Brain–machine interface for eye movements

Arnulf B. A. Graf[1] and Richard A. Andersen[1]

Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA 91125

Contributed by Richard A. Andersen, October 24, 2014 (sent for review July 29, 2014; reviewed by Apostolos P. Georgopoulos and Ranulfo Romo)

A number of studies in tetraplegic humans and healthy nonhuman primates (NHPs) have shown that neuronal activity from reach-related cortical areas can be used to predict reach intentions using brain–machine interfaces (BMIs) and therefore assist tetraplegic patients by controlling external devices (e.g., robotic limbs and computer cursors). However, to our knowledge, there have been no studies that have applied BMIs to eye movement areas to decode intended eye movements. In this study, we recorded the activity from populations of neurons from the lateral intraparietal area (LIP), a cortical node in the NHP saccade system. Eye movement plans were predicted in real time using Bayesian inference from small ensembles of LIP neurons without the animal making an eye movement. Learning, defined as an increase in the prediction accuracy, occurred at the level of neuronal ensembles, particularly for difficult predictions. Population learning had two components: an update of the parameters of the BMI based on its history and a change in the responses of individual neurons. These results provide strong evidence that the responses of neuronal ensembles can be shaped with respect to a cost function, here the prediction accuracy of the BMI. Furthermore, eye movement plans could be decoded without the animals emitting any actual eye movements and could be used to control the position of a cursor on a computer screen. These findings show that BMIs for eye movements are promising aids for assisting paralyzed patients.

learning | lateral intraparietal area | brain–machine interface | eye movements | saccades

Brain–machine interfaces (BMIs) have been successfully used to predict reaches and arm movements (1–7). However, little effort has been concentrated on building a BMI based on eye movements. This gap is surprising because the motor and neuronal mechanisms of eye movements are very well understood and arguably simpler than those of arm movements. Specifically, eye movements are rapid and ballistic. The lateral intraparietal cortex (LIP) is ideally suited to be the site for a BMI based on eye movements (8). LIP neurons are known to encode eye movement plans, among other signals such as eye position (9–16). We recently showed that eye movement plans can be accurately predicted from the responses of populations of LIP neurons using Bayesian inference (16). The aim of the present study was twofold. First, a BMI was used with small neuronal ensembles of LIP neurons to predict, in real time, eye movement plans without the animals actually making eye movements. Second, the BMI application induced learning-related changes in the saccade system. Learning can produce changes in reach areas, but how learning-related changes occur at the level of LIP neuronal ensembles is still unclear (17, 18).

Here, we show that the intended eye movement activity can be used to accurately position a cursor on a computer screen. These results suggest that an eye movement BMI can be used as a prosthetic to assist locked-in patients who cannot produce eye movements. Moreover, such an eye movement BMI can also be used to assist tetraplegic persons to decode intended limb movements by providing an extra channel of target position information (19). Learning, defined as an increase in the prediction accuracy, occurred at the level of neuronal ensembles, particularly for difficult predictions. The population learning had two components: an update of the parameters of the BMI based on its history and a change in the responses of individual neurons. These results provide strong evidence that the responses of neuronal ensembles can be shaped with respect to a cost function, which here is the prediction accuracy of the BMI. Such learning adds additional support for the utility of an eye movement BMI based on LIP activity.

Results

Instructed and Brain Control Saccade Plans. We recorded from the LIP of two adult male monkeys, targeting a different hemisphere in each animal. We used a microdrive with five individually adjustable microelectrodes to simultaneously record from small populations of neurons. In all, 40 populations of neurons were obtained, totaling 278 neurons (on average, 7 neurons per population). We used an eye movement-based BMI to study learning-related changes in the neuronal populations. This BMI was defined by two sequential phases. First, the parameters of the BMI were learned offline by inferring instructed planned saccades from the persistent activity of populations of LIP neurons (20). Second, eye movement plans were predicted from the population activity in real time without overt behavior. The same population sample was used in both phases of the task, and each recorded session yielded a different sample of neuronal population activity.

During the instructed trials, a real-time behavioral control unit instructed the visual display and monitored the position of one eye. These data were combined with the recorded neuronal data in the data acquisition unit and saved for offline analysis (Fig. 1*A*). The animals performed delayed memory saccades (*Materials and Methods*). On completion of fixation, a saccade to one of the eight randomly chosen neighboring locations was cued. The animals had to remember the instructed saccade and form a plan to move their eyes. On extinction of the fixation cue, the saccade to the cued location was executed (Fig. 1*B*). In this research, we exclusively examined the memory period (red box in Fig. 1*B*) when the saccade plan was made; an analysis across task epochs was previously

Significance

We developed a brain–machine interface (BMI) that records single cell activity from populations of neurons for decoding planned eye movements (i.e., saccades) without the animals executing them. The recordings were made from the lateral intraparietal area, an important cortical node in the primate saccade system. To our knowledge, this is the first BMI using neuronal populations to predict eye movements in a brain control task. Learning occurred during brain control trials, improving prediction performance by changing the response properties of single neurons. An eye movement BMI can assist patients locked-in from stroke or neurodegenerative disease who are unable to make any movements including eye movements.

Author contributions: A.B.A.G. and R.A.A. designed research; A.B.A.G. performed research; A.B.A.G. analyzed data; and A.B.A.G. and R.A.A. wrote the paper.

