Published in final edited form as:

J Neural Eng. 2014 February ; 11(1): 016004.

Intention Estimation in Brain Machine Interfaces

Joline M. Fan^{1,2}, Paul Nuyujukian^{1,3}, Jonathan C. Kao⁴, Cynthia A. Chestek^{4,5}, Stephen I. Ryu^{4,6}, and Krishna V. Shenoy^{1,4,5,7}

¹Department of Bioengineering, Stanford University, Stanford, CA

²School of Medicine, University of California San Francisco, San Francisco, CA

³School of Medicine, Stanford University, Stanford, CA

⁴Department of Electrical Engineering, Stanford University, Stanford, CA

⁵Stanford Institute for Neuro-Innovation and Translational Neuroscience, Stanford, CA

⁶Department of Neurosurgery, Palo Alto Medical Foundation, Palo Alto, CA

⁷Department of Neurobiology, Stanford University, Stanford, CA

Abstract

Objective—The objective of this work was to quantitatively investigate the mechanisms underlying the performance gains of the recently reported ‘recalibrated feedback intention-trained Kalman Filter’ (ReFIT-KF).

Approach—This was accomplished by designing variants of the ReFIT-KF algorithm and evaluating training and online data to understand the neural basis of this improvement. We focused on assessing the contribution of two training set innovations of the ReFIT-KF algorithm: intention estimation and the two-stage training paradigm.

Main Results—Within the two-stage training paradigm, we found that intention estimation independently increased target acquisition rates by 37% and 59%, respectively, across two monkeys implanted with multiunit intracortical arrays. Intention estimation improved performance by enhancing the tuning properties and the mutual information between the kinematic and neural training data. Furthermore, intention estimation led to fewer shifts in channel tuning between the training set and online control, suggesting that less adaptation was required during online control. Retraining the decoder with online BMI training data also reduced shifts in tuning, suggesting a benefit of training a decoder in the same behavioral context; however, retraining also led to slower online decode velocities. Finally, we demonstrated that one- and two-stage training paradigms performed comparably when intention estimation is applied.

Significance—These findings highlight the utility of intention estimation in reducing the need of adaptive strategies and improving the online performance of BMIs, helping to guide future BMI design decisions.

shenoy@stanford.edu.

AMS classification scheme numbers: 93B51, 93B05, 93D21, 92C20

Introduction

Brain-machine interfaces (BMIs) are medical systems that translate neural activity from the brain into control signals that guide prosthetic devices. BMIs may ultimately offer disabled patients a way to interact with the environment, including restoring the ability to conduct activities of daily living. Intra-cortical BMIs operate by measuring neural activity, such as action potentials, and mapping these signals to relevant control signals, such as muscle activation, force, or end effector kinematics, i.e. the position or velocity of a prosthetic arm or a computer cursor on a screen (for recent reviews see e.g., [1–8]). In the past few years, there has been considerable effort devoted to making BMI technologies more clinically viable, including an ongoing FDA phase-I clinical trial (e.g., [9–14]) and efforts to move beyond the need for multi-wire cables connected to the subject by developing electronic circuits to wirelessly transmit neural signals (for recent reviews see e.g., [15, 16]). However, despite these advances, the performance of the algorithm that maps neural activity to kinematics (i.e., the decoder) remains an important limitation to the more wide spread use and efficacy of BMIs (e.g., [6,17,18]).

One approach to improving the performance of closed-loop BMIs over time is through behavioral learning and adaptation. In these paradigms, the decoding algorithm can be arbitrarily assigned [19, 20] or loosely correlated to native arm control [21, 22]. Over time, subjects can improve their performance of fixed decoders through adaptive strategies, which correlate to changes in neural tuning [23–25]. After a sufficient amount of time, observed changes in the neural data between native arm movement and BMI control stabilize [26]. Such shifts in neural tuning may be attributed to inherent changes in the neural representation or to suboptimal ‘out-of-the-box’ decoding quality, requiring the subject to adopt a new control strategy [20].

Another approach is to develop algorithms that mimic the neural-to-kinematic biological mapping as closely as possible, which are termed biomimetic decoders [27, 28]. Here, the goal is to minimize the need for behavioral learning and adaptation by building a decoder whose control strategy is similar to that of native arm movement, with the aim of maximizing ‘out-of-the-box’ performance. A central assumption enabling these decoders is that neural firing characteristics should not change significantly from training to online testing under BMI control if the decoder has high predictive power and is controlled in a manner similar to that of the native limb. As we will demonstrate, shifts in neural tuning between training and testing sets can therefore be used as an indicator that a decoder is relying more heavily on an adaptive strategy rather than a good biomimetic fit. It is important to note that while the basic neuroscience question of how cortical neural activity relates mechanistically to movements (kinematics) (e.g., [29–32]), muscles (e.g., [33–35]), and internal cortical dynamics (e.g., [7, 36–39]) is still an open and contentious question (e.g., [40–43]), biomimetic controllers seek merely to provide a net neural-to-kinematics description or mapping that does not deviate dramatically from experimental observations of neural-to-kinematic relationships in the relevant workspaces. Strategies of adaptation and biomimetic fitting can be used in combination, particularly when the biomimetic decoding strategies does not well mimic native arm control.

Recently we reported the design and characterization of a high-performing algorithm and training procedure, termed the ‘recalibrated feedback intention-trained Kalman filter’ (ReFIT-KF) [44–48], which significantly improved performance of the highest-performance algorithm (i.e., the velocity based Kalman filter (VKF), [10, 49]). ReFIT-KF involves a two-step training procedure, in which an initial decoder is trained off of native arm reaches and subsequently used for a second online training session. A second decoder is then generated from the online training data augmented by intention estimation modifications [48]. While the various contributions of each innovation of ReFIT-KF have been detailed in a previous publication [48], the specific mechanisms related to the quantitative impact of intention estimation modifications on the resulting fit and new decoder have not yet been investigated, and may span adaptive and biomimetic control strategies.

In this study, we investigated the neural mechanisms contributing to the performance gains of ReFIT-KF. We asked whether the key to high performance of ReFIT-KF is the retraining step, the intention estimation training set modifications, or a combination of the two. For instance, the two-stage training procedure may be important for adaptive processes, such as normalizing contextual changes between native arm reaching and BMI control [20]. Alternatively, the modifications to the training data may enable the neural data to be more precisely correlated to the kinematic data (i.e., the intended kinematics), thereby leading to a more biomimetic decoder. As will be further explained, training set modifications are applied to the second stage of training in the two-stage training procedure. By comparing the performance of various algorithms derived from the ReFIT-KF training philosophy, we demonstrate that intention estimation modifications are largely responsible for the improvement in overall performance. We investigate the effects intention estimation modifications have on the training data and utilize the shifts in neural tuning between the training and testing set as a tool to assess the degree of adaption required for online decoding. Finally, we show that intention estimation modifications can be directly applied to native arm reaching training data to generate a one-stage decoder that performs comparably to ReFIT-KF. It is important to note that the ReFIT-KF algorithm also features a positional feedback signal, which we do not examine in this study.

Materials and Methods

Behavioral Task

All protocols and experimental methods were approved by the Stanford University Institutional Animal Care and Use Committee (IACUC). Two adult rhesus macaques (J and L) were trained to make point-to-point reaches towards targets presented in a 3-D environment (MSMS, MDDE, e.g., [50]). While the monkeys reached in a 3-D environment to targets rendered on a 2-D fronto-parallel plane, and thus primarily exhibited 2-D reaches, the monkeys controlled a virtual cursor moving strictly in the 2-D fronto-parallel plane. A 3-D percept was enabled by a stereographic display using fused mirrors and separate LCD monitors, which refreshed at 120 Hz. Hand position was recorded using an infrared reflective bead taped to the monkeys’ index fingers and a Polaris tracking system (Northern Digital, Ontario, Canada), sampling at 60 Hz.

