

Current Biology, Volume 23

Supplemental Information

Volitional Control of Neural Activity

Relies on the Natural Motor Repertoire

Eun Jung Hwang, Paul M. Bailey, and Richard A. Andersen

Supplemental Inventory

Supplemental Figures

Figure S1, related to Figure 2

Figure S2, related to Figure 3

Figure S3, related to Figure 4

Figure S4, related to Figure 5

Supplemental Results, related to Figures 3, S3, and S4

Supplemental Experimental Procedures

Supplemental References

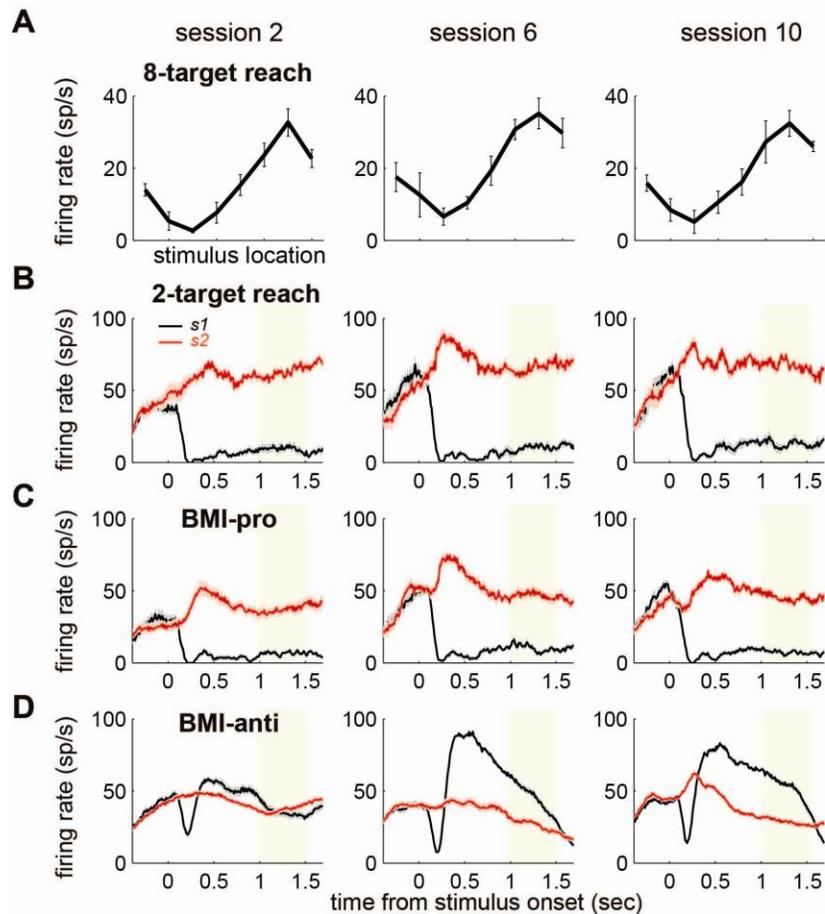


Figure S1. Long-Term Training of a Single Parietal Reach Region Neuron, Related to Figure 1

A. The firing rate (mean \pm SD) of the direct neuron for each of 8 target locations in the reach task of sessions 2, 6, and 10 of monkey Y. The mean correlation coefficient between tuning curves is 0.99.

B-D. The temporal dynamics of the firing rate (mean \pm SEM) of the direct neuron for two stimulus locations in the reach, BMI-pro, and BMI-anti tasks. The firing rate in the yellow region governs reward in the BMI tasks.