Reviewers: A.P.G., University of Minnesota Medical School; and R.R., Universidad Nacional Autonóma de México.

The authors declare no conflict of interest.

[1]To whom correspondence may be addressed. Email: graf@vis.caltech.edu or andersen@vis.caltech.edu.

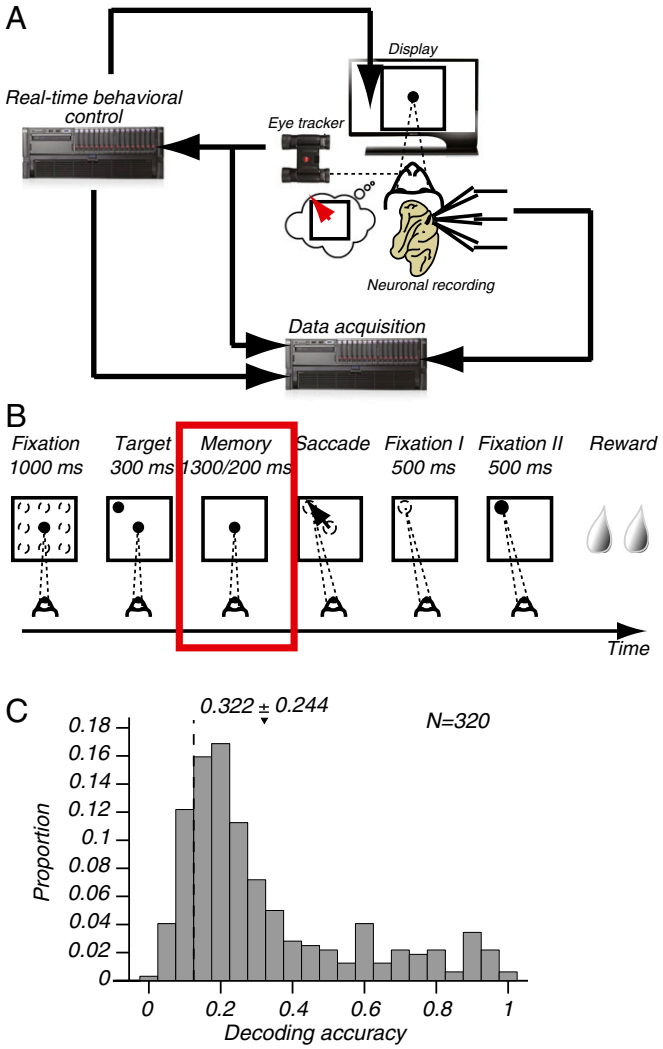

Fig. 1. Instructed trials. (*A*) Schematic of the experimental setup. A real-time behavioral control unit operated the visual display and monitored eye position. Eye position and neural signals were collected by a data acquisition unit and analyzed offline. (*B*) Epochs of the task for one trial. After initial fixation, one of eight possible target locations was briefly cued. After a variable delay, the fixation extinguished instructing a saccade to the remembered location of the target. Once the saccade was made the position, the target was reilluminated followed by a reward on correct trials. (*C*) Distribution of the decoding accuracy across the 40 populations and eight saccade directions during the memory epoch (red box in *B*). Eighty-three percent of saccade directions were decoded above the 1/8 chance level (dotted line), and the mean and SE of the decoding accuracy was 0.322 ± 0.244.

presented for a different dataset (16). The spike count of each neuron in a population across the memory period was used as the neuronal population response. The BMI was trained offline on these data to predict the planned saccade direction given the neuronal population response. To address population coding in LIP, we used Bayesian inference (16, 21–24). The activity of the neurons was assumed to be independent, and the response statistics were described by a truncated Gaussian distribution (*Materials and Methods*). To evaluate the training of the BMI, we computed in Fig. 1*C* the distribution of the saccade direction prediction accuracies across the populations and saccade directions. Similar to previous findings (16), the saccade direction was accurately predicted from the neuronal responses: in 83% of cases, the prediction accuracy was above chance. The parameters of the BMI obtained during the instructed saccades were thus an accurate description of the inference of eye movements from responses of small populations of neurons.

These instructed trials were followed daily by brain control trials. In the brain control trials, the BMI predicted in real time the saccade plans from the neuronal population response without the monkey making any overt eye movements. The experimental setup is an evolution of the one used in the instructed trials, except that here the BMI makes online predictions about the eye movement plan. Once the estimated saccade direction was generated by the BMI, it was relayed to the real-time behavioral controller (Fig. 2*A*). The animal maintained fixation for the entire trial and was rewarded when the planned eye movement to the remembered target location was accurately predicted in real time based solely on the neuronal population response and without an actual eye movement (Fig. 2*B*). If the prediction was incorrect, the animal was shown the incorrectly decoded plan and did not receive a reward. By providing this feedback, we hypothesized that the animal could learn to adjust the neuronal responses to maximize the prediction accuracy of the BMI without performing any overt behavior.

