NIH-PA Author Manuscript

Volitional Control of Neural Activity Relies on the Natural Motor Repertoire

Eun Jung Hwang[1], **Paul M. Bailey**[1,2], and **Richard A. Andersen**[1]

[1]Division of Biology, California Institute of Technology, Pasadena, CA 91125, USA

[2]Beth Israel Deaconess Medical Center, Boston, MA 02215, USA

Summary

Background—The results from recent brain machine interface (BMI) studies suggest that it may be more efficient to use simple arbitrary relationships between individual neuron activity and BMI movements than the complex relationship observed between neuron activity and natural movements. This idea is based on the assumption that individual neurons can be conditioned independently regardless of their natural movement association.

Results—We tested this assumption in the parietal reach region (PRR), an important candidate area for BMIs in which neurons encode the target location for reaching movements. Monkeys could learn to elicit arbitrarily assigned activity patterns, but the seemingly arbitrary patterns always belonged to the response set for natural reaching movements. Moreover, neurons that are free from conditioning showed correlated responses with the conditioned neurons as if they encoded common reach targets. Thus, learning was accomplished by finding reach targets (intrinsic variable of PRR neurons) for which the natural response of reach planning could approximate the arbitrary patterns.

Conclusions—Our results suggest that animals learn to volitionally control single neuron activity in PRR by preferentially exploring and exploiting their natural movement repertoire. Thus, for optimal performance, BMIs utilizing neural signals in PRR should harness, not disregard, the activity patterns in the natural sensorimotor repertoire.

Introduction

With brain machine interfaces (BMIs), the neural activity directly controls a machine (e.g., prosthetic arms for paralyzed patients) via decoders that translate the neural activity to movements of the machine. Subjects can learn to control BMIs, sometimes even for arbitrarily determined decoding rules[1–6]. Moritz et al.[1] showed that monkeys learned to volitionally control the activity of any two neurons in the primary motor cortex (M1) for a BMI in which the activation of one neuron stimulated their paralyzed wrist flexor, while the activation of the other stimulated the wrist extensor. Based on this finding, it was proposed that the use of decoders to implement simple arbitrary rules between individual neuron activity and BMI movements may be a more efficient approach than implementing the complex rules observed between the neural activity and natural movements[1, 7].

Corresponding Author: Eun Jung Hwang, Ph.D., Division of Biology, Mail Code 216-76, California Institute of Technology, Pasadena, CA 91125, USA, Tel: 626.395.8337, Fax: 626.795.2397, eunjung@caltech.edu.

In a related study, Jarosiewicz et al.[3] examined the learning mechanism in a BMI task in which a subset of the M1 neurons that were used for the decoder was decoded incorrectly to produce a visuomotor rotation between the desired and the decoded cursor movements. As the monkey learned to offset the visuomotor rotation, preferred directions (PDs) shifted in the direction of the visuomotor rotation across the correctly and incorrectly decoded neurons. However, the incorrectly decoded neurons showed a slightly larger shift in their PDs. A subsequent modeling study[8] showed that these results could be replicated by a single learning mechanism in which the activation of each neuron is updated to a newly explored value whenever the explored value produced a BMI output associated with a larger reward. A key feature of this model is that the explorative signal is randomly and independently assigned to each neuron and that individual neurons independently adapt according to their own activation-reward experience. Such a learning mechanism that facilitates the independent adaptation of individual neurons will be hereafter referred to as "individual-neuron".

Although an individual-neuron mechanism could elegantly reproduce the observed neural changes, the same group originally suggested an alternative, equally viable learning mechanism: the slightly larger change for the incorrectly decoded subset reflects individual-neuron learning, but the dominant global shift of the PDs reflects the behavioral strategy of re-aiming to counter the applied rotation. A cognitive strategy of manipulating an intrinsic variable of natural movements, such as target direction, prevents independent adaptation of individual neurons because the strategy influences a global network of neurons that are sensitive to the manipulated variable. We will hereafter refer to this learning mechanism as "intrinsic-variable". Thus far, it is unclear whether different mechanisms co-exist, if there is a preference for one mechanism over another, or if such a preference changes depending on the circumstances.

Elucidating the predominant forms of learning can help to build optimal BMI decoders[9]. If intrinsic-variable learning predominates, decoders implementing simple arbitrary rules, as suggested by some studies, would not be optimal because learning arbitrary patterns is not guaranteed due to the limited repertoire of activity patterns associated with natural movements. In contrast, if individual-neuron learning predominates, animals would learn to produce virtually any arbitrary activity pattern through the independent adaptation of individual neurons, and thus decoders implementing simple arbitrary rules might indeed be efficacious. At the extreme, each neuron could be individually trained, and essentially no decoder would be required at all.

Elucidating the learning mechanisms could address the even larger issue of whether the brain is so plastic that any area can be trained to operate a BMI. If there are few or no limits on learning BMI tasks across the cortex, as suggested by individual-neuron learning, then there is no need to select particular areas for particular types of BMIs.

A straightforward way to distinguish between individual-neuron versus intrinsic-variable learning is to test whether BMI subjects can learn to produce neural activity patterns that cannot be associated with any natural movement. However, this approach has the shortcoming of requiring complete knowledge of the parameters that are encoded in a brain area to determine the activity patterns that cannot be learned.

A less direct but efficient approach is to examine the behavior of neurons that are observed but are not used for decoding. These neurons are referred to as "untrained neurons," as opposed to "trained neurons," which are used for decoding and thus directly contribute to the BMI output. When intrinsic-variable learning is a possible solution (i.e., a cognitive strategy of manipulating intrinsic variables can produce the appropriate activity pattern),

intrinsic-variable learning would influence untrained neurons in a predictable way based on the cognitive strategy. In contrast, under individual-neuron learning, the activity changes of untrained neurons would be negligible on average because the activity of untrained neurons, which is explored independently from the activity of trained neurons, has no systematic relationship with reward.

Consider the following thought-experiment to better understand this approach. Suppose that trained and untrained neurons respond identically in a reach task: both neurons fire more spikes when reaches are made towards a stimulus on the right than on the left. Now consider an arbitrary BMI decoder rule: the trained neuron must fire more spikes to move the cursor to the left stimulus than to the right, opposite to the reach task. The individual-neuron mechanism allows the trained neuron to adapt to this decoder rule. Intrinsic-variable learning is also possible in this case because the cognitive strategy of planning a reach in the opposite direction of the stimulus (anti-reach) is a viable solution. Thus, both mechanisms can account for the trained neuron to produce the appropriate activity pattern, flipping its tuning preference for the two stimulus locations. However, the untrained neuron will behave differently between the two mechanisms. Under intrinsic-variable learning, via anti-reach planning, both the trained and the untrained neuron would flip their preferred stimulus locations. In contrast, under individual-neuron learning, the untrained neuron would not flip its preferred stimulus because untrained neurons are not reinforced in any consistent way during their independent activity explorations and thus, their net activity change should be near zero.