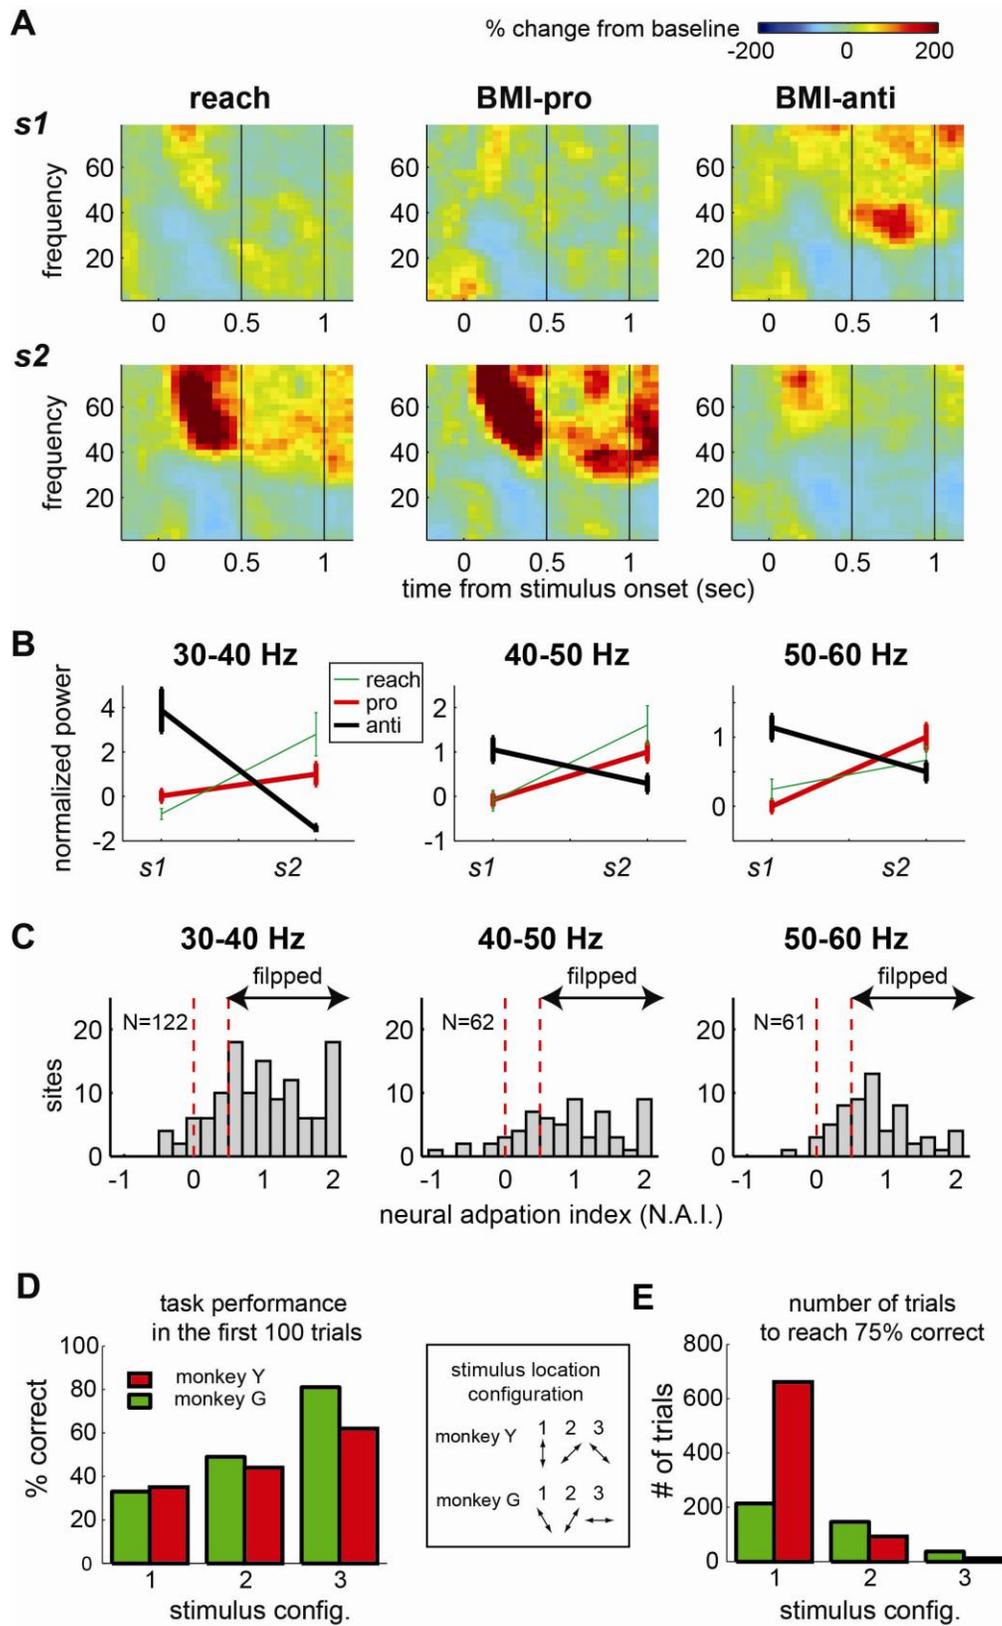


Figure S2.

Figure S2. Further Evidence for Intrinsic-Variable Learning: LFPs Flip the Preferred Stimulus in the BMI-Anti Task (A-C), and Learning Is Generalized in the BMI-Anti Task (D and E), Related to Figure 3

- A.** The LFP power spectrograms, normalized to the % change from the baseline (0.5 s interval before stimulus onset) for two stimuli. The firing rate of the trained neuron in Figure 3A in the interval between the two vertical lines governs reward in the BMI tasks.
- B.** The modulation of the LFP spectral power (mean \pm SEM) by the stimulus during the specified time interval in three frequency bands.
- C.** Distributions of the neural adaptation index for the spectral power across all available LFP sites. The median was 0.97, 0.91, and 0.81 respectively.
- D.** The % correct of the first 100 trials (spanning 13 ± 2.8 (SD) minutes) for each stimulus location pair in the first session when the pair was first presented. The number for stimulus pair represents the training order among them.
- E.** The number of trials required to reach 75% correct.

