Article

Neuron

Intrinsic Variable Learning for Brain-Machine Interface Control by Human Anterior Intraparietal Cortex

Highlights

- AIP neurons learn to modulate their activity to compensate for errors in BMI tasks
- Changes in the neural activity reflect a cognitive re-adaptation mechanism
- AIP fails to compensate for errors when novel neural activity patterns are required
- Learning in AIP is constrained by the pre-existing neuronal structure

Authors

Sofia Sakellaridi, Vassilios N. Christopoulos, Tyson Aflalo, ..., Debra Ouellette, Nader Pouratian, Richard A. Andersen

Correspondence

sofia@caltech.edu

In Brief

Sakellaridi, Christopoulos, et al. studied the learning mechanism in human AIP using brain-machine interface paradigms. They found that changes in neural activity during learning reflect cognitive re-adaptation mechanisms. When cognitive strategies were not adequate for compensation, AIP failed to learn.

Sakellaridi et al., 2019, Neuron *102*, 694–705
May 8, 2019
https://doi.org/10.1016/j.neuron.2019.02.012

CellPress

CellPress

Neuron
Article

Intrinsic Variable Learning for Brain-Machine Interface Control by Human Anterior Intraparietal Cortex

Sofia Sakellaridi,[1,2,5,6,*] Vassilios N. Christopoulos,[1,2,5] Tyson Aflalo,[1,2] Kelsie W. Pejsa,[1,2] Emily R. Rosario,[3] Debra Ouellette,[3] Nader Pouratian,[4] and Richard A. Andersen[1,2]

[1]Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA 91125, USA
[2]Tianqiao and Chrissy Chen Brain-Machine Interface Center, Chen Institute for Neuroscience, California Institute of Technology, Pasadena, CA 91125, USA
[3]Casa Colina Hospital and Centers for Healthcare, Pomona, CA 91767, USA
[4]Department of Neurosurgery, Interdepartmental Program in Neuroscience, and Brain Research Institute, David Geffen School of Medicine at UCLA, Los Angeles, CA 90095, USA
[5]These authors contributed equally
[6]Lead Contact
*Correspondence: sofia@caltech.edu
https://doi.org/10.1016/j.neuron.2019.02.012

SUMMARY

Although animal studies provided significant insights in understanding the neural basis of learning and adaptation, they often cannot dissociate between different learning mechanisms due to the lack of verbal communication. To overcome this limitation, we examined the mechanisms of learning and its limits in a human intracortical brain-machine interface (BMI) paradigm. A tetraplegic participant controlled a 2D computer cursor by modulating single-neuron activity in the anterior intraparietal area (AIP). By perturbing the neuron-to-movement mapping, the participant learned to modulate the activity of the recorded neurons to solve the perturbations by adopting a target re-aiming strategy. However, when no cognitive strategies were adequate to produce correct responses, AIP failed to adapt to perturbations. These findings suggest that learning is constrained by the pre-existing neuronal structure, although it is possible that AIP needs more training time to learn to generate novel activity patterns when cognitive re-adaptation fails to solve the perturbations.

INTRODUCTION

The brain-machine interface (BMI) is an emerging technology that translates brain activity into motor actions. Besides being a powerful tool that enables people with paralysis to control external devices with their thoughts, the BMI has also shed light on the neural representation of cognitive functions, such as learning and motor adaptation (Armenta Salas and Helms Tillery, 2016; Carmena, 2013; Chase et al., 2012; Ganguly and Carmena, 2010; Golub et al., 2018; Hwang et al., 2013; Koralek et al., 2012; Sadtler et al., 2014; Wander et al., 2013). The main advantage of BMIs for studying learning is that brain plasticity associated with learning and adaptation can be explored within a particular node in a network, without taking into account the biomechanics of the limb or networks at lower levels, such as the spinal cord (Carmena et al., 2003). Although recent BMI studies have showed that behavioral performance improves with practice, there is no strong consensus on the nature of learning and its constraints in humans (Aflalo et al., 2015; Collinger et al., 2013; Hochberg et al., 2006, 2012; Truccolo et al., 2008; Wang et al., 2013). The two prevailing mechanisms that have been considered to account for changes in neuronal activity associated with learning are the individual-neuron mechanism and the intrinsic-variable mechanism. The first one posits that entirely new patterns of neural activity can be learned by individual neurons or populations of neurons (Cerf et al., 2010; Chapin et al., 1999; Fetz, 1969, 2007; Fetz and Baker, 1973; Fetz and Finocchio, 1971; Gage et al., 2005). In groundbreaking research, Fetz and colleagues provided evidence that non-human primates (NHPs) could modulate single motor cortex neurons or even pairs of neurons in opposite directions (Fetz, 1969, 2007; Fetz and Baker, 1973). Moritz and colleagues further showed that pairs of motor cortex neurons could arbitrarily change their tuning to control stimulation of antagonist muscles in NHP limbs in a limb movement task (Moritz et al., 2008). Although these findings suggest that certain brain areas can generate novel and arbitrary patterns of activity, an equally plausible interpretation is that animals generate only existing patterns of activity within new contexts (Chase et al., 2010; Golub et al., 2018; Hwang et al., 2013; Jarosiewicz et al., 2008). These patterns of activity may be constrained by the pre-existing structure of the network, and animals learn to manipulate an intrinsic variable of the natural movement, such as target direction, to generate activity patterns that comply with the new environmental demands (Jarosiewicz et al., 2008; Paz et al., 2003; Paz and Vaadia, 2004; Wise et al., 1998). This mechanism prevents independent adaptation of individual neurons, because the cognitive strategy

influences a global network of neurons that are sensitive to the intrinsic variable. Therefore, it is still unclear whether changes in neural activity that accompany learning reflect individual-neuron or intrinsic-variable learning.

To dissociate between these two learning mechanisms, we need to explore whether brain areas are capable of generating activity patterns that cannot be associated with any possible natural movement. The lack of any introspective description of strategies in animal BMI models of learning make the dissociation between the two mechanisms at best correlative and inferred. For instance, given that animals cannot be verbally instructed on which effector to use to perform intended movements (e.g., thinking of using only the right wrist) in a BMI task, it would not be possible to rule out that animals learn to produce the correct responses by using cognitive strategies to explore portions of the intrinsic manifold that were not identified in the baseline portions of the task (e.g., generating unexplored patterns by thinking about wiggling the toes) rather than generating truly novel patterns of activity. These limitations reveal the importance of using human subjects, who can be instructed to follow specific rules during BMI tasks and can report whether cognitive strategies were adopted to produce the correct responses. In the current study, a human participant with tetraplegia (C3-C4 complete lesion; 6 years post-injury) was implanted with a microelectrode array in the anterior intraparietal area (AIP) and learned to modulate the activity of a *single* neuron to control a 2D cursor in a goal-decoding BMI center-out task with two peripheral stimuli. The participant was instructed to attempt intended movements using only the right wrist and to not switch effectors during the task. She was also instructed to verbally report the direction of the intended movement after each trial in addition to any other strategies she employed to solve the task. Only the start and the end position of the cursor were presented to the participant. In a BMI-pro task, the neuron's intrinsic coding for desired direction of wrist movements was used to decode the desired goal. In a BMI-anti task, the neuron was trained to switch its activity so that it was active for desired movements opposite the direction of the goal stimulus. Besides the verbal feedback from the participant, we also looked at the behavior of neurons that were not used for decoding but were found to be selective for wrist movement and therefore were part of the pre-existing structure that plans wrist movements. These neurons are referred to as untrained neurons, as opposed to trained neurons that are used for decoding and, therefore, directly contribute to the BMI output. If the untrained neurons do not change their turning in the BMI-anti task, then the individual neuron mechanism predominates; if they flip their preferred directions, then intrinsic variable learning predominates. In favor of the intrinsic variable mechanism, we found that both trained and untrained neurons changed their tuning in the BMI-anti task. This change reflected an explicit strategic choice of the subject, as the participant reported that she solved the task by attempting intended wrist movements in the direction opposite to the stimulus (i.e., re-aiming strategy).

We further tested whether the learning mechanism in AIP varies with the complexity of the BMI task. In a second BMI perturbation experiment, the participant had to control the cursor by simultaneously modulating the activity of two trained neurons. Similar to the first experiment, the activity patterns of the trained and untrained neurons, as well as the verbal feedback from the participant, provide evidence in favor of the intrinsic variable learning mechanism. These findings suggest that the intrinsic variable learning mechanism predominates in AIP regardless of the task complexity. However, it could be argued that these tasks favor the intrinsic variable learning over the individual neuron learning, because there is always a cognitive solution that can produce the correct responses. We designed a BMI perturbation task in which no cognitive strategy could produce the neural activity patterns necessary to solve the task. In this case, an individual neuron mechanism was required. The participant could not produce the correct responses, indicating that AIP neurons were not capable of generating novel and arbitrary patterns of activity to solve the task. This suggests that either AIP neuron can only engage pre-existing structures of the neural network for short-term learning (about 1 h) or a much more extensive period of training may be required if AIP neurons can learn to generate entirely new neuronal activity patterns.