To make these predictions, the BMI sampled the neuronal population response in real time, in our case every 100 ms, by computing spike counts of each neuron across a temporal interval that started with the extinction of the target and ended at

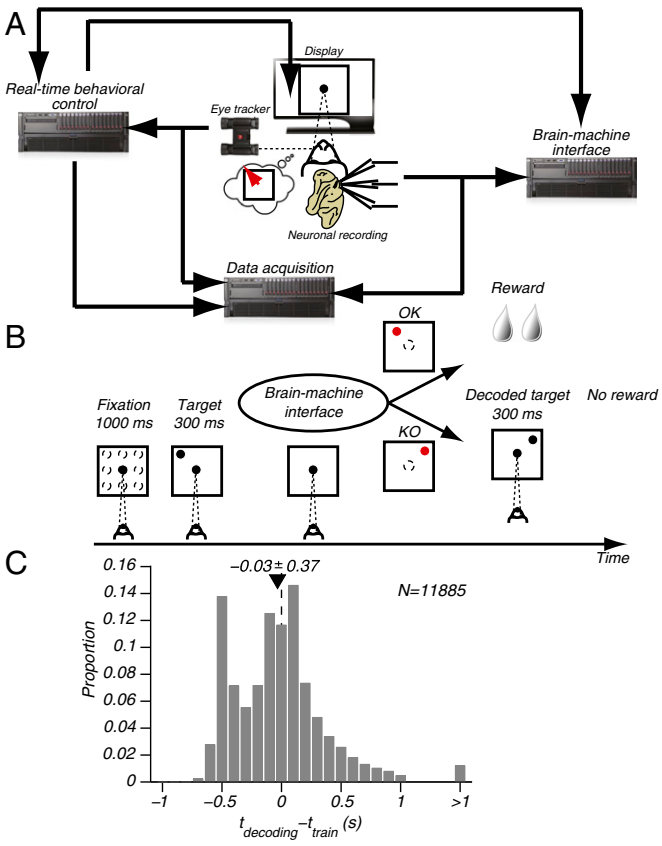

Fig. 2. Brain control trials. (*A*) Schematic of the experimental setup. The setup was similar to that used for instructed trials (Fig.1*A*), except a BMI unit estimated the planned movement and provided feedback to the behavioral control unit. (*B*) Epochs of the task. The trials were similar to the instructed trials (Fig. 1*B*) except after the fixation target was extinguished, the animal did not move his eyes and received feedback of the decoded location. If the location was correct and the animal did not move his eyes, he received a reward. (*C*) Distribution across all brain control trials of the difference between the online decoding time (self-paced BMI) and the training time from the instructed trials. The mean and SE were −0.03 ± 0.37.

the time the saccade direction estimate was generated. The BMI generated, for each sampling time, a saccade direction estimate by updating the Bayesian decoder derived from the instructed trials. By using the history of these estimates within a trial, the BMI evaluated whether an estimate was stable (*Materials and Methods*). When stability was achieved, the estimate was considered final and passed on to the real-time behavioral control unit (Fig. 2*A*). As such, one defining feature of our BMI is that it is self-paced: rather than using the same timescale for training the BMI in the instructed trials, the BMI during the brain control trials found the optimal timescale for each online estimate. This increases the robustness of the BMI to drifts in neuronal responses and cell isolation. The distribution of the differences between the prediction timescales from the training and brain control trials is shown in Fig. 2*C*. This distribution is strongly peaked around zero, indicating that the BMI automatically determined a decoding timescale that was on average similar to the training timescale. However, the heavy tails of the distribution also indicate that the best prediction accuracy was often achieved on timescales different from training, even on shorter timescales. This result suggests a change in the neuronal population response, which may underlie population learning.

The second defining feature of our BMI was that its parameters were dynamically updated (*Materials and Methods*). At the beginning of the brain control trials, the parameters of the BMI were derived from the spike counts of each neuron in the memory period and the instructed saccade plans. For each estimate generated by the BMI, the spike count and cued saccade were used to update the database of spike counts and saccade directions. The parameters of the BMI were then recomputed based on this updated database. The updated BMI made predictions in the following brain control trial. Akin to recursive Bayesian estimation, the decoder was updated by each new trial, thus empirically learning from the data.

Population Learning. For each brain control trial, the BMI had access to the entire trial history, namely all of the instructed trials and all of the previous brain control trials; it could learn from its history. To quantify this learning, we computed the change in prediction accuracy during the brain control trials (*Materials and Methods*). Fig. 3 illustrates, for one population and one saccade direction, the increase in decoding accuracy during the brain control trials (black area in the figure). We defined this increase in the decoding accuracy as population learning (25) and its neuronal correlates (26). To assess how much of this effect was due to empirically updating the decoding parameters, an analysis was performed without history. In this case, the decoder was only trained on the instructed trials and was tested for each brain control trial separately, thus corresponding to a BMI with frozen parameters (red area in Fig. 3). The decoding accuracy decreased when omitting trial history, suggesting unlearning. To quantify population learning, a learning selectivity index (LSI) was computed by comparing the last quartile (*lq*) and the first quartile (*fq*) of the decoding accuracy during the brain control trials

$$LSI = \frac{lq - fq}{lq + fq}.$$

The LSI varied between −1 (strong decrease in decoding accuracy, i.e., strong unlearning) and 1 (strong increase in decoding accuracy, i.e., strong learning). We computed the LSI of each population for the saccade direction that had the most brain control trials and compared the LSI for both learning mechanisms in Fig. 4*A*. We found that, in 78% of cases, learning with history was stronger that learning without history (dots above the diagonal). The LSI for learning with history was also positive in 78% of cases, showing a strong intraday learning progression.