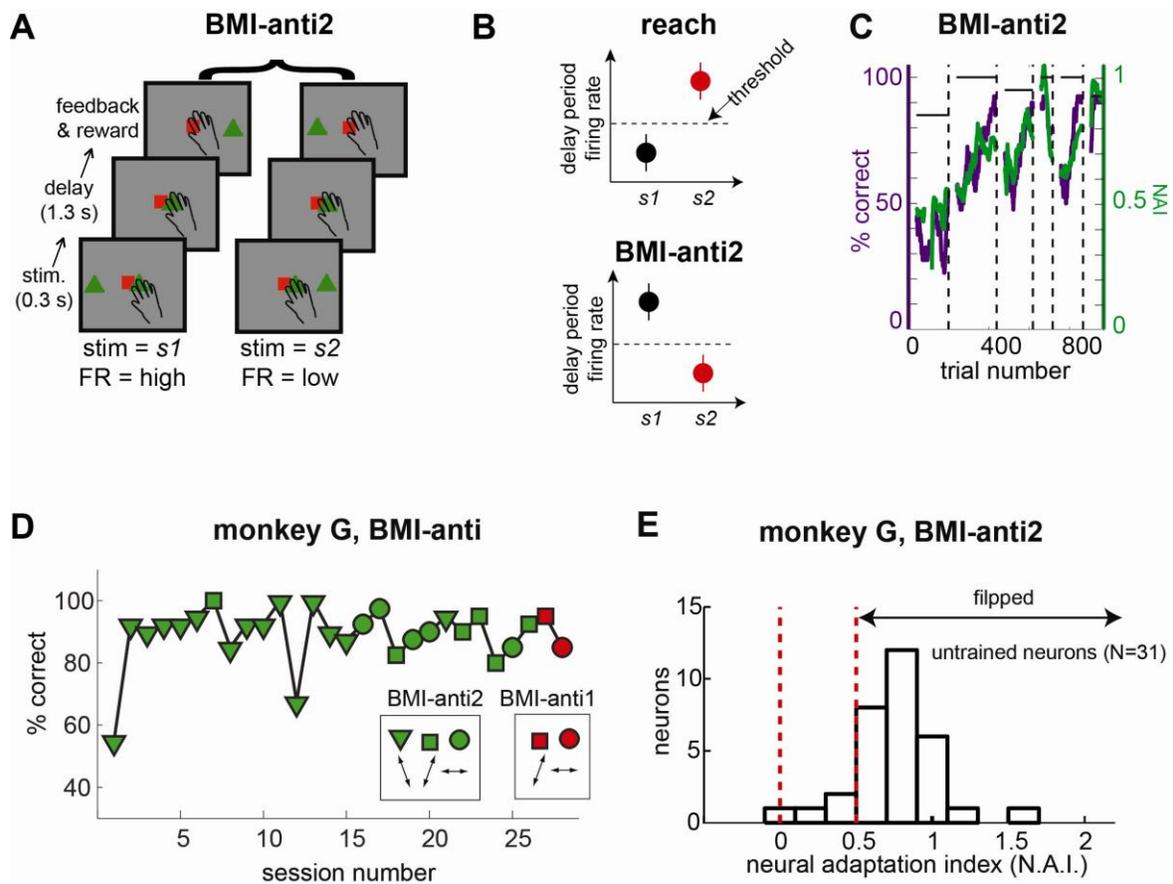


Figure S3. Further Evidence for Intrinsic-Variable Learning: Task Modification Facilitates the Discovery of a Successful Cognitive Strategy, Related to Figure 4

A. The temporal event sequence in successful trials for two stimulus locations in the BMI-anti2 task.

B. The activity pattern of a hypothetical neuron for successful trials in the reach versus BMI-anti2 task.

C. The % correct (purple) and neural adaptation index (green) in the BMI-anti2 task training sessions from Monkey G. The dashed vertical lines indicate the end of each of the 6 daily sessions. The horizontal bars indicate the % correct in the BMI-pro task in the corresponding sessions.

D. The peak performance in each of 26 BMI-anti2 and 2 BMI-anti1 task sessions for monkey G. The different symbols indicate different stimulus pairs. The configuration for each stimulus pair is illustrated in the inset.

E. The distribution of the neural adaptation index of untrained neurons (N=31) for all successful trials of BMI-anti2 task training from monkey G.

