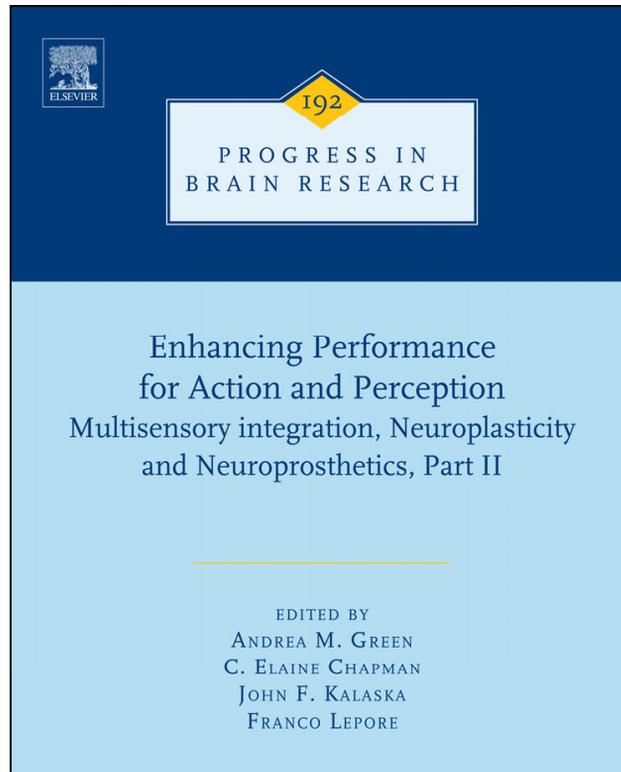


**Provided for non-commercial research and educational use only.
Not for reproduction, distribution or commercial use.**

This chapter was originally published in the book *Progress in Brain Research*, Vol. 192, published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues who know you, and providing a copy to your institution's administrator.



All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>

From: Steven M. Chase and Andrew B. Schwartz, Inference from populations: going beyond models. In Andrea M. Green, C. Elaine Chapman, John F. Kalaska, Franco Lepore, editors: *Progress in Brain Research*, Vol. 192, Amsterdam: The Netherlands, 2011, pp. 103-112.

ISBN: 978-0-444-53355-5

© Copyright 2011 Elsevier B.V.

Elsevier

CHAPTER 7

Inference from populations: going beyond models

Steven M. Chase and Andrew B. Schwartz*

Department of Neurobiology and the Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, Pennsylvania, USA

Abstract: How are abstract signals, like intent, represented in neural populations? By creating a direct link between neural activity and behavior, brain–computer interfaces (BCIs) can help answer this question. Early instantiations of these devices sought mainly to mimic arm movements: by building models of arm tuning for the neurons, desired arm movements could be read out and used to control various prosthetic devices. However, as the functionality of these devices increases, a more general approach that relies less on endogenous control signals may be required. Here we review some of the current, model-based approaches for finding volitional control signals for spiking-based BCIs, and present some new approaches for finding control signals without resorting to parametric models of neural activity.

Keywords: brain–machine interface; motor cortex; decoding; cosine tuning.

Spiking-based brain–computer interfaces (BCIs) map the activity of dozens to hundreds of recorded neurons to the control of some device. This device could be something fairly simple, like a spelling tool that can help locked-in patients communicate with the outside world (Musallam et al., 2004; Santhanam et al., 2006), or it could be something more complex, like a cursor on a computer screen (Hochberg et al., 2006; Kennedy et al., 2000; Mulliken et al., 2008; Serruya et al., 2002; Taylor et al., 2002), a robotic arm (Chapin

et al., 1999; Velliste et al., 2008; Wessberg et al., 2000), or even a muscle stimulator (Moritz et al., 2008; Pohlmeier et al., 2009), that could help paralyzed patients regain the ability to move. In each case, the goal of these devices is the same: translate the intent of the user, encoded in trains of action potentials, into the desired action of the device.

At first blush, this seems quite simple. If the user wants the computer cursor to move upward, he simply imagines the cursor moving upward: somewhere in the brain, this process of imagination causes changes in the firing rates of neurons, and the BCI needs only to decode these changes into upward cursor movement. The reality,

*Corresponding author.

Tel.: +412-383-7021; Fax: +412-383-5460

E-mail: abs21@pitt.edu

however, is probably quite a bit different. Because the neural processes related to imagination and imagery are not well understood, they are hard to tap into with any degree of reliability. Further, the process of imagining typing the letter “p” seems fundamentally different from the process of imagining “grab the cup.” Does this imply that a BCI that controls a spelling device must use a different set of neurons than a BCI that controls a robotic hand?

The answer, of course, is no. Most BCIs that rely on single-neuron recordings extract them from the primary motor cortex (e.g., Hochberg et al., 2006; Serruya et al., 2002; Taylor et al., 2002; Velliste et al., 2008). Since this area is a major source of “output” from the brain, it is a logical place to find signals that might be used to control an external device. The idea is that by tapping into endogenous control signals related to some observable, quantifiable behavior, like arm movement, the subject could eventually learn to associate those behaviors with certain actions of the device. This has proven to be a successful tactic: signals from the proximal arm area of primary motor cortex have successfully been used to control computer cursors moving in two (Carmena et al., 2003; Hochberg et al., 2006) and three (Jarosiewicz et al., 2008; Taylor et al., 2002) dimensions, and a monkey was even successfully trained to feed itself with a 4 degree of freedom robotic arm (Velliste et al., 2008). However, the human hand alone contains ~ 20 independently controllable degrees-of-freedom. To further complicate things, the neural representation of hand shape is only beginning to be worked out (see, e.g., Hendrix et al., 2009; Lemon, 1993; Vargas-Irwin et al., 2010). There is still a long way to go before prosthetic arms can replace the capabilities of a lost limb.