A center-out task was performed, prompting the cursor-based acquisition of targets that alternated between the middle of a workspace and the perimeter of either an 8 cm or 12 cm circle. Successful trials required that the monkey hold the cursor within the target acquisition boundaries (4.8 cm by 4.8 cm) for 500 ms and acquire the target within a maximum trial time of 5 seconds. A liquid reward was provided for each successful trial.

All experiments were performed in an A-B-A fashion with 250 center-out trials per block. This setup ensured that factors such as monkey motivation and channels' contents were the same across different decoders within a given experimental day. Two experimental days for Monkey J (2012-07-26, 2012-07-27) and for Monkey L (2011-07-21, 2011-07-22) were used for the comparison between ReFIT-KF and Re-KF (see *The ReFIT algorithm* section). Across two experimental days, Monkey J completed 1222 center-out ReFIT-KF trials and 1208 center-out Re-KF trials; Monkey L completed 590 center-out ReFIT-KF trials and 550 center-out Re-KF trials. t-tests were conducted to evaluate if the distribution of last acquire times across different decoders were statistically significant ($p < 0.05$).

Four experimental days for Monkey J (2011-05-12, 2011-05-13, 2011-05-16, 2011-05-17) and three experimental days for Monkey L (2011-04-07, 2011-04-08, 2011-04-15) were used for the comparison between ReFIT-KF and FIT-KF (see *The ReFIT algorithm* section). Monkey J completed 1376 center-out ReFIT-KF trials and 2381 center-out FIT-KF trials; Monkey L completed 1091 center-out ReFIT-KF trials and 1417 center-out FIT-KF trials. For comparing the performance and tuning characteristics of FIT-KF versus KF, three experimental days for Monkey J (2011-05-19, 2011-05-20, 2011-05-24) and three experimental days for Monkey L (2011-04-18, 2011-04-20, 2011-04-21) were used. Monkey J completed 1300 center-out FIT-KF trials and 1115 center-out KF trials; Monkey L completed 724 center-out FIT-KF trials and 449 center-out KF trials. Again, t-tests were conducted to evaluate if the performance across different decoders, as measured by last acquire times, is statistically significant ($p < 0.05$).

Electrophysiology

Monkeys J and L were implanted with 96-electrode Utah arrays (400 μm electrode separation and 1 mm penetration depth; Blackrock Microsystems Inc., Salt Lake City, UT) in primary motor cortex (M1) and/or dorsal premotor cortex (PMd) contralateral to the reaching arm. In Monkey J, an M1 array and a PMd array (192 total channels) were implanted on 24 August 2009. In Monkey L, an array straddling the PMd/M1 border was implanted on 22 January 2008. Data for Monkey J were collected across 69 sessions from 2 September 2010 to 7 July 2012. Data for Monkey L were collected across 63 sessions from 1 July 2010 to 22 July 2011.

During an experimental session, neural signals measured from the electrode arrays were initially processed using a Cerebus recording system (Blackrock Microsystems Inc., Salt Lake City, UT). Spiking activity was aggregated by setting a -4.5 root mean square threshold on each channel without any spike sorting. This has been shown to produce a stable source of action-potential based neural signals [51,52]. Behavioral task control and neural decodes ran on separate PCs using a Simulink / xPC platform (Mathworks, Natick, MA), yielding a low communication latency of 3 ms.

The ReFIT-KF algorithm

The ReFIT-KF algorithm entails a two stage-training paradigm [44, 48]. In the first stage of training, the subject completes 500 trials of a 12 cm, center-out task with 8 targets using active arm and hand movements. The kinematics of the cursor movements and the simultaneous neural recordings are used to build a ‘first-pass decoder,’ such as a velocity or position-velocity based Kalman filter (KF). In the second stage of training, the subject completes an additional 500 trials of a 12 cm, center-out task using the KF. The number of trials between the first and second stages were matched to equalize the sampling and variability of target selection in the data. The online training data, comprised of the simultaneously recorded kinematic and neural data during online KF control, are used to build a ‘final decoder.’ If intention estimation is applied to the online training data, this final decoder is precisely the ReFIT-KF filter [44,48]. If intention estimation is not applied, the final decoder is what we term the ‘recalibrated Kalman filter’ (Re-KF).

In this study, ‘intention estimation’ refers to training set modifications, in which the cursor velocities in the training set are rotated to point towards the direction of the target and their magnitudes are set to zero during the periods the cursor is successfully held at the target [48]. It is important to note that the final retrained decoder receives no further target information, and does not make use of prior knowledge regarding target location as some algorithms do for separate reasons [53–57]. These intention estimation modifications, inferred by the experimenter, are based on the assumption that the monkey intends to move to the target at all times and intends to cease moving when holding on the target to select it. For all decoders, neural data were binned in non-overlapping 50 ms intervals. This has been shown previously to be long enough to provide reliable estimates of neural firing rates, but short enough to not introduce deleteriously large time-lags into the BMI closed-loop control [44,47,48,58].

If intention estimation is applied to the initial arm reaching data, the final one-stage decoder is termed the ‘feedback-intention-trained Kalman filter’ (FIT-KF). Here, the initial arm reaching training task is designed to allow the kinematics data to more closely model the distribution of the kinematic data obtained in the second training stage of the two-stage training paradigm. Specifically, we used the training task where a target is randomly placed in a 16 by 16 cm² workspace as opposed to the original 12 cm center-out task, in order to better sample the workspace area. In addition, we eliminated slow velocities at 1) the beginning of the reach (150 ms for Monkey L, and 250 ms for Monkey J) and 2) after the monkey reaches the target but before the monkey successfully holds the target. This is done in order to more closely match the distribution of velocity magnitudes seen in the two-stage training method.

Preferred Direction

Preferred directions (PDs) were calculated in two different ways. First, when characterizing the directional tuning of a voltage-thresholded electrode channel during online control, the PDs were determined based on assigning spike counts to the prompted direction of movement. We assume that within a time period (150 ms to 550 ms after target onset for each trial) the monkey intends to move toward the prompted target and neural activity thus

best reflects this intention. Second, when characterizing the directional tuning of the training data, the PDs are viewed from the perspective of the decoder and are based on assigning spike counts to cursor kinematics in 50 ms intervals and applying intention estimation modifications as indicated. In these cases, the tuning of the training data is determined after applying intention estimation modifications in order to fully capture the experimenter's "best-guess" of the neural tuning for building a high performing biomimetic decoder.

In both cases, the PD was calculated by fitting mean spike counts per direction (0 to 2π in intervals of $\pi/4$) with a cosine of frequency 2π , and determining the phase minimizing the least squares error of the sinusoidal fit [29]. Only channels that were cosine tuned in all control modalities ($R^2 > 0.5$) for an individual day were considered in the analysis characterizing changes in directional tuning (30 - 45 channels for Monkey L and 75-85 channels for Monkey J).

Statistically significant changes in PDs were determined using a bootstrap analysis. Firing rates were resampled 1000 times per direction to obtain a distribution of PDs for each channel and condition (i.e. Hand, KF, Re-KF, ReFIT-KF), and the means of these distribution were subtracted in order to simulate the null hypothesis of identical means. Sets of PDs were then randomly drawn from the zero-mean distributions, and the PDs were subtracted to form a control distribution to simulate the null hypothesis. To form the distribution of shifts in PD between two conditions, 1000 PDs were randomly sampled from one distribution and subtracted from a random sample of PDs from the second distribution. If the absolute mean shift in PD across two conditions for a given channel was greater than 95% of the shifts based on noise, it is said to be significant. F-tests were conducted to evaluate if the variance in the distribution of shifts in PDs between different training sets and their respective online control was statistically significant ($p < 0.05$).