Based on this rationale, we investigated the BMI learning mechanism in the parietal reach region (PRR), an important candidate area to provide control signals for BMIs[6, 10–12]. PRR neurons primarily encode the planned reach target location in visual coordinates[13–16]. Thus, if intrinsic-variable learning occurs, the main cognitive strategy to change PRR neuronal activity would involve manipulations of the reach target location, such as target re-aiming. We observed that untrained neuron activity in BMI tasks was correlated with trained neuron activity, similar to reach tasks, indicating a cognitive strategy obtained by intrinsic-variable learning. Thus, intrinsic-variable learning predominated in PRR, suggesting that not all brain areas or all arbitrary decoders can be trained to operate a BMI.

Results

In the first study, examining the BMI learning mechanism in PRR, we tested the thought-experiment described in the Introduction using macaque monkeys. Each experimental session consisted of 3 task blocks in the following order: reach, BMI-pro, and BMI-anti1 (Fig. 1A). The individual reach trials consisted of 3 epochs: stimulus, delay, and reachreward. During the stimulus period (0.3 s), one of two diametrically opposing locations (stimulus 1 or 2) was randomly illuminated. The delay period (approximately 1.3 s) was initiated with the stimulus offset and ended with a "go" signal. During the reach-reward period, the monkeys made a reach and received a juice reward if the reach was made to the previously presented stimulus location. The BMI trials followed a similar sequence, replacing the reach-reward period with a feedback-reward period. During the delay period, the cursor feedback location was decoded from the firing rate of the trained neuron. If the firing rate conformed to the stimulus-response rule of a given BMI task (i.e., successful trials), the feedback cursor was placed at the same location as the stimulus cue, and the monkeys were rewarded during the feedback-reward period. Otherwise, the cursor feedback was placed opposite the stimulus and no reward was given.

The following stimulus-response rules were used in the BMI tasks (Fig. 1B). If the firing rate of the trained neuron was lower for stimulus 1 than for stimulus 2 in the reach task, the

BMI-pro task rule was that lower firing rates for stimulus 1 than for 2 would result in a reward. The firing rate threshold, dividing the high and low firing rates, was computed using the maximum-likelihood classifier (*Experimental Procedures*). The BMI-anti task enforced the opposite rule of higher firing rates for stimulus 1 than for 2, similar to the thought-experiment described in the Introduction. This rule forces the trained neuron to flip its preferred stimulus, associated with higher firing rates, between the BMI-pro and BMI-anti tasks. Notably, although we use the term "stimulus-response," the delay period activity, is not a sensory response but rather reflects the monkey's movement plan[17, 18]. Neither of the monkeys had been exposed to any target re-aiming task, such as the anti-reach task, until the BMI-anti task block was performed in this study. Thus, the monkeys were not biased in advance to favor a target re-aiming strategy over individual-neuron learning. We first describe the findings from monkey Y, followed by those of monkey G.

Trained neurons flip their tuning in the BMI-anti1 task

Monkey Y performed the first 10 experimental sessions, each on different days, with the same pair of stimuli. The same trained neuron was used across sessions 2–10 (Fig. S1). The activity of this trained neuron in early, intermediate, and late sessions is shown in Figures 2A–B. In the BMI-pro task, as expected from the stimulus-response rule consistent with the tuning property in the familiar natural reach task, the firing rate of the trained neuron was properly discriminated between the two stimuli from the earliest session for that neuron (Fig. 2A). In contrast, the firing rate in the BMI-anti1 task was indiscriminate between the two stimuli in the earliest session, and only gradually became more discriminate with the opposite pattern from the BMI-pro task (Fig. 2B). The tuning of the trained neuron for the two stimuli differed between the two BMI tasks, as measured using the neural adaptation index (NAI; *Experimental Procedures*). If the tuning did not change, the index was 0. If the tuning for the two stimuli perfectly flipped, the index was 1. Consistent with the firing rate histograms in Figures 2A–B, the NAI gradually increased towards 1 in parallel with task performance accuracy across sessions 2–10 (Fig. 2C).

After the first 10 sessions with one stimulus pair, monkey Y performed 18 more sessions, up to 2 per day, with 3 additional pairs of stimuli (Fig. 2D). The additional 18 sessions used different sets of neurons. For two of the new stimulus pairs, the task performance accuracy was over 80%, even in their first sessions. For the remaining new stimulus pair, performance accuracy gradually increased with training, similar to the original stimulus pair. The average peak performance of monkey Y in the BMI-anti1 task across all 28 sessions was 77±16.5% (mean±s.d.), and the average NAI was 0.85±0.288.

Untrained neurons also flip their tuning in the BMI-anti task1

To determine which mechanism was primarily responsible for the BMI-anti1 task learning, we examined the activity of the untrained neurons. Previously, we discussed how intrinsic-variable learning and not individual-neuron learning would drive untrained neurons to flip their preferred stimulus, similar to the trained neuron. Figure 3A displays the activity of a trained neuron and 3 untrained neurons simultaneously recorded in a typical session for monkey Y (session 23). The activity in the decoding window clearly changed, flipping the preferred stimulus between the two BMI tasks for both trained and untrained neurons (see Supplemental Information for details about the temporal dynamics of neuronal activity).

To quantify the changes in activity of the untrained neurons associated with BMI-anti1 task learning, we computed their NAI for successful trials in which the trained neuron flipped the firing rate for two stimuli. An NAI greater than 0.5 indicates that the preferred stimulus was flipped, while an NAI less than 0.5 indicates that it was not. Thus, under intrinsic-variable learning, the NAI of untrained neurons for successful trials would be greater than 0.5, while

the NAI for unsuccessful trials would be less than 0.5. In contrast, under individual-neuron learning, the tuning of untrained neurons would not change and thus the NAI would be near zero for both successful and unsuccessful trials. The majority (74 of 124) of untrained neurons exhibited an NAI greater than 0.5 for successful trials as they flipped their preferred stimulus, and the median index (0.70) was significantly greater than 0.5 (Wilcoxon signed-rank test, $p<1e-11$) (Fig. 3B). By comparison, the median NAI of the same untrained neurons for unsuccessful trials in which the trained neuron did not flip its firing rate was 0.38, which was significantly smaller than 0.5 (Wilcoxon signed-rank test, $p<1e-7$). These results are consistent with the intrinsic-variable learning hypothesis.

Trained and untrained neurons fluctuate their activity together irrespective of performance level

Monkey G performed the BMI-anti1 task in 12 sessions, each on different days, with the same stimulus pair, using different trained neurons. In contrast to monkey Y, the performance of monkey G showed initial improvement and became saturated at relatively low levels, even after performing 12 sessions with more than 5000 trials (Fig. 4A). The initial improvement from approximately 0% to 50% suggests that the monkey stopped planning pro-reaches, which would have resulted in a performance accuracy near zero. However, the average performance accuracy and NAI of the last 1000 trials were only 57% and 0.58, respectively, indicating no further improvement. Despite the lack of improvement, if the monkey had pursued intrinsic-variable learning, not only trained but also untrained neurons would have flipped their preferred stimulus in successful trials. However, if the monkey had pursued individual-neuron learning, then untrained neurons would not have flipped their preferred stimulus. Therefore, we examined the NAI of untrained neurons in successful trials. The majority of untrained neurons (22 of 32) had an NAI greater than 0.5, and the median index (0.64) was significantly greater than 0.5 (Wilcoxon signed-rank test, $p<1e-6$) (Fig. 4B). The same analysis of the untrained neurons in monkey Y during the first 7 sessions, over which the average performance accuracy was 57%, produced similar results: the majority (17 of 26) had an NAI greater than 0.5, and the median index (0.66) was significantly greater than 0.5 (Wilcoxon signed-rank test, $p<1e-4$). This result indicates that both monkeys pursued intrinsic-variable learning in the BMI-anti1 task, regardless of their performance level.