RESULTS

Volitional Control of Single Neurons in a Goal-Decoding BMI Task

We designed a BMI center-out task in which a participant with tetraplegia intended wrist movements toward one of two stimuli presented at diametrically opposing locations (i.e., up-down; Figure 1A). Each session consisted of three task blocks in the following order: BMI-calibration; BMI-pro; and BMI-anti. BMI-calibration trials started with a peripheral cue randomly presented for 0.8 s at one of the two potential locations on the screen. Following the cue offset, a delay period of 1.5 s ensued. After the go signal, the participant had 2 s to attempt a wrist movement toward the direction of the stimulus location. The only verbal instruction given prior to the beginning of each session was "attempt a wrist movement to the stimulus location after the go signal." The average firing rate activity of the recorded neurons after the go signal in the BMI-calibration was used to identify a single neuron that best decodes the stimulus location—i.e., a neuron that fires more spikes when intended wrist movements are made to one stimulus over the other. We refer to this neuron as trained, because it was used for decoding and therefore directly contributed to the BMI output.

The BMI-pro task had a similar sequence to the calibration, with an extra feedback period following the go period. The stimulus feedback location was decoded from the firing rate of the trained neuron during the go period, according to the following stimulus response rule. If the firing rate of the trained neuron was higher for stimulus S_1 than for stimulus S_2 in the calibration, the BMI-pro rule was that higher firing rates for stimulus S_1 than for S_2 would result in correct trials—i.e., the participant should increase the activity of the neuron in response to S_1 and suppress the activity of the neuron in response to S_2 (Figure 1B). We used linear discriminant analysis (LDA) to identify the firing rate threshold that divides the high and the low firing rates (see STAR Methods section for more details). If the firing rate conformed to the stimulus response rule, the stimulus feedback location was placed at the same location as the stimulus location

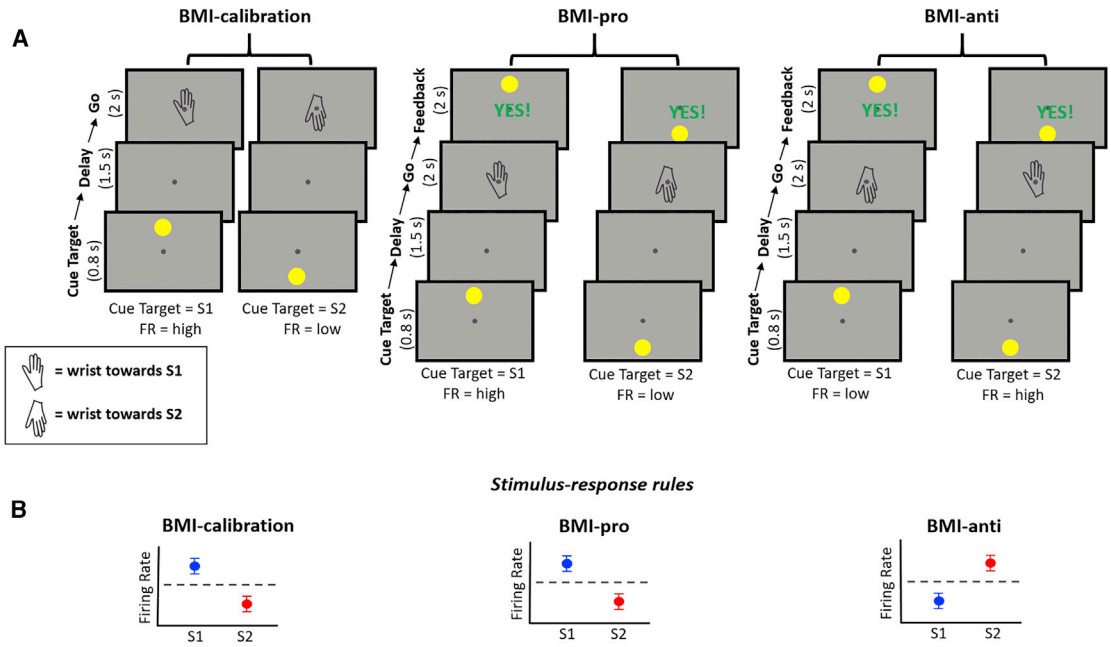

Figure 1. Task Event Sequences and Stimulus-Response Rules of the Brain-Machine Interface Learning Experiment with a Single Neuron

(A) A graphical representation of the brain-machine interface (BMI) learning paradigm with a single neuron.

(B) The stimulus response rules for the BMI-calibration (left), the BMI-pro task (middle), and BMI-anti task (right). Each panel illustrates the firing rate (mean ± SEM) of a hypothetical trained neuron for successful trials in each of the three tasks.

in the cue target period and a "yes" message was displayed on the screen, indicating a successful trial. Otherwise, the stimulus feedback location was placed opposite to the cue stimulus and a "no" message was presented, indicating an unsuccessful trial. The BMI-anti task enforced the opposite stimulus response rule of lower firing rate for stimulus S_1 than for S_2. This rule forced the trained neuron to flip its preferred stimulus associated with higher firing rates, between the BMI-pro and the BMI-anti trials. In the BMI-anti task, the stimulus feedback location was placed opposite to the intended wrist direction for successful trials and at the same location otherwise.

The participant performed 8 experimental sessions in total, each on different days. For each session, the neuron that best discriminated between the two stimuli was selected as the "trained" neuron. We did not attempt to maintain the same neuron across sessions, in part, because we found that neurons recorded on a channel would change from day to day. The activities of two single trained neurons recorded in the 1st and 3rd sessions are illustrated in Figures 2A and 2B, respectively. In the BMI-pro task, the firing rates of the trained neurons were properly discriminated between the two stimuli in both sessions (Figures 2A and 2B, left column), because the participant intended movements in the direction of the displayed stimulus. However, in the BMI-anti task, the trained neuron of the 1st session was not correctly discriminated, indicating that the participant was not able to modulate the activity of the trained neuron to conform to the new stimulus response rule (Figure 2A, middle and right panels). According to the verbal report, the participant was still mostly performing intended wrist movements to the direction of the presented stimulus, even though she was also exploring some new strategies, such as intending faster wrist movements. During the 3rd session, and particularly at the end of the 1st block of the BMI-anti task, the participant found and reported that the best strategy to solve the task was to intend wrist movements in a direction opposite to the stimulus goal (anti-wrist movements). At this time, the firing rate of the trained neuron was able to properly discriminate by producing the opposite pattern from the BMI-pro task (Figure 2B, right panel). To confirm that the activity of the trained neuron was modulated by anti-wrist movements, we looked at the activity of the untrained neurons—i.e., neurons that do not participate in the BMI control but are part of the pre-existing structure that plans the wrist movements. If the trained neuron modulates its activity independently from the pre-existing structure (i.e., individual-neuron mechanism), the untrained neurons will not flip their preferred stimulus, similar to the trained neuron. On the other hand, if the participant learns a new cognitive strategy (i.e., attempting anti-wrist movements), the untrained neurons will also flip their preferred stimulus (i.e., intrinsic variable learning mechanism). Figures 3A and 3B depict the firing rates of one trained and two untrained neurons, respectively, from a BMI-pro (top row) and BMI-anti (bottom row) task from a typical session (i.e., neurons recorded from a single day). Consistent with intrinsic variable learning, the trained and untrained neurons flipped their preferred stimuli in the BMI-anti task, verifying that the participant aimed in the opposite direction to the stimulus. The average performance in the BMI-pro task across all 8 sessions was 0.89 ± 0.11 (mean ± SD). For the BMI-anti task, in the first 2 sessions and the first block of the 3rd session, the average

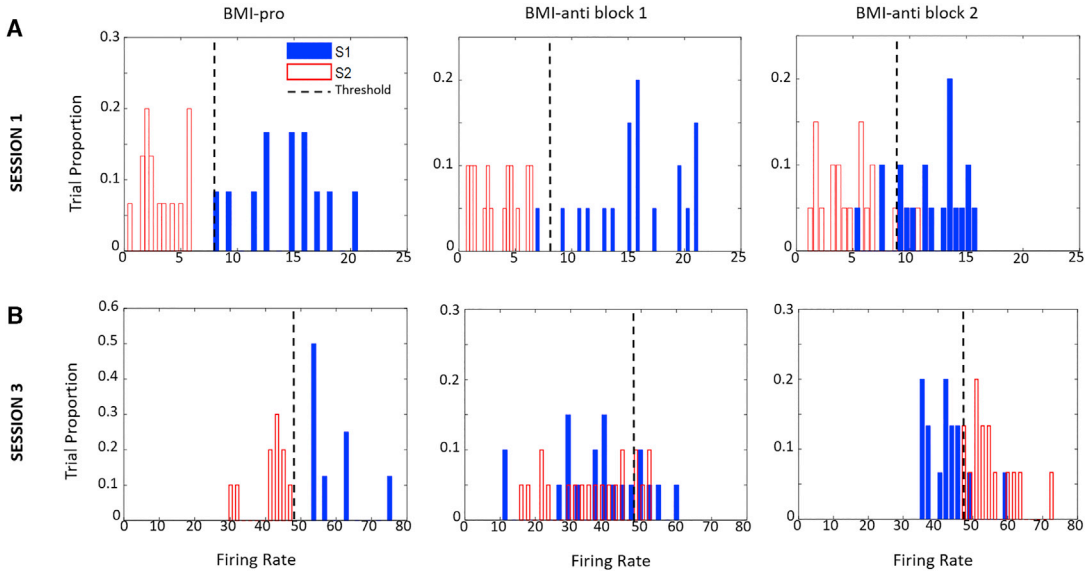

Figure 2. Firing Rate Distributions of Single Neurons in BMI-Pro and BMI-Anti Tasks

(A and B) The firing-rate distributions of two different single trained neurons for each of the two stimuli in the BMI-pro and BMI-anti task blocks in (A) session 1 and (B) session 3 for both correct and incorrect trials.

performance was 0.21 ± 0.16. After adopting the anti-wrist movement strategy, the performance increased to 0.84 ± 0.09 (the last block of the 3rd session + the last 5 sessions). Moreover, 91 untrained cells (out of a total of 281 recorded untrained cells across the last 5 sessions and the last block of the 3rd session) showed activity that discriminated between S_1 and S_2 during the go period of the BMI-pro task. Eighty of those 91 cells (~88%) flipped their preferred stimulus between BMI-pro and BMI-anti tasks during the go period (two-tailed t test $p < 0.05$; Figure 3C). The average activity of the untrained neurons, whose activity flipped between the BMI-pro and BMI-anti task, did not change between these two tasks (two-tailed t test $p > 0.05$ for all the untrained neurons that flipped their preferred stimulus). The remaining 11 units (~12%) did not flip their preferred stimulus.