Information Profiles

Two metrics were used for neuron ranking and evaluating information density for a given channel. The first metric is based on mutual information, as shown in Equations 1–3. Here, $p(x)$ denotes the probability that a channel fires x spikes in a given 50 ms bin width, X represents all such encountered x , M is the number of reach directions, and y_j is the reach direction, which can either be the prompted direction, cursor direction, or cursor direction modified by intention estimation. Then, $p(x|y_i)$ denotes the probability that a channel fires x spikes in a given 50 ms bin when reaching in direction y_i . We calculate mutual information between binned spike counts and reach direction. We use the prompted reach direction for analyzing the ordered shifts in PD, and the cursor direction for analyzing the benefits of intention estimation. For each channel, the entropy, $H(X)$, and the entropy conditioned on reach direction, $H(X|Y)$, are calculated to determine the mutual information, $I(X; Y)$. This statistical metric captures how informative neural firing is of reach direction. We note that a sufficient amount of data was used to estimate the probability distributions such that a bias-corrected version of mutual information (e.g. [59]) was not required.

$$H(X) = - \sum_{x \in \chi} p(x) \log(p(x)) \quad (1)$$

$$H(X|Y) = - \sum_{j=1}^M p(y_j) \sum_{x \in \chi} p(x|y_j) \log(p(x|y_j)) \quad (2)$$

$$I(X;Y) = H(X) - H(X|Y) \quad (3)$$

The second information metric relies on the extent of direction tuning of the channel, as characterized by the modulation over variance ratio (MDV). This metric has been shown to lead to similar cell rankings as model-dependent ranking methods [60]. For the modulation over variance ratio, the mean and variance of the firing rate, binned across the movement direction of the cursor (with or without intention estimation), are calculated. For each channel, the minimum mean firing rate is subtracted from the maximum mean firing rate, which is termed the ‘modulation’. Because some channels may have large modulation but may be extremely noisy and variant, the modulation is normalised by the mean variance of the firing rates, across all directions, which gives rise to the MDV metric.

Results

Online Performance

In order to investigate the mechanism by which ReFIT-KF improves performance, we assessed the individual contribution of intention estimation and the two-stage training procedure. To evaluate the effects of intention estimation, we compared the online performance of two decoders that differ only by whether intention estimation was applied or not: the recalibrated Kalman filter (Re-KF) and the ReFIT-KF. In these experiments and analyses, neither decoders incorporate positional feedback into the algorithm.

Figure 1 illustrates the two-stage training paradigm by which these two decoders are generated. In the first stage, the subject completes a reaching task under hand control. The cursor kinematics and simultaneous neural activity, which comprise the ‘training data,’ are used to build a ‘first-pass BMI decoder’ (KF) for online control. In the second stage of training, the subject completes an additional reaching task using the KF. Training data from the online KF task is then used to build the ‘final BMI decoder’ (i.e. Re-KF or ReFIT-KF). The bolded flow diagram in Fig. 1 describes the training framework of the ReFIT-KF algorithm in which intention estimation modifications are applied to the KF training data. The Re-KF decoder is generated from the same two-stage training procedure; however, intention estimation modifications are not applied in the second stage of training.

Decoder performance metrics are obtained from a center-out reaching task in which the subject moves the cursor from the center of a workspace to peripheral targets and holds the cursor on the target for 500 ms to activate the selection. The subject must then move the cursor back to the center of the workspace for the next selection. In Fig. 2, the online closed-loop BMI performance of KF, Re-KF, and ReFIT-KF are shown, along with the natural hand movement for comparison, from a 12 cm center-out reaching task.

In Fig. 2a, c, the mean distance of a reach to a target is plotted as a function of duration of the reach for each decoder. The thicker line later in time along each curve indicates the ‘dial-

in-time'. This is the average time that passes before the monkey successfully holds the target, after first reaching the target. The 500 ms hold time itself is not shown or included in this dial-in-time calculation. The dial-in-time can be used as a measure of the fine control that enables the monkey to narrow in on a target. The time preceding the dial-in-time is the average time required to move from the center of the workspace to a peripheral target. This period is more related to speed and the global directionality of a movement. The overall acquire time, defined as the time of target presentation until the time of final target acquisition (not including the 500 ms hold time), for each decoder is plotted in Fig. 2b, d. The time below the white lines denotes the average time to first reach the target. Additional time is often needed to achieve final target acquisition due to behaviors such as target overshoot; the time above the white line denotes the dial-in-time.

Across two experimental days per subject, ReFIT-KF reached targets 37% faster than Re-KF in Monkey L ($p < 0.001$, t-test) and 59% faster in Monkey J ($p < 0.001$), as determined by normalizing the difference in last acquisition times by the Re-KF last acquisition time. Acquisition time was improved primarily due to a decrease in the dial-in time. Here, both models were trained on a two-stage training paradigm but ReFIT-KF incorporated intention estimation modifications in its training set, suggesting that intention estimation modifications independently improves performance.

The contribution of the retraining step in the two stage training paradigm is less clear when evaluating performance. To evaluate the benefit of retraining on performance, we compared the performance of KF with Re-KF, which is built using the same decoding algorithm as KF but from online KF training data. In Fig. 2, Re-KF is shown to have slower acquisition times as compared to KF, suggesting that retraining leads to slower velocities; this will be addressed in the discussion section.

Importantly, we note that retraining with intention estimation (ReFIT-KF, Fig. 2, gold) leads to a finer degree of online control than simply retraining alone (Re-KF, Fig. 2, blue). In the following two sections, we investigate how intention estimation modifications alter the training data characteristics; specifically, we examine its effects on the second stage of training in the two-stage training paradigm (i.e. as in ReFIT-KF). We later show how intention estimation can also be applied to a one-stage training paradigm.

Effects of Intention Estimation Modification on Training Data

Intention estimation modifications alter training data by associating neural firing rates with rotated kinematics that are presumably better reflected in the neural data. To evaluate the effects of intention estimation modifications on the training data, we examined the changes in the tuning curve characteristics (i.e. modulation and per-channel variance) of the training data before and after intention estimation modifications were applied to the data.

As exemplified in Fig. 3a and b, intention estimation modifications are shown to alter tuning curve characteristics, as seen by comparing the light grey tuning curve, in which intention estimations are not applied, to the black tuning curve, in which they are applied. The tuning curves of Fig. 3a and b exemplify an increase in modulation by 1.36 and 1.25 fold, respectively, and a decrease in firing rate variability across directions by 0.88 and 0.83 fold,

respectively, when intention estimation is applied versus not applied. Error bars in the figure denote standard error. Modulation per channel is defined by the difference in the minimum and maximum firing rates of the fitted tuning curve; the per-channel firing rate variance is the variance of the firing rates in each direction, averaged across all directions, for a given channel.