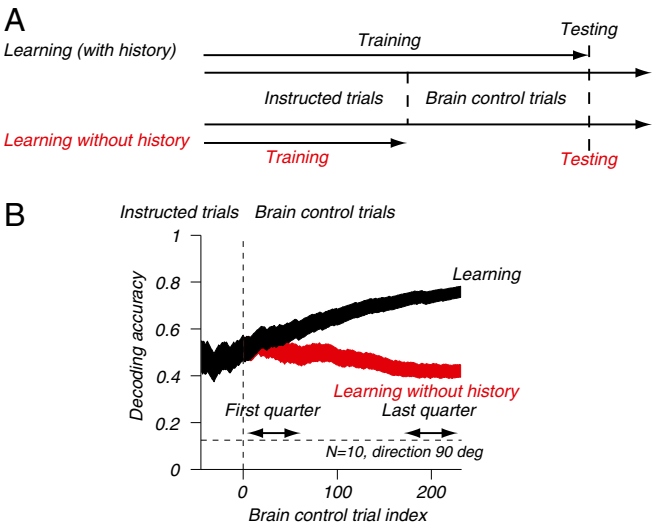

Fig. 3. Population learning. (*A*) Schematic of the two learning regimes: learning with history in which the BMI was trained on all of the previous trials (instructed and previous brain control trials) and learning without history in which the BMI was trained using solely the instructed trials. (*B*) Decoding accuracy as a function of the progression of brain control trials for both learning regimes for one population of 10 neurons and one saccade direction. The horizontal dotted line represents chance (1/8), and the vertical line represents the transition between instructed and brain control trials. The first and last quartiles of the decoding accuracy, used to compute the learning selectivity index, are also indicated.

However, in the absence of history, only 43% of LSIs were positive, indicating that unlearning frequently occurred when ignoring trial history.

Next, we asked how the LSI depended on the prediction accuracy at the beginning of the brain control trials (Fig. 4 *B* and *C*). When using trial history, the LSIs were dependent on the initial decoding accuracy: they were largest for low initial decoding accuracies and decreased regularly as the decoding accuracy increased (Fig. 4*B*). These findings suggest that learning effects are most prominent for difficult behaviors, which here are the saccade directions that were poorly predicted. The LSIs, when ignoring trial history, were largely independent of the decoding accuracy (Fig. 4*C*). Thus, unlearning can occur independently of the decoding accuracy. Taken together, these findings show that learning mainly occurs when there is something to learn, e.g., a difficult task with a low decoding accuracy. Unlearning, however, can occur independently of the difficulty of the task. In other words, the prediction accuracy can increase (i.e., learning) only when it is low to begin with, and it can decrease (i.e., unlearning) independently of whether it is initially high or low.

The population learning and unlearning that we documented depends on the mechanisms used by the BMI (parameters computed with or without trial history). Moreover, the fact that learning and unlearning occurred suggests that populations of neurons in LIP exhibit learning-related changes that allowed them to better use a BMI, in concordance with previous findings in different animal models and tasks (27, 28). By providing feedback, the animals were able to learn to use the BMI by increasing the accuracy of eye movement predictions.

Contributions of Individual Neurons to Population Learning. We examined whether learning to control the BMI resulted in changes in the activity at the level of single neurons. LIP neurons are known to have monomodal tuning curves with respect to planned saccade directions. To quantify the change in tuning, we computed the difference between the preferred direction of each

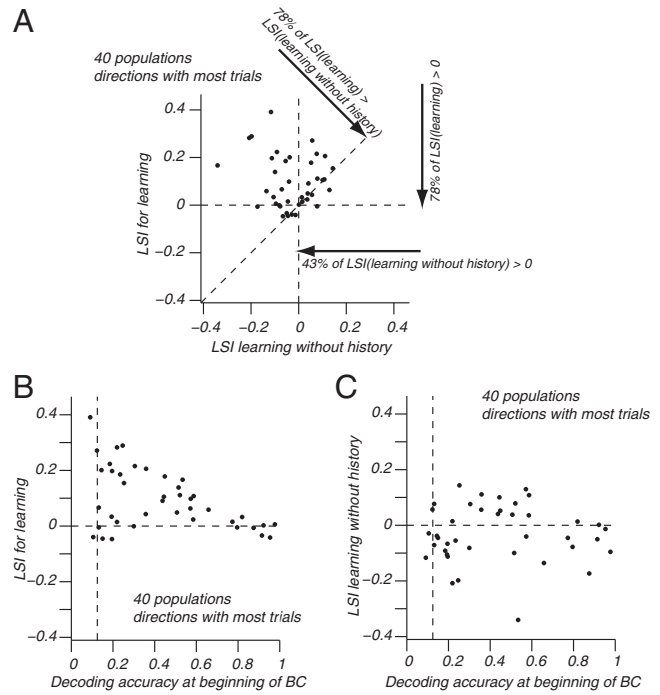

Fig. 4. Population learning quantified by the LSI of the saccade direction with most trials for each population. (*A*) Comparison of learning with and without history. (*B*) Relation between the LSI with history and the decoding accuracy at the beginning of the brain control (BC) trials. The vertical dotted line indicates chance (1/8). (*C*) Same as in *B*, except for the LSI without history.