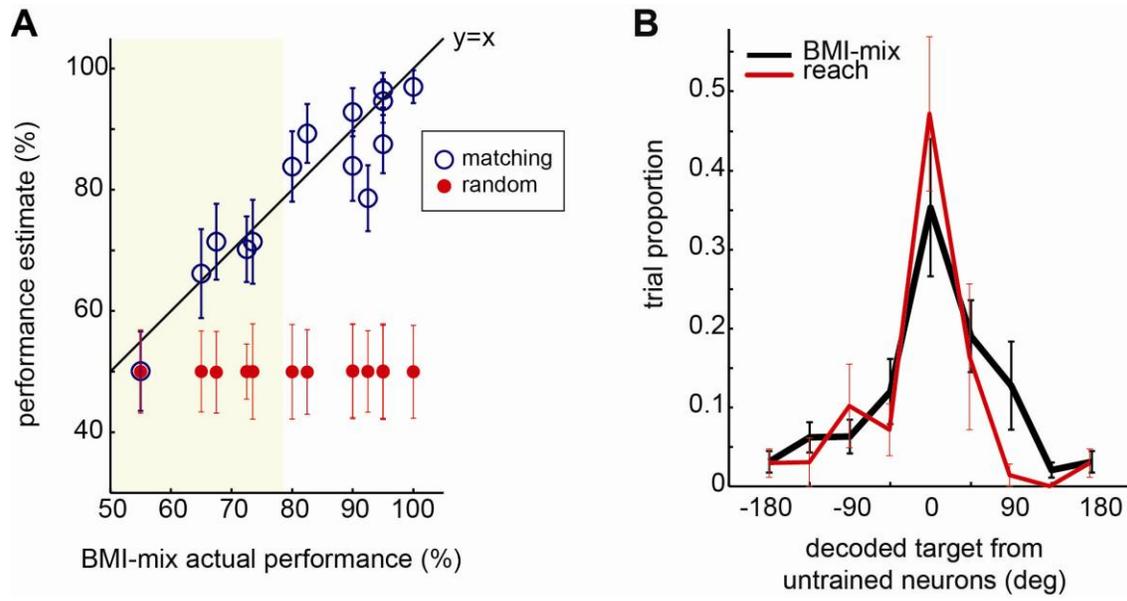


Figure S4. The BMI-Mix Task Activity Resembles the Reach Planning Activity of Matching Targets, Related Figure 5

A. The actual versus estimated performance (mean \pm SD of 10000 repetitions) using firing rate samples from the reach planning activity for the matching and random targets, respectively (each session yielded one blue and one red circle). The yellow box contains 5 sessions with actual performance $<80\%$.

B. Both trained and untrained neurons encode the same targets during the BMI-mix task in 5 sessions with performance $<80\%$. The black line is the histogram of the difference between the targets decoded from the trained and untrained neurons in the BMI-mix task when the target decoded from the trained neurons is a matching target (mean \pm SEM, across 5 BMI-mix sessions). The peak at zero indicates that when trained neurons encode a matching target, untrained neurons also encode the same matching target. The red line is the same but for the 2-target reach task when the target decoded from the trained neurons is the reach target.

Supplemental Results

The Temporal Dynamics of Neuronal Activity in the BMI-Anti Task, Related to Figure 3

Figure 3A shows the temporal dynamics of the firing rates for a trained neuron and 3 untrained neurons in a BMI-anti1 session. These temporal dynamics closely resemble the population activity in the parietal reach region when monkeys plan an anti reach[1]. In contrast to the anti reach task, however, tuning tended to revert back after the decoding window period was completed in the BMI-anti1 task. The task design might explain this phenomenon. In the BMI-anti1 task, the feedback cursor was presented at the same location as the stimulus cue for successful trials. Thus, it is possible that after the decoding window closed, the monkey might have quickly formed a default plan for a reach to the expected feedback location, similar to the pro reach plan. Alternatively, the animal could have shifted attention to the appearance of the feedback target. Consistent with this idea that the anticipated feedback location influenced the activity following the decoding window, the tuning did not revert back in the BMI-anti2 task, where the feedback cursor was presented opposite the stimulus for successful trials.

Further Evidence for Intrinsic-Variable Learning: Task Modification Can Facilitate the Discovery of a Successful Cognitive Strategy, Related to Figure S3

Figure 4 shows that the behavior of untrained neurons in the BMI-anti1 task was consistent with the intrinsic-variable learning hypothesis even for monkey G whose task performance was low. If monkey G indeed pursued intrinsic-variable learning in the BMI-anti1 task, then facilitating the discovery of a successful cognitive strategy might help learning. We confirmed this idea using the following strategy. Monkey G was tested in a slightly modified form of the BMI-anti1 task, BMI-anti2, in which the stimulus-response rule was the same as the BMI-anti1 task but the opposite feedback cursor policy was employed (**Figure S3A-B**). If the monkey planned an anti reach, then the feedback cursor appeared opposite the stimulus and the monkey received a reward. If the monkey planned a pro reach, then the feedback cursor appeared at the stimulus location and the monkey did not receive a reward. We conjectured that this modification, which provided veridical feedback with respect to the planned reach direction, might facilitate the formation of a cognitive association between the intention of the monkey and the feedback, and thus might more explicitly guide the monkey to develop the strategy of planning anti reaches. In fact, unlike in the BMI-anti1 task, monkey G achieved a stable high performance level in the BMI-anti2 task within the first 6 sessions, each on different days (**Figure S3C**), suggesting that the feedback modification, in fact, facilitated the learning of the same stimulus-response rule.