The effects of intention estimation on the change in modulation and variance were examined across 60 days of training data for Monkey J (2011-06-10 to 2010-09-02) and 55 days of training data for Monkey L (2011-05-19 to 2010-07-01). Changes in the tuning curves for each channel were averaged across all days for each monkey. Histograms of the number of channels displaying changes in modulation and variance when intention estimation was applied to the training data are shown in Fig 3c, d for Monkey J (blue) and Monkey L (red). On average, intention estimation modifications increased the modulation by 1.54 ± 0.12 fold (mean $\pm$ std, significantly different than mean of 1, $p < 0.001$, t-test, Monkey J) and 1.38 ± 0.16 fold ($p < 0.001$, Monkey L), as calculated by the ratio of the modulation with intention estimation relative to that without intention estimation. Intention estimation also decreased the per-channel variance across reach directions by 0.90 ± 0.03 (mean $\pm$ std, significantly different than mean of 1, $p < 0.001$, t-test, Monkey J) and 0.95 ± 0.05 ($p < 0.001$), as calculated by the ratio of per-channel variance with intention estimation relative to that without intention estimation. Intention estimation modifications reclassify neural firing rates with different rotated kinematics. These modifications are shown here to enhance tuning characteristics by increasing the modulation and decreasing the per-channel variance of the tuning curve.

We further evaluate how the information content changes when intention estimation modifications are applied to the training data by using the metrics of modulation over variance (MDV) and mutual information (MI). The MDV metric is determined by the ratio of the modulation over the per-channel firing rate variance. The ratios are then averaged over all viable channels and experimental days. Across the same 60 days of training data for Monkey J (2010-09-02 to 2011-06-10) and 55 days of training data for Monkey L (2010-07-01 to 2011-05-19), intention estimation modifications to the data are found to increase the MDV by a factor of 1.60 ± 0.16 (mean $\pm$ std, statistically significant from mean of 1, $p < 0.001$, t-test, Monkey J) and 1.34 ± 0.20 ($p < 0.001$, Monkey L).

The MI metric evaluates the information shared between the neural firing rates and kinematic training data (see methods). Applying intention estimation to the data is found to increase the MI by a factor of 1.57 ± 0.21 (mean $\pm$ std, statistically significant from mean of 1, $p < 0.001$, t-test, Monkey J) and 1.54 ± 0.22 ($p < 0.001$, Monkey L). This suggests that intention estimation plays a critical role in making the training data more harmonious by enabling neural firing rates to be associated with kinematics that are better reflected in the neural data. Ultimately, intention estimation leads to an increased correlation between neural and kinematic training data, improving the online performance of the decoder. To reiterate, the above intention estimation modification analyses were performed offline, and no online retraining occurred before these tuning curve and mutual information improvements were found.

Preferred Directions

Given that intention estimation can alter the training data, we evaluated how the tuning characteristics of the training data matches that of online cursor control. If the decoder captures the natural relationship between neural firing and kinematics precisely, few new adaptive strategies would be required for high controllability, as the decoder becomes controlled in a manner similar to that during the training period. As a result, one would expect fewer shifts in the tuning properties between the training data and the online closed-loop BMI data.

In Fig. 4, we evaluated the shifts in preferred direction (PD) of the neural data for three training sets and their respective online performance, in order to probe the components of the two stage paradigm. We measured the shifts in PD for each step illustrated in Fig. 1 (i.e. stage i, iia, iib) to assess the degree to which retraining and intention estimation modifications can help minimize changes in tuning. Fig. 4a shows histograms of the shifts in PD between the training and testing data; red indicates PDs that are statistically different between the training and testing data. The green line depicts the distribution of change in PD as a result of noise. A wider distribution of shifts in PD indicates that the neural characteristics changed notably between the training set and online control sessions. Fewer shifts in the PD suggest that the neural characteristics of the training data are predictive of online control neural characteristics and are thus preserved during online decoder control. The shifts in PD for all viable channels, ordered from the most to least informative channels as measured by MI, are plotted in Fig. 4b. Figure 5 is a replica of Fig. 4 but in Monkey L.

As shown in Fig. 4, we first evaluated tuning shifts between training data derived from native hand control versus the first pass online control, KF, as in stage i. Over two experimental days per monkey, 75% of channels in Monkey J and 58% in Monkey L had significant changes in PD. These shifts in the PDs may be attributed to altered online neural characteristics that compensate for changes in behavioral context or for poor decoder quality. If these shifts were attributable to changes in the behavioral context (i.e. native arm control versus BMI control), we would expect fewer shifts in PD if the same decoder were built from BMI controlled training data, i.e. stage iia. On the other hand, if the shifts in PD are attributable to adaptation of subjects control strategy which result from poor decoder performance, we would expect fewer shifts in PD if the decoder were more biomimetic decoder as when intention estimation modifications are applied, i.e. stage iib.

In the second column of Fig. 4, we illustrate the shifts in PD between the online training model versus the final retrained online model, Re-KF, as in stage iia. Over the same experimental days, 67% of channels in Monkey J and 50% in Monkey L have significant changes in PD; furthermore, the variance in the distribution of PD shifts between stages i and iia decreased by 1.32 fold ($p < 0.001$, F-test, Monkey J) and 1.26 fold ($p = 0.05$, Monkey L) from data aggregated across two experimental days. These findings indicate that there is a benefit to retraining the decoder on a training set of the same behavioral context.

In the third column of Fig. 4, we demonstrate that adding intention estimation modifications to the training data can further minimize the shifts in PD. Over two experimental days, 34% of channels in Monkey J and 35% in Monkey L have significant shifts in PD, fewer than the

prior stages. The variance in the distribution of PD shifts between stage i to stage iib decreased by 3.62 fold ($p < 0.001$, F-test, Monkey J) and 1.55 ($p < 0.001$, Monkey L) across two experimental days. Furthermore, the variance in the PD shifts between stage iia and iib decreased by 2.75 fold ($p < 0.001$, F-test, Monkey J) and 1.23 ($p = 0.08$, Monkey L).

In Fig. 4b, we note that less contributory channels tend to have larger shifts in PD as compared to more contributory channels. In particular, the mean absolute shift in PD in the more contributory channels (i.e. the first half of channels in order of contribution by MI, channels 1-43) was smaller than that in the less contributory channels (i.e. the second half of channels in order of contribution, channels 44-85). For instance, in stage i of Fig. 4b, the mean shift in PD for the more contributory channels is 30.3° as compared to 37.6° in the less contributory channels. In stage iia, the mean shift in PD is 16.5° in the more contributory channels versus 21.1° in the less contributory channels, and in stage iib, 5.4° versus 9.2° , respectively. In stage i of Fig. 5b, we note a similar trend such that the mean shift in PD for more contributory channels is 17.3° as compared to 36.6° in the less contributory channels. In stage iia, the mean shift in PD is 17.9° versus 25.5° , and in stage iib, 8.6° versus 16.9° , respectively.

The combined effects of retraining and intention estimation on the shifts in PD were evaluated across 60 days of training data for Monkey J (2010-09-02 to 2011-06-10) and 55 days with Monkey L (2010-07-01 to 2011-05-19). On each experimental day, the ReFIT algorithm was built and performed online. Just as before, the shifts in PD for stage i were compared to that of stage iib (i.e. ReFIT-KF). We note that averaged across 60 days for Monkey J, 71% of channels significantly changed their PD during stage i as compared to 34% of channels during stage iib; furthermore, the variance in shifts in PD decreased by 3.00 fold ($p < 0.001$, F-test) between stage i and stage iib. Averaged across 55 days in Monkey L, 51% of channels significantly changed their PD during stage i as compared to 27% of channels in stage iib; the variance in shifts in PD decreased by 2.16 fold ($p < 0.001$) between stage i and stage iib.

Developing a one-stage decoder

In the previous section, we showed that both retraining and intention estimation reduces the neural shifts between training data and online control; however, we also noted that retraining reduced online cursor velocities. We therefore explored the feasibility of applying intention estimation to a one-stage training paradigm in order to maintain high speeds and achieve good control. Intention estimation modifications in the context of hand movements corrects for intrinsic movement variability and noisiness relative to cortical neural activity (e.g., [61–63]); furthermore, it infers that when the subject is trying to stop, he is commanding a zero velocity.