Learning Mechanisms and Task Complexity

Although the single-neuron BMI experiment provides evidence in favor of the intrinsic variable learning mechanism, it could be hypothesized that individual neuron learning would be more likely to be pursued for more cognitively complex tasks. In a second experiment, we explored whether the preferred learning mechanism varies with the task complexity. We designed a BMI-fsb (fsb, feasible) task in which the cursor location was controlled by the activity of two trained neurons at the same time. We refer to this task as feasible, because it is still possible for the participant to solve it by adopting a new cognitive strategy. Each session consisted of three different task blocks in the following order: BMI-calibration; BMI-pro; and BMI-fsb (Figure 4A). At each stage, there were three epochs, cue target, delay, and go, similar to the single-neuron BMI experiment. In the BMI-calibration, the participant attempted wrist movements to eight radially arranged targets. Based on the tuning properties of two trained neurons in calibration, we selected a pair of stimulus locations (out of $\binom{8}{2} = 28$ possible pairs of stimuli) to be presented in the BMI-pro and the BMI-fsb tasks. The linear boundary in the BMI-pro task was the one that *best* separated the two firing-rate clusters, each formed by the activity of the two trained neurons during the go period, for each stimulus in the BMI-calibration. On the other hand, we selected the linear boundary for the BMI-fsb task under two constraints: (1) it provides about 50% classification accuracy for both stimuli and (2) it best separates the two clusters of another pair of targets. Therefore, the participant could either alter the cognitive strategy (i.e., re-aiming to directions away from the presented stimuli) or produce new activity patterns through independent adaptation of the trained neurons to match with the new stimulus response rule. Both the BMI-pro and BMI-fsb tasks had the same sequence as the BMI-calibration with an extra feedback period following the go signal. A simplified, conceptual illustration of the stimulus response rules in the two BMI tasks for the pair of stimuli (1, 5) is shown in Figure 4B. The linear threshold (black continuous line) in the BMI-pro task accurately separates the two firing-rate clusters. On the other hand, the BMI-fsb threshold (black discontinuous line) divides the two clusters such that half of the trials lie above and half below the threshold for both stimuli. However, it accurately separates the clusters for intended wrist movements toward the direction of targets 3 and 7 presented in the BMI-calibration. Therefore, the participant could either attempt wrist movements to the direction of 3 and 7 in response to stimuli 1 and 5, respectively, or she has to learn to generate activity patterns that conform to the stimulus response rule through independent adaptation of the two trained neurons.

The participant performed 15 experimental sessions in total, each on different days, with different pairs of stimuli and different

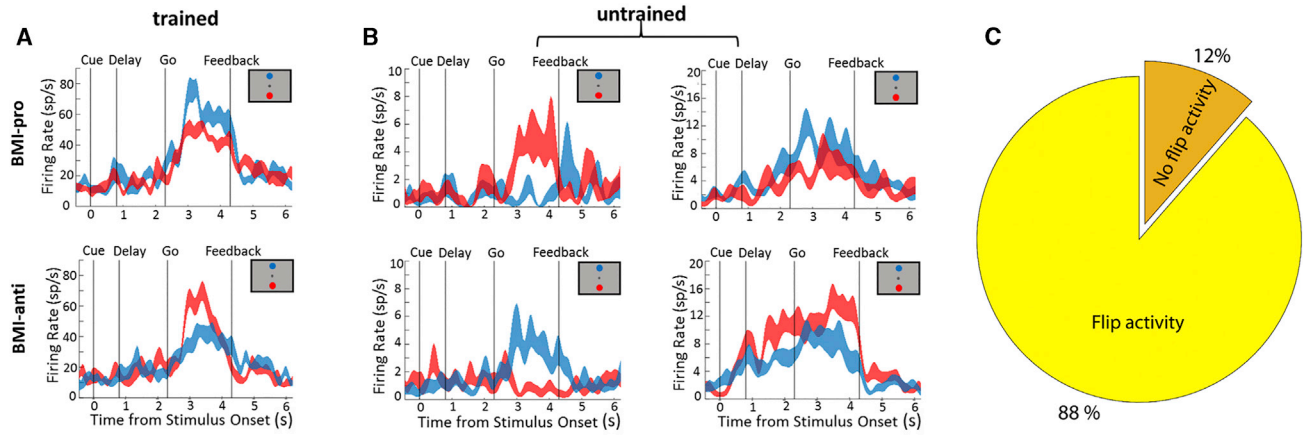

Figure 3. Trained and Untrained Neurons Flip Their Preferred Stimuli in the BMI-Anti Task

(A) The temporal dynamics of the firing rates for a trained AIP neuron that directly contributed to the BMI output in a typical session.

(B) Similar to (A) but for two untrained AIP neurons.

(C) The percentage of untrained neurons that flip their activity between BMI-pro and BMI-anti tasks after the participant adopted the anti-wrist movement strategy.

linear threshold boundaries that ranged in direction between 0° (i.e., horizontal line) and 180°. The average direction of the threshold in the BMI-pro task was 88.8° ± 59.2° (mean ± SD) and in the BMI-fsb task was 86.2° ± 55.8° across all 15 trials. The average angle between the BMI-pro and the BMI-fsb threshold was 95° ± 28.4°. In the BMI-pro task, the participant immediately produced the appropriate activity patterns for both trained neurons, achieving an average performance of 0.88 ± 0.12 (mean ± SD) across all 15 sessions. Within a timescale of a session, she also learned to generate activity patterns that comply with the stimulus response rule in the BMI-fsb task, achieving 0.78 ± 0.21 decode performance across all 15 sessions. A characteristic example is illustrated in Figure 5, in which the pair of stimuli (1, 4) was selected based on the tuning properties of the two trained neurons in BMI-calibration (Figure 5A). According to the stimulus response rule in the BMI-pro task, the averaged activity of the two trained neurons should fall above the linear threshold (black line) for stimulus 4 and below the threshold for stimulus 1 for successful trials. The task performance accuracy was over 95% in the BMI-pro task, in which intended wrist movements were planned to the direction of the presented stimuli (Figure 5B). In the BMI-fsb task, the linear threshold was selected so that it substantially deteriorated the performance for intended wrist movements to the presented stimuli (black discontinuous trace; Figure 5A). However, the participant learned to generate rule-complying activity patterns (Figure 5C), achieving performance greater than 90%. One way to modulate the neural activity of the trained neurons was to aim in different directions. We examined the behavior of the trained neurons during the BMI-calibration to determine whether planning movements to any of the tested directions could have produced activity patterns that comply with the stimulus response rule in the BMI-fsb task. The activity for stimulus 1 in the BMI-fsb task matched the activity for intended wrist movement toward stimulus 3, and the activity for stimulus 4 matched the activity for intended wrist movements to stimulus 7 (Figure 5D). This is consistent with the verbal feedback provided by the participant: she reported that she was mostly aiming in the direction of stimulus 3 (cyan) in response to stimulus 1 and toward stimulus 7 (red) in response to stimulus 4 (Figure 6A).

These findings suggest that the BMI-fsb task was solved by re-aiming intended wrist movements—a form of intrinsic variable learning. To confirm this hypothesis, we compared the activity of the two trained neurons with 20 untrained neurons, whose activity was significantly discriminated between the two stimuli, for each successful BMI-fsb trial. We used a nearest-neighbor decoding algorithm to infer the stimulus that the neuronal population encoded. The decoding algorithm selected the stimulus associated with the ensemble activity in the BMI-calibration that was closest to the ensemble activity of a given BMI-fsb trial in terms of the Euclidean distance. Figure 6B depicts the eight-target decoding results from the BMI-fsb task across all successful trials (correspondence matrix). The x and y axes represent the stimulus decoded by the two trained and untrained neurons, respectively. The correspondence matrix therefore allows us to visualize the degree of correspondence between the stimulus specified by the trained and untrained units. The two bright spots for stimuli 3 and 7 indicate that they were encoded concurrently for the trained and the untrained neurons in most of the trials. These two stimuli were the same two best-matching stimuli as inferred by the activity pattern (Figure 5D) and reported by the participant (Figure 6A). A more detailed analysis of how the activity distribution of the untrained neurons is modulated to counter the perturbation of the stimulus response rule in this particular example is presented in Figure S1. Overall, in about 50% of the BMI-fsb trials across all sessions, both trained and untrained neurons encoded the same matching target (49.6% ± 3.93%; mean ± SEM; across all 15 sessions; Figure 6C, blue curve). Similarly, in about 50% of the BMI calibration trials, both trained and untrained neurons encoded the same cued target (48.4% ± 4.95% trials across all 15 sessions; Figure 6C, red curve).