neuron, i.e., the direction eliciting the strongest response, and the cued direction (Fig. 5*A*). The preferred direction was calculated from responses during either the instructed or the brain control trials (Fig. 5*A*, gray and black curves, respectively). The preferred direction was determined using a circular Gaussian interpolation through the tuning curve, and only neurons with a good fit ($r^2 > 0.7$) were considered. We then computed the distribution of the difference between the preferred and cued directions across these neurons (Fig. 5*B*). This distribution for the brain control trials had a smaller mean, SE, and circular variance than for the instructed trials. This finding shows that individual neurons change their response properties in the brain control trials to shift their tuning curves toward the cued direction. The cued direction thus acted as an attractor during the brain control trials.

We finally asked whether the changes at the level of single neurons were reflected in population learning as quantified by the LSI (29). One neuron in each population was selected for the strongest correlation between the LSI for learning (with history) and difference in spike counts between the instructed and brain control trials (Fig. 6*A*). The five dark dots in Fig. 6*A* represent five saccade directions in which one neuron had a strong correlation between the LSI and the response difference. The gray dots represent the directions for the neuron in each population with the largest correlation, for a total of 225 directions with sufficient number of trials (10 or more). The findings show that for each population there was at least one neuron that had a change in response between instructed and brain control trials that was reflected in the LSI. We suggest that these changes in response were underlying the shifts in the tuning curves studied in Fig. 5. Given that the LSI is a metric quantifying population learning, we then studied the correlation coefficient between the LSI and response difference for all of the neurons. Fig. 6*B* illustrates how the correlation coefficients vary within each population of neurons ordered by decreasing correlation, with the dark line corresponding to the population of nine neurons that included the neuron highlighted with the dark dots in Fig. 6*A*. The curves are decreasing, suggesting that all neurons in a population do not contribute to the same extent to population learning. More specifically, the population learning is sparse: only a few neurons change their response properties to increase the prediction accuracy of the neuronal population code during learning. As such, the changes in the neuronal population code can be traced back to a small subset of neurons that dominate population learning.

Discussion

The present study demonstrates, for the first time, the successful application of an eye movement-based BMI using LIP activity. The populations of neurons in LIP learned to rapidly shift their response characteristics to increase the prediction accuracy of oculomotor plans. This learning-related change was strongest for eye movement plans that were difficult to predict, showing that learning effects are strongest for conditions "in which there is something to learn." Conversely, unlearning occurred independently of the difficulty of the condition, suggesting that "it is always possible to unlearn." A decoder that learned with trial history outperformed a decoder agnostic to trial history, and the tuning of individual neurons also changed during learning. Therefore, the BMI learned by adjusting its parameters, and individual neurons learned to shift their response properties to better use the BMI. These two types of learning occurred simultaneously in our BMI and were successfully untangled and examined.

Previous studies have shown that there are two general categories of learning mechanism for single neurons in BMIs (18). For an individual neuron mechanism, each neuron is separately manipulated to produce a particular response pattern independent of the activity of its neighbors. For an intrinsic variable mechanism, the natural repertoire of response patterns is cognitively explored and affects not only the neuron being recorded from but

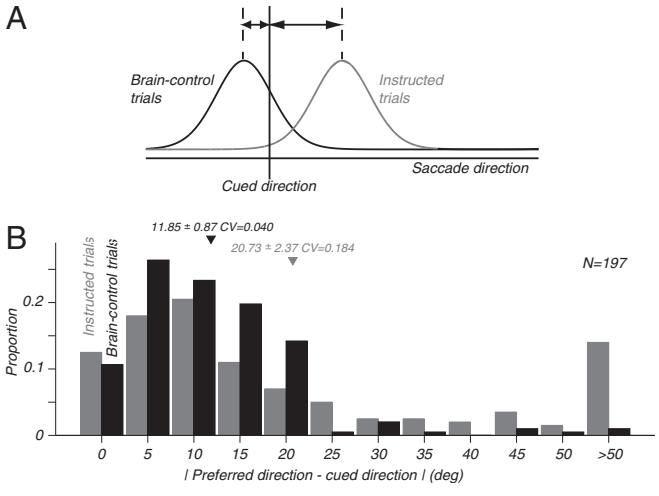

Fig. 5. Shifts in the preferred saccade direction of individual neurons between the instructed and the brain control trials. (*A*) Schematic of tuning curves and preferred directions of a model neuron during the instructed and brain control trials. The shift in tuning is quantified by the angular difference between the preferred direction and the cued saccade direction. (*B*) Distributions of the difference between the preferred and the cued saccade directions for the instructed (gray) and the brain control (black) trials. These distributions were obtained using the 197 neurons that had a good fit with a circular Gaussian function ($r^2 > 0.7$). The mean, SE, and circular variance of both distributions are indicated.