Further evidence for the intrinsic-variable learning hypothesis

If intrinsic-variable learning is pursued, then facilitating the discovery of a successful cognitive strategy might help learning. We confirmed this idea using 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 (Fig. S3A–B). Unlike in the BMI-anti1 task, monkey G achieved a stable high performance level in the BMI-anti2 task within the first 6 sessions, and showed consistent learning afterwards (Fig. S3C–E; Supplemental Information).

Further supporting the intrinsic-variable hypothesis, the preferred stimulus also flipped for local field potentials in the BMI-anti tasks, which reflect the average activity of the local neural ensemble comprising mostly untrained neurons (Fig. S2A–C). Another finding suggestive of intrinsic-variable learning is that the monkeys generalized the BMI-anti task learning to different stimulus pairs and different neurons (Fig. S2D–E). Individual-neuron learning cannot explain this generalization because each time a new trained neuron is used or new stimuli are introduced, new neural-activity-explorations are necessary, which would require a similar amount of training across different neurons or stimuli. In contrast, intrinsic-variable learning with a cognitive strategy, such as planning the anti-reach, can be generalized across different neurons and stimuli[19].

The relationship between BMI task complexity and the preferred learning mechanism: BMI-mix task

One might wonder whether a preferred learning mechanism would vary depending on the complexity of a cognitive solution for the task. In other words, would individual-neuron learning be more likely to be pursued if the task becomes more cognitively complex to solve? To address this question, we conducted a second study using a BMI-mix task in which the reward was contingent on two neurons and the two neurons were specifically reinforced to respond independently. Only monkey Y performed the BMI-mix task. Fourteen different sessions were recorded on 14 different days with different sets of neurons (2 trained and 4±1.0 untrained neurons per session).

Before each session, the monkey performed a reach task with 8 equidistant targets around a fixation point. Based on the tuning properties of the two trained neurons in the 8-target reach task, a pair of diametrically opposing stimulus locations was selected. Subsequently, with the selected pair of stimuli, each session proceeded with the reach, BMI-pro, and BMImix task blocks, in the same way as the BMI-anti study but with new stimulus-response rules.

The following stimulus-response rules were applied. The 2D space, each axis representing the firing rate of each trained neuron during the delay period, was divided into two regions by a linear boundary (Fig. 5A; *Experimental Procedures*). To be successful, the delay period activity of the trained neurons should fall on the correct side of the boundary, which differed between the two stimuli. The linear boundary for the BMI-pro task was the one that best separated the two firing rate clusters, each formed by the delay period activity for each stimulus in the reach task. Thus, in the BMI-pro task, if the monkey planned a reach to the stimulus, the delay period activity of the trained neurons would conform to the stimulus-response rule. The linear boundary for the BMI-mix task was the one that best separated the two clusters formed by swapping the two stimuli for the reach task activity of the second trained neuron. Thus, the BMI-mix task was part BMI-pro, as the first trained neuron must respond as in the reach task, and part BMI-anti, as the second trained neuron must respond in the opposite way from the reach task.

In the BMI-pro task, the monkey immediately produced the appropriate activity pattern for both trained neurons, achieving 97±3.0% peak performance on average across the sessions (Fig. 5B). For the BMI-mix task, in 9 of 14 sessions, the monkey remarkably learned to produce rule-complying activity patterns, achieving performance accuracy greater than 80% (91±6.4%). Thus, in these BMI-mix sessions, the first trained neuron maintained the same preferred stimulus as the reach task, while the second trained neuron flipped its preferred stimulus (Fig. 5C).

Intrinsic-variable learning is possible in the BMI-mix task

If the newly emerged activity pattern of the trained neurons in these high performance sessions could not be produced by encoding their intrinsic variable (i.e., reach target location), then we can rule out the possibility of intrinsic-variable learning. Thus, in each session, we examined the activity pattern of the trained neurons during the 8-target reach task to determine whether planning reaches to any of the 8 targets could have produced the new activity pattern in the BMI-mix task. For the example session, the activity for stimulus 1 in the BMI-mix task matched the reach-planning activity for the target at 45° counterclockwise from the stimulus, while the activity for stimulus 2 matched the reach-planning activity for the target at 135° counterclockwise from the stimulus (Fig. 5D). Similarly, in all other sessions we found reach targets for which reach-planning activity matched the activity in the BMI-mix task (Fig. S4A; Supplemental Information). Thus, the

new 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.

Untrained neurons indicate intrinsic-variable learning in the BMI-mix task

If an intrinsic-variable mechanism indeed underlay the BMI-mix task learning, not only the trained neurons but also the untrained neurons would encode the same matching targets in successful trials. To address this possibility, we compared the target location encoded by the trained versus untrained neurons for each successful BMI-mix trial. The target that any neural ensemble encoded was inferred using the nearest neighbor-decoding algorithm, which selected the target associated with the ensemble activity in the 8-target reach task that was closest to the ensemble activity of a given BMI-mix trial in terms of Mahalanobis distance. Thus, the decoded target varied among the 8 target locations. Figure 5E shows the 8-target decoding result for the example BMI-mix session in Figures 5A–D. The two bright squares on the diagonal indicate that the trained and untrained neurons concurrently encoded the two specific targets most frequently. Intriguingly, the two specific targets were the same two best matching targets inferred from the activity pattern of the trained neurons as previously described. In 37% of the trials in which the trained neurons encoded a matching target, the untrained neurons also encoded the same matching target in this example session (30±5.1% across all 9 sessions; Fig. 5F). This apparently low number, due to the limited decoder accuracy, is expected as shown in the 2-target reach task: in 28±6.0% trials in which the trained neurons encoded the reach target, the untrained neurons also encode the same target. Even in the 5 sessions with performance < 80%, the trained and untrained neurons encoded the same targets in the BMI-mix task (Fig. S4B). These results suggest that the monkey achieved success in the BMI-mix task by planning reaches to matching targets, a form of intrinsic-variable learning.

Discussion

The repertoire of natural movement-associated activity in paralyzed patients

Our results show that the brain, at least in PRR, explores an existing repertoire of movement-associated activity patterns to control BMIs. This constraint raises a question of how rich a repertoire paralyzed patients can have. BMIs based on PRR are conceived to be cognitive prostheses for which the discrete target location of movements, a cognitive variable, is decoded from the neural activity[10–12]. The representation of cognitive variables is most likely intact, even after long-term paralysis, although it has not been directly tested in PRR[20]. Thus, it is expected that paralyzed patients can readily control PRR-based BMIs as long as the decoder is tuned to reinforce the neural activity patterns observed while the patients vary their intended movement targets.