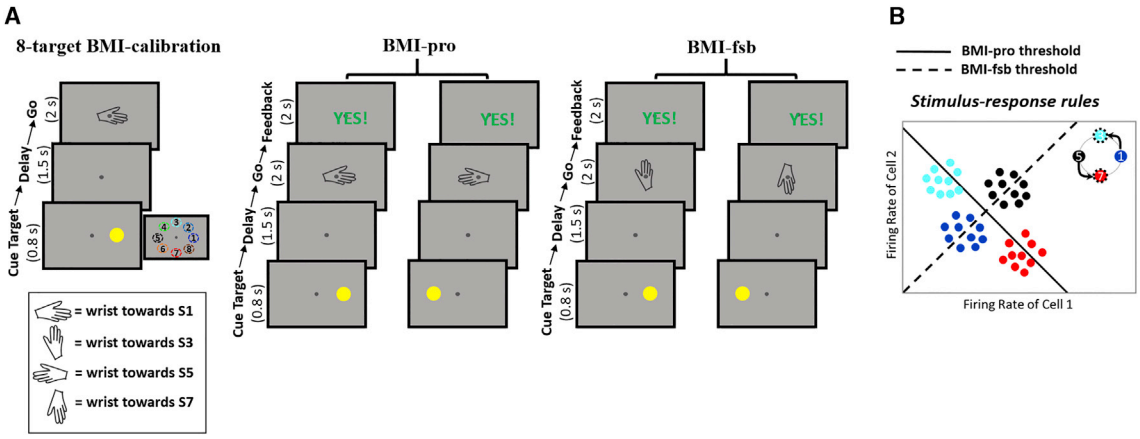

Figure 4. Task Event Sequences and Stimulus-Response Rules of the Brain-Machine Interface Learning Experiment with Two Neurons

(A) A graphical representation of the BMI learning paradigm with two neurons.

(B) A simplified, conceptual illustration of the stimulus response rules for the BMI-pro and BMI-fsb trials for the pair of stimuli 1 and 5. Points correspond to the hypothetical firing rates of the two trained neurons in response to four stimuli (color coded) in the BMI-calibration.

Learning Mechanisms in BMI Tasks with No Feasible Cognitive Strategies

So far, we explored the learning mechanisms for BMI tasks in which there was always a feasible intrinsic variable remapping strategy to compensate for perturbations. An interesting question is what would happen if no cognitive strategy was adequate to solve the BMI tasks. To address this question, we designed a BMI-nfsb task (nfsb, non-feasible) in which there was no strategy that could produce rule-complying activity patterns. We selected a linear threshold that performed poorly on classifying the firing-rate clusters of the two trained neurons for at least one of the two stimuli presented in the BMI-nfsb task. Ideally, we looked for a linear boundary that results in chance performance (50%) regardless of the intended movement direction. We used evolutionary algorithms (EAs) to identify such a linear threshold (see STAR Methods section for more details). EAs are adaptive heuristic algorithms that mimic the process of natural selection (Bandyopadhyay et al., 1995; de la Fraga and Coello Coello, 2011; Goldberg, 1989). As such, they represent a principled exploitation of a random search to solve complex optimization problems. The results show that, in the timescale of a session (i.e., about 1 h), neurons could not generate activity patterns that conform to the stimulus response rule in the BMI-nfsb task.

A characteristic example is depicted in Figure 7. Based on the tuning properties of the two trained neurons from the BMI-calibration (Figure 7A), we selected the pair of stimuli 2 and 7 in the BMI-pro and BMI-nfsb tasks. This pair was selected because there is a linear threshold (black solid line), named BMI-pro threshold, which accurately separates the two firing-rate clusters of the trained neurons—i.e., the firing rates are above this linear threshold for intended wrist movements to the stimulus 7 and below it for intended movements to the stimulus 2. We used this threshold in the BMI-pro task. On the other hand, the EAs generated a linear threshold for the BMI-nfsb task (black discontinuous line), in which the activity of the two trained neurons should fall within the right and the left side of this boundary in response to the stimuli 2 and 7, respectively. Figure 7B depicts the percentage of the calibration trials that fall into the right (blue) and the left (red) side of the BMI-nfsb threshold for intended wrist movements to all possible 8 targets. Notice that the threshold poorly separates the firing-rate clusters of the two trained neurons for intended wrist movements to the seven out of eight possible targets—except for target 2, in which the activity of the two trained neurons of all calibration trials falls into the right side of the threshold. On the other hand, the highest percentage (~65%) of trials that fall into the left side of the threshold corresponds to intended movements toward target 1. Therefore, when stimulus 7 is presented in the BMI-nfsb task, the best performance that the participant could achieve is about 65% by aiming toward the direction where target 1 was located in the BMI-calibration. This suggests that there is no cognitive strategy that can produce high performance accuracy in response to stimulus 7 in the BMI-nfsb task. Consequently, an individual neuron mechanism is required to compensate for changes on the stimulus response rule. On the other hand, when stimulus 2 is presented, the best performance (100%) could be achieved for intended wrist movements toward target 2, suggesting that the participant should not modify her strategy.

We found that, in the BMI-pro task, the participant immediately produced appropriate activity patterns for the two trained neurons to solve the task, achieving performance accuracy over 90% (Figure 7C). On the other hand, the two neurons could not learn to generate rule-complying activity patterns for trials in which stimulus 7 was presented in the BMI-nfsb task, achieving performance accuracy of about 50% (Figure 7D). The participant reported that, although she was re-aiming to different directions to compensate for the applied perturbation, none of them seemed to offset the perturbation and improve her performance when stimulus 7 was presented. Overall, we ran 6 sessions on different days, in which there was always a feasible intrinsic variable remapping strategy to acquire one stimulus (i.e., cognitively solvable stimulus), but not the other (i.e., cognitively non-solvable stimulus), after perturbing the stimulus response rule (i.e., BMI-nfsb task). Note that both stimuli could be achieved with

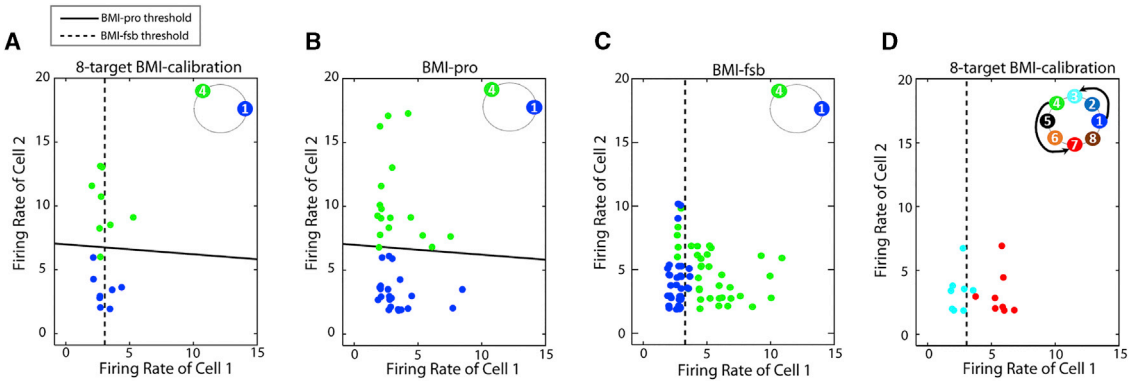

Figure 5. Evidence of Intrinsic-Variable Learning in the BMI-fsb Task

(A–C) The firing rates for the two trained neurons in response to the stimuli 1 and 4 (colored coded) in (A) the 8-target BMI-calibration, (B) BMI-pro, and (C) BMI-fsb tasks. The black solid and dashed lines represent the linear thresholds used for decoding in the BMI-pro and BMI-fsb tasks, respectively.

(D) The firing rates of the same two trained neurons in response to targets 3 (cyan) and 7 (red) in the BMI-calibration (the arrows indicate the relationship between the stimulus and matched target locations). The two stimuli (cyan and red) correspond to the matched target locations.

high performance in the BMI-pro task by aiming directly to their locations. We measured the performance as the number of trials in which the activity of the two trained neurons fell on the correct side of the boundary. The participant was able to produce the appropriate activity pattern for the pair of trained neurons in the BMI-pro and BMI-nfsb tasks for the cognitively solvable stimulus, achieving performance of 94.03% ± 2.65% (mean ± SEM) in the BMI-pro and 87.22% ± 3.0791% in the BMI-fsb tasks across 6 sessions (no significant differences in the performance between the two tasks for the cognitively solvable stimulus; two-tailed t test $p = 0.1324$; Figure 7E). However, she could not learn to produce rule-complying activity patterns for the cognitively non-solvable stimulus in the BMI-nfsb task, and her performance was significantly deteriorated from 88.90% ± 3.93% in the BMI-pro task to 50.0% ± 5.55% in the BMI-nfsb task across 6 sessions (two-tailed t test $p < 0.001$; Figure 7E). We interpret these findings as evidence that AIP neurons cannot learn to produce novel and arbitrary patterns of activity to comply with the stimulus response rule in the BMI-nfsb task. These results suggest that either there is a constraint on learning imposed by the existing neural structure of AIP or more extensive training is required to generate entirely novel neuronal activity patterns.