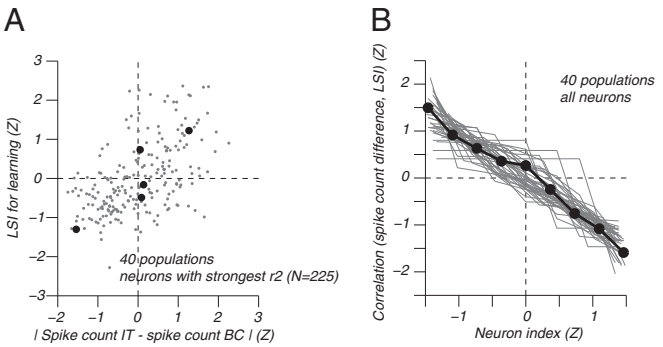

Fig. 6. Sparseness of population learning (with history). (*A*) Relation between the LSI and the difference in spike count between instructed and brain control trials (Z score on both axes) across the 225 saccades that had 10 or more trials. For each of the 40 populations, the directions (gray dots) corresponding to the neuron that had the strongest correlation between its LSI and the spike count difference are represented. The five directions for one such neuron are represented by the black dots, and the corresponding correlation was $r^2 = 0.88$. (*B*) Correlation of the LSI and the spike count difference for all of the neurons within a population (neuron index), the neurons being ordered by decreasing correlation (Z score on both axes). The gray lines represent the relation between this correlation as a function of the neuronal index. The dark line represents the neurons of the population from which the example neuron in *A* (black dots) was drawn.

also the more global pattern of activity of the cortical area. In the parietal reach region, located near LIP but encoding reaches instead of saccades, the activity of the untrained neurons supports the intrinsic variable mechanism (18). In the current study, either mechanism could produce the shifting of response properties observed in LIP. Because all neurons participated in the training, the results cannot differentiate between these two learning mechanisms.

Our study has laid the foundations of a BMI based on eye movements that uses neuronal responses to predict eye movement plans. BMIs are poised to have significant social/engineering and medical applications. In the former case, they will usher a novel area of interaction between man and machine. In the latter, they will help patients with limb or oculomotor paralysis regain interaction with their environment (5, 30–32). Although BMIs based on arm movements, reaches, and grasping are well established (2–4, 6, 7) and have led to successful clinical applications (5), the current study presents the first BMI for eye movements using populations of neurons. Such a BMI can work per se in isolation or as a complement to an arm movement BMI (19). An eye movement-based BMI would be highly advantageous for patients who are completely paralyzed and cannot move their eyes; that is, patients who are locked in due to strokes or neurodegenerative disease. Such a BMI can also assist tetraplegic patients by providing a second movement modality to specify the goals of intended manual movements. Using the neuronal correlate of eye movements instead of externally recorded physical eye movements significantly enhances the ease of use of the interface as it does not require additional eye tracking equipment. Finally the goal-oriented nature of signals for ballistic saccades is well tailored for BMIs that are primarily concerned with the end point of desired movements.

Materials and Methods

Instructed and Brain Control Task. The instructed trials have been described previously (16). Briefly, the animals made delayed memory saccades starting on a 3 × 3 Cartesian grid with nodes separated by 8° and ending on one of the eight neighboring nodes (Fig. 1*B*). Daily, after the instructed trial sessions, the animals performed brain control trials in which they maintained fixation during the entire trial (Fig. 2*B*). Although no overt behavior was performed, the animals were rewarded when the planned eye movement to the remembered target location was accurately predicted in real time based solely on the neuronal population response. If a wrong target was predicted, this target was displayed as the animal maintained fixation to provide an error signal. The brain control trials allowed us to correlate the neuronal plasticity at the population level with the accuracy of the predicted eye movement plan without the animal initiating a movement. Specifically, the animals maintained fixation for 1,000 ms. A target was then flashed in the periphery for 300 ms. While still maintaining fixation, the neuronal responses were decoded in real time. When the decoder reached a sufficient confidence on its prediction (see below), an estimated target location was generated. If the estimate matched the cued target location, the fixation was extinguished and the animal was rewarded. If the estimate differed from the cued target, the animal had to maintain fixation while the estimated target location was flashed for 300 ms. The fixation was then extinguished, and no reward was delivered. The animals thus had to fixate longer and received no reward. Also, by presenting the decoded target, the animals received visual feedback to learn to use the neuronal population activity to generate accurate estimates. In other words, the animals had to learn to use the decoder based on a given neuronal population sample.

The cued saccades were on a grid as described in ref. 16. However, for each brain control session, we had the animal plan saccades specifically tailored for the neuronal sample at hand. Because the brain control task is intrinsically a saccade direction estimation task, the focus was on saccade direction estimation and started with center-out brain control saccades in all eight directions. We then determined the best-predicted saccade directions and cued the animal only on these. This often resulted in a two or three alternative forced choice experiment across saccade directions contralateral to the recorded brain hemisphere. If the animal was successful at these best brain control saccades, the animal was finally cued to plan these saccades across the 3 × 3 grid. This behavioral strategy allowed us to (*i*) maintain the animal's motivation and (*ii*) focus on the oculomotor behaviors best represented by our neuronal population sample. We thus had a vastly different number of trials per condition, resulting in nonuniform priors for saccade direction and presaccade eye position in a Bayesian inference framework (see below). The eye position and saccade direction were independent behavioral variables by design (16), and the animals only received feedback about the saccade direction. All studies were carried out in compliance with the guidelines of the Caltech Institutional Animal Care and Use Committee and the National Institutes of Health *Guide for the Care and Use of Laboratory Animals* (33).

