



HHS Public Access

Author manuscript

J Neural Eng. Author manuscript; available in PMC 2018 January 21.

Published in final edited form as:

J Neural Eng. 2016 June ; 13(3): 036021. doi:10.1088/1741-2560/13/3/036021.

A four-dimensional virtual hand brain-machine interface using active dimension selection

Adam G. Rouse

Department of Neuroscience, University of Rochester, Rochester, NY. Department of Neurology, University of Rochester, Rochester, NY. Department of Biomedical Engineering, University of Rochester, Rochester, NY

Abstract

Objective—Brain-machine interfaces (BMI) traditionally rely on a fixed, linear transformation from neural signals to an output state-space. In this study, the assumption that a BMI must control a fixed, orthogonal basis set was challenged and a novel active dimension selection (ADS) decoder was explored.

Approach—ADS utilizes a two stage decoder by using neural signals to both i) select an active dimension being controlled and ii) control the velocity along the selected dimension. ADS decoding was tested in a monkey using 16 single units from premotor and primary motor cortex to successfully control a virtual hand avatar to move to eight different postures.

Main Results—Following training with the ADS decoder to control 2, 3, and then 4 dimensions, each emulating a grasp shape of the hand, performance reached 93% correct with a bit rate of 2.4 bits/s for eight targets. Selection of eight targets using ADS control was more efficient, as measured by bit rate, than either full four-dimensional control or computer assisted one-dimensional control.

Significance—ADS decoding allows a user to quickly and efficiently select different hand postures. This novel decoding scheme represents a potential method to reduce the complexity of high-dimension BMI control of the hand.

Keywords

brain-machine interface; brain-computer interface; motor cortex; neuroprosthetics; hand shaping; grasping; finger movements

1. Introduction

In recent years, brain-machine interfaces (BMI) have increased in complexity, which is typically defined by the number of degrees of freedom (DOF) being controlled. A cursor on a display is an example of a two DOF system [1]. The horizontal and vertical components of the cursor are controlled independently, based on the recorded neural signals. Translation of the hand can be described as a three DOF system in standard Cartesian coordinates, both in

neurophysiologic studies [2] and in BMI implementations [3–5]. In addition to controlling translation, a current goal is to create BMIs that control currently available robotic hands [6–9] that replicate many of the natural actions of the human hand. Creating BMIs that can replicate natural control of the 22 DOFs of the human hand remains challenging [10].

The 22 DOFS of the hand often do not move independently. A large fraction of the variance in complex, multi-joint movements can be mathematically reduced to a small number of principal components [11–13]. Such principal components or “synergies” have been used to define the basis set of motion for robotic design [14,15], have been applied in decoding and reconstructing movements from recorded neural data [16,17], and have been implemented as DOFs to be used in real-time BMI control [18]. Current BMIs thus often have designs that combine subsets of joints of the hand into fixed synergies [10]. Rather than controlling all the DOFs of the hand independently, an underlying BMI state-space with fewer control dimensions is defined. The BMI state-space is decoded from the incoming neural signals and then output in a fixed manner to the controlled DOFs (e.g. robotic joints, stimulated muscles, or displayed virtual features) that make up the physical plant [19].

While current BMI implementations rely on fixed state-spaces with all output dimensions continuously and equally controlled, there are alternative parameterizations of the hand that consider discrete classes of movement that may be more effective. Napier [20] described distinct patterns of prehensile movements, such as precision and power grips. Arbib and colleagues [21] proposed the idea of virtual fingers where the number of digits that oppose the thumb could be selected and moved together as a “virtual finger” to create varying sized grasps. Other neuroprosthetics, like electromyographic (EMG) controlled prosthetic arms, have implemented switching algorithms [22–24] that allow the user to first select a type of grasp and then execute it. In reality, the natural musculoskeletal system of the hand does not necessarily require fixed, orthogonal synergies, and likely is controlled by more complex neural encoding [25]. Yet alternative encoding schemes for BMIs controlled by cortical neurons remain relatively unexplored.

Several methods have been employed in current BMI controlled devices to improve performance, though at the cost of completely autonomous user control. Some BMI designs use a priori knowledge of the desired movement to mix an automatically generated signal with the user’s neural signals, creating movements that more closely replicate natural movements while limiting user flexibility and depending more on autonomous, external sensor-controlled robotics [26–30]. Selective control of certain dimensions has been used as a training technique to allow sequential learning of the different DOFs [5,10,31,32], but in these cases the dimension(s) to be controlled in a given trial are specified by the designer/experimenter. Does such selective control require an external agent, or might the user’s neural signals both specify the dimension(s) to be controlled and control the movement along those dimension(s)?

Here, a novel BMI was designed to use such active dimension selection (ADS) to control a four dimensional basis that mimics four common human grasps. This approach combines the ideas of classes of grasps and selective control of movement. This novel ADS decoder uses the same set of recorded neural signals in a two-step process where with each time step: i)

the dimension being controlled was selected, and ii) the velocity along the chosen dimension was actively controlled; meanwhile, the three remaining dimensions were passively held in a neutral posture. ADS allows the same neural signals to accurately control a selected subset of dimensions important to the current task while offloading control of other dimensions to an automated process. While this limits flexibility in the range of movements that can be performed compared to when all dimensions can be controlled simultaneously, many movements including common grasps can be defined with a single movement dimension. The present study demonstrates that ADS allows these reduced dimension movements to be performed reliably and efficiently.

2. Methods

2.1. Virtual hands

Two 20-DOF (4 DOF for each of the 5 digits) virtual hands were created using the MSMS software package [33]. Four reduced dimensions were designed to emulate the flexion/extension movement of four common grasps observed in humans: power (all five digits), pinch (thumb and index finger), tripod (thumb, index, and middle fingers), and opposition (thumb and little finger). In contrast to many commonly defined joint synergies, these four dimensions were not defined to be orthogonal in the 20-DOF joint space. Rather they were defined with the goal of creating grasping movements most useful to a subject when constrained to move along only a single dimension at a time.

These four grasp dimensions then were used as a Cartesian state-space, defining the task environment. This four-dimensional state-space completely defined and reliably mapped to the position of the 20 DOF virtual hands. In the state-space, the value for each dimension (power, pinch, tripod, and opposition) ranged from -1 (maximum extension) to +1 (maximum flexion). The current state-space values were then translated to the 20 DOF virtual hand by summing the amount of flexion/extension for each DOF based on which grasp dimensions included motion of each DOF (e.g. total index flexion/extension at the metacarpophalangeal and inter-phalangeal joints was the sum of the power, pinch, and tripod values in the state-space; little finger flexion/extension was the sum of the power and opposition values).

The task display (Fig. 1A) consisted of two virtual hands: an avatar hand on the left and a target hand on the right. The target hand was presented stationary at up to eight different target postures that corresponded to eight different locations in the 4-dimensional state-space (Fig. 1B). These locations were at either +0.667 (flexion) or -0.667 (extension) along one of the four grasp dimensions, with all other dimensions set to 0. The monkey's task then was to move the avatar hand to match the target posture. The avatar hand was deemed to have matched the target hand posture when the state-space coordinates of the avatar were greater than +0.5 (flexion) or less than -0.5 (extension) in the target dimension and between -0.167 and +0.167 in all other dimensions.

In addition to the avatar and target virtual hands, the background color on each half of the display also varied, providing additional feedback on the current state-space location of the avatar hand and the target hand posture. The state-space was mapped to a Hue-Saturation-

Value color space with hue representing the dimension (radial direction) of movement and saturation representing the amount of flexion/extension (distance from the origin). Since the hue/saturation color wheel is two dimensional, the four dimensional state-space was collapsed by summing the locations of the four state-space values when the four dimensional axes are defined to be 45° apart as shown in Fig. 1B.