In this chapter, we review some of the approaches taken toward extracting potential control signals for BCI devices, and argue that to achieve the next generation of prosthetic devices we will need a new approach that does not rely on assumptions about how arm

movements might actually be encoded. By going beyond these models, we can use the neural population responses to infer the existence of potential control signals, without making particular reference to what those control signals might be. We postulate that an understanding of these control signals will not only facilitate the design of high functioning prosthetic devices, but will also help answer basic questions about the neural correlates of intent.

Prostheses based on arm movements

In the 1980s, Georgopoulos and colleagues released a series of papers detailing how populations of neurons in the proximal arm area of motor cortex might represent arm movements (Georgopoulos et al., 1982; Schwartz et al., 1988), and further how this information could be read out from simultaneously recorded neurons (Georgopoulos et al., 1986, 1988; Kettner et al., 1988). These papers detailed an approach to neural decoding which has become fairly common in the BCI field: first perform a series of experiments with natural arm movements to build an encoding model detailing how firing rates depend on arm movements, then invert this encoding model and use decoded arm movements to drive a remote effector (e.g., a computer cursor; Taylor et al., 2002). Alternately, specification of the encoding model could be skipped entirely, and instead the arm movement data could be used to train an algorithm to extract kinematic data directly from the spike trains (e.g., Mulliken et al., 2008; Wessberg et al., 2000). In either case, this approach relies extensively on arm movement data to calibrate the decoding device.

Apart from the obvious clinical problem of using arm movement data to train a device that is supposed to replace the functionality of a lost arm, this approach may not be ideal because it assumes the control signals appropriate for an arm are the same as the signals appropriate for the device. A number of studies now suggest that

this is not the case. First, it has been observed that as the subjects learn to modulate their cortical signals to control the cursor directly, they stop moving their arms (Carmena et al., 2003; Chapin et al., 1999; Taylor et al., 2002). This indicates that there must be at least *some* change in the neural signal that is specific to brain-control. Second, tuning curves recorded during the arm control session have been found to be different from the tuning curves recorded during the brain-control session (Carmena et al., 2003; Taylor et al., 2002); in fact, these tuning curve differences tend to increase with training (Ganguly and Carmena, 2009; Taylor et al., 2002). Taylor and colleagues investigated this issue in detail by recording from neurons in the primary motor cortex of monkeys while they performed center-out reaching movements under both hand- and brain-control. On the first day, they found that neurons changed their preferred directions between the two sessions by an average of $\sim 65^\circ$, and these tuning curve differences increased as the subject trained with the BCI (Fig. 1). This change in tuning is indicative of a brain-control specific change in the control signals.

It is still an open question as to *why* the tuning differs between hand-control and brain-control. It could be that the motor cortex is sensitive to the differences in the dynamic properties of the remote effector and the limb. Another possibility is that the neural activity is influenced by the proprioceptive feedback, which changes between the hand-control and brain-control tasks because the subjects are no longer moving their arms. In fact, Hatsopoulos and colleagues have shown that when the arm is passively moved in concert with a brain-controlled cursor, decoding performance improves (Suminski et al., 2010). Since they calibrated their decoder with arm movement data, this could be an indication that the tuning curves change less between hand-control and brain-control when the arm is kept moving in the brain-control condition. An interesting implication of this work is that it suggests that providing proper proprioceptive feedback to

patients might improve their ability to use a neural prosthetic device. To date, however, providing realistic proprioceptive feedback to patients remains a challenge.

Prostheses based on cursor imagery

Another, related approach does not rely on arm movements themselves, but rather on motor imagery. In work with human subjects, Hochberg and colleagues found that many cells in primary motor cortex could be driven by imagined movements, such as imagined wrist or elbow flexion (2006). Presumably, the neural tuning to these “imagined” movements is similar to the tuning these neurons would have had to natural arm movements, and studies in monkeys comparing the neural tuning during active movement with that during passive observation seem to bear this out (Tkach et al., 2007). Electroencephalographic signals recorded from humans have shown that the modulations induced by motor imagery have a similar spatial distribution to the modulations produced by actual movement, although they are weaker. Providing visual feedback about the modulation, however, can increase the magnitude of the modulations to be the same as or greater than with overt movement (Miller et al., 2010). While these findings allow researchers to overcome the clinical problem of relying on arm movement data to build their decoder, as with the arm movement data there is evidence that these imagery-based neural tunings may still not be capturing all of the neural activity that could be used to drive a prosthetic device (Wahnoun et al., 2006).

Several researchers have explored iterative approaches to try to uncover the right set of control signals. Although details can vary from lab to lab, an example of the procedure used by Schwartz and colleagues is as follows (Chase et al., 2009; Fraser et al., 2009; Jarosiewicz et al., 2008). First, the decoding parameters are initialized, by assigning values at random to every neuron. Alternate approaches involve initializing