Population Recordings. We used an identical recording setup as described in ref. 16. Briefly, we targeted the LIP area using a microdrive of five individually adjustable microelectrodes. This multielectrode technique allowed recording from different empirical samples of population activity across recording sessions. The same population sample was kept between the instructed and the brain control trials (verified online by analysis of the waveform shapes). In addition to previously mentioned reasons (16), the recording of multiple neurons simultaneously was necessary to obtain the best possible neuronal sampling, thereby increasing the heterogeneity of the neuronal ensembles (24) and maximizing the prediction accuracy of the BMI.

Inferring Saccade Plans from Neuronal Populations. We predicted eye movement plans from the neuronal population response in LIP using Bayesian inference (16, 22–24, 34–43). The neuronal population response was the spike count of each neuron over a given temporal window (see below). The likelihood function $p(\mathbf{r}|\theta)$ was determined first: the probability to obtain a population response r for a given saccade direction θ, evaluated across all θ. The N neurons were assumed to be independent

$$p(\mathbf{r}|\theta) = \prod_{i=1}^{N} p(r_i|\theta).$$

The response statistics of each neuron were approximated using a truncated Gaussian distribution over positive integers

$$p(r_i|\theta) = \frac{G(r_i, \mu_i, \sigma_i, \theta)}{\int_0^\infty G(r, \mu_i, \sigma_i, \theta)dr} \quad \text{where}$$

$$G(r, \mu, \sigma, \theta) = \frac{1}{\sqrt{2\pi\sigma(\theta)^2}} \exp\left\{-\frac{[r-\mu(\theta)]^2}{2\sigma(\theta)^2}\right\}.$$

This model required the empirical derivation of the mean $\mu(\theta)$ and the variance $\sigma(\theta)^2$ of the spike count for each saccade direction. The probability

to obtain a saccade direction given a behavior $p(\theta|\mathbf{r})$—the posterior distribution—is then computed using Bayes' theorem

$$p(\theta|\mathbf{r}) = \frac{p(\mathbf{r}|\theta)p(\theta)}{p(\mathbf{r})},$$

where $p(\theta)$ is the prior function determined by the number of occurrences of the saccade direction as instructed by the task and $p(\mathbf{r})$ is the partition function. The maximum of the posterior distribution was chosen as the estimate of the saccade direction corresponding to a neuronal population activity (MAP estimation). The prediction accuracy was finally defined by the proportion of veridical estimates evaluated across all trials.

For the instructed trials, the neuronal population response was the spike count of each neuron across a forward temporal window starting 100 ms after the beginning of the memory epoch to avoid the responses corresponding to the visual cue bleeding into the memory epoch. The end of the temporal window was defined by the longest memory duration that was common to all randomized memory epochs across trials. The posterior was evaluated using a leave-one-out cross-validation scheme to avoid overfitting and assure a good generalization (44). Specifically, the parameters of the posterior were built on a subset of trials, and the posterior was evaluated on another nonoverlapping subset.

For the brain control trials, the neuronal population response was the spike count of each neuron across a backward temporal window ending on the time the saccade direction estimate was generated. The length of the window was identical to the length used for the instructed trials, namely the longest common memory duration, to avoid artifacts based on inferring behaviors from spike counts collected over temporal windows of different length. The posterior here was computed without cross-validation because the parameters of the posterior were derived from a novel trial that had never been seen before and was thus different from the brain control trials used to compute the posterior. It was computed for learning with memory using all of the previous trials and for learning without memory using only the instructed trials.

BMI. In the brain control trials, we sampled neuronal responses every 100 ms (our definition of real time). Using the neuronal population response as defined above, the BMI generated a saccade direction estimate every 100 ms based on the collected neuronal data until it reached a confidence level. The confidence metric was the stability of the decoder's estimates: when the decoder generated the same estimate for a consecutive 400 ms (a timescale we empirically found to be sufficiently long to yield steady-state estimates), this estimate was deemed to be stable and hence accepted. If no stable estimate was found within 3 s, the decoder timed out, and the cued target location was used as an incorrect estimate. The BMI was thus self-paced. Furthermore, the BMI was allowed to learn by incrementally updating its neuronal response history. The database of spike counts from the instructed and preceding brain control trials was updated by the spike count used to generate the online direction estimate. The Bayesian decoding model needed only to be updated for the cued condition, thus allowing for a fast update of the decoder to satisfy the constraints imposed by a real-time BMI. This recursive Bayesian framework corresponds to learning with feedback because the spike count database is updated using the cued saccade direction and not the estimated one because the latter may possibly make the estimates diverge. Finally, the saccade direction estimate was computed across the 3×3 nodes of the grid because saccade direction and eye position were, by construction, independent variables. As such, the predictions of the BMI generalized across different initial eye positions.

ACKNOWLEDGMENTS. We thank M. Yanike for helpful comments on the manuscript, K. Pejsa for animal care, and V. Shcherbatyuk and T. Yao for technical and administrative assistance. This research was supported by National Institutes of Health Research Grants EY005522, EY013337, and EY007492 and the Boswell Foundation.