The monkey's task then consisted of moving the avatar hand in two, three, or four state-space dimensions to match one of the four, six, or eight target hand postures, respectively. The sequence of events in a single trial is illustrated in Fig. 1A. At the start of each trial, the controlled avatar hand and target hand both were automatically centered to a neutral posture (0, 0, 0, 0). At the cue time point, the target hand changed abruptly to a pseudorandomly selected target posture e.g. (+0.667, 0, 0, 0). For 300 ms immediately following the cue, the avatar hand remained fixed at the neutral posture, preventing unwanted movement that would have resulted from spontaneous neural activity that had occurred before the cue. After this 300 ms the hand was free to move through the 2, 3, or 4-dimensional state-space under brain control (described below). The trial ended when one of the target postures was contacted (matched the target boundaries defined above), or else after 5 s timed out. If the correct posture was achieved, the monkey was given a water reward. If the trial was incorrect, either because one of the wrong target postures was contacted or because the 5 s timed out, the same target was repeated up to two more times to encourage the monkey to not skip difficult targets.

2.2. Neural Recordings

A male Rhesus monkey weighing approximately 10 kg was used as the subject in the present study. All procedures for the care and use of nonhuman primates followed the Guide for the Care and Use of Laboratory Animals and were approved by the University Committee on Animal Resources at the University of Rochester, Rochester, New York. Floating microelectrode arrays (MicroProbes, Inc), each with 16 recording electrodes had been implanted in the monkey's left hemisphere using procedures described in detail previously [34]. The present study was conducted using the electrodes on eight such arrays: four in primary motor cortex (M1) and four in ventral premotor cortex (PMv). The electrode lengths ranged from 1.5–7.5 mm, such that neurons were recorded from the hemispheric surface to the depth of the anterior bank of the central sulcus (M1) and the posterior bank of the arcuate sulcus (PMv). The electrodes in M1 used for BMI control spanned a distance of approximately 17 mm along the central sulcus and were centered on the region of M1 from which movement of the digits were evoked with intracortical microstimulation. The electrodes in PMv used for control spanned a distance of approximately 9mm along the inferior limb of the arcuate sulcus. Sixteen of the 128 recorded channels were selected based on the best discriminable spiking units. Eight of the channels were from the arrays implanted in M1 and eight from the arrays in PMv; no distinction between the two areas was made in the present experiment and analyses. Online spike sorting was performed to select the most prominent spiking unit on each channel (Plexon, Inc). The same 16 channels were used from day to day. Occasionally, a unit was lost, which required switching channels at the start of the day. No switching was performed within sessions.

2.3. BMI Decoding

Each unit's firing rate was estimated by convolving the spike train with a 500 ms Gaussian window (center = -250 ms, $\sigma = 125$ ms). The estimated firing rate was then square root transformed to better normalize the variance [35]. To create zero-centered neural control signals, a running average of each unit's square root transformed firing rate was calculated as an exponential moving average updated after each trial with the average from the new trial added with a weight of $\frac{1}{50}$ to $\frac{49}{50}$ of the previous average. The average then was subtracted from the instantaneous square-root transformed firing rate to generate the normalized instantaneous firing rate for each unit. The BMI system sampled neural firing rates and updated joint positions at 100 Hz.

2.3.1 Active Dimension Selection (ADS) decoder—The ADS decoder relies on the same underlying linear relationship used in classical linear, velocity decoders that decode state-space velocities from the normalized instantaneous firing rates. However, the ADS decoder does not use the neural signals to control all state-space velocities simultaneously. Rather, the ADS decoder selects only a single dimension that is being moved by the neural signals while the other dimensions are held at or returned to a neutral posture.

Like a classic linear decoder, the ADS decoder used decoding weights (W) for each dimension that are a simple linear transformation from firing rate to state-space velocity.

$$\dot{x}_{4 \times 1} = W_{4 \times 16} N_{16 \times 1}, \quad \text{Eqn. 1}$$

Here, the instantaneous firing rates of the 16 recorded units (N) is multiplied by a weight matrix (W) that transforms the firing rates to velocities (\dot{x}) in the four-dimensional state-space.

The novel component of the ADS decoder is the manner in which these state-space velocities are used. Rather than all four velocity estimates simultaneously controlling the virtual hand, only the largest magnitude velocity signal moves the hand along that single dimension of the state-space.

$$x_{t,i} = x_{t-1,i} + W_{i \times 16} N_{t,16 \times 1}, \quad i \text{ such that } \dot{x}_i = \max\{|\dot{x}_{1..4}|\}, \quad \text{Eqn. 2}$$

The new position along dimension i in the state-space becomes the previous position plus the velocity signal along dimension i based on the instantaneous firing rates (N) at time step t . The ADS decoder selects the i^{th} dimension based on that dimension having the largest magnitude velocity signal. For the other dimensions, the state-space positions are not actively controlled by the neural signals. Instead, the ADS decoder returns the other dimensions in the state-space to zero, corresponding to a neutral hand posture along those dimensions.

$$x_{t,j} = \alpha x_{t-1,j}, \text{ for all } j \neq i, \quad \text{Eqn. 3}$$

where $\alpha = 0.926$

For all dimensions j other than the active dimension i , the previous position is multiplied by a constant (α) to approximate an exponential decay back to zero with a time constant of approximately 130 ms.

If at any 10 ms time step a different dimension attained the largest velocity signal, the ADS decoder selected this dimension as the new active dimension, and the previously active dimension decayed back towards zero. The four state-space variables from a single, example trial all have been plotted as a function of time in Fig. 2. In this trial, the monkey first extended the avatar hand along a wrong dimension (light grey) before the brain signals corrected and flexed the avatar hand along the correct dimension (blue) to the desired hand posture.

2.3.2 Decoding Weights Updating—As shown in Eqn. 1, decoding weights (W) were used to linearly transform the 16 instantaneous firing rates (N) to velocities (\dot{x}) in the d dimensions of the state-space. Following each session, the data was organized and analyzed to generate decoding weights for the following session. The input data for generating the weights consisted of the average firing rate for the 16 recorded units for each trial during the previous session. For each unit, the average firing rate was the average of the normalized instantaneous firing rate from 300 ms after the cue was presented until the trial ended by contacting one of the target hand postures or timing out after 5 s. These average firing rates thus represent a single value for each n^{th} unit for each k^{th} trial and were arranged in an $N=16 \times K$ trials matrix.

The weighting matrix (W) was generated for each session using one of two methods: i) uniform weights, or ii) regressed weights (illustrated in Figure 3). The uniform weights method was used for all sessions during the first 90 days. For day 91 and all subsequent sessions, regressed weights were used. The weights were estimated with the following two procedures:

- i. Uniform weights were originally generated on the first day for each output dimension by randomly assigning +1 or -1 to each of the 16 neurons, with the constraints that: a) for each dimension, 8 neurons were assigned a weight of +1 and 8 neurons were assigned -1; and b) the weight vectors for all dimensions were mutually orthogonal. The weights were updated for each session from the previous recording session by assigning weights based on the dimension and unit that showed the most modulation. This was determined by creating a receiver operator characteristic (ROC) curve for each dimension [36]. The ROC curve was constructed by comparing the fraction of flexion trials and extension trials along a given dimension with firing rates for a given neuron less than a series of firing rate thresholds. The area between the ROC curve and the 45° no-discrimination line was then used to determine the difference in modulation for

each neuron for flexion versus extension trials for a given dimension. The neurons that increased their firing for flexion compared to extension were given a weight of +1 and those with extension greater than flexion were given a weight of -1. The largest modulated neurons for a given dimension based on the area of the ROC curve were assigned first with each new +1 or -1 weight to a given unit and dimension being assigned contingent on meeting constraints a and b above. Weights were assigned using this greedy approach until no longer possible and the remaining weights were assigned +1 or -1 as necessary given the constraints.

- ii. After day 91, regression weights were generated using ordinary least squares to predict a target matrix, Y , from the neural data, N . Each column of Y represented a trial with d rows corresponding to the d dimensions and +1 indicating a flexion target or -1 indicating an extension target in the given dimension. As an example in the 4D case, the Y matrix:

$$Y = \begin{bmatrix} +1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & +1 \end{bmatrix}$$

has three columns to represent the targets for three trials: the first trial with a 1st dimension flexion target; the next trial, a 2nd dimension extension target; and the third trial, 4th dimension flexion target. Linear least squares regression was then performed to generate the best weight matrix (W) to predict the target matrix (Y) from the neural data (N).

$$W_{d \times 16} \times N_{16 \times k} = Y_{d \times k}$$

To maximize the separation between dimensions and encourage the generation of neural signals that were distinct for different targets, the resulting weight matrix, W , was transformed such that all rows were orthogonal. To transform the weight vectors for all dimensions by similar amounts, pairs of weight vectors were symmetrically rotated to be orthogonal to one another. Such rotations were iterated until the weight vectors for all dimensions were mutually orthogonal.