DISCUSSION

General

Although a large body of evidence suggests that both humans and animals are able to learn a variety of new tasks and adapt to unpredictable changes of the environment, to our knowledge, there have been no systematic studies of single-neuron learning in humans. The prevailing hypotheses suggest that neurons learn new tasks by either generating entirely new patterns of activity, independently of the pre-existing neuronal structure (individual neuron learning; Cerf et al., 2010; Chapin et al., 1999; Fetz, 1969, 2007; Fetz and Baker, 1973; Fetz and Finocchio, 1971; Gage et al., 2005) or they combine existing patterns of activity that are already part of the natural repertoire of the neural population (intrinsic variable learning; Golub et al., 2018; Hwang et al., 2013; Paz et al., 2003; Paz and Vaadia, 2004; Wise et al., 1998). However, most of these studies have used laboratory animal models, making it challenging (or even impossible) to distinguish between these two mechanisms without verbal instruction and feedback from the participants. To overcome this issue, we designed BMI center-out intended wrist movement tasks, where a human participant with tetraplegia controls a 2D computer cursor by modulating the activity of AIP neurons. The participant was instructed to use only the right wrist. Verbal feedback reports were obtained on every trial about the strategy that was followed to solve the BMI tasks. By perturbing the mapping between the neural activity and the cursor location, we provide evidence in favor of the intrinsic variable learning mechanism. In particular, we found that the participant learned to volitionally control neuronal activity in AIP by re-aiming in directions that offset the perturbation. Importantly, when there were no cognitive strategies that could solve the BMI tasks, AIP neurons could not generate novel patterns of activity, independent of the pre-existing neuronal structure. These findings suggest that the pre-existing neuronal structure constrains learning, and this may explain why learning a new task is easier when it is related to skills that we already have.

From Animal- to Human-BMI: Human Participants Help to Dissociate the Two Learning Mechanisms

Animal BMI studies have significantly contributed in understanding the mechanisms underlying learning and adaptation (Golub et al., 2016). Their main advantage is that we can directly compare neural activity between manual control and BMI control, enabling insights on the cortical reorganization during the process of learning neuroprosthetic control (Carmena et al., 2003; Chapin et al., 1999; Ganguly et al., 2011; Vyas et al., 2018; Zacksenhouse et al., 2007). For instance, a growing body of research indicates that, when switching from manual control to neuroprosthetic control, the brain enters into a novel

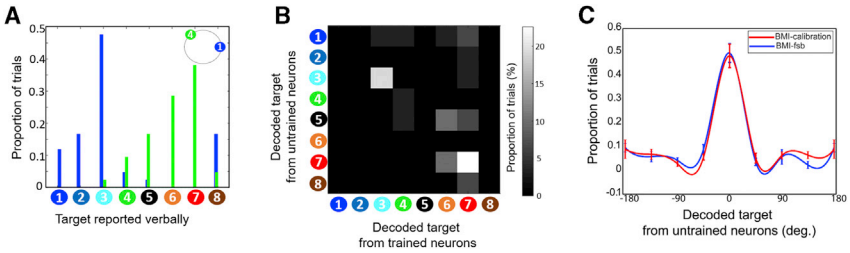

Figure 6. Trained and Untrained Neurons Encode the Same Matching Targets in the BMI-fsb Task

(A) The proportion of trials for the eight targets that the participant verbally reported to have attempted a wrist movement in response to stimuli 1 and 4 during the BMI-fsb task in the example session described in Figure 5.

(B) Correspondence matrix for the BMI-fsb session presented in Figure 5. It shows the proportion of trials for the eight targets decoded by the trained versus untrained neurons. The x axis represents the target decoded by the trained neurons and the y axis the target decoded by the untrained neurons in the example session described in Figure 5.

(C) Probability distribution (blue trace) across 15 sessions of the difference between the targets decoded by the trained and untrained neurons in the BMI-fsb task when the target decoded from the trained neurons was a matching target. The red trace describes the probability distribution from 15 sessions of the difference between the targets decoded from the trained and untrained neurons in the 8-target BMI-calibration when the target decoded by the trained neurons was the BMI-calibration target. The peak at zero corresponds to the case in which the trained and the untrained neurons encode the same matching target. Error bars represent SEM, and red and blue lines show the spline interpolation to the data.

state, producing new representations in the untrained neurons (Clancy et al., 2014; Ganguly et al., 2011; So et al., 2012). Animal models also provide a valuable tool to gain additional insights into the mechanism of neuroplasticity in learning during BMI by disrupting neural activity using invasive techniques, such as optogenetics (Gulati et al., 2017). Similar to our work, previous BMI studies in NHPs perturbed the neuron-to-movement map and provided evidence that animals learn to volitionally control the activity of the recorded neurons to adapt to perturbations (Golub et al., 2018; Hwang et al., 2013; Jarosiewicz et al., 2008). Consistent with our findings, Hwang et al. (2013) showed that NHPs can alter the pattern of activity of a trained posterior parietal cortex (PPC) neuron in a BMI task; however, the simultaneously recorded untrained neurons also changed their activity patterns in a predictable manner consistent with the intrinsic variable learning mechanism (Hwang et al., 2013). The activity of the untrained PPC neurons maintained consistent strength of tuning after the animals adapted to external perturbations, as we also found in the untrained AIP neurons. Additionally, in a neural population M1 study, after dimensionality reduction, NHPs were trained to make BMI movements within or out of the manifold of intrinsic dimensions in M1 (Sadtler et al., 2014). Within the manifold, perturbations could be solved by adopting new cognitive strategies, whereas out-of-the-manifold perturbations require the animals to generate novel neuronal activity patterns. The results showed that the animals could solve the within the manifold perturbations in a single day by re-associating existing activity patterns with different intended movements (a form of intrinsic variable learning; Golub et al., 2018). However, the out-of-the-manifold perturbations could not be solved. These findings are consistent with easier intrinsic variable (i.e., within the manifold) rather than individual neuron (i.e., outside the manifold) learning mechanisms.

Although providing important insights, it could be argued that NHP studies cannot conclusively determine the type of learning mechanisms employed to overcome external perturbations. This is because NHPs cannot be asked what strategy they follow to solve a task and they cannot be instructed to constrain how they solve the task. A constrained strategy is important to dissociate between learning mechanisms, especially when recording from cortical regions with high-dimensional mixed representations. Recent electrophysiological studies reported that neurons in NHP prefrontal cortex, rodent PPC, and human AIP are tuned to mixtures of task-related variables (Fusi et al., 2016; Raposo et al., 2014; Rigotti et al., 2013; Zhang et al., 2017). Therefore, a BMI-nfsb task does not necessarily require implicit generation of novel activation profiles. Instead, it is possible that the participant could solve the task by attempting movements using other effectors or cognitive strategies (e.g., finger, arm, or foot movements). Given that NHPs and rodents are non-verbal animals, they can explore and select other intrinsic variables to solve the BMI tasks. In our study, we constrained the participant to attempt intended movements using only her right wrist. The participant was allowed to explore different cognitive strategies for improving her performance, but she had to always perform the intended actions by using her right wrist. By verbally instructing the participant to use a constrained strategy, we restrict the available interpretations to intrinsic variable or individual neuron learning mechanisms.

Verbal reporting is also important to interpret the behavior of the untrained neurons—i.e., neurons that do not directly contribute to the BMI output but are part of the pre-existing neuronal structure that plans the wrist movements. According to our findings, both trained and untrained neurons fluctuate their activity together and in a predictable way based on the cognitive strategy used. Although this is a form of intrinsic variable learning, it could be equivalently interpreted as individual neuron learning that affects both the trained and the untrained neurons. In other words, cortical plasticity may result in new patterns of activity in both trained and untrained neurons, even when individual neuron learning occurs. However, the participant reported a target re-aiming strategy, indicating that the changes of the neuronal activity in the trained and untrained neurons reflect an intrinsic variable learning mechanism. Overall, verbally instructed strategies and self-reports are major advantages, because they can better reveal how individuals learn new tasks and compensate for perturbations. However, it could be argued that the participant may not be capable of providing an accurate description of her strategies, because she might be unaware of changes in her mental strategy when trying to compensate for the perturbation. Although this is possible, whatever changes in her strategy may have occurred proved incapable of