Learning mechanisms in other areas of the brain

We do not claim that intrinsic-variable learning must predominate in all brain areas nor that intrinsic-variable learning must occur only in PRR. Intrinsic-variable learning does seem to play a significant role in the frontal eye field (FEF). When a monkey volitionally controlled the activity of a FEF neuron, the increased activity shifted the spatial attention of the monkey (intrinsic variable of FEF neurons) to the response field of the neuron, indicating that the monkey learned to produce reward-associated activity by directing spatial attention to specific locations[21–24].

The learning mechanism that plays a dominant role in M1, the primary target of BMI studies, remains unknown. Ganguly et al.[25] reported that the tuning modulation depth of trained neurons became sharper and the modulation depth of untrained neurons became shallower as the monkey became better at controlling the cursor (driven by the activity of

the trained neurons in M1). However, this result alone cannot support one mechanism over another because it is unknown whether or not the newly emerged activity belongs to the natural sensorimotor repertoire. A critical test is to examine the behavior of untrained neurons in tasks for which intrinsic-variable learning is a viable solution. For instance, the task in the study of Jarosiewicz et al. is solvable, at least in large part, by target re-aiming. If untrained neurons in M1 were observed in that task and had shifted their PDs in the same direction as all of their trained neurons, then intrinsic-variable learning in M1 would be strongly supported.

Differences between M1 and PRR

Understanding the difference in intrinsic variables between M1 and PRR might provide a useful insight into the open question of area-dependent learning mechanisms. The issue of what M1 neurons intrinsically encode has been contentious, as M1 neurons are sensitive to a wide range of movement parameters, such as position, velocity, force, and torque, and their tuning properties are highly heterogeneous in terms of kinematic versus kinetic features, joint versus extrinsic coordinates, etc.[26–30]. At the extreme, Churchland et al. [31] proposed that the preparatory activity of M1 neurons exist not to represent specific movement features but to initialize a dynamical system whose evolution will produce movement activity. According to this view, there are no intrinsic variables in M1.

Although movement parameter coding in the parietal cortex has not been examined as extensively as in M1, a few studies that directly compared neuronal activity between the parietal area 5 and M1 observed that area 5 neurons were much less sensitive to kinetic variables, such as torque and force, than M1 neurons[32, 33]. PRR appears to represent movements at an even more abstract level than area 5, as it encodes the static reach target more strongly than the dynamically changing reach direction during movement compared to area 5[16]. Furthermore, PRR encodes the spatial goal locations largely in visual coordinates[13, 34, 35], suggesting that the spatial reference frame used by PRR is simpler than M1. These differences between M1 and PRR might lead to different learning mechanisms. Future studies directly comparing the intrinsic variables and the dominant learning mechanisms among different brain areas will provide valuable information on the effective design strategies for BMIs in each area[9].

Limitations of the current study

The current study cannot rule out the possibility that individual-neuron learning may play a role under different experimental settings (e.g., longer training periods than we tested, or different decoders which the intrinsic-variable mechanism cannot possibly learn). Testing decoders that the intrinsic-variable mechanism cannot learn is an interesting topic for future study although it is a formidable task given that such an experiment requires complete knowledge of the intrinsic response repertoire of a cortical area.

We examined learning involving up to 2 trained neurons, fewer than many practical BMI applications would use. An important topic for future study is the examination of learning involving more trained neurons, i.e., more degrees of freedom.

Finally, our premise that untrained neurons would not exhibit activity change on average if neurons can adjust their activity independently might appear too strong an idealization of individual-neuron learning. However, some studies indeed proposed BMI learning models that meet this strict premise[8, 36]. Moreover, the strict premise is useful to test the emerging view that reinforcing arbitrary mappings between neural activity patterns and movements is an efficacious approach to facilitating BMI learning. Less strict premises

inevitably limit the capacity of individual-neuron learning such that not all arbitrary activity patterns are producible, making it less distinguishable from intrinsic-variable learning.

Experimental Procedures

The California Institute of Technology Institutional Animal Care and Use Committee approved the animal procedures used in this study, which were performed in accordance with NIH guidelines. Details of the behavioral tasks and neural recording procedures are described in the Supplemental Information.

Neural adaptation index (NAI)

The NAI was computed as follows: $\{1 - (FR_{anti,s1} - FR_{anti,s2})/(FR_{pro,s1} - FR_{pro,s2})\}/2$, where $FR_{anti,s1}$ denotes the mean firing rate in the delay period following stimulus 1 in the BMI-anti task, and the same notation applies to the other variables.

Linear discriminants in the BMI tasks

For the BMI-pro task using one trained neuron, the threshold dividing the high and low firing rates was computed as the maximum-likelihood classifier under the assumptions of uniform prior and Poisson distribution as follows: $x = (M_1 - M_2)/\log(M_1 / M_2)$, where x is the trained neuron firing rate and M_i is the mean firing rate for stimulus i during the reach task. The same threshold was used for the BMI-anti task, but the stimuli associated with the high versus low firing rates were flipped.

For the BMI-pro task using two trained neurons, the linear boundary dividing the two firing rate zones was computed as the maximum-likelihood classifier under the assumptions of uniform prior and independent Poisson distributions as follows: $y = \{\log(M_{21} / M_{11})/\log(M_{12}/M_{22})\} \cdot x + \{M_{11} + M_{12} - M_{21} - M_{22}\}/\log(M_{12} / M_{22})$, where y is the firing rate of the first neuron, x for the second neuron, and M_{ij} for the mean firing rate of neuron j for stimulus i during the reach task.

For the BMI-mix task, a linear boundary was computed in the same way as the BMI-pro task boundary, except the mean firing rates of the second neuron were flipped between the two stimuli as follows: $y = \{\log(M_{21} / M_{11})/\log(M_{22}/M_{12})\} \cdot x + \{M_{11} + M_{22} - M_{21} - M_{12}\}/\log(M_{22} / M_{12})$.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

1. Moritz CT, Perlmutter SI, Fetz EE. Direct control of paralysed muscles by cortical neurons. Nature. 2008; 456:639–642. [PubMed: 18923392]

2. Ganguly K, Carmena JM. Emergence of a stable cortical map for neuroprosthetic control. PLoS Biol. 2009; 7:e1000153. [PubMed: 19621062]

3. Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. Functional network reorganization during learning in a brain-computer interface paradigm. Proc. Natl. Acad. Sci. 2008; 105:19486–19491. [PubMed: 19047633]

4. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature. 2006; 442:164–171. [PubMed: 16838014]

5. Lebedev MA, Carmena JM, O'Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, Nicolelis MAL. Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. J. Neurosci. 2005; 25:4681–4693. [PubMed: 15888644]

6. Mulliken GH, Musallam S, Andersen RA. Decoding trajectories from posterior parietal cortex ensembles. J. Neurosci. 2008; 28:12913–12926. [PubMed: 19036985]

7. Fetz EE. Volitional control of neural activity: implications for brain-computer interfaces. J. Physiol. 2007; 579:571–579. [PubMed: 17234689]

8. Legenstein R, Chase SM, Schwartz AB, Maass W. A reward-modulated Hebbian learning rule can explain experimentally observed network reorganization in a brain control task. J. Neurosci. 2010; 30:8400–8410. [PubMed: 20573887]

9. Green AM, Kalaska JF. Learning to move machines with the mind. Trends in Neurosciences. 2010 In Press, Corrected Proof.

10. Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA. Cognitive control signals for neural prosthetics. Science. 2004; 305:258–262. [PubMed: 15247483]

11. Andersen RA, Hwang EJ, Mulliken GH. Cognitive neural prosthetics. Annu Rev Psychol. 2010; 61:169–190. C161–C163. [PubMed: 19575625]

12. Hwang, E.; Andersen, R. Cognitively driven brain machine control using neural signals in the parietal reach region; Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE; 2010. p. 3329-3332.

13. Pesaran B, Nelson MJ, Andersen RA. Dorsal premotor neurons encode the relative position of the hand, eye, goal during reach planning. Neuron. 2006; 51:125. [PubMed: 16815337]

14. Buneo CA, Jarvis MR, Batista AP, Andersen RA. Direct visuomotor transformations for reaching. 2002; 416:632.

15. Snyder LH, Batista AP, Andersen RA. Coding of intention in the posterior parietal cortex. Nature. 1997; 386:167–170. [PubMed: 9062187]

16. Mulliken GH, Musallam S, Andersen RA. Forward estimation of movement state in posterior parietal cortex. Proc. Natl Acad. Sci. U. 2008; 105:8170–8177. [PubMed: 18499800]

17. Gail A, Andersen RA. Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J. Neurosci. 2006; 26:9376–9384. [PubMed: 16971521]

18. Hwang EJ, Andersen RA. Spiking and LFP activity in PRR during symbolically instructed reaches. J Neurophysiol. 2011

19. Imamizu H, Uno Y, Kawato M. Internal representations of the motor apparatus: Implications from generalization in visuomotor learning. J. Exp. Psychol. 1995; 21:1174–1198.

20. Truccolo W, Friehs GM, Donoghue JP, Hochberg LR. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. J. Neurosci. 2008; 28:1163–1178. [PubMed: 18234894]

21. Moore T, Armstrong KM. Selective gating of visual signals by microstimulation of frontal cortex. Nature. 2003; 421:370–373. [PubMed: 12540901]

22. Schall JD. On the role of frontal eye field in guiding attention and saccades. Vision Research. 2004; 44:1453–1467. [PubMed: 15066404]

23. Schafer RJ, Moore T. Selective attention from voluntary control of neurons in prefrontal cortex. Science. 2011; 332:1568–1571. [PubMed: 21617042]

24. Buschman TJ, Miller EK. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science. 2007; 315:1860–1862. [PubMed: 17395832]

25. Ganguly K, Dimitrov DF, Wallis JD, Carmena JM. Reversible large-scale modification of cortical networks during neuroprosthetic control. Nat Neurosci. 2011; 14:662–667. [PubMed: 21499255]

26. Kalaska, JF. From intention to action: motor cortex and the control of reaching movements. In: Sternad, D., editor. Progress in Motor Control. Vol. Volume 629. Springer US: 2009. p. 139-178.

27. Ashe J, Georgopoulos AP. Movement Parameters and Neural Activity in Motor Cortex and Area 5. Cereb. Cortex. 1994; 4:590–600. [PubMed: 7703686]

28. Rokni U, Richardson AG, Bizzi E, Seung HS. Motor learning with unstable neural representations. Neuron. 2007; 54:653–666. [PubMed: 17521576]

29. Hatsopoulos, Nicholas G, Suminski, Aaron J. Sensing with the Motor Cortex. Neuron. 2011; 72:477–487. [PubMed: 22078507]

30. Cherian A, Krucoff MO, Miller LE. Motor cortical prediction of EMG: Evidence that a kinetic brain machine interface may be robust across altered movement dynamics. J. Neurophysiol. 2011

31. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? Neuron. 2010; 68:387–400. [PubMed: 21040842]

32. Hamel-Pâquet C, Sergio LE, Kalaska JF. Parietal Area 5 Activity Does Not Reflect the Differential Time-Course of Motor Output Kinetics During Arm-Reaching and Isometric-Force Tasks. J. Neurophysiol. 2006; 95:3353–3370. [PubMed: 16481461]

33. Kalaska JF, Cohen DA, Prud'homme M, Hyde ML. Parietal area 5 neuronal activity encodes movement kinematics, not movement dynamics. Exp. Brain. Res. 1990; 80:351–364. [PubMed: 2113482]

34. Wu W, Hatsopoulos N. Evidence against a single coordinate system representation in the motor cortex. Exp. Brain. Res. 2006; 175:197–210. [PubMed: 16775704]

35. Batista AP, Buneo CA, Snyder LH, Andersen RA. Reach plans in eyecentered coordinates. Science. 1999; 285:257–260. [PubMed: 10398603]

36. Heliot R, Ganguly K, Jimenez J, Carmena JM. Learning in Closed-Loop Brain Machine Interfaces: Modeling and Experimental Validation. Systems, Man, Cybernetics, Part B: Cybernetics, IEEE Transactions on 40. 2010:1387–1397.