After the first 6 sessions with one stimulus pair, monkey G performed 20 more sessions of the BMI-anti task2, up to 3 per day, with 2 additional stimulus pairs, and then performed 2 more sessions of the BMI-anti task1 to test generalization (**Figure S3D**). The average peak performance and the average NAI in the BMI-anti task2 across all sessions were $89 \pm 10.0\%$ and 0.95 ± 0.156 , respectively (**Figure S3D**). Furthermore, for successful trials, the majority of untrained neurons (27 of 31) showed an NAI greater than 0.5, and the median index (0.78) was significantly greater than 0.5 (Wilcoxon signed-rank test, $p < 1e-5$) (**Figure S3E**). Therefore, untrained neurons also showed intrinsic-variable learning in the modified task.

The BMI-Mix Task Activity Closely Resembles the Reach-Planning Activity for the Matching Targets, Related to Figure S4A

For the example shown in Figure 5A-D, the mean activity for stimulus 1 in the BMI-mix task matched the mean reach-planning activity for the target at 45° counterclockwise from the stimulus, while the mean activity for stimulus 2 matched the mean reach-planning activity for the target at 135° counterclockwise from the stimulus. Likewise, in all other sessions, the reach target that matched the activity for each stimulus in the BMI-mix task was identified based on the mean activity. We further assessed how well the reach-planning activity for the matching targets resembled the activity in the BMI-mix task using the following strategy.

If the monkey indeed planned reaches to the matching targets, the BMI-mix task performance should be accurately estimated from the reach-planning activity of the matching targets. Thus, as a performance estimate, we computed the proportion of the randomly sampled firing rates from the reach-planning activity of the matching targets that conformed to the BMI-mix stimulus-response rule. For the particular example shown in Figures 5A-D, the performance estimate was 94%, similar to the actual performance of 95%. We estimated the chance level performance using the same analysis but with randomly selected targets instead of the matching targets. The performance estimate using random targets was 51%. In all 9 sessions with performance accuracy greater than 80%, the performance estimated using the matching targets was not significantly different from the actual performance (**Figure S4A**; paired *t*-test; $t(8)=0.86$, $p>0.4$), while the performance estimated using random targets was significantly different from the actual performance (paired *t*-test; $t(8)=16.7$, $p<1e-6$). Statistical tests using all 14 sessions produced similar results. These results clearly show that the newly emerged activity patterns in the BMI-mix task could be elicited through target re-aiming, i.e., planning reaches to matching targets transformed from stimuli locations.

Supplemental Experimental Procedures

Behavioral Tasks

The monkeys sat in a dark room in front of an LCD monitor mounted behind a touch-sensitive screen. Eye position was tracked with an infrared eye tracker (ISCAN). The monkeys maintained eye fixation throughout the trial in both reach and BMI tasks. All trials began as the monkeys fixated on a central square and touched a central triangle (**Figure 1A**). After a 0.5-s hold period, a triangular stimulus was presented at 10 cm eccentricity from the central fixation point for 0.3 s (stimulus period). The stimulus location was randomly alternated between two diametrically opposing locations, except for the 8-target reach task in which the location was alternated among 8 equidistant locations. A delay period of 1.3 ± 0.08 (SD) s followed the stimulus offset and ended with a “go” signal (the extinction of the central triangle). In the reach tasks, the monkeys were required to make center-out reaches to the remembered stimulus location after the “go” signal. Seven to ten successful reaches to each location were completed.