On occasional days when a channel's unit was lost for the subsequent recording session, the unit of another channel was selected by the investigator and was assigned the previous unit's weights.

2.4. Training

The monkey was initially trained with a joystick to become familiar with the avatar and target hands and the sequential events of a behavioral trial. The horizontal and vertical positions of the joystick were mapped to the velocity of power and pinch grasps, respectively. After the monkey could reliably move the avatar hand using the joystick to the correct position when a target hand was presented, the monkey was switched to brain control. Since the joystick allowed only two DOF control, brain control was first trained on

the same two dimensions. The first time the monkey experienced movements of the virtual hand along the third and fourth dimensions was when these dimensions were added during brain control.

Early brain-control training used two methods of computer assistance. First, a constant velocity bias in the correct state-space direction was added to the brain controlled velocity signal to increase the probability of movement toward the correct target. Second, error detection was turned off to allow the monkey to contact the correct target after first contacting an incorrect target in the same trial. These two assistance measures were adjusted daily by the experimenter with the goal of the monkey being rewarded on approximately 75% of trials. As performance improved, both forms of assistance were gradually eliminated so that the avatar hand was controlled by only recorded brain signals.

2.5. Experiments

All initial training was done with the ADS decoder. Once performance plateaued with the ADS decoder, two experiments were performed. The first experiment added catch trials during some recording sessions where the BMI decoder was switched from the ADS decoder to two other decoders. The second experiment used only the ADS decoder but varied the feedback display to examine how well and which parts of the visual display were being used by the monkey during the BMI task.

2.5.1. Experiment 1) ADS vs. CDS vs. Full decoding—To measure the information available and precision achieved with similar neural signals, in some sessions two alternative decoding schemes were used in randomly intermixed catch trials while still using the ADS decoder in most trials. In each case, the decoding weight matrix (W) was the same across decoders. The first was full dimensional control (Full) where all dimensions were controlled simultaneously. As described by Eqn. 1, for Full control the decoding weights (W) were matrix multiplied by the instantaneous firing rates for the 16 units (N) to produce the simultaneous velocity signals along all dimensions. These velocity signals controlled all hand dimensions simultaneously with the virtual hand moving freely within the state-space. Thus, Full decoding used the same decoding weights as ADS but did not incorporate an actively selected dimension (Eqn. 2 and 3 above). The second was a one-dimensional computerized dimension selection (CDS) control in which the computer automatically selected the active dimension and the neural signals were responsible only for controlling the velocity of the hand along that dimension with the appropriate row of decoding weights in W . All other dimensions were fixed in a neutral position. The Full and CDS trials each comprised 20% of randomly selected trials and the ADS decoder was used in the remaining 60% of trials.

2.5.2. Experiment 2) Varying display feedback—To examine which parts of the visual display were being used by the monkey, three sessions were performed with blocks of three modified display conditions in addition to a block with the regular display. These sessions were performed with ADS decoding only. The three modified blocks were: i) only hands – only the target and avatar hands were displayed against a grey background instead of the color shifting background, ii) only colors – only the constant-color target background

and shifting-color avatar background were displayed with neither the target or avatar hand, and iii) no avatar, only target – the target hand with color background was displayed but the avatar hand and its shifting-color background were not.

2.6. Performance metrics

Five different metrics were used to measure performance. Performance was often worse early in the session as the monkey and system adjusted to new decoding weights as well as later in the day as the monkey became less motivated; therefore, only ordinal trials 101–700 of approximately 1000–1400 total trials per day were used in assessing performance. Also, since the monkey was required to repeat incorrectly performed trials up to 3 times and thus knew the desired target at the start of the subsequent trial(s), only trials that followed a correct trial were used in all five performance metrics.

1. *Success rate.* The percentage of all trials successfully completed by the monkey. The denominator includes correct trials, trials where an incorrect target was contacted, and trials where the 5 s timeout was reached.
2. *Percentage timeouts.* The percentage of all trials where neither a correct or incorrect target was contacted but rather the movement time reached 5 s.
3. *Percentage correct.* The percentage of target postures correctly contacted. To allow for exact assessment of chance performance, only trials where one of the targets was contacted were included in the denominator. Timeout trials, where no targets were contacted within 5 s, were not included. Thus, chance levels for 4, 6, and 8 possible targets were 25, 16.7, and 12.5%.
4. *Mean movement time.* The mean movement time was the total time from cue to trial end (either target contact or 5 s trial timeout) for all trials divided by the number of trials.
5. *Bit rate.* The bits per trial was calculated as:

$$\text{Bits per trial} = \log_2(N) + p \cdot \log_2(N) + (1-p) \cdot \log_2\left(\frac{1-p}{N-1}\right)$$

where N is the number of possible targets and p is the percentage correct [37,38]. The bits per trial was then divided by the mean movement time to generate the bit rate in bits per second. The bit rate is a useful metric for comparing BMI performance between experiments, and represents the number of useful bits of information transmitted when making selections with an output device (e.g. answering yes/no questions, spelling words). The bits per trial is based on the percentage correct since an output device would receive the transmitted information of a correct or incorrect trial but simply wait and ignore timeout trials. The mean movement time used to calculate bit rate does include the time of timeout trials to reflect the increased transmittal time.

3. Results

3.1 Learning and performance with Active Dimension Selection

The monkey was trained on two-dimensional (2D) followed by three- (3D) and then four-dimensional (4D) control. Each dimension had two target postures (flexion and extension), providing 4, 6, and 8 targets for 2D, 3D, and 4D, respectively. The learning time course for training (white regions) and blocks of experiment 1 (grey regions) is plotted in Fig. 4. On Day 1, the monkey was first switched from joystick control to 2D brain-control (blue markers), using uniform decoding weights (W) for all channels. During early recording sessions computer assistance aided the monkey in contacting the correct target, and the success rate shown here was estimated off-line by simulating the avatar hand motion that would have been produced by the recorded neural activity had it acted solely through ADS brain-control with no computer assistance (open blue circles). Because the initial weights were arbitrary, performance started close to chance (25% for 4 targets) and steadily increased. After ~3 weeks, computer assistance was discontinued. The monkey then performed the 2D task entirely with ADS brain-control (filled blue circles), continuing to improve until plateauing at approximately 80%. Next, experiment 1 with CDS and Full-control catch trials were included for several sessions (grey region).

Following 2D, a 3rd dimension was introduced on day 63 (red markers). Initial training again included computer assistance (open red circles). After advancing to total brain-control (filled red circles), on day 91 the decoding weights (W) were switched from constant uniform weights to weights updated daily based on linear regression (vertical black line), resulting in a sudden increase in success rate from 73% to 89%. Thereafter, the monkey continued to improve in 3D control. Catch trial experiments were again performed. Finally, the 4th dimension was added, initially with computer assistance (open yellow circles) and then progressing to total brain-control (filled yellow circles).

The experiments then were repeated with 2D, 4D, and 3D control. Although with the initial introduction of each added dimension 20–30 recording sessions were needed for ADS performance to plateau, when 2D, 4D, and 3D control were repeated, ADS returned to a high level in less than five sessions. Furthermore, when using uniform weights success rate plateaued at ~ 80%, but after switching to regression weights, the ADS success rate for 2D-, 3D-, and 4D-control, all reached greater than 95% correct during some sessions. On days with catch trials, ADS control tended to have slightly lower success rates. In comparison, CDS-control catch trials were nearly 100% correct, while Full-control catch trials showed substantially lower performance that decreased with the number of controlled dimensions.