Highlights

1. Monkeys learned to volitionally elicit arbitrarily assigned neural activity patterns

2. These seemingly arbitrary activity patterns can be elicited by planning natural movements

3. Untrained neurons indicated the same natural movements as trained neurons

4. Animals volitionally control neural activity by exploiting a natural motor repertoire
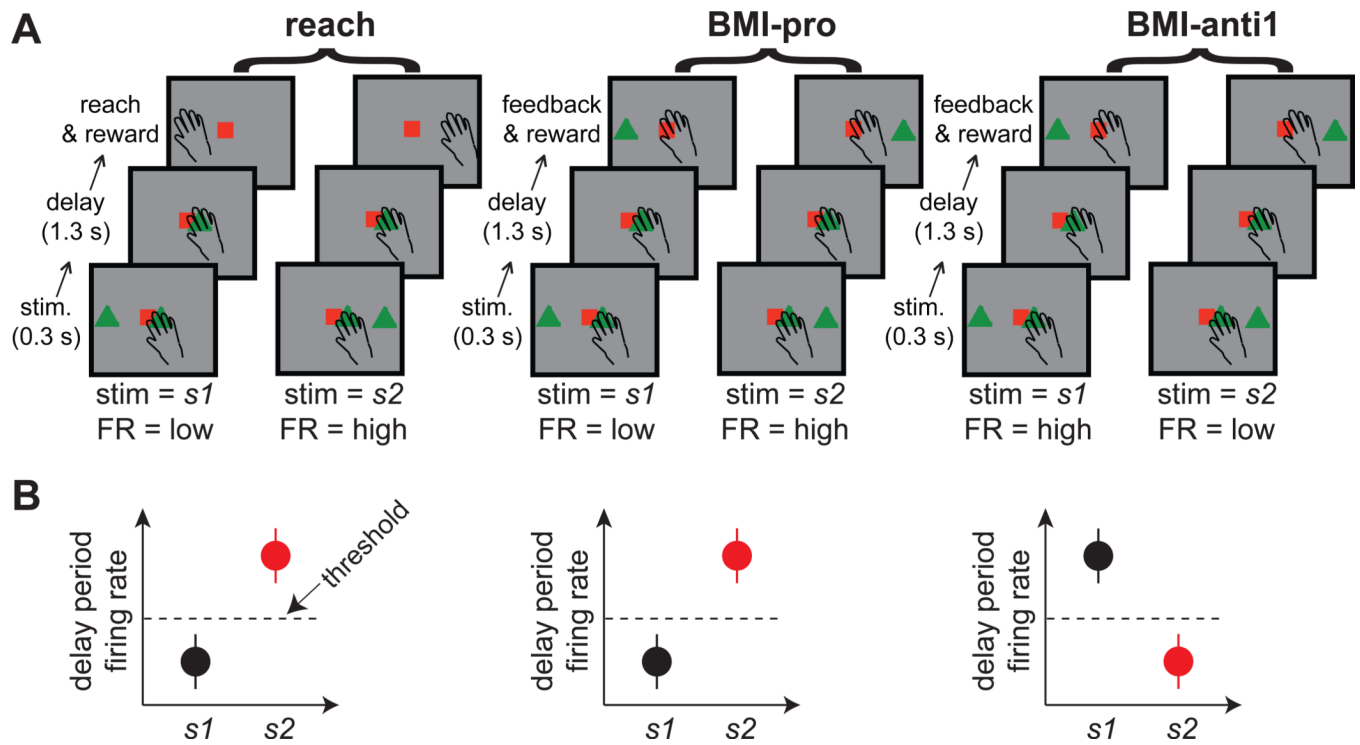
Figure 1. The task event sequence and the stimulus-response rule
A. The temporal event sequence in successful trials for two stimulus locations in the reach, BMI-pro, BMI-anti1 tasks. **B.** The activity pattern of a hypothetical neuron for successful trials in each of the three tasks.
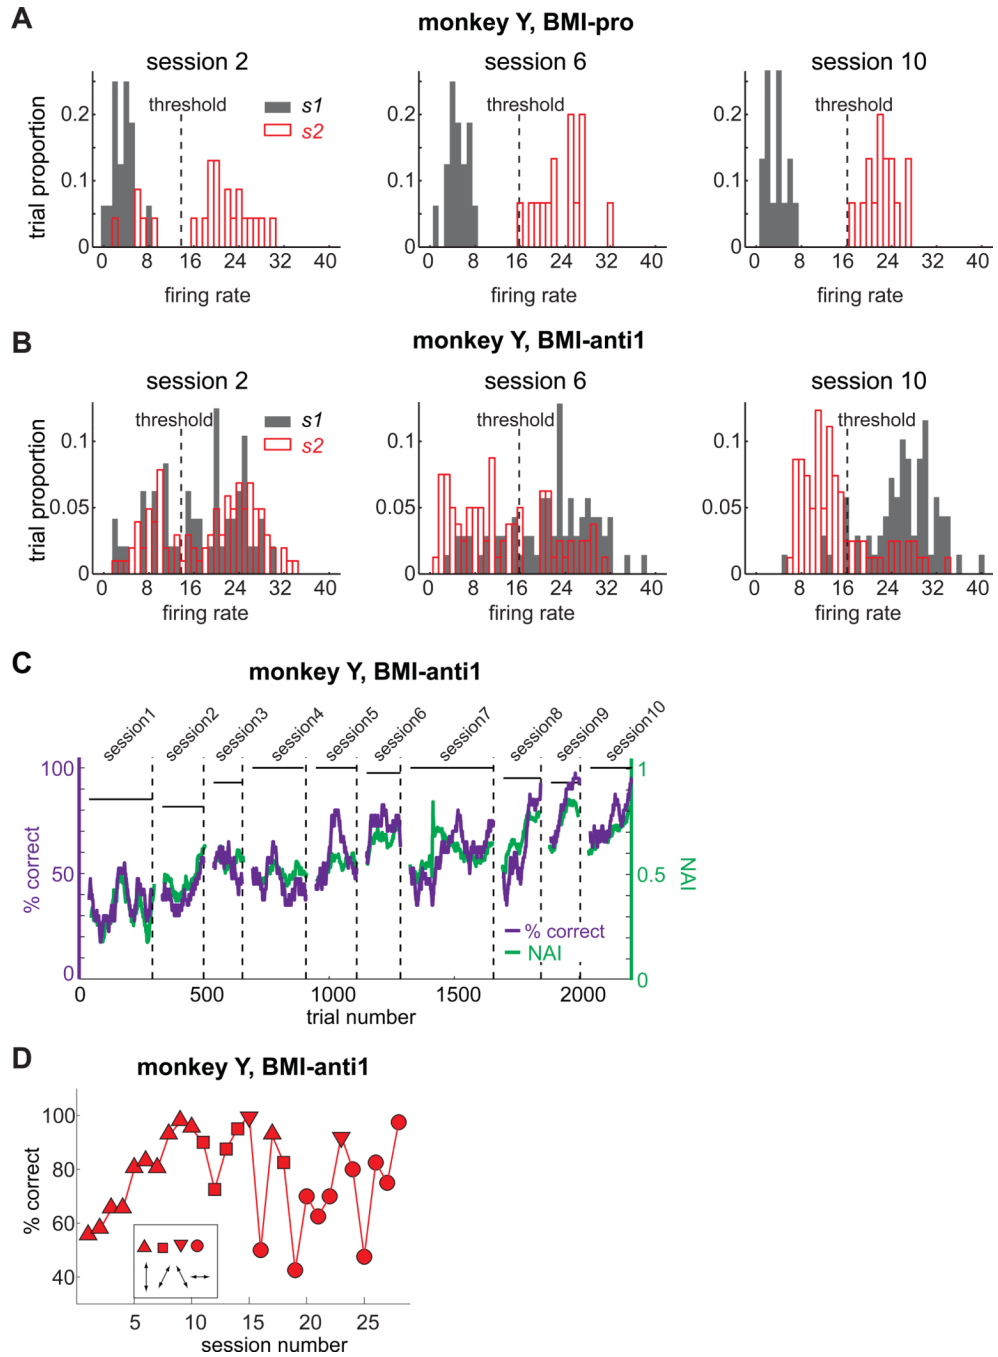
Figure 2. Monkey Y learned the BMI-anti1 task with long-term training over 10 days
A. The firing rate distributions of a single neuron for each of the two stimuli in the BMI-pro task of sessions 2, 6 and 10, recorded from monkey Y. **B.