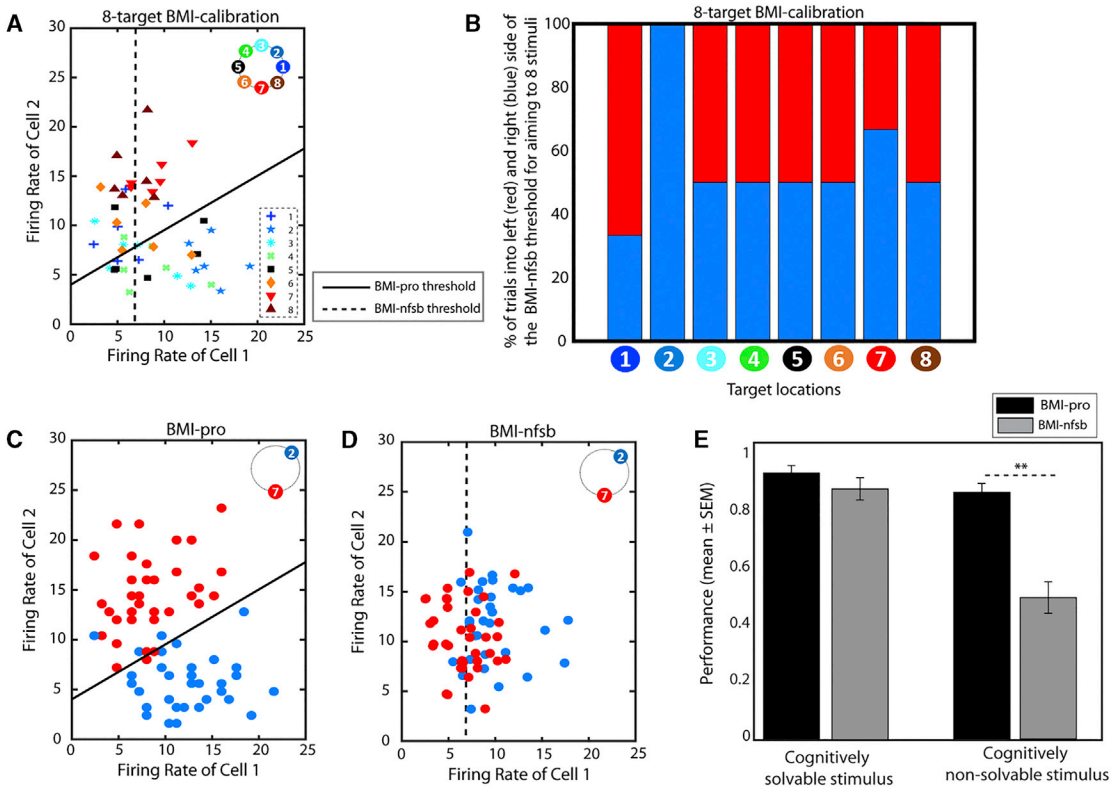

Figure 7. AIP Fails to Generate Novel Patterns of Activity in the BMI-nfsb Task

(A) The firing rates for the 2 trained neurons in response to the eight stimuli (color and shape coded) in the 8-target BMI-calibration. The pair of stimuli 2 and 7 was selected in the BMI-pro and BMI-nfsb tasks. The black solid and dashed lines represent the linear boundaries used for decoding in the BMI-pro and BMI-nfsb tasks, respectively.

(B) Percentage of trials that fall into the left (red) and right (blue) side of the BMI-nfsb task threshold in response to the eight stimuli in the 8-target BMI-calibration.

(C) The firing rates for the two trained neurons in response to the stimuli 2 and 7 in the BMI-pro. To be successful, the activity of the two trained neurons should fall below and above the BMI-pro threshold in response to stimuli 2 and 7, respectively.

(D) Similar to (C) but for the BMI-nfsb task. To be successful, the activity of the two trained neurons should fall on the right and left side of the BMI-nfsb threshold in response to stimuli 2 and 7, respectively.

(E) Proportion of successful trials (i.e., performance) across 6 sessions. Successful trials were achieved when the activity of the trained neurons during the go period fell on the correct side of the boundary. Participant achieved performance accuracy 94.03% ± 2.65% (mean ± SEM) in the BMI-pro and 87.22% ± 3.0791% in the BMI-nfsd across 6 sessions (no significant difference in the performance; two-tailed t test; $p = 0.1324$) for the cognitively solvable stimulus—i.e., stimulus that could be acquired by adopting a cognitive strategy. However, the performance was significantly deteriorated from 88.90% ± 3.93% in the BMI-pro to 50.0% ± 5.55% in the BMI-nfsb (two-tailed t test $p < 0.001$) across the 6 sessions for the cognitively non-solvable stimulus—i.e., stimulus that could not be acquired by adopting any cognitive strategy. Instead, an individual neuron mechanism was required to generate novel patterns of activity and was not accomplished during each daily session.

generating novel activity patterns that comply with the stimulus response rule in the BMI-nfsb task. Further, from this perspective, an inability of the participants to explain how a solution is achieved itself becomes an interesting finding, as may be the case for individual neuron learning.

When Cognitive Strategies Are Not Enough to Solve the BMI Tasks

One major limitation in most of the previous studies is that they assessed the neural mechanisms of learning in different brain areas by using BMI tasks that favor one hypothesis over the other. For instance, it could be argued that perturbations that can be solved by adopting new cognitive strategies, such as re-aiming to different directions, favor the intrinsic variable learning over the individual neuron learning mechanism. Hence, it is likely that trained neurons could modulate their activity independently from the pre-existing structure of the brain network only when cognitive strategies fail to solve the BMI tasks. To test this hypothesis, we designed a BMI task in which there was no cognitive strategy that can produce neural activity patterns that conform to the stimulus response rule. In this case, an individual neuron mechanism is required. Because the participant was instructed to use only the right wrist for intending movements, successful outcomes would indicate that AIP neurons are capable of generating arbitrary patterns of activity that do not belong to the response set of the BMI movements. The results show that AIP neurons could not generate novel and arbitrary patterns of neural activity, at least within a single session, to

solve the BMI perturbation task. This result suggests that either the individual neuron mechanism is a slow learning process—and therefore extensive training is required before neurons can generate novel and arbitrary activity patterns—or AIP cannot elicit arbitrary activity patterns.

Constraints on Learning

Although our findings suggest that learning is constrained by the neural structure of the implanted area, it does not imply that intrinsic variable learning is the only mechanism for learning in AIP or other brain areas. Instead, it could be hypothesized that learning in BMI tasks, in which cognitive strategies can produce rule-complying activity patterns, involves fast timescale mechanisms that underlie adaptation (Salinas 2004). On the other hand, when no cognitive strategies are adequate to solve the BMI tasks, learning may engage mechanisms required for acquiring new skills (Picard et al., 2013; Rioult-Pedotti et al., 2000), and therefore, people may need an extensive period of training to improve their performance (Sadtler et al., 2014). Additionally, it is possible that the individual neuron learning mechanism exists in other brain areas, such as the primary motor cortex (M1) (Ganguly et al., 2011). For instance, neural recordings in M1 from NHPs during brain control showed that the effects were subtle on untrained neuron activity, usually consisting of a reduction in gain rather than changes in tuning direction (Ganguly et al., 2011). Noticeably, the gain reduction of the untrained neurons after learning brain control of a cursor seems to be at odds with our findings—i.e., no significant gain modulation of the untrained neurons after learning to adapt to perturbations in BMI tasks. However, Ganguly et al. (2011) observed the down-modulation of the untrained neurons when comparing manual control with brain control. On the other hand, the comparison in our study was between two BMI tasks—i.e., BMI-pro and BMI-anti—and therefore it is possible that the gain of the untrained neurons was already reduced in the BMI-pro task—a hypothesis that could not be tested due to lack of manual control data from the participant. Furthermore, Jarosiewicz et al. (2008) found that, although the dominant response to a rotation of a subset of the neural population was compensation based on re-aiming (a form of intrinsic-variable learning that affects all of the neurons in the population equally), re-weighting and re-mapping compensation (forms of individual neuron learning in which the tuning curves of a subset of neurons in the population are modulated through a plastic mechanism) were also present in the neural responses (Chase et al., 2010; Jarosiewicz et al., 2008). Although these findings favor some degree of short-term individual neuron learning, the lack of verbal feedback cannot conclusively support one mechanism over the other for the reasons that we explained above.

Where Does Intrinsic Variable Learning Occur?

We showed that intrinsic variable learning is the predominant strategy in AIP to compensate for perturbations in BMI tasks. However, this does not necessarily mean that learning occurs within AIP. In fact, it is likely that learning emerges in other brain regions and then is transferred to AIP. An intrinsic variable learning mechanism, such as re-aiming to directions that offset the perturbation, should be implemented by brain areas that remap the actual stimulus location to intended motor actions. Most of the recorded AIP neurons in our study encode only the selected motor action (see example units in Figures 3A and 3B)—i.e., activity increases only after the go signal. One candidate area that transforms sensory information (i.e., target location) into motor actions (i.e., intended arm movements) is the human homolog of the macaque parietal reach region (PRR). A study in our lab showed that PRR neurons encode the task rule before the appearance of the stimulus location (Gail and Andersen, 2006). By training NHPs to perform pro- and anti-reach movements and presenting the task rule (i.e., pro- versus anti-reaches) prior to stimulus onset, Gail and Andersen (2006) reported that PRR neurons integrate cue location with the task rule to encode the desired action plan. Therefore, one hypothesis is that human PRR combines the stimulus location with the stimulus response rule (i.e., intended movements to the stimulus location versus re-aiming to different directions) to generate the appropriate action plans. Then, this information is transferred to AIP and other motor regions to specify the characteristics of the action plans. This hypothesis is also supported by other studies showing that PRR neurons encode both the potential and the selected reach plans in contrast to other regions within PPC, such as area 5d (Cui and Andersen, 2011) and the primary motor cortex (Cisek and Kalaska, 2005), which encode only the selected reach plan. Overall, it is likely that intrinsic variable involving spatial remapping of reach targets occurs in PRR before it is transferred to AIP. However, we do not have recording data from human PRR, and therefore, further work is required to validate this hypothesis.