Using the novel ADS decoder, the monkey thus successfully controlled the virtual hand to achieve the instructed target hand posture with a high level of accuracy (Supplementary Video 1 and 2). This performance was quantified by analyzing all days after switching to regressed weights (Day 91) on the training days when only ADS control was used with normal feedback and without computer assistance (White regions with filled circles in Fig. 4). Because performance tended to improve during the first few days of training when switched to a different number of dimensions and performance also tended to decrease during and following the Full and CDS catch trial sessions, the best and most stable

performance tended to be the late training sessions with ADS-only decoding preceding catch trial sessions. To examine these late training sessions, the final three sessions before initiating catch trial sessions were also analyzed separately (3 sessions in white area immediately preceding grey regions). For 4D, the ADS decoder was utilized twice after switching to regressed weights so there is a total of six late training sessions. The 3D ADS decoder was also repeated twice but since the second instance had only two sessions preceding catch trial sessions there were only five available late training sessions. The 2D ADS decoder was used for only one series of sessions using regressed weights so there are only three late training sessions. As shown in Fig. 5A, using ADS the monkey successfully contacted the correct target with an average of 89% correct even for the 4D task. When analyzing only the late training sessions immediately before switching to Full/CDS/ADS sessions (larger circles), the average was 93% correct for the 4D task. The average movement times during these late training sessions were 0.86, 1.16, and 1.03 s for 2D, 3D, or 4D, respectively (Fig. 5D). Because larger numbers of targets resulted in only a modest decrease in accuracy with similar movement times, the bit rate was highest for the 8-target, 4D task, at 2.4 bits/s (Fig. 5E).

Fig. 6 illustrates the neural activity during movements to the 8 different targets in a single 4D session. The activity of a typical single unit is plotted for all trials in Fig. 6A. This neuron's firing decreased in anticipation of the cue and then became strongly modulated ~250 ms after the cue appeared. This neuron in general fired more for extension movements of the avatar hand, except in the thumb-little finger opposition dimension where it fired more for flexion. Fig. 6B shows the mean firing rates across all trials to each of the 8 targets for each of the 16 units used for BMI control in this session. Rather than being modulated for only a few targets, most units were modulated to varying degrees for all 8 targets. This represents a key benefit of ADS. If units are assigned to control all dimensions simultaneously, some units must either have small modulation to maintain near zero velocity along certain dimensions or units must cancel out in those dimensions. By using active dimension selection, all units can be modulated corresponding to the currently selected dimension's weights, irrespective of tuning along other dimensions.

3.2 Experiment 1) Comparing ADS with other decoders

To better quantify the effect of active dimension selection, sessions were conducted with Full (Supplementary Video 3) and CDS (Supplementary Video 4) control catch trials (grey regions in Fig. 4). The full-control condition required the neural signals to control all dimensions simultaneously whereas in CDS control the neural signals only controlled flexion/extension along the correct dimension, selected automatically by the computer. Example brain controlled trajectories in the 4D state-space are shown in Fig. 7 for trials of power grip flexion. Although actually four-dimensional, for purposes of illustration the trajectories of the four state-space variables are plotted here in two separate two-dimensional Cartesian coordinate spaces. Using the ADS decoder (Fig. 7A), trajectories are rendered more efficient by restricting movement to a single dimension at a time. Fig. 7B illustrates that the CDS decoder automatically restricted movement entirely to only the power flexion/extension dimension. In contrast, Fig. 7C shows how Full control allowed movement within the entire state-space, making precise movement to the correct target extremely difficult.

Performance in the ADS, CDS, and Full-control trials for the final five 2D, 3D, and 4D catch-trial sessions in Fig. 8. With Full control, although the monkey had a success rate of 93% for 4 targets in the 2D case, the success rate dropped in the 3D and 4D case to 45% and 22% (Fig. 8A), respectively, largely the result of 39% and 72% of trials ending in 5 s timeouts. In the remaining trials a target was contacted, but fewer than 80% of targets contacted were correct (Fig. 8B). In contrast, with either CDS or ADS, less than 1% of trials ended with timeout errors, and thus the success rate and percentage correct were equivalent. With CDS, the monkey had near perfect performance selecting from the two targets along the computer-selected dimension in the 2D, 3D, and 4D cases. With ADS, success rate was slightly lower than with CDS, but much higher than with Full control in both the 3D and 4D cases. Percentage correct with ADS was not different from CDS in 2D but significantly lower in the 3D and 4D cases (χ^2 test, $p < 0.05$) with ADS being at 99%, 86%, and 95% for 2D, 3D, and 4D, respectively. Percentage correct with ADS was significantly higher than Full-control in all cases (χ^2 test, $p < 0.05$). Mean movement times were relatively constant as a function of the number of task dimensions with either ADS or CDS, but increased with more task dimensions under Full control, primarily as the result of more 5 s timeouts. For ADS, the mean movement times were 0.76, 0.98, and 0.93 s for 2D, 3D, and 4D, respectively.

The bit rate is a combination of percentage correct and mean movement time. With either ADS or Full control, the monkey selected from 4, 6, or 8 targets which corresponds to 2, 2.58, and 3 of bits per trial, respectively. With CDS, the computer limited the monkey to 2 targets, or 1 bit per trial, whether working in 2D, 3D, or 4D. Movement times with CDS and ADS were similar while the bits per trial were more in the ADS case. The longer movement times and lower percentage correct with Full control lowered the bit rate. Consequently, ADS provided higher bit rates than either CDS or Full-control.

3.3 Experiment 2) Hand shape versus background color of the feedback display

The complete task display conveyed the brain-controlled and target variables using both virtual hand shape and changing background color. The relative importance of hand shape versus background color was determined by reducing the information displayed during three 4D sessions conducted at the end of the study following the initial training and Experiment 1 shown in Fig. 4. In each of these three sessions, ADS decoding was used to control the virtual hand in blocks of trials with three different types of reduced feedback. First, the color shifting background was replaced with a neutral grey background so only the displayed avatar and target hands provided information about the task. Second, the avatar and target hands were removed and only the two colored backgrounds—a target color rectangle and a feedback color shifting rectangle—indicated the brain-controlled BMI scene. Third, the brain-controlled avatar hand and color background were removed leaving only the target hand and color displayed, requiring the animal to modulate its neural signals with no visual feedback.

Fig. 9 compares the monkey's performance averaged across the three sessions with the regular display (dark grey) and with each of the three reduced display modes. With only the virtual hands displayed on a constant grey background (light grey), success rate was slightly

worse than with full visual feedback (89% vs. 73%). In contrast, without the two virtual hands, using only colors (white), the monkey performed at chance (11%). Finally, with no avatar hand or color feedback, viewing only the target (medium grey), the monkey's performance (45%) was reduced compared to viewing only hands, though better than when viewing only colors. The success rate in all 4 conditions were statistically different from one another ($p < 0.001$, pairwise comparisons using χ^2 test). Due to the nature of ADS decoding, the percentage of timeouts was small for all display conditions (Fig. 9B). There were no observed timeout errors with the regular display or the only hands conditions. A small percentage of only colors (4.5%) and only target (3.8%) trials ended in timeouts. Consequently the percentage correct of completed trials was identical or nearly identical to the success rate in all cases. The mean movement time only slightly increased from the regular display (1.01 s) to the only hand (1.10 s) condition, but increased significantly for the only colors (1.76 s) and only target (1.72 s) conditions ($p < 0.05$, bootstrapped confidence intervals). Combining percentage correct and mean movement time reduced the bit rate from 2.2 to 1.3 bits/s when the regular display was reduced to only hands but fell to 0 when only colors were presented as the monkey was performing at chance. Hence the monkey relied more on the virtual hands than the background colors. The no avatar, only target condition resulted in a low bit rate of 0.3 bits/s albeit above chance performance. Thus neural modulation and resulting control were improved with visual feedback and were not simply a ballistic neural trajectory generated after the target was displayed.

4. Discussion

The ADS decoder provided multi-DOF BMI control of an avatar hand that enabled the monkey to quickly and accurately select the presented target postures. The bit rate of 2.4 bits/s for eight targets in four dimensions using an avatar hand compares favorably to many standard 2D center-out BMI task bit rates of 1 bit/s or less [39]. A more recent study reported a bit rate of 3–4 bits/s [40] in a standard 2D center out task although with additional processing using Kalman filtering with online BMI optimization and more spiking channels. Here, BMI control was achieved with simple linear decoding weights, using no Kalman filtering or other advanced statistical processing, no online updating within a session, and only 16 spiking units. For simple target selection tasks, rather than controlling a cursor on a 2D screen, future consideration should be given to BMI schemes that use a higher number of output dimensions but reduced movement within those dimensions.