As shown in Fig. 6, a FIT-KF decoder was built by applying intention estimation modifications to a new hand training set. A new hand training set comprised of reaching to randomly placed targets, as described in the methods, was developed to approximately mimic the broader spectrum of kinematics the first-pass decoder obtains during the two-stage training paradigm, as compared to kinematics of the native arm training set. We compare two training paradigms in which both decoders are trained from hand data but one

paradigm also includes intention estimation. To maximize performance, we also incorporated the feedback control innovation in the comparison of both algorithms unlike the previous set of experiments. We, however, do not explicitly tease apart the contribution of feedback control in this study.

As shown in Fig. 7, FIT-KF performed comparably to ReFIT-KF. Furthermore, applying intention estimation modifications to the hand training data increased the rate of successfully acquiring targets by 30% in Monkey L ($p < 0.001$, t-test) and 50% in Monkey J ($p < 0.0001$), as measured by normalizing the difference in acquire times between KF and FIT-KF by the acquisition time of KF. In this figure, performance metrics were based on an 8 cm center-out reaching task, in which ReFIT-KF and FIT-KF were directly compared in an A-B-A fashion (see Methods) across four experimental days for Monkey J (2011-05-12, 2011-05-13, 2011-05-16, 2011-05-17) and three experimental days for Monkey L (2011-04-07, 2011-04-08, 2011-04-15). In addition, KF and FIT-KF were directly compared across three experimental days for Monkey J (2011-05-19, 2011-05-20, 2011-05-24) and three experimental days for Monkey L (2011-04-18, 2011-04-20, 2011-04-21). Direct comparisons of FIT-KF, ReFIT-KF and KF were collapsed together to give rise to the performance metrics of Fig. 7.

As before, intention estimation applied to hand reaching training data is shown to reduce the shifts in PD between the training and online testing data. In Fig 8 and 9, a histogram of the shifts in PD are shown between hand training data and online KF control (stage ia) on the left, and between hand training data with intention estimation modifications and online KF control (stage ib) on the right. We note that averaged across the three experimental days for Monkey L and Monkey J, 71% (Monkey L) and 66% (Monkey J) of channels had shifts in PD that were significant between hand training data without intention estimation and online testing. When intention estimation modifications were applied to the hand training data, 34% (Monkey L) and 52% (Monkey J) of channels had significant shifts in PD, demonstrating that fewer channels had significant shifts in tuning between the training and testing set. Furthermore, the variance in shifts in PD decreased by 1.42 fold ($p < 0.001$, t-test, Monkey L) and 1.19 fold ($p = 0.011$, Monkey J) when intention estimation modifications were applied.

As before, the mean shift in PD was also compared between the more versus less contributory channels. For instance, in stage ia of Fig. 8b, the mean shift in PD for the more contributory channels (i.e. first half of channels in order of contribution by MI, channels 1 to 37) is 23.8° , as compared to 29.4° in the less contributory channels (channels 38 to 75). In stage ib of Fig. 8b, the mean shift in PD was 12.7° versus 17.7° , respectively. In stage ia of Fig 9, the mean shift in PD was 24.7° for the more contributory channels and 29.4° for the less contributory channels, and in stage ib of Fig 9b, 19.8° versus 22.4° , respectively. These findings reaffirm the importance of intention estimation in not only the two-stage training paradigm, but also the one-stage paradigm, allowing for higher online performance and a corresponding decrease in PD shifts.

Discussion

In this study, we investigated two components of the high performing ReFIT-KF decoder [48]: retraining and intention estimation. We show that intention estimation directly improved decoder performance. We demonstrate that intention estimation can augment the tuning curve characteristics of the training data used to fit a decoder by enhancing the modulation and reducing the per-channel variance. These changes demonstrate that the intention estimation modifications enable neural firing rates to be better associated with kinematics reflected in the neural data. Without intention estimation, firing rates may be sub-optimally binned with respect to their intended kinematics, thereby decreasing the correlations between each channel's firing rate and reach kinematics. We demonstrate that the modified training set ultimately leads to improvement in the online performance of the decoder.

When intention estimation is applied, we note that the distribution of shifts in PD between the training and tests sets decreased, which suggests that the subject was able to control the decoder online with minimal adaptation or change in control strategy. The decrease in PD shifts with intention estimation demonstrates that algorithmic modifications can alter the tuning shifts, and that modifying the training data can lead to a more biomimetic online control strategy. As shown in Fig. 2 and 7, intention estimation is shown to improve performance primarily by its reduction in dial-in-time, highlighting its effect on controllability. As noted in the methods, the intention estimation modifications used in this study involve two modifications to the training data: rotating the cursor velocity toward the target and setting the velocity magnitude during the period the cursor is successfully held at the target to zero. As shown in a previous paper [48], the modification of rotating the cursor velocity toward the target more significantly improved offline decode performance as compared to the scaling the velocity magnitude at the target. As we used threshold crossings and did not spike sort individual channels, we note that this study does not address how single neuron preferred directions are affected.

We furthermore observe that retraining can be helpful as it also reduces the shifts in PD, presumably because the training and testing data are of the same behavioral modality. However, from Fig. 2 showing the comparison of KF and Re-KF, we observe an additional undesired effect of retraining in which overall decoded velocities tend to decrease, leading to longer acquisition times. Because the retrained decoder relies on online cursor kinematics, the performance of the retrained decoder is dependent on that of first-pass decoder. In this sense, errors in decoding can be compounded through each retraining iteration; for example, if the decoding algorithm tends to produce slower velocities, then the decoder will slow down through retraining iterations. Slower retrained decoders may arise if the online training data has a larger fraction of slower velocities, e.g. if there is an increased amount of time spent dialing into the target during online cursor control. Alternatively, when interim decoders have poor online performance, the online cursor kinematics may at any time be associated incorrectly with the subject's intention and relative cortical activity. As a consequence, instances of high velocity kinematics occur for a more diverse set of firing rates, causing the correlation between high velocity kinematics and neural data to decrease and resulting in a slower decoder.

Retraining may be helpful when put in tandem with intention estimation modifications, as these modifications may correct for errors in the first-pass training data between the subject's intention (and related neural activity) and cursor movement due to poor first-pass decode performance. Without correction, these errors may be propagated, leading to slower decoders, as discussed. Given the observation, however, that retraining without intention estimation modifications slow down overall acquisition times, we assessed the viability of directly applying intention estimation to the first stage of training. Because FIT-KF and ReFIT-KF performed similarly, whereas ReFIT-KF performs superiorly to Re-KF, we highlight the importance of intention estimation over retraining in the ReFIT-KF paradigm when using this animal model [64]. The performance gains observed when intention estimation modifications were directly applied to arm reaching data furthermore suggest that there is sufficient variability in the natural reaches of the subject for a given movement intention (e.g., [61–63]). In this setting, intention estimation modifications appear to correct for intrinsic movement variability relative to the neural activity.

It is important to note that many of the channels that exhibit small changes in PD tend to contribute significantly to the decoder, as shown in Fig. 4b and 5b. There are two explanations for this observation. First, it is possible that channels that are more informative and contributory are consistent in their tuning over different contexts. Second, because more of the informative channels are weighted heavily in the decoder, cursor movement is best achieved by adapting a control strategy that modulates contributory channels in a predictable way that matches the training paradigm. Therefore, the PDs of these contributory channels would be expected to vary less [65]. Nevertheless, we find that tuning shifts in even largely non-contributory channels reflects reduction in the performance of the decoder.