In the BMI tasks, within 0.1 s after the “go” signal, we displayed a feedback cursor at one of the two locations based on the delay period firing rate of the trained neuron(s) (**Figure 1B**). The decoding window in the delay period varied across BMI-anti task sessions, 1.0-1.5 s from stimulus onset in the first 17 sessions, 0.9-1.4 s in 1 session, and 0.5-1.0 s in 10 sessions for monkey Y, and 0.9-1.4 s in 28 sessions for monkey G. The decoding window was fixed at 0.5-1.0 s from the stimulus onset in the BMI-mix task. If the stimulus-response rule was satisfied (**Figure 1B**), the monkeys received a juice reward with a high pitch acoustic tone after maintaining the central touch and eye fixation for an additional 0.3 s after the feedback; otherwise, the trial was aborted with a low-pitch acoustic tone. Due to the constraint of maintaining the touch fixation for an additional 0.4 s after the “go” signal, the monkeys stopped making actual reaches in response to the “go” signal in the BMI trials.

Monkey Y had previous experience in BMI-pro tasks. However, monkey G performed the BMI-pro task for the first time in the BMI-pro task block in the first session of this study. The BMI-pro task in each session continued until the percent correct for 40 consecutive trials first exceeded 70% (44 ± 7.5 trials per session). The BMI-anti task continued as long as the monkeys continued to be engaged in the task and isolation of the trained neuron was maintained (190 ± 75 trials per session). In some sessions, however, the task was terminated once the monkeys’ performance saturated at a near perfect level. The same criteria were applied to the BMI-mix task (305 ± 237.6 trials per session).

Each BMI-anti experimental session consisted of three task blocks: reach, BMI-pro, and BMI-anti. During the early days of training, we conducted only a single session on each day. However, once the monkey became proficient at the BMI-anti task and rapidly achieved a high performance level within each session, we occasionally conducted up to 3 sessions on a single day. Each session used different sets of neurons, except for sessions 2-10 for monkey Y, in which the same neuron was used as the trained neuron.

Note that, in the BMI-anti and BMI-mix tasks, the central hand targets and visual cues were indicated in a different color from the BMI-pro task (blue versus green for monkey Y, magenta versus green for monkey G) to show the contextual change.

Neural Adaptation Index

If the preferred stimulus did not change and the tuning modulation depth increased, then the NAI was less than 0. If the preferred stimulus was the same and the tuning depth decreased, the NAI was calculated between 0 and 0.5. If the preferred stimulus flipped and the tuning depth decreased, the NAI was between 0.5 and 1. If the preferred stimulus flipped and the tuning depth increased, the NAI was greater than 1. NAI usually ranges between 0 and 1. However, the NAI was unreliable for signals with a small tuning depth in the BMI-pro task, which is the denominator of the formula. Unlike trained neurons, untrained neurons could be insensitive to the stimulus locations in the BMI-pro task, resulting in nonsensical values. Thus, when computing NAIs for untrained neurons, we included neurons only if their tuning depth in the BMI-pro task was at least 1 spike/s.

Neural Recording

The monkey was implanted with a head holder and a recording chamber housing a 16-channel semi-chronic microdrive (Neuralynx, Bozeman, MT). Sixteen electrodes were strategically placed over the intra parietal sulcus (IPS), guided by structural magnetic resonance images so that most of them would be located in the parietal reach region, i.e., the medial bank of IPS. The electrodes were spread over a $2 \times 6.5 \text{ mm}^2$ area along the IPS on the brain surface, but they became more densely populated once they entered the brain due to the oblique incidence angles required to hit the IPS in the center of the chamber. Across sessions, the electrode locations were adjusted in depth by small amounts to improve the quality of unit isolations or to acquire a new set of neurons. A commercial 16-channel neural signal recording system (Plexon MAP, Dallas, TX) was used to record and store neural signals. LFPs and spikes were separated using hardware band-pass filters in a preamplifier (LFP: 3.3-88 Hz and spike: 154 Hz – 8.8 kHz). We performed online spike detection and sorting using a commercial software package (Plexon Rasputin). The spike count computation and target decoding were customized using real-time MATLAB codes.

Supplemental References

1. Gail, A., and Andersen, R.A. (2006). Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. *J. Neurosci.* 26, 9376-9384.