Active dimension selection leverages the fact that complex movements like grasping are not random but predictable and coordinated across joints. BMIs are limited because the number of recorded neurons is relatively small compared to the entire neural population of the natural motor system. With limited information, performance can be improved when good inferences can be made about the intended movements. Better statistical priors improve prediction. In reality, movements do not uniformly span even the reduced space of low-dimensional principal component descriptions. Although not emphasized, hand grasps plotted in the first two principal components in Figure 7 of Santello et al. [11] show a distinctive “L” pattern in the PC space rather than being uniformly distributed. Thus, a large fraction of the hand grasp shapes can be defined by specifying points along the two axes defined by the “L,” without needing to cover the entire two-dimensional space of the first

two PCs. The present results illustrate the advantages of constraining movements to such principal axes within a state-space.

Using neural signals to select the active dimension effectively transformed BMI control from a four-dimensional task into four one-dimensional tasks. This permitted the ADS decoder to select targets quickly and correctly, with only a slight performance decrement compared to one-dimensional, CDS control. Conversely, Full control became increasingly difficult as the number of dimensions increased, with many trials having no target contacted within the 5 s time limit. Using catch trials with no cue as to control type (ADS vs. CDS vs. Full control), the information and precision possible with identical neural signals was measured accurately. Although the number of neural signals and their noise limit any BMI implementation, by ignoring the noise in unwanted dimensions ADS improves on classic Full-dimensional decoders for performing a specific set of BMI-controlled movements.

Feedback and neural adaptation are critical for BMI applications. Although initial training as well as the feedback experiment (Experiment 2) was performed using only ADS decoding, one can speculate about the differences in the effect of training and feedback that might be expected with ADS decoding compared to traditional full-dimensional decoding. BMI performance is determined by how well users can generate neural signals that are easily decodable with little overlap between different conditions. Training in the ADS condition allows the user to develop neural activity that has strong modulation along the active dimension being controlled with no need to precisely modulate other dimensions of control, provided the proper dimension has been selected. Additionally, determining through visual feedback which dimensions to flex or extend can be difficult in an increasingly high dimensional space. The feedback is likely less ambiguous in the ADS environment where the different classes or selected dimension result in distinct hand movements and only one dimension must be monitored at any given time. These factors suggest that ADS could improve learning rates and performance compared to traditional full-dimensional BMIs. Further work will be needed to directly compare the effects of training and feedback on performance with ADS versus full-dimensional decoders.

The proof-of-principle ADS avatar hand presented here was limited in two aspects that could likely be improved in the future. First, in the current implementation, the movement was limited to only one dimension out of the larger, four dimensional state-space. For instance, the power grasp could not be performed with varying degrees of flexion at individual digits as might be needed for an irregularly shaped rock versus a spherical ball. With more recorded neurons and more complex models, active dimension selection could be expanded to enable the BMI to actively control two or three dimensions at a time in a state-space with an even higher number of dimensions. The optimal ratio of actively controlled to total state-space dimensions has yet to be examined. Second, the current decoding scheme was also relatively limited in the precision of movement along a given dimension. The experiment was designed and the monkey was trained to select different target postures as quickly as possible with no required hold time. The largest magnitude velocity signal was used to select the active dimension, making it difficult to hold a constant, neutral posture. This could potentially be remedied by adding a state or dimension that holds the hand stationary, as well as by improving the decoding of precise, low-velocity movements.

Active dimension selection is a novel technique that allows the neural signals to select accurate control of a subset of dimensions important to the task, while offloading control of other dimensions to a more automated process. This approach allows grasp dimensions to be defined that are not required to be orthogonal in joint space, and thereby generates efficient trajectories along a single grasp dimension. In contrast, other approaches define dimensions of grasping based on either physical degrees of freedom of a prosthetic [10] or orthogonal principal components of hand movement [18]. Such approaches require simultaneous control of each added dimension to increase the available movement repertoire; transitioning to these added dimensions can lead to inconsistency and frustration [10]. The ADS approach allows for the possibility of adding novel hand movements with less interference from previously learned movements. Further work will be needed to explore how to balance the flexibility of simultaneous dimension control versus efficient movement paths of ADS to best select and generate useful movements.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The author would like to thank Marc H. Schieber for providing laboratory resource support, experimental time with the animal used in this study, and assistance with editing this manuscript. Additionally, the author would like to thank Matthew Hazelett and Samuel Brown for their assistance in software development. Finally, the author would like to thank Drs. Gerald E. Loeb and Rahman Davoodi as well as the rest of the Medical Device Development Facility at the University of Southern California for their development of the MSMS software package. This work was funded by the National Institute of Health (NIH) through National Institute of Neurological Disorders and Stroke grants R01NS065902 and R01NS079664 to MHS.

References

1. Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP. Instant neural control of a movement signal. *Nature*. 2002; 416:141–2. [PubMed: 11894084]
2. Georgopoulos AP, Schwartz AB, Kettner RE. Neuronal Population Coding of Movement Direction. *Science*. 1986; 233:1416–9. [PubMed: 3749885]
3. Taylor DM, Helms Tillery SI, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science*. 2002; 296:1829–32. [PubMed: 12052948]
4. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. *Nature*. 2008; 453:1098–101. [PubMed: 18509337]
5. 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*. 2012; 485:372–5. [PubMed: 22596161]
6. Johannes MS, Bigelow JD, Burck JM, Harshbarger SD, Kozlowski MV. An Overview of the Developmental Process for the Modular Prosthetic Limb. *Johns Hopkins APL Tech Dig*. 2011; 30:207–16.
7. Resnik L, Klinger SL, Etter K. The DEKA Arm: Its features, functionality, and evolution during the Veterans Affairs Study to optimize the DEKA Arm. *Prosthet Orthot Int*. 2013; 38:492–504. [PubMed: 24150930]
8. Hutchinson DT. The Quest for the Bionic Arm. *J Am Acad Orthop Surg*. 2014; 22:346–51. [PubMed: 24860130]
9. Bennett DA, Dalley SA, Truex D, Goldfarb M. A Multigrasp Hand Prosthesis for Providing Precision and Conformal Grasps. *IEEE/ASME Trans Mechatronics*. 2015; 20:1697–704.