Conclusions

This study is the first systematic examination of single neuron learning in a human BMI paradigm, which is an important advance considering previous studies used laboratory animal models for eventual application to humans. We recorded neural activity from a tetraplegic human participant implanted with a 96-channel microelectrode array in AIP to test two potential mechanisms of learning in BMI tasks. By perturbing the mapping between the neural activity and the cursor location, we found that the participant learned to volitionally control single neuron activity in AIP by using existing patterns of activity that are already part of the natural repertoire of the neuronal populations. These findings suggest that the ability to learn new BMI tasks is limited by the pre-existing structure of the brain network. It remains to be determined whether AIP and/or other brain areas in humans can learn to generate novel and arbitrary patterns of activity over longer timescales.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral setup

SUPPLEMENTAL INFORMATION

Supplemental Information can be found with this article online at https://doi.org/10.1016/j.neuron.2019.02.012.

ACKNOWLEDGMENTS

This work was supported by the National Institute of Health (5R01EY01554512), the Tianqiao and Chrissy Chen Brain-Machine Interface Center at Caltech, the Boswell Foundation, and the Swartz Foundation. The authors would also like to thank subject N.S. for participating in the studies and Viktor Scherbatyuk for technical assistance.

AUTHOR CONTRIBUTIONS

S.S., V.N.C., T.A., and R.A.A. designed the study. S.S., V.N.C., and T.A. developed the experimental tasks. S.S. and V.N.C. collected the data. S.S., V.N.C., and T.A. analyzed the results. S.S., V.N.C., T.A., and R.A.A. interpreted results and wrote the paper. E.R.R. provided experimental facilities, administrative assistance, and coordination with Casa Colina Hospital and Centers for Healthcare. K.W.P. provided administrative assistance. D.O. provided onsite assistance during experimental sessions. N.P. performed the surgery implanting the recording arrays in subject N.S.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 17, 2018
Revised: November 5, 2018
Accepted: February 6, 2019
Published: March 7, 2019

REFERENCES

Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., Shanfield, K., Hayes-Jackson, S., Aisen, M., Heck, C., et al. (2015). Neurophysiology. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. Science *348*, 906–910.

Armenta Salas, M., and Helms Tillery, S.I. (2016). Uniform and non-uniform perturbations in brain-machine interface task elicit similar neural strategies. Front. Syst. Neurosci. *10*, 70.

Bandyopadhyay, S., Murthy, C.A., and Pal, S.K. (1995). Pattern classification with genetic algorithms. Pattern Recognit. Lett. *16*, 801–808.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. *10*, 433–436.

Carmena, J.M. (2013). Advances in neuroprosthetic learning and control. PLoS Biol. *11*, e1001561.

Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biol. *1*, E42.

Cerf, M., Thiruvengadam, N., Mormann, F., Kraskov, A., Quiroga, R.Q., Koch, C., and Fried, I. (2010). On-line, voluntary control of human temporal lobe neurons. Nature *467*, 1104–1108.

Chapin, J.K., Moxon, K.A., Markowitz, R.S., and Nicolelis, M.A. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. Nat. Neurosci. *2*, 664–670.

Chase, S.M., Schwartz, A.B., and Kass, R.E. (2010). Latent inputs improve estimates of neural encoding in motor cortex. J. Neurosci. *30*, 13873–13882.

Chase, S.M., Kass, R.E., and Schwartz, A.B. (2012). Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. J. Neurophysiol. *108*, 624–644.

Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron *45*, 801–814.

Clancy, K.B., Koralek, A.C., Costa, R.M., Feldman, D.E., and Carmena, J.M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. Nat. Neurosci. *17*, 807–809.

Collinger, J.L., Wodlinger, B., Downey, J.E., Wang, W., Tyler-Kabara, E.C., Weber, D.J., McMorland, A.J., Velliste, M., Boninger, M.L., and Schwartz, A.B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. Lancet *381*, 557–564.

Cui, H., and Andersen, R.A. (2011). Different representations of potential and selected motor plans by distinct parietal areas. J. Neurosci. *31*, 18130–18136.

de la Fraga, L.G., and Coello Coello, C.A. (2011). A review of applications of evolutionary algorithms in pattern recognition. In Pattern Recognition, Machine Intelligence and Biometrics, P.S.P. Wang, ed. (Springer).

Fetz, E.E. (1969). Operant conditioning of cortical unit activity. Science *163*, 955–958.

Fetz, E.E. (2007). Volitional control of neural activity: implications for brain-computer interfaces. J. Physiol. *579*, 571–579.

Fetz, E.E., and Baker, M.A. (1973). Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. J. Neurophysiol. *36*, 179–204.

Fetz, E.E., and Finocchio, D.V. (1971). Operant conditioning of specific patterns of neural and muscular activity. Science *174*, 431–435.

Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. Curr. Opin. Neurobiol. *37*, 66–74.

Gage, G.J., Ludwig, K.A., Otto, K.J., Ionides, E.L., and Kipke, D.R. (2005). Naive coadaptive cortical control. J. Neural Eng. *2*, 52–63.

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

Ganguly, K., and Carmena, J.M. (2010). Neural correlates of skill acquisition with a cortical brain-machine interface. J. Mot. Behav. *42*, 355–360.

Ganguly, K., Dimitrov, D.F., Wallis, J.D., and Carmena, J.M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. Nat. Neurosci. *14*, 662–667.

Goldberg, D.E. (1989). Genetic Algorithms in Search (Optimization, and Machine Learning Addison-Wesley Professional).

Golub, M.D., Chase, S.M., Batista, A.P., and Yu, B.M. (2016). Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. Curr. Opin. Neurobiol. *37*, 53–58.

Golub, M.D., Sadtler, P.T., Oby, E.R., Quick, K.M., Ryu, S.I., Tyler-Kabara, E.C., Batista, A.P., Chase, S.M., and Yu, B.M. (2018). Learning by neural reassociation. Nat. Neurosci. *21*, 607–616.

Gulati, T., Guo, L., Ramanathan, D.S., Bodepudi, A., and Ganguly, K. (2017). Neural reactivations during sleep determine network credit assignment. Nat. Neurosci. *20*, 1277–1284.

Hochberg, L.R., Serruya, M.D., Friehs, G.M., Mukand, J.A., Saleh, M., Caplan, A.H., Branner, A., Chen, D., Penn, R.D., and Donoghue, J.P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature *442*, 164–171.

Hochberg, L.R., Bacher, D., Jarosiewicz, B., Masse, N.Y., Simeral, J.D., Vogel, J., Haddadin, S., Liu, J., Cash, S.S., van der Smagt, P., and Donoghue, J.P.

(2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature *485*, 372–375.

Hwang, E.J., Bailey, P.M., and Andersen, R.A. (2013). Volitional control of neural activity relies on the natural motor repertoire. Curr. Biol. *23*, 353–361.

Jarosiewicz, B., Chase, S.M., Fraser, G.W., Velliste, M., Kass, R.E., and Schwartz, A.B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. Proc. Natl. Acad. Sci. USA *105*, 19486–19491.

Koralek, A.C., Jin, X., Long, J.D., 2nd, Costa, R.M., and Carmena, J.M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. Nature *483*, 331–335.

Moritz, C.T., Perlmutter, S.I., and Fetz, E.E. (2008). Direct control of paralysed muscles by cortical neurons. Nature *456*, 639–642.

Paz, R., and Vaadia, E. (2004). Learning-induced improvement in encoding and decoding of specific movement directions by neurons in the primary motor cortex. PLoS Biol. *2*, E45.

Paz, R., Boraud, T., Natan, C., Bergman, H., and Vaadia, E. (2003). Preparatory activity in motor cortex reflects learning of local visuomotor skills. Nat. Neurosci. *6*, 882–890.

Picard, N., Matsuzaka, Y., and Strick, P.L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in M1. Nat. Neurosci. *16*, 1340–1347.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. *17*, 1784–1792.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Rioult-Pedotti, M.-S., Friedman, D., and Donoghue, J.P. (2000). Learning-induced LTP in neocortex. Science *290*, 533–536.

Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426.

Salinas, E. (2004). Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. J. Neurosci. *24*, 1113–1118.

So, K., Ganguly, K., Jimenez, J., Gastpar, M.C., and Carmena, J.M. (2012). Redundant information encoding in primary motor cortex during natural and prosthetic motor control. J. Comput. Neurosci. *32*, 555–561.

Truccolo, W., Friehs, G.M., Donoghue, J.P., and Hochberg, L.R. (2008). Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. J. Neurosci. *28*, 1163–1178.

Vyas, S., Even-Chen, N., Stavisky, S.D., Ryu, S.I., Nuyujukian, P., and Shenoy, K.V. (2018). Neural population dynamics underlying motor learning transfer. Neuron *97*, 1177–1186.e3.

Wander, J.D., Blakely, T., Miller, K.J., Weaver, K.E., Johnson, L.A., Olson, J.D., Fetz, E.E., Rao, R.P., and Ojemann, J.G. (2013). Distributed cortical adaptation during learning of a brain-computer interface task. Proc. Natl. Acad. Sci. USA *110*, 10818–10823.

Wang, W., Collinger, J.L., Degenhart, A.D., Tyler-Kabara, E.C., Schwartz, A.B., Moran, D.W., Weber, D.J., Wodlinger, B., Vinjamuri, R.K., Ashmore, R.C., et al. (2013). An electrocorticographic brain interface in an individual with tetraplegia. PLoS ONE *8*, e55344.

Wise, S.P., Moody, S.L., Blomstrom, K.J., and Mitz, A.R. (1998). Changes in motor cortical activity during visuomotor adaptation. Exp. Brain Res. *121*, 285–299.

Zacksenhouse, M., Lebedev, M.A., Carmena, J.M., O'Doherty, J.E., Henriquez, C., and Nicolelis, M.A. (2007). Cortical modulations increase in early sessions with brain-machine interface. PLoS ONE *2*, e619.