** The firing rate distributions of the same single neuron in the BMI-anti task. **C.** The % correct and neural adaptation index (NAI) in the BMI-anti1 task from the first 10 sessions. The dashed vertical lines indicate the end of each session. The horizontal bars indicate the % correct in the BMI-pro task in the corresponding sessions. **D.** The peak performance in each of the 28 BMI-anti1 task sessions for monkey Y. The different symbols indicate different stimulus pairs. The configuration for each stimulus pair is illustrated in the inset. See also Figure S1.
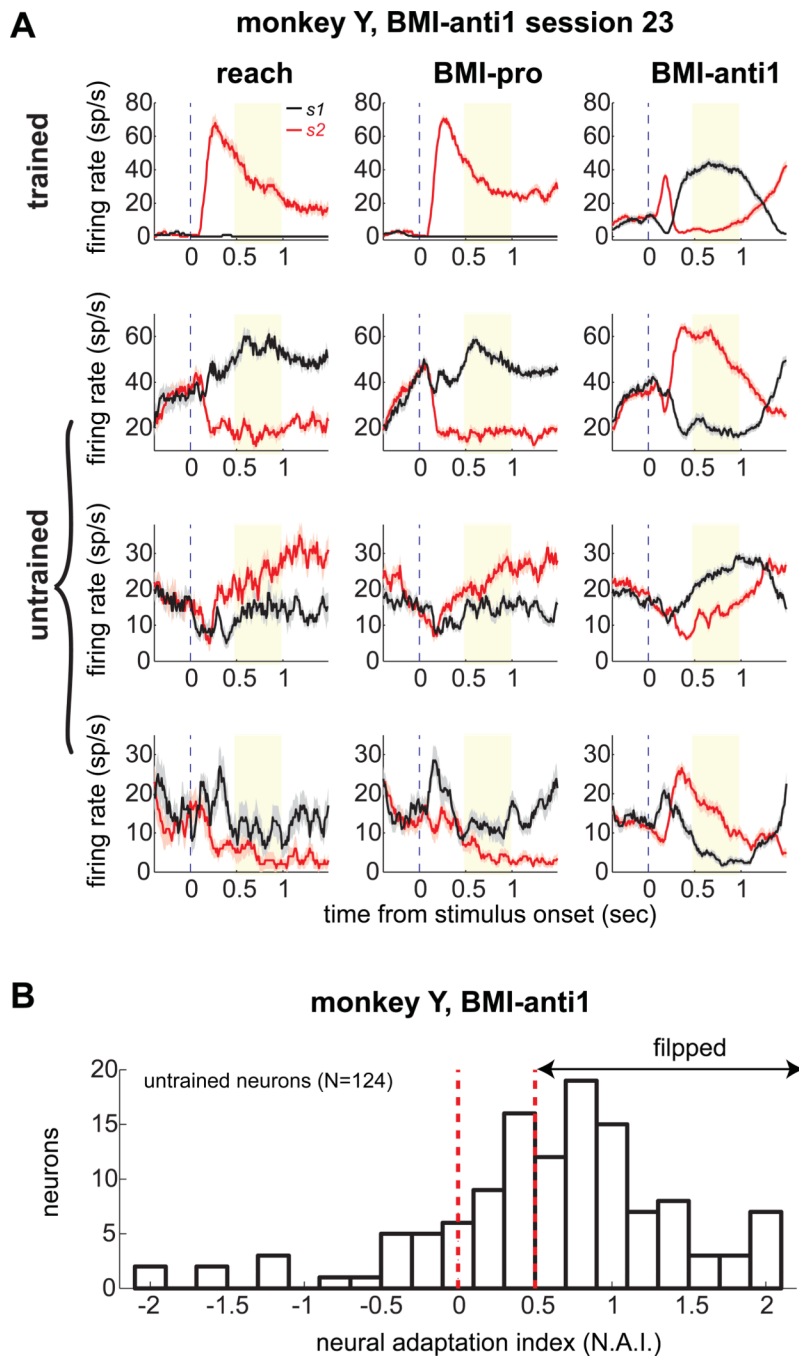
Figure 3. Untrained neurons flip the preferred stimulus in the BMI-anti task1
A. The temporal dynamics of the firing rates (mean±s.e.m.) for a trained neuron that directly contributed to the BMI output and 3 untrained neurons in session 23 for monkey Y. **B.** The distribution of the neural adaptation index of untrained neurons (N=124) in monkey Y for all successful trials.
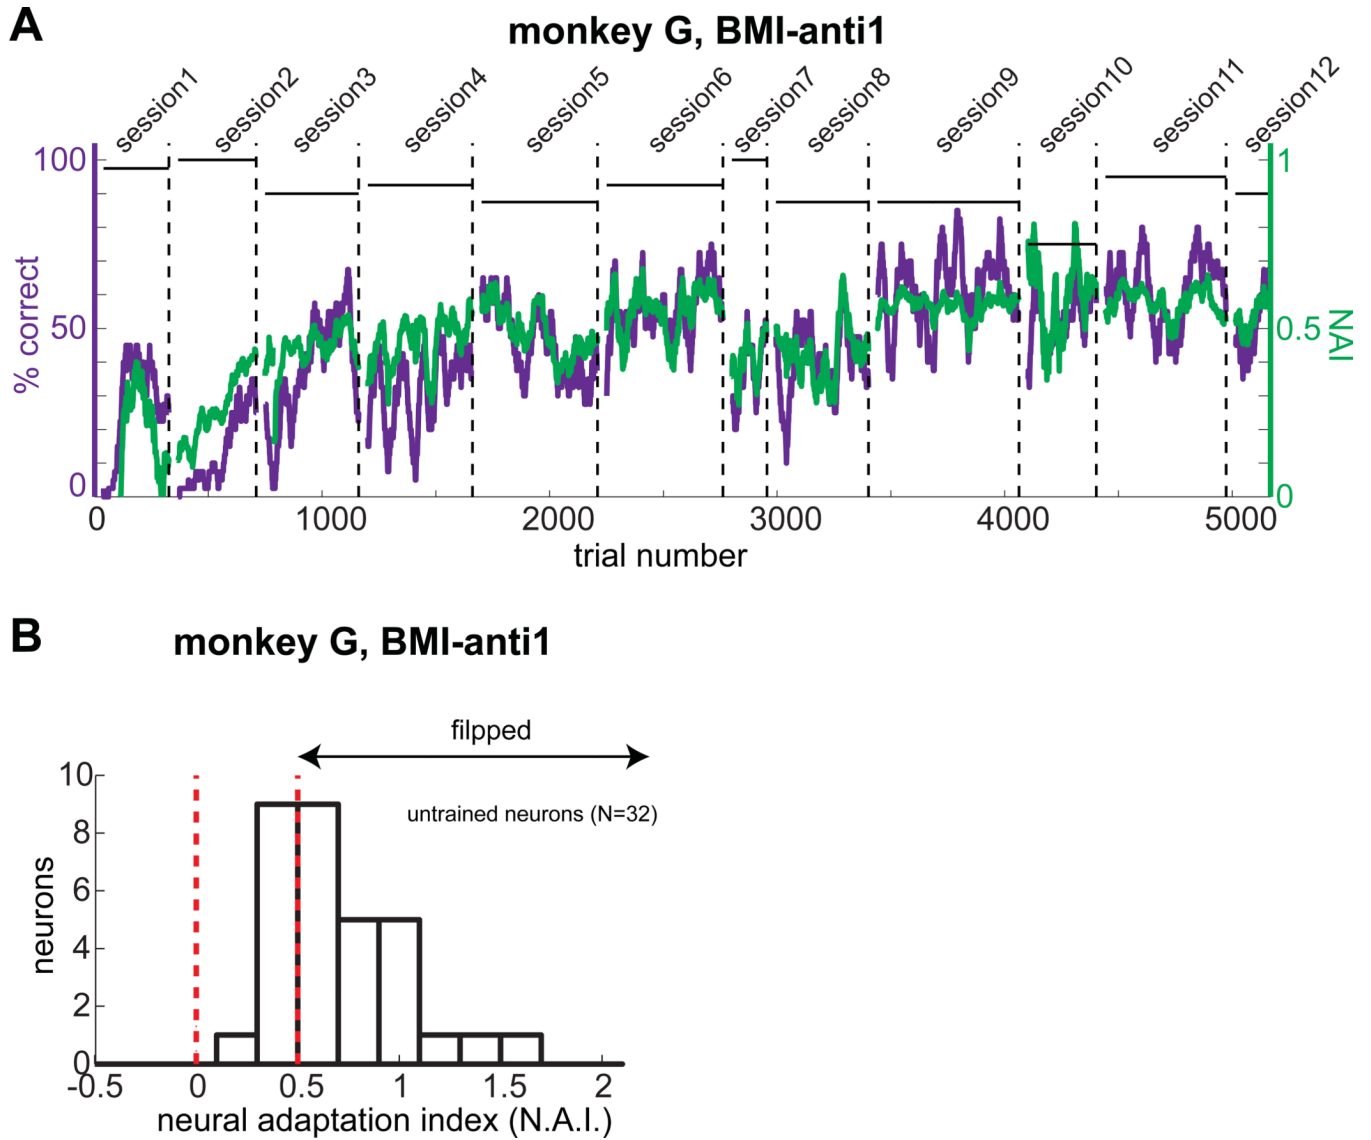
A

