ganglia in habit formation. *Nat Rev Neurosci* 7:464–476. [CrossRef Medline](#)