Zhang, C.Y., Aflalo, T., Revechkis, B., Rosario, E.R., Ouellette, D., Pouratian, N., and Andersen, R.A. (2017). Partially mixed selectivity in human posterior parietal association cortex. Neuron *95*, 697–708.e4.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
MATLAB R2016a	Mathworks	RRID:SCR_001622; https://www.mathworks.com/
Psychophysics toolbox PTB3	N/A	RRID:SCR_002881; http://psychtoolbox.org/
Other		
Neuroport Recording System, Utah array implant	Blackrock Microsystems	https://blackrockmicro.com/

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Sofia Sakellaridi (sofia@caltech.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

The participant of this study was N.S., a 60 years old female tetraplegic patient, who has a complete spinal cord injury at cervical level C3-C4 (implanted 6 years post injury) with paralysis of all limbs. All procedures were approved by the California Institute of Technology, University of California Los Angeles, and Casa Colina Center for Rehabilitation Internal Review Boards. Informed consent was obtained from NS after the nature of the study and the possible risks were explained. Study sessions were performed at most 3 days/week at Casa Colina Center for Rehabilitation.

METHOD DETAILS

Behavioral setup

All the tasks were performed with N.S. seated in her motorized wheelchair in front of a 27-inch LCD monitor in a lit room. The monitor was positioned so that the screen occupied approximately 40 degrees of visual angle. Stimulus presentation was controlled using the psychophysics toolbox for MATLAB (Mathworks) (Brainard, 1997).

Implant methodology and physiological recordings

Subject N.S. was implanted with two 96-channel Neuroport arrays. One array was implanted on the surface of the Superior Parietal Lobule (putative human BA5) and one at the junction of the intraparietal sulcus with the postcentral sulcus (putative human AIP), two areas that are specialized in reaching and grasping movements, respectively. The implantation locations were determined based on preoperative fMRI (Aflalo et al., 2015). NS performed an imagined hand reaching and grasping task during a functional magnetic resonance imaging (fMRI) scan. A Siemens Trio 3T scanner at the University of California, Los Angeles Semel Institute for Neuroscience & Human Behavior Imaging Center was used for scanning. Anatomical sequences were acquired using a T1-weighted MP-RAGE sequence with the following parameters (TR = 1590 ms; TE = 2.7ms; fov = 176 × 256 × 256 mm; 1 mm isotropic voxels). The parameters for the functional scans were as follows: T2*-weighted single-shot echo-planar acquisition sequence (TR = 2000 ms; slice thickness = 3 mm; in-plane resolution = 3x3mm; TE = 30ms; flip angle = 80; fov = 192 × 192mm; matrix size = 64 x64; 33 slices (no-gap) oriented 20 degrees relative to ACPC line). Surface reconstruction of the cortex was performed using the Freesurfer software (http://surfer.nmr.mgh.harvard.edu/). NS was asked to imagine or attempt different reaching and grasping movements with visual and attentional controls in order to localize limb and hand areas in SPL and AIP. The BOLD responses from the reaching and the grasping tasks were used to determine the implantation sites for the surgery. The surgery was performed at the UCLA Medical Center on August 26, 2014. A single craniotomy was made over SPL and AIP of the dominant hemisphere and the two arrays were implanted in the hand/limb regions of SPL and AIP. The AIP array was placed at Talairach coordinate [-36 lateral, 48 posterior, 53 superior]. More information about the implantation method can be found in a previous study from our lab (Aflalo et al., 2015). Although a different subject participated, the same procedure was followed to implant and record the neural data.

Data collection

The data were collected 2 years after implanting the recording arrays, over a period of about 1 year (including pauses) in 2-3 study sessions per week. Each session lasted about 1h. Neural activity was amplified, digitized, and recorded at 30 kHz from each recording

electrode using the Neuroport neural signal processor (NSP). Unsorted action potential threshold crossings of −4.5 times the root-mean-square of the high-pass filtered (250Hz) full bandwidth signal was used to register action potentials for on-line analysis. Clearly isolatable single-units were sorted using a window discriminator approach when possible. For offline analysis single and multiunit activity was sorted using mixture of Gaussian modeling of the waveforms detected by threshold crossings (for further information regarding the data collection see (Aflalo et al., 2015)). Due to a lack of single units in SPL, only data from AIP was analyzed in the current study.

Task procedure

In the single-neuron experiment, each session consisted of three primary task blocks: BMI-calibration, BMI-pro, BMI-anti. The first 4 sessions were performed in separate blocks, in which one block of BMI-pro trials preceded two blocks of BMI-anti trials. In the last 4 sessions, blocks of BMI-pro trials were randomly interleaved with blocks of BMI-anti trials. NS ran 20 calibration trials (10 per stimulus), 48.88 ± 2.8 (MEAN ± SEM) BMI-pro trials and 61.25 ± 10.92 BMI-anti trials per session. In the two-neuron experiments, each session consisted of three primary blocks in the following order: BMI-calibration, BMI-pro and BMI-fsb (for a pair of feasible stimuli) or BMI-nfsb (for one feasible and one non-feasible stimulus). NS ran 48 calibration trials (6 trials per stimulus), 32.67 ± 1.18 BMI-pro trials, 97.33 ± 5.89 BMI-fsb trials and 72 BMI-nfsb trials per session. In between blocks, NS rested for about 1 min. Prior to experimental trials, NS was instructed to control the 2D location of the cursor on the screen using only her right wrist and no other body parts. In particular, the only instruction given to NS was "attempt a wrist movement with your right hand after the go-signal." We also instructed NS after each trial to verbally report the direction of the intended movement based on points on a clock (e.g., intended movement at 12 o'clock).

QUANTIFICATION AND STATISTICAL ANALYSIS

All analyses were performed using MATLAB 2016a.

Linear discriminants in the BMI tasks with one neuron

In the single-neuron experiment, we used linear discriminant analysis (LDA) in the BMI-calibration to identify the linear threshold that divides the high and low firing rates under the assumption of a uniform prior distribution. In particular, the classifier took as input a vector comprised of the number of spikes occurring after the go-signal for the trained neuron. The following assumptions were made for the classification models: (a) the prior probability across the classes was uniform, (b) the conditional probability distribution of each feature on any given class was normal, (c) only the mean firing rates differ for each class (the covariance of the normal distributions were the same for each class), and (d) the firing rates of each input are independent (covariance of the normal distribution was normal). We evaluated the classification accuracy offline using the calibration data based on leave-one out cross-validation. The linear threshold from the LDA classification was used both in the BMI-pro and the BMI-anti tasks, with the difference that the stimuli associated with high versus low firing rate was flipped in the BMI-anti task. Classification accuracy for the on-line decoding in the BMI-pro and BMI-anti tasks was based on subject's performance during the testing session.

Linear discriminants in the BMI tasks with two neurons

Similarly, in the BMI-pro task with two trained neurons, the linear boundary that divides the two-firing rate clusters was found using LDA under the assumption of a uniform prior and independent distribution. The linear boundary for the BMI-pro task was the one that best separated the two-firing rate clusters acquired during the BMI-calibration– i.e., ideally 100% classification accuracy for both stimuli. For the BMI-fsb task, we selected a linear threshold based on the classification accuracy in the BMI-calibration of each possible pair of stimuli, except the one used in the BMI-pro task. In particular, we computed all 27 linear thresholds for each pair of stimuli in the BMI-calibration that was not presented in the BMI-pro task and selected one under the following constraints: i) it poorly classifies the two firing-rate clusters for the presented pair of stimuli in the BMI-pro task and ii) it best separates the firing-rate clusters formed by intended wrist movements to another pair of stimuli.

Finally, for the BMI-nfsb task, the linear threshold was selected using evolutionary algorithms (EAs) (Bandyopadhyay et al., 1995; de la Fraga and Coello Coello, 2011; Goldberg, 1989). The criterion (fitness function) that was used to generate the linear boundary was given as follows:

$$f = \sum_{i=1}^{8} r_i, \qquad \text{where} \quad r_i = \begin{cases} \dfrac{group_1^i}{group_1^i + group_2^i}, & if \;\; group_1^i < group_2^i \\ \dfrac{group_2^i}{group_1^i + group_2^i}, & if \;\; group_2^i < group_1^i \end{cases}$$

Where $group_x^i$ is the number of trials for the stimulus i in the BMI-calibration in which the firing rate of the two trained neurons falls within the one (x = 1) or the other (x = 2) side of the linear boundary. Note that r_i takes values between 0 (i.e., best classification – all trials are classified into one of the groups) to 0.5 (chance classification – half of the trials are classified into one of the group and the rest of them to the other). The linear boundary was given by maximizing the fitness function f – i.e., $max(f) = \sum_{i=1}^{8} 0.5 = 4$.

Each chromosome in the EA population consists of 4 genes, which correspond to the 2-dimensional points (x_1, y_1) and (x_2, y_2) that define the linear threshold. The parameters of the EA algorithm were as follows: Population = 100 (i.e., number of chromosomes), Generations = 100 (i.e., total number of iterations), probability of crossover = 0.8, probability of mutations = 0.2.

DATA AND SOFTWARE AVAILABILITY

Data and MATLAB analysis scripts available upon request from Sofia Sakellaridi (sofia@caltech.edu).

ADDITIONAL RESOURCES

This study was conducted as part of NIH clinical trial NCT01958086.