10. Wodlinger B, Downey JE, Tyler-Kabara EC, Schwartz AB, Boninger ML, Collinger JL. Ten-dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations. *J Neural Eng.* 2015; 12:016011. [PubMed: 25514320]
11. Santello M, Flanders M, Soechting JF. Postural hand synergies for tool use. *J Neurosci.* 1998; 18:10105–15. [PubMed: 9822764]
12. Mason CR, Theverapperuma LS, Hendrix CM, Ebner TJ. Monkey hand postural synergies during reach-to-grasp in the absence of vision of the hand and object. *J Neurophysiol.* 2004; 91:2826–37. [PubMed: 14762155]
13. Todorov E, Ghahramani Z. Analysis of the synergies underlying complex hand manipulation. *Conf Proc IEEE Eng Med Biol Soc.* 2004; 6:4637–40. [PubMed: 17271341]
14. Ciocarlie M, Goldfeder C, Allen P. Dimensionality reduction for hand-independent dexterous robotic grasping 2007. *IEEE/RSJ Int Conf Intell Robot Syst.* 2007; 20:3270–5.
15. Matrone GC, Cipriani C, Secco EL, Magenes G, Carrozza MC. Principal components analysis based control of a multi-DoF underactuated prosthetic hand. *J Neuroeng Rehabil.* 2010:7–16. [PubMed: 20149260]
16. Holdefer RN, Miller LE. Primary motor cortical neurons encode functional muscle synergies. *Exp brain Res.* 2002; 146:233–43. [PubMed: 12195525]
17. Artemiadis PK, Shakhnarovich G, Vargas-irwin CE, Donoghue JP, Black MJ. Decoding grasp aperture from motor-cortical population activity. *Proc IEEE EMBS Conf Neural Eng.* 2007:518–21.
18. Vinjamuri R, Weber DJ, Mao Z-H, Collinger JL, Degenhart AD, Kelly JW, Boninger ML, Tyler-Kabara EC, Wang W. Toward synergy-based brain-machine interfaces. *IEEE Trans Inf Technol Biomed.* 2011; 15:726–36. [PubMed: 21708506]
19. Velliste M, McMorland aJC, Diril E, Clanton ST, Schwartz AB. State-space control of prosthetic hand shape. *Conf Proc IEEE Eng Med Biol Soc.* 2012; 2012:964–7. [PubMed: 23366054]
20. Napier JR. The prehensile movements of the human hand. *J Bone Jt Surg - Br.* 1956; 38:902–13.
21. Arbib, MA., Iberall, T., Lyons, D. Coordinated Control Programs for Movements of the Hand. In: Goodwin, AW., Darian-Smith, I., editors. *Hand Function and the Neocortex (Experimental Brain Research Series)*. Heidelberg: Springer-Verlag; 1985. p. 111-29.
22. Ajiboye AB, Weir RF, Heckathorne CW, Childress DS. Neurofuzzy logic as a control algorithm for an externally powered multifunctional hand prosthesis. *Proc 2002 MyoElectric Control Prosthetics Symp.* 2002:21–4.
23. Kent BA, Lavery J, Engeberg ED. Anthropomorphic Control of a Dexterous Artificial Hand via Task Dependent Temporally Synchronized Synergies. *J Bionic Eng.* 2014; 11:236–48.
24. Parker P, Englehart K, Hudgins B. Myoelectric signal processing for control of powered limb prostheses. *J Electromyogr Kinesiol.* 2006; 16:541–8. [PubMed: 17045489]
25. Rouse AG, Schieber MH. Advancing brain-machine interfaces: moving beyond linear state space models. *Front Syst Neurosci.* 2015; 9:1–13. [PubMed: 25709570]
26. Kim HK, Biggs SJ, Schloerb DW, Carmena JM, Lebedev MA, Nicolelis MAL, Srinivasan MA. Continuous shared control for stabilizing reaching and grasping with brain-machine interfaces. *IEEE Trans Biomed Eng.* 2006; 53:1164–73. [PubMed: 16761843]
27. Markovic M, Dosen S, Cipriani C, Popovic D, Farina D. Stereovision and augmented reality for closed-loop control of grasping in hand prostheses. *J Neural Eng.* 2014; 11:046001. [PubMed: 24891493]
28. Muelling, K., Venkatraman, A., Valois, J., Downey, JE., Weiss, J., Javdani, S., Hebert, M., Schwartz, AB., Collinger, JL., Bagnell, JA. Autonomy Infused Teleoperation with Application to BCI Manipulation. *Proceedings of Robotics: Science and Systems; Rome, Italy.* 2015.
29. McMullen DP, Hotson G, Katyal KD, Wester Ba, Fifer MS, McGee TG, Harris A, Johannes MS, Vogelstein RJ, Ravitz AD, Anderson WS, Thakor NV, Crone NE. Demonstration of a semi-autonomous hybrid brain-machine interface using human intracranial EEG, eye tracking, and computer vision to control a robotic upper limb prosthetic. *IEEE Trans Neural Syst Rehabil Eng.* 2014; 22:784–96. [PubMed: 24760914]

30. Hotson G, Smith R, Rouse A, Schieber MH, Thakor N, Wester B. High Precision Neural Decoding of Complex Movement Trajectories using Recursive Bayesian Estimation with Dynamic Movement Primitives. *IEEE Robot Autom Lett.* 2016:1–8.
31. 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 Biol.* 2003; 1:193–208.
32. Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJC, Velliste M, Boninger ML, Schwartz AB. High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet.* 2013; 381:557–64. [PubMed: 23253623]
33. Davoodi R, Loeb GE. Real-time animation software for customized training to use motor prosthetic systems. *IEEE Trans Neural Syst Rehabil Eng.* 2012; 20:134–42. [PubMed: 22186964]
34. Mollazadeh M, Aggarwal V, Davidson AG, Law AJ, Thakor NV, Schieber MH. Spatiotemporal variation of multiple neurophysiological signals in the primary motor cortex during dexterous reach-to-grasp movements. *J Neurosci.* 2011; 31:15531–43. [PubMed: 22031899]
35. Ashe J, Georgopoulos aP. Movement parameters and neural activity in motor cortex and area 5. *Cereb Cortex.* 1994; 6:590–600.
36. Green, D., Swets, J. *Signal Detection Theory and Psychophysics.* New York: Wiley & Sons, Inc; 1966.
37. Pierce, JR. *An Introduction to Information Theory.* Dover; 1980.
38. McFarland DJ, Sarnacki WA, Wolpaw JR. Brain – computer interface (BCI) operation: optimizing information transfer rates. *Biol Psychol.* 2003; 63:237–51. [PubMed: 12853169]
39. Tehovnik EJ, Woods LC, Slocum WM. Transfer of information by BMI. *Neuroscience.* 2013; 255:134–46. [PubMed: 24120558]
40. Nuyujukian P, Fan JM, Kao JC, Ryu SI, Shenoy KV. A high-performance keyboard neural prosthesis enabled by task optimization. *IEEE Trans Biomed Eng.* 2014; 62:21–9. [PubMed: 25203982]

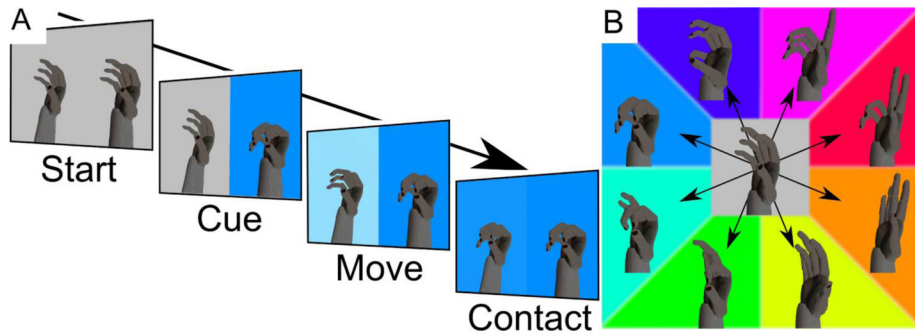


Figure 1.

Virtual hand grasping task. A) The task display featured a target hand posture and background color on the right and the brain controlled avatar hand with a colored background (representing the current location in the state-space) on the left. Each trial consisted of i) Start both – the target hand and avatar hand were returned to a neutral posture, ii) Cue – the target hand was moved abruptly to one of eight possible hand postures while the avatar hand remained fixed for 300 ms, iii) Move – the avatar hand moved under brain control, the example video frame shows the avatar hand (left) midway to the flexed power grasp target posture (right); the less saturated blue (left) also instructs the monkey more flexion is necessary to match the fully saturated blue (right), and iv) Contact – the avatar hand is within the flexed power target boundaries. B) The eight different target hand postures along the four different grasp dimensions are shown schematically, collapsed into a two-dimensional arrangement. Each of the four grasp dimension has target flexion and extension hand postures (with associated colors) – precision **pinch** (green/magenta), three-finger **tripod** (cyan/red), **power** grasp (blue/orange), and thumb and little finger **opposition** (indigo/yellow). These four dimensions were mapped to the 20 DOFs of the virtual hand.




Figure 2.

State-space variables with ADS decoding for a single trial of power grasp flexion. The four state-space variables are plotted on a single axis. Dotted horizontal lines represent the flexion (+0.5) and extension (-0.5) targets that must be crossed along the single target dimension. Dashed horizontal lines represent the region around zero (-0.167 to +0.167) that the other state dimension values must be within to complete a successful target contact. The power grasp dimension (blue) was required to cross the flexion target line at +0.5. Another dimension (light grey) was initially active, and slightly extended before neural control switched to select the correct dimension, after which the ADS decoder returned this dimension to zero. The other two dimensions (black and dark grey) never deviated from zero. N.B. Scores for different dimensions have been offset slightly from zero for display purposes.