While we show that tuning curve shifts can be reduced by generating decoders that are more biomimetic, adaptive strategies may also be important to further improve decode performance. In particular, shifts in tuning may differ in degree depending on not only the algorithm employed but also the animal model used and the training methodologies (e.g., passive observation versus reaching, [64]). When neural properties shift as a result of a context change or over long period of time, co-adaption algorithms [23, 66] or a two-stage training procedure would be critical for adjusting to the shifts in neural tuning. In a clinical setting, a patient may necessarily require a different initial training paradigm, such as observation or imagination, and adaptive measures due to larger contextual changes may be necessary.

In this paper we demonstrate that intention estimation modifications provide a way to develop high-performance out-of-the-box decoders (i.e., FIT-KF and ReFIT-KF). In doing so, we assumed that mimicking the native neural-to-arm mapping would result in a good initial decoder. In particular, we focus on improving the directional controllability of the cursor to be more similar to the control strategies of the native arm. Improving the biomimetic nature of cursor speeds is not investigated in this study and may require nonlinear decoders for biomimetic design. However, it remains unclear what the ideal final control strategy of a BMI is and whether it might be biomimetic strategy or an entirely different strategy. For example, other groups have observed subjects adopting new control strategies during BMI control, such as reducing arm movements, despite being trained with

arm movements [21, 22]. In instances like this, a combination of biomimetic and adaptive control strategies would be beneficial for seeding the decoder at a level of high performance and then relying on adaptation if a different control strategy is desired. Unlike the minimal tuning shifts observed in this paper when employing a biomimetic decoder, large, stable shifts in tuning may be seen when there are explicit changes in contexts or control modalities [26]. Yet, we have provided here an understanding for how certain algorithmic and training modifications, such as with intention estimation, may be able to denoise and improve the training data. As a result, tuning shifts can be decreased based on algorithmic changes, which is useful if the subject desires an out-of-the-box high-performance (biomimetic) controller.

Performance continues to be a challenge for the clinical viability of BMIs. With an improved understanding of how intention estimation modification affects the training set data and, ultimately, fit of a decoder, it will become important to develop training paradigms that better elicit and predict patient intentions. It is important to again highlight that intention estimation modifications are applied offline during the fitting of the decoder. Knowledge of the target position is required only for the initial training period; for clinical applications, high-performance cursor control independent of known target position could therefore be achieved after the initial training period. Encouragingly, the ReFIT-KF and Velocity-KF performance results, including ReFIT outperforming Velocity-KF, have been translated to a person with paralysis (amyotrophic lateral sclerosis, or ALS) as part of the BrainGate multisite FDA phase-I clinical trial [67, 68]. Another study has also reported that a closed-loop based KF, incorporating a variant of the kinematic vector rotation retraining, has translated to and shown performance benefit beyond an otherwise matched open-loop KF in two people with paralysis (pontine stroke, and ALS) as part of the BrainGate clinical trial [14]. By continuing to understand and utilize biomimetic and adaptive strategies, the performance of BMIs should continue to increase which is important for BMI clinical viability.

Acknowledgments

We thank M Mazariegos and J Aguayo for surgical assistance and veterinary care, B Oskotsky for computational support, and S Eiseensee, B Davis, and E Castaneda for administrative and program management support. This work was supported by National Science Foundation graduate research fellowships (JMF, JCK, CAC), Stanford Medical Scientist Training Program (PN), Soros Fellowship (PN), the Christopher and Dana Reeve Foundation (SIR, KVS), Stanford University Graduate Fellowship (JMF, CAC), and the following awards to KVS: Burroughs Wellcome Fund Career Awards in the Biomedical Sciences, DARPA Revolutionizing Prosthetics program contract N66001-06-C-8005, DARPA REPAIR contract N66001-10-C-2010, NIH NINDS CRCNS R01-NS054283, NIH Bioengineering Research Grant R01-NS064318, NIH EUREKA Award R01-NS066311, and NIH Directors Pioneer Award 1DP1OD006409.

References

- [1]. Hatsopoulos NG, Donoghue JP. The science of neural interface systems. *Annu Rev Neurosci.* 2009; 32:249–266. [PubMed: 19400719]
- [2]. Scherberger H. Neural control of motor prostheses. *Curr Opin Neurobiol.* 2009; 19:629–633. [PubMed: 19896364]
- [3]. Nicolelis MAL, Lebedev MA. Principles of neural ensemble physiology underlying the operation of brain-machine interfaces. *Nat Rev Neurosci.* 2009; 10:530–540. [PubMed: 19543222]

- [4]. Andersen RA, Hwang EJ, Mulliken GH. Cognitive neural prosthetics. *Annu Rev Psychol.* 2010; 61:169–190. [PubMed: 19575625]
- [5]. del J, Millan R, Carmena JM. Invasive or noninvasive: understanding brain-machine interface technology. *IEEE Eng Med Biol Mag.* 2010; 29:16–22. [PubMed: 20209672]
- [6]. Green AM, Kalaska JF. Learning to move machines with the mind. *Trends in Neurosci.* 2010 in press.
- [7]. Shenoy KV, Kaufman MT, Sahani M, Churchland MM. A dynamical systems view of motor preparation: implications for neural prosthetic system design. *Prog Brain Res.* 2011; 192:33–58. [PubMed: 21763517]
- [8]. Homer, Mark L.; Nurmikko, Arto V.; Donoghue, John P.; Hochberg, Leigh R. Sensors and decoding for intracortical brain computer interfaces. *Annual Review of Biomedical Engineering.* 2013; 15(1):383–405. PMID: 23862678.
- [9]. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature.* 2006; 442:164–171. [PubMed: 16838014]
- [10]. Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Black MJ. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J Neuroeng.* 2008; 5:455–476.
- [11]. Simeral JD, Kim S-P, Black MJ, Donoghue JP, Hochberg LR. Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *J Neural Eng.* 2011; 8:025027. [PubMed: 21436513]
- [12]. Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Friehs GM, Black MJ. Point-and-click cursor control with an intracortical neural interface system by humans with tetraplegia. *IEEE Trans Neural Syst Rehabil Eng.* 2011; 19:193–203. [PubMed: 21278024]
- [13]. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, Haddadin S, Liu J, Cash SS, van der Smagt P, Donoghue JP. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature.* May; 2012 485(7398):372–5. [PubMed: 22596161]
- [14]. Jarosiewicz, Beata; Masse, Nicolas Y.; Bacher, Daniel; Cash, Sydney S.; Eskandar, Emad; Friehs, Gerhard; Donoghue, John P.; Hochberg, Leigh R. Advantages of closed-loop calibration in intracortical brain–computer interfaces for people with tetraplegia. *Journal of Neural Engineering.* 2013; 10(4):046012. [PubMed: 23838067]
- [15]. Nurmikko AV, Donoghue JP, Hochberg L, Patterson WR, Song Y-K, Bull C, Borton D, Laiwalla F, Park S, Ming Y, Aceros J. Listening to brain microcircuits for interfacing with external world – progress in wireless implantable microelectronic neuroengineering devices. *Proceedings of IEEE.* 2010; 98:375.
- [16]. Gilja V, Chestek CA, Nuyujukian P, Foster J, Shenoy KV. Autonomous head-mounted electrophysiology systems for freely-behaving primates. *Curr Opin Neurobio.* 2010; 20:676–686.
- [17]. Ryu SI, Shenoy KV. Human cortical prostheses: Lost in translation? *Neurosurg Focus.* 2009; 27:E5. [PubMed: 19569893]
- [18]. Gilja V, Chestek CA, Diester I, Henderson JM, Deisseroth K, Shenoy KV. Challenges and opportunities for next generation intra-cortically based neural prostheses. *IEEE Trans Biomed Eng.* 2011 doi: 10.1109/TBME.2011.2107553:in press2.
- [19]. Moritz CT, Perlmutter SI, Fetz EE. Direct control of paralysed muscles by cortical neurons. *Nature.* 2008; 456:639–642. [PubMed: 18923392]
- [20]. Chase SM, Schwartz AB, Kass RE. Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural Netw.* 2009; 22:1203–1213. [PubMed: 19502004]
- [21]. Carmena JM, Lebedev MA, Crist RE, O’Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, Nicolelis MAL. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology.* 2003; 1:193–208.
- [22]. Lebedev MA, Carmena JM, O’Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, Nicolelis MAL. Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *Journal of Neuroscience.* 2005; 25:4681–4693. [PubMed: 15888644]