1. Georgopoulos AP, Taira M, Lukashin A (1993) Cognitive neurophysiology of the motor cortex. *Science* 260(5104):47–52.
2. Wessberg J, et al. (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408(6810):361–365.
3. Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA (2004) Cognitive control signals for neural prosthetics. *Science* 305(5681):258–262.
4. Lebedev MA, et al. (2005) Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *J Neurosci* 25(19): 4681–4693.
5. Hochberg LR, et al. (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442(7099):164–171.
6. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB (2008) Cortical control of a prosthetic arm for self-feeding. *Nature* 453(7198):1098–1101.
7. O'Doherty JE, et al. (2011) Active tactile exploration using a brain-machine-brain interface. *Nature* 479(7372):228–231.
8. Andersen RA, Burdick JW, Musallam S, Pesaran B, Cham JG (2004) Cognitive neural prosthetics. *Trends Cogn Sci* 8(11):486–493.
9. Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230(4724):456–458.
10. Andersen RA, Essick GK, Siegel RM (1987) Neurons of area 7 activated by both visual stimuli and oculomotor behavior. *Exp Brain Res* 67(2):316–322.
11. Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70(1):216–220.
12. Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L (1990) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10(4):1176–1196.
13. Schall JD, Thompson KG (1999) Neural selection and control of visually guided eye movements. *Annu Rev Neurosci* 22:241–259.
14. Goldberg ME, Bisley J, Powell KD, Gottlieb J, Kusunoki M (2002) The role of the lateral intraparietal area of the monkey in the generation of saccades and visuospatial attention. *Ann N Y Acad Sci* 956:205–215.
15. Costa VD, Averbeck BB (2014) Looking into the future. *eLife* 3:e03146.
16. Graf ABA, Andersen RA (2014) Inferring eye position from populations of lateral intraparietal neurons. *eLife* 3:e02813.
17. Moritz CT, Fetz EE (2011) Volitional control of single cortical neurons in a brain-machine interface. *J Neural Eng* 8(2):025017.
18. Hwang EJ, Bailey PM, Andersen RA (2013) Volitional control of neural activity relies on the natural motor repertoire. *Curr Biol* 23(5):353–361.
19. Pesaran B, Andersen RA (2010) Prosthetic devices and methods and systems related thereto. US Patent 7,797,040.
20. Brody CD, Romo R, Kepecs A (2003) Basic mechanisms for graded persistent activity: Discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol* 13(2):204–211.
21. Dayan P, Abbott LF (2001) *Theoretical Neuroscience* (MIT Press, Cambridge, MA).
22. Jaynes ET (2003) *Probability Theory: The Logic of Science* (Cambridge Univ Press, Cambridge, UK).
23. Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Annu Rev Neurosci* 26:381–410.
24. Graf ABA, Kohn A, Jazayeri M, Movshon JA (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat Neurosci* 14(2):239–245.
25. Wolpert DM, Diedrichsen J, Flanagan JR (2011) Principles of sensorimotor learning. *Nat Rev Neurosci* 12(12):739–751.
26. Gilbert CD, Sigman M, Crist RE (2001) The neural basis of perceptual learning. *Neuron* 31(5):681–697.
27. Koralek AC, Jin X, Long JD, 2nd, Costa RM, Carmena JM (2012) Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483(7389): 331–335.
28. Koralek AC, Costa RM, Carmena JM (2013) Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* 79(5):865–872.
29. Graf ABA (2014) From neuronal populations to behavior: A computational journey. *Front Comput Neurosci* 8:81.
30. Nicolelis MA (2003) Brain-machine interfaces to restore motor function and probe neural circuits. *Nat Rev Neurosci* 4(5):417–422.
31. Donoghue JP (2008) Bridging the brain to the world: A perspective on neural interface systems. *Neuron* 60(3):511–521.
32. Leuthardt EC, Schalk G, Roland J, Rouse A, Moran DW (2009) Evolution of brain-computer interfaces: Going beyond classic motor physiology. *Neurosurg Focus* 27(1):E4.
33. Committee for the Update of the Guide for the Care and Use of Laboratory Animals (2011) *Guide for the Care and Use of Laboratory Animals* (National Academies Press, Washington, DC), 8th Ed.
34. Foldiak P (1993) The ideal homunculus: Statistical inference from neural population responses. *Computation and Neural Systems*, eds Eeckman FH, Bower JB (Kluwer Academic Publishers, Norwell, MA), pp 55–60.
35. Seung HS, Sompolinsky H (1993) Simple models for reading neuronal population codes. *Proc Natl Acad Sci USA* 90(22):10749–10753.
36. Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. *J Comput Neurosci* 1(1-2):89–107.
37. Sanger TD (1996) Probability density estimation for the interpretation of neural population codes. *J Neurophysiol* 76(4):2790–2793.
38. Oram MW, Földiák P, Perrett DI, Sengpiel F (1998) The 'Ideal Homunculus': Decoding neural population signals. *Trends Neurosci* 21(6):259–265.
39. Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11(1):91–101.
40. Sanger TD (2003) Neural population codes. *Curr Opin Neurobiol* 13(2):238–249.
41. Averbeck BB, Lee D (2004) Coding and transmission of information by neural ensembles. *Trends Neurosci* 27(4):225–230.
42. Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nat Neurosci* 7(5):456–461.
43. Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9(11):1432–1438.
44. Duda R, Hart P, Stork D (2001) *Pattern Classification* (John Wiley & Sons, New York).