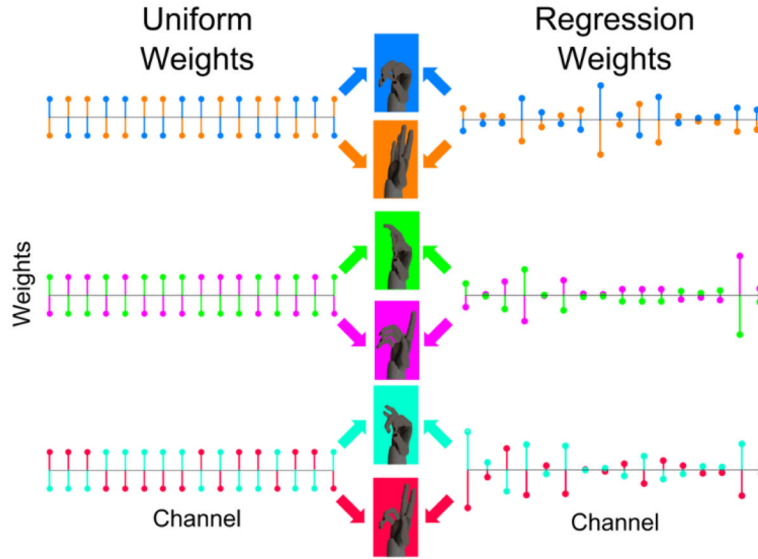


Figure 3. Schematic of channel weights for the 3D ADS decoder. The uniform weights on the left are either +1 or -1 for all 16 channels. Each dimension (power – blue/orange, pinch – green/magenta, tripod – cyan/red) has 8 positive weights and eight negative weights. If a weight is +1 (e.g. blue up and orange down in the power grasp dimension), then flexion of the given grasp is facilitated when the unit has a firing rate above baseline and extension is facilitated when the firing rate is below baseline. Conversely, units with -1 weights (e.g. orange up and blue down) facilitate extension with increased firing and flexion with decreased firing. The sum across all 16 units of the unit’s weight multiplied by its firing rate determines the net velocity of flexion/extension. In ADS control, only the dimension with the maximum absolute signal moves while the other dimensions are passively fixed. The regressed weights on the right are similarly used by the decoder, but can be different in their magnitude as determined by linear least squares regression.

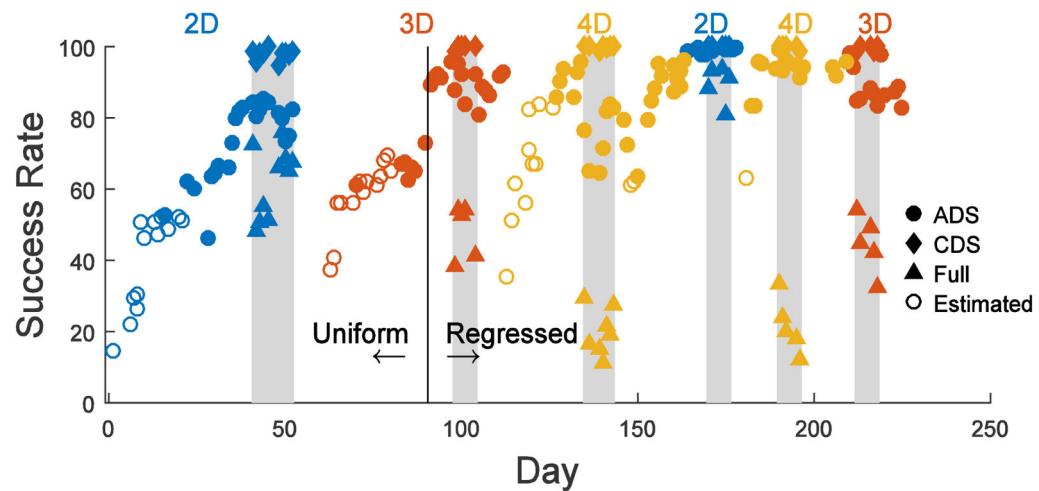


Figure 4.

Success rate across recording sessions during training (white regions) and experiment 1 (grey regions). The monkey was trained to control first the 2D, 3D, and then 4D virtual hand. Following this initial training, 2D, 4D, and 3D control were tested again. Open circles represent days when computer assistance was used, and the virtual hand was only partially controlled by the brain. For these days, the success rate has been estimated with post-hoc analysis of how the recorded neural signals would have performed without computer assistance. Initially, the decoding weights were uniform across all 16 channels. On day 91 (vertical black line), the uniform weights were replaced with regressed weights based on the previous day's data. The grey regions are days when CDS (diamond) and full (triangle) control catch trials were performed in addition to standard trials performed with the ADS (circle) decoder.

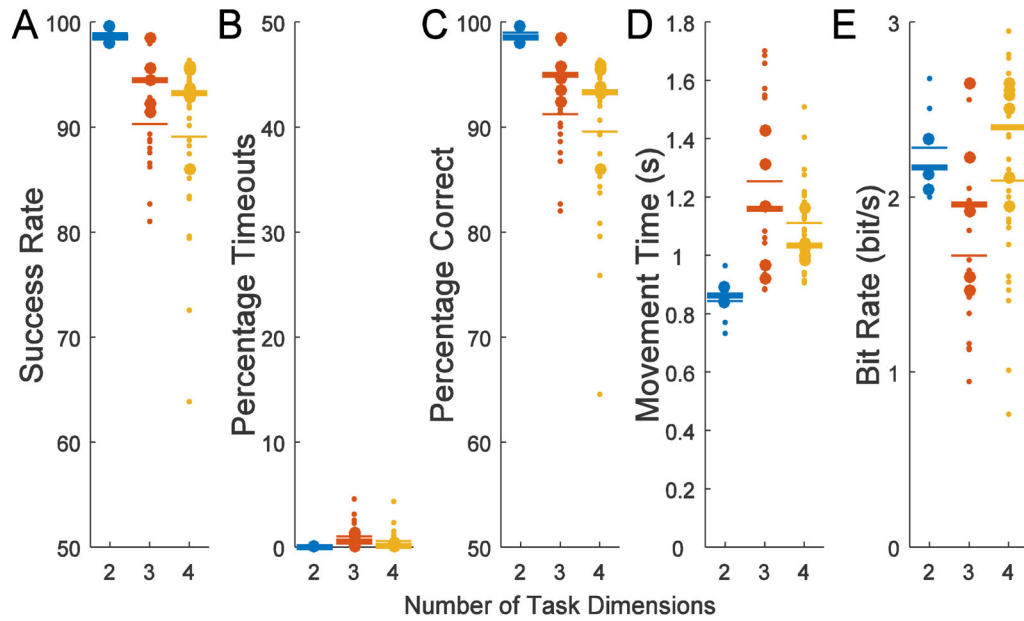


Figure 5.

Performance for ADS virtual hand control in ADS-only training sessions as measured by A) success rate, B) percentage timeouts C) percentage correct, D) mean movement time, and E) bit rate. Success rate is the percentage of all trials completed successfully. The percentage correct includes only trials completed and not trials that timed out. The mean movement time includes all trials. The bit rate is a combination of the percentage correct and mean movement times. The larger circles represent the late training sessions as defined as those immediately preceding switching to the Full/CDS/ADS experiment (3 sessions preceding grey regions in Fig. 4). For 2D, there was only one set of three sessions while for 3D and 4D there were two sets of three sessions. However, for 3D there are only five data points since the Full/CDS/ADS experiment was started once with only two preceding all-ADS sessions. The thinner horizontal lines are means across all ADS training sessions with regressed weights while bold horizontal lines are means for larger circles representing late training sessions immediately preceding Full/CDS/ADS experiment sessions.

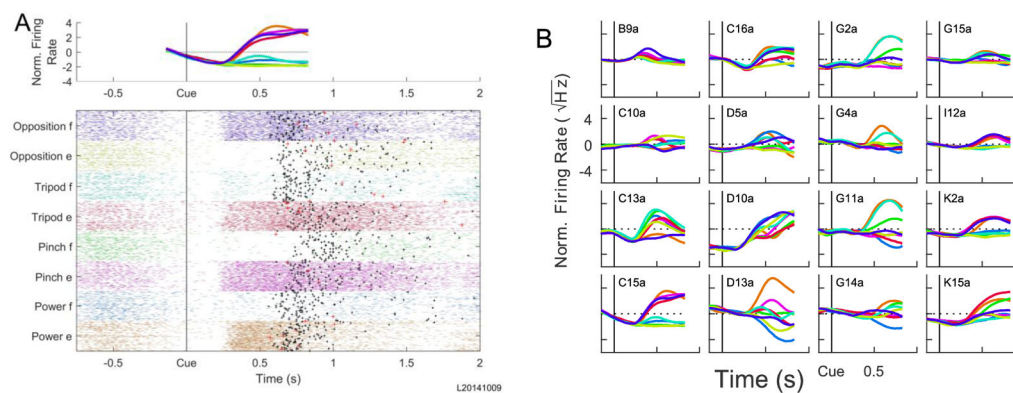


Figure 6.