- [23]. Taylor DM, Tillery S. I. Helms, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science*. 2002; 296:1829–1832. [PubMed: 12052948]
- [24]. Mulliken GH, Musallam S, Andersen RA. Decoding trajectories from posterior parietal cortex ensembles. *J Neurosci*. 2008; 28:12913–12926. [PubMed: 19036985]
- [25]. Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. Functional network reorganization during learning in a brain-computer interface paradigm. *Proc Natl Acad Sci U S A*. 2008; 105:19486–19491. [PubMed: 19047633]
- [26]. Ganguly K, Carmena JM. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*. 2009; 7:e1000153. [PubMed: 19621062]
- [27]. Fagg AH, Hatsopoulos NG, de Lafuente V, Moxon KA, Nemati S, Rebesco JM, Romo R, Solla SA, Reimer J, Tkach D, Pohlmeier EA, Miller LE. Biomimetic brain machine interfaces for the control of movement. *J Neurosci*. 2007; 27:11842–11846. [PubMed: 17978021]
- [28]. Ethier C, Oby ER, Bauman MJ, Miller LE. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature*. May; 2012 485(7398):368–71. [PubMed: 22522928]
- [29]. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci*. 1982; 2:1527–1537. [PubMed: 7143039]
- [30]. Georgopoulos AP, Schwartz AB, Kettner. Neuronal population coding of movement direction. *Science*. 1986; 233:1416–1419. [PubMed: 3749885]
- [31]. Schwartz AB. Direct cortical representation of drawing. *Science*. 1994; 265:540–542. [PubMed: 8036499]
- [32]. Pearce TM, Moran DW. Strategy-dependent encoding of planned arm movements in the dorsal strategy-dependent encoding of planned arm movements in the dorsal premotor cortex. *Science*. 2012 (Scienceexpress, in press).
- [33]. Kalaska JF, Cohen DA, Hyde ML, Prudhomme M. A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *J Neurosci*. 1989; 9:2080–2102. [PubMed: 2723767]
- [34]. Scott SH, Kalaska JF. Reaching movements with similar hand paths but different arm orientations. i. activity of individual cells in motor cortex. *J Neurophysiol*. 1997; 77:826–852. [PubMed: 9065853]
- [35]. Scott SH, Sergio LE, Kalaska JF. Reaching movements with similar hand paths but different arm orientations. ii. activity of individual cells in dorsal premotor cortex and parietal area 5. *J Neurophysiol*. 1997; 78:2413–2426. [PubMed: 9356393]
- [36]. Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, Shenoy KV. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci*. 2010; 13:369–378. [PubMed: 20173745]
- [37]. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*. 2010; 68:387–400. [PubMed: 21040842]
- [38]. Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. Single-trial neural correlates of arm movement preparation. *Neuron*. 2011; 71:555–64. [PubMed: 21835350]
- [39]. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. Neural population dynamics during reaching. *Nature*. 2012; 487:51–6. [PubMed: 22722855]
- [40]. Scott SH. Inconvenient truths about neural processing in primary motor cortex. *J Physiol*. 2008; 586:1217–1224. [PubMed: 18187462]
- [41]. Kalaska JF. From intention to action: motor cortex and the control of reaching movements. *Adv Exp Med Biol*. 2009; 629:139–178. [PubMed: 19227499]
- [42]. Graziano MSA. New insights into motor cortex. *Neuron*. 2011; 71:387–8. [PubMed: 21835337]
- [43]. Graziano MSA. Cables vs. networks: old and new views on the function of motor cortex. *J Physiol*. 2011; 589:2439. [PubMed: 21572141]

- [44]. Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Ryu SI, Shenoy KV. High-performance continuous neural cursor control enabled by a feedback control perspective. *Frontiers in Neuroscience (COSYNE)*. 2010 page doi: 10.3389/conf.fnins.2010.03.00249.
- [45]. Gilja, V.; Nuyujukian, P.; Chestek, CA.; Cunningham, JP.; Yu, BM.; Ryu, SI.; Shenoy, KV. Neuroscience Meeting Planner. Society for Neuroscience; San Diego, CA: 2010. A high-performance continuous cortically-controlled prosthesis enabled by feedback control design. Program No. 20.6Online
- [46]. Nuyujukian, P.; Gilja, V.; Chestek, CA.; Cunningham, JP.; Fan, JM.; Yu, BM.; Ryu, SI.; Shenoy, KV. Neuroscience Meeting Planner. Society for Neuroscience; San Diego, CA: 2010. Generalization and robustness of a continuous cortically-controlled prosthesis enabled by feedback control design. Program No. 20.7Online
- [47]. Nuyujukian, P.; Kao, J.; Fan, JM.; Stavisky, S.; Ryu, SI.; Shenoy, KV. *Frontiers in Neuroscience. Conference Abstract: Computational and Systems Neuroscience (COSYNE)*. Salt Lake City, UT: 2012. A high-performance, robust brain-machine interface without retraining; p. III-65.
- [48]. Gilja, V.; Nuyujukian, P.; Chestek, CA.; Cunningham, JP.; Yu, BM.; Fan, JM.; Ryu, SI.; Shenoy, KV. A brain machine interface control algorithm designed from a feedback control perspective. *Proc. Of the 34th Annual International Conference IEEE EMBS*; San Diego, CA. 2012. In press
- [49]. Sussillo D, Nuyujukian P, Fan JM, Kao JC, Stavisky SD, Ryu SI, Shenoy KV. A recurrent neural network for closed-loop intracortical brain-machine interface decoders. *J Neural Eng*. 2012; 9(2): 026027. [PubMed: 22427488]
- [50]. Davoodi R, Loeb GE. MSMS software for VR simulations of neural prostheses and patient training and rehabilitation. *Stud Health Technol Inform*. 2011; 163:156–62. [PubMed: 21335781]
- [51]. Fraser GW, Chase SM, Whitford A, Schwartz AB. Control of a brain-computer interface without spike sorting. *J Neural Eng*. Oct.2009 6:055004. [PubMed: 19721186]
- [52]. Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy KV. Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *J Neural Eng*. 2011; 045005:8.
- [53]. Kemere CT, Shenoy KV, Meng TH. Model-based neural decoding of reaching movements: a maximum likelihood approach. *IEEE Trans Biomed Eng*. 2004; 51:925–932. [PubMed: 15188860]
- [54]. Yu BM, Kemere C, Santhanam G, Afshar A, Ryu SI, Meng TH, Sahani M, Shenoy KV. Mixture of trajectory models for neural decoding of goal-directed movements. *J Neurophysiol*. 2007; 97:3763–3780. [PubMed: 17329627]
- [55]. Srinivasan L, Eden UT, Mitter SK, Brown EN. General-purpose filter design for neural prosthetic devices. *J Neurophysiol*. 2007; 98:2456–2475. [PubMed: 17522167]
- [56]. Kulkarni J, Paninski L. Efficient analytic computational methods for state-space decoding of goal-directed movements. *IEEE Signal Processing Magazine*. 2008; 25:78–86.
- [57]. Lawhern V, Hatsopoulos NG, Wu W. Coupling time decoding and trajectory decoding using a target-included model in the motor cortex. *Neurocomputing*. Apr.2012 82:117–126. [PubMed: 22379284]
- [58]. Cunningham JP, Nuyujukian P, Gilja V, Chestek CA, Ryu SI, Shenoy KV. A closed-loop human simulator for investigating the role of feedback-control in brain-machine interfaces. *J Neurophysiol*. 2010 doi:10.1152/jn.00503.2010, in press.
- [59]. Milekovic, Tomislav; Ball, Tonio; Schulze-Bonhage, Andreas; Aertsen, Ad; Mehring, Carsten. Detection of error related neuronal responses recorded by electrocorticography in humans during continuous movements. *PLoS ONE*. 2013; 8(2):e55235. [PubMed: 23383315]
- [60]. Sanchez JC, Carmena JM, Lebedev MA, Nicolelis MAL, Harris JG, Principe JC. Ascertaining the importance of neurons to develop better brain-machine interfaces. *IEEE Trans Biomed Eng*. 2004; 51:943–53. [PubMed: 15188862]
- [61]. Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008; 9:292–303. [PubMed: 18319728]
- [62]. van Beers RJ, Haggard P, Wolpert DM. The role of execution noise in movement variability. *J Neurophysiol*. 2004; 91:1050–1063. [PubMed: 14561687]