monkey G, BMI-anti1

B

monkey G, BMI-anti1

Figure 4. Trained and untrained neurons fluctuate their activity together regardless of performance level

A. The % correct and neural adaptation index in the BMI-anti1 task training sessions from Monkey G. The dashed vertical lines indicate the end of each of the 12 daily sessions. The horizontal bars indicate the % correct in the BMI-pro task in the corresponding sessions. **B.** The distribution of the neural adaptation index of untrained neurons (N=32) for all successful trials of BMI-anti1 task training from monkey G.
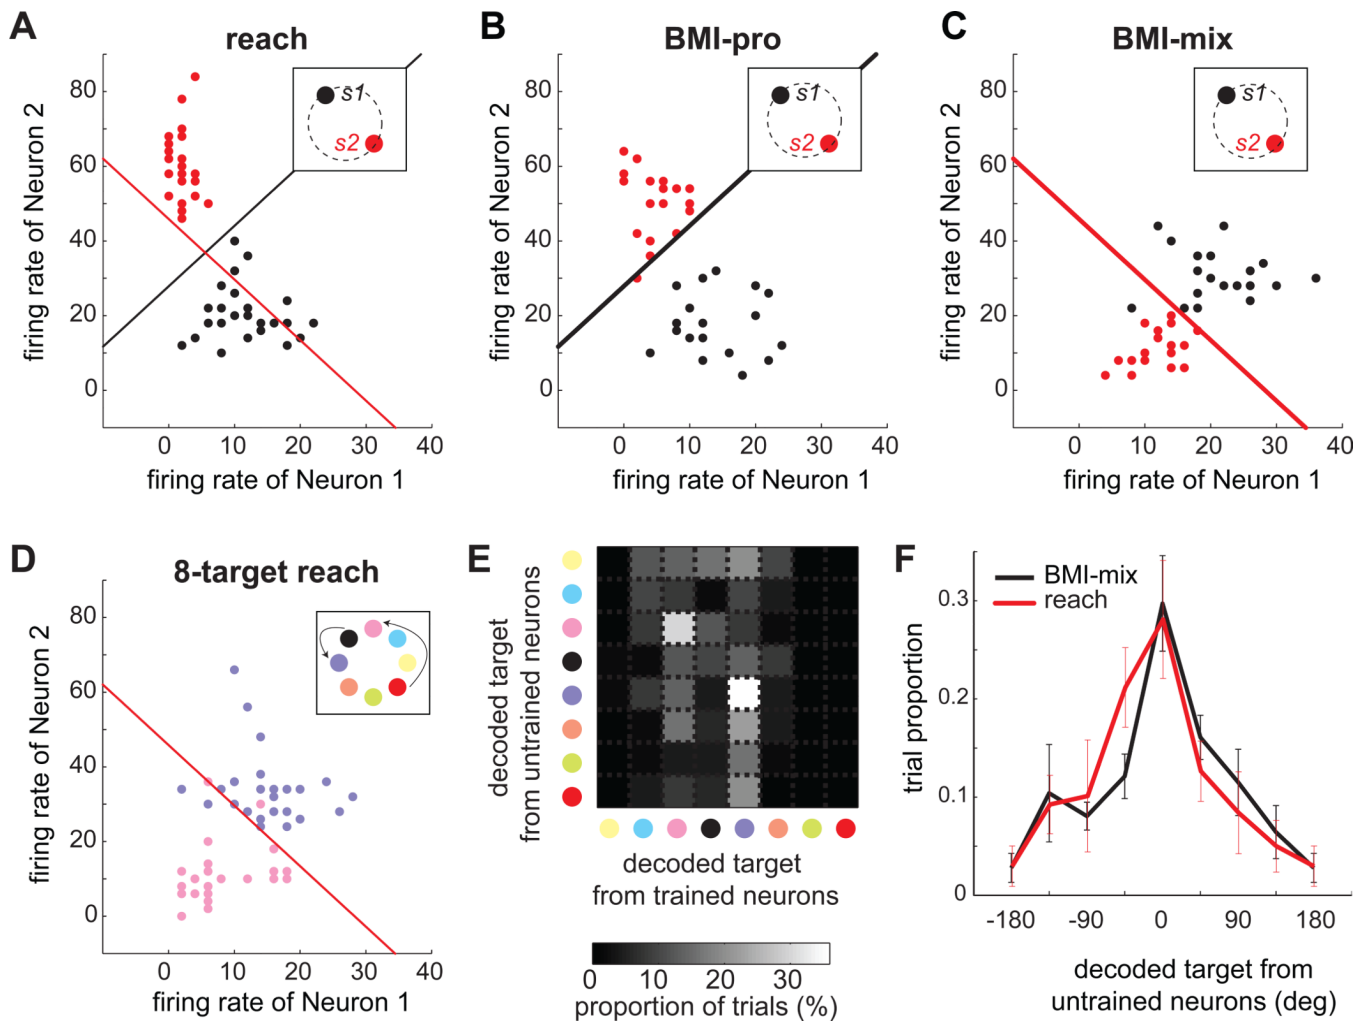
Figure 5. Intrinsic-variable learning is not task-dependent
A–C. The firing rates for the two trained neurons in response to the two stimuli (color coded) in the reach, BMI-pro, and BMI-mix tasks from an example session. The black and red lines represent the linear boundaries used for decoding in the BMI-pro and BMI-mix tasks, respectively. The insets show the location of the two stimuli. **D.** The firing rates of the same two neurons for two reach targets (purple and pink in the inset; the arrows indicate the relationship between the stimulus and matched target locations) in the 8-target reach task. The matching target for stimulus 1 was rotated 45° counterclockwise relative to the actual stimulus location, while the matching target for stimulus 2 was rotated 135° counterclockwise. **E.** The x-axis represents the target decoded from the trained neurons and the y-axis from the untrained neurons in the example session shown in **C**. The proportion of trials is shown in grayscale. **F.** The black line shows the probability distribution 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 ±s.e.m., across 9 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 shows the probability distribution of the difference between the targets decoded from the trained and untrained neurons in the 2-target reach task when the target decoded from the trained neurons is the reach target.