Individual unit firing rates. A) A raster plot is shown of all observed action potentials for trials from a single 8 target, 4D session for an example unit (C15a). Trials are sorted by the eight target hand postures and aligned at the Cue (vertical black line). Black dots indicate when the correct target was contacted while larger red crosses indicate contact with an incorrect target. The mean normalized firing rate calculated by the online control algorithm as a function of time for the eight different trial types is plotted on top. B) The mean normalized firing rate for all 16 units used for control in the same session.

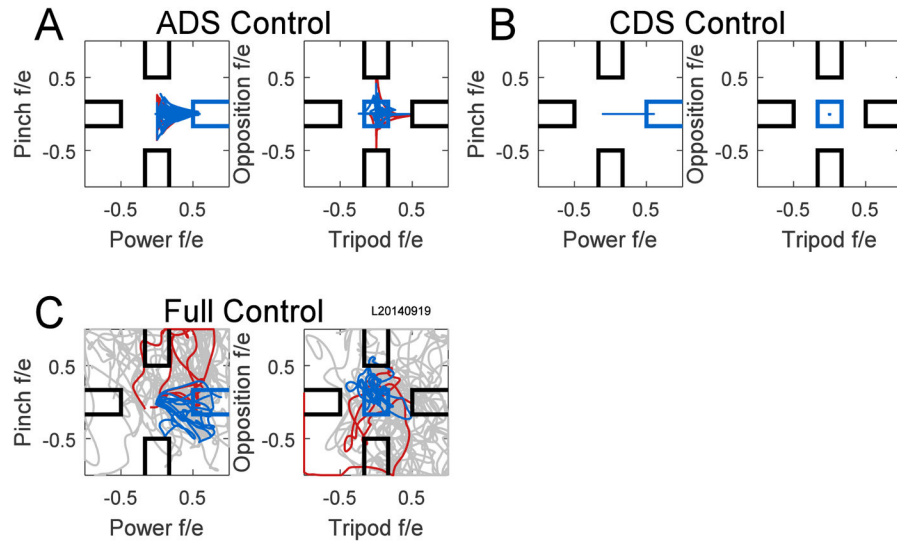


Figure 7.

State-space variable trajectories for power grasp flexion trials in a 4D session with catch trials. A) All ADS trials are shown from a single recording session. The four dimensions of the state-space are shown in two separate two-dimensional plots. Power and pinch are the abscissa and ordinate of the left plot; tripod and opposition are the abscissa and ordinate of the right plot, respectively. Trajectories of correct trials are blue, while trajectories that contacted incorrect targets are red. B) CDS control catch trials. The power grasp dimension was selected automatically by the computer and the neural signals only controlled movement in the power dimension. All CDS trials in the session were correct. C) Full control catch trials. Trajectories from correct trials are blue; from incorrect trials, red; and from trials that timed out after 5 s, grey. f/e, flexion(+)/extension(-).

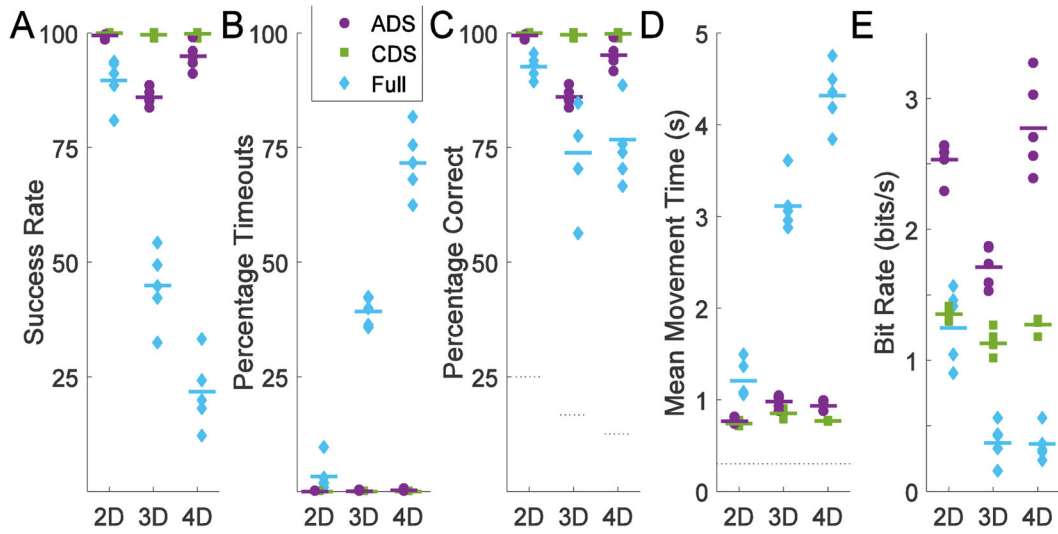


Figure 8. Performance comparison of Experiment 1. The final five recording sessions in which the control algorithm was switched between ADS, CDS, and Full control were analyzed for each number of dimensions. A) Success rate incorporate error trials, including trials with 5 s timeouts. B) The percentage timeouts is percentage of trials with 5 s timeouts. C) The percentage correct only includes trials where a target was contacted. Grey lines show chance levels for 4, 6, and 8 targets. D) Mean movement times again incorporate error trials, including trials with 5 s timeouts. The grey line indicates the 300 ms before the avatar was allowed to move. E) Bit rate combines mean movement time and percentage correct into a single metric of maximum communication capacity based on information theory.

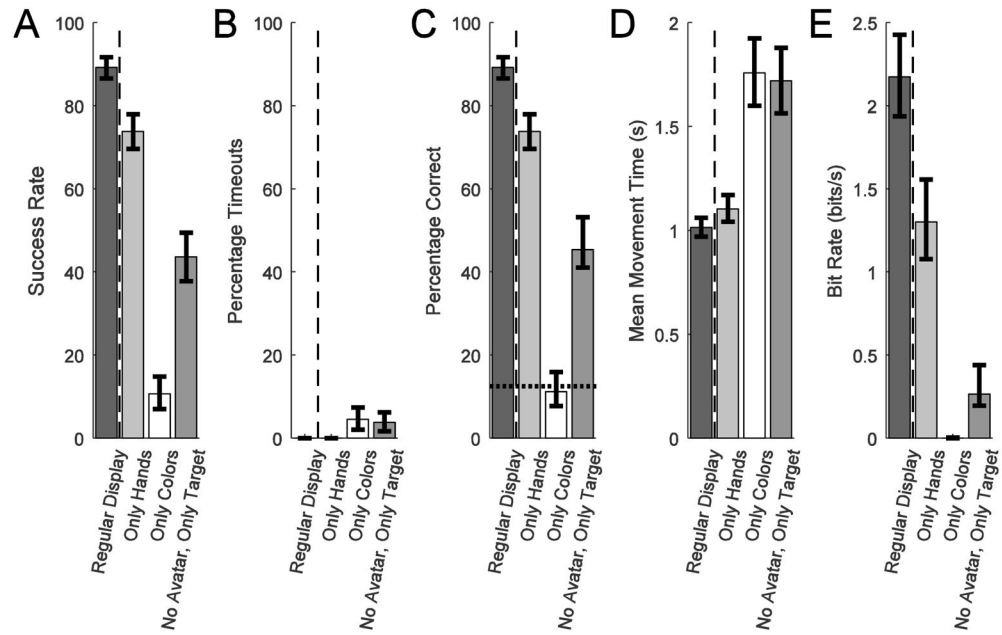


Figure 9.

Performance for three different reduced display conditions compared to the regular task display using ADS decoding. In the Only Hands condition, the target and avatar hands were displayed, but the background color was a static grey. In the Only Colors condition, the target color and shifting avatar color were displayed, but not the target or avatar hands. In the No Avatar, Only Target condition, the target hand and its associated color were displayed, but not the avatar hand or color. The averages were calculated across approximately 450 trials in each condition pooled from three sessions. The 95% confidence intervals are shown using binomial proportions for the percentages, bootstrapping for the movement times, and combined best and worst percentage correct and movement times for the bit rate.