- [63]. Churchland MM, Afshar A, Shenoy KV. A central source of movement variability. *Neuron*. 2006; 52:1085–1096. [PubMed: 17178410]
- [64]. Nuyujukian P, Fan JM, Gilja V, Kalanithi PS, Chestek CA, Shenoy KV. Monkey models for brain-machine interfaces: the need for maintaining diversity. *Conf Proc IEEE Eng Med Biol Soc*. 2011:1301–5. [PubMed: 22254555]
- [65]. Ganguly, Karunesh; Dimitrov, Dragan F.; Wallis, Jonathan D.; Carmena, Jose M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat Neurosci*. May; 2011 14(5):662–667. [PubMed: 21499255]
- [66]. Osborne AL, Dangi S, Moorman HG, Carmena JM. Closed-loop decoder adaptation on intermediate time-scales facilitates rapid bmi performance improvements. *IEEE Transactions in Neural Systems and Rehabilitation Engineering*. 2012; 20:468–477.
- [67]. Gilja, V.; Pandarinath, C.; Blabe, CH.; Hochberg, LR.; Shenoy, KV.; Henderson, JM. Neuroscience Meeting Planner. Society for Neuroscience; Washington, DC: 2013. Design and application of high performance intracortical brain computer interface for a person with amyotrophic lateral sclerosis.
- [68]. Simeral, JD.; Bacher, D.; Sarma, AA.; Schmansky, NJ.; Stavisky, SD.; Jarosiewicz, B.; Milekovic, T.; Rosler, DM.; Pandarinath, C.; Cornwell, AS.; Henderson, JM.; Shenoy, KV.; Kirsch, RF.; Donoghue, JP.; Hochberg, LR. Neuroscience Meeting Planner. Society for Neuroscience; Washington, DC: 2013. Evolution of the braingate real-time brain-computer interface (bci) platform for individuals with tetraplegia or limb loss.

Figure 1.

Schematic of the two-stage training paradigms. In the first step of training, the monkey performs a 2D center out and back reaching task with the native arm. [i] A first pass BMI model is built using the recorded neural data and the kinematic data from the arm movements. In the second step of training, the monkey performs the same 2D center out and back reaching task, while using the first pass BMI model (KF). A second set of training data is collected during this online run to build the final BMI model. This model-build is carried out [iia] without training set modifications (Re-KF) and [iib] with intention estimation modifications (IEMs) of the training set (ReFIT-KF). The bolded training paradigm indicates the method of training used in the ReFIT-KF algorithm [48]. The blue boxes schematize the subject's online cursor control activity; the black boxes schematize the training data, which is derived from the online cursor control.

Figure 2.

Effect of intention estimation modifications on the performance of three decoders. (a) The average distance from the current BMI cursor location to the target for center-out reaches. The thick line indicates the dial-in-time. Data from Monkey L, 2011-07-21 and 2011-07-22. (b) Bar graph of the average acquire time required to successfully obtain the target for Monkey L. The white line indicates the average time the subject reaches the target. (c) Same as in (a) but for Monkey J, 2012-07-26, 2012-07-27. (d) Same as in (b) but in Monkey J.

Figure 3.

Effects of intention estimation on tuning characteristics. (a) Tuning curve characteristics using training data from two representative channels with (black) and without (grey) intention estimation for channel 41, Monkey J, 2011-06-10. Error bars denote standard error. (b) Same as (a) but for channel 76, Monkey J, 2011-06-10. (c) Histogram of the ratio of the tuning curve modulation when intention estimation modifications are applied over not applied across all viable channels for Monkey J (blue, averaged data across 60 days) and Monkey L (red, averaged data across 55 days). (d) Histogram of the ratio of per-channel firing rate variance when intention estimation is applied over not applied across all viable channels for Monkey J (blue, averaged data across 60 days) and Monkey L (red, averaged data across 55 days).

Figure 4.

Tuning shifts between training sets and respective online testing sets for each component of the two-stage training paradigm. (a) Histograms of the shifts in PD between the native hand training data and the online first pass decoder, KF (stage i); first-pass decoder training set and the online final decoder data, Re-KF (stage iia); the first-pass decoder training set with intention estimations modifications (IEMs) and the online final decoder data, ReFIT-KF (stage iib). Red bars indicate statistically significant deviations, and the green line indicates the bootstrapped distribution of shifts due to noise. (b) Shifts in PD per channel ordered from the most (channel 1) to least contributory channels (channel 85) for stage i, iia, and iib. The channels that have statistically significant shifts in PD are colored in red. Data from Monkey J, 2012-07-26.

Figure 5.

Tuning shifts between training sets and respective online testing sets, as in Fig 4. Data from Monkey L, 2011-07-21.

Figure 6.

Schematic of the one-stage training paradigms. A final BMI control model is built from the native hand control training data without [ia] and with intention estimation modifications (IEMs) [ib]

Figure 7.

Comparison of the online performance of one and two-stage decoders with intention estimation. (a) The average distance to target is plotted for center-out reaches using hand control (green), KF (blue), ReFIT-KF (yellow), and FIT-KF (red). Data from Monkey L, aggregated across 6 experimental days. (b) The average acquire times to successfully hold a target are plotted for the four control modalities, where the white line indicates the first time the subject reaches the target without necessarily holding it in Monkey L (c) Same comparison as (a) but in Monkey J, data aggregated across 7 experimental days. (d) Same as (b) but in Monkey J.

Figure 8.

Tuning shifts between training and online testing sets for a one-stage training paradigm. (a) Histograms of the shifts in PD between the hand training data versus the online first pass decoder, KF (stage ia), and between the hand training data with intention estimation modifications versus the online first pass decoder, FIT (stage ib). (b) Shifts in PD ordered from most to least contributory channels. Data from Monkey L, 2011-04-21.

Figure 9. Tuning shifts between training sets and respective online testing sets, as in Fig 8. Data from Monkey J, 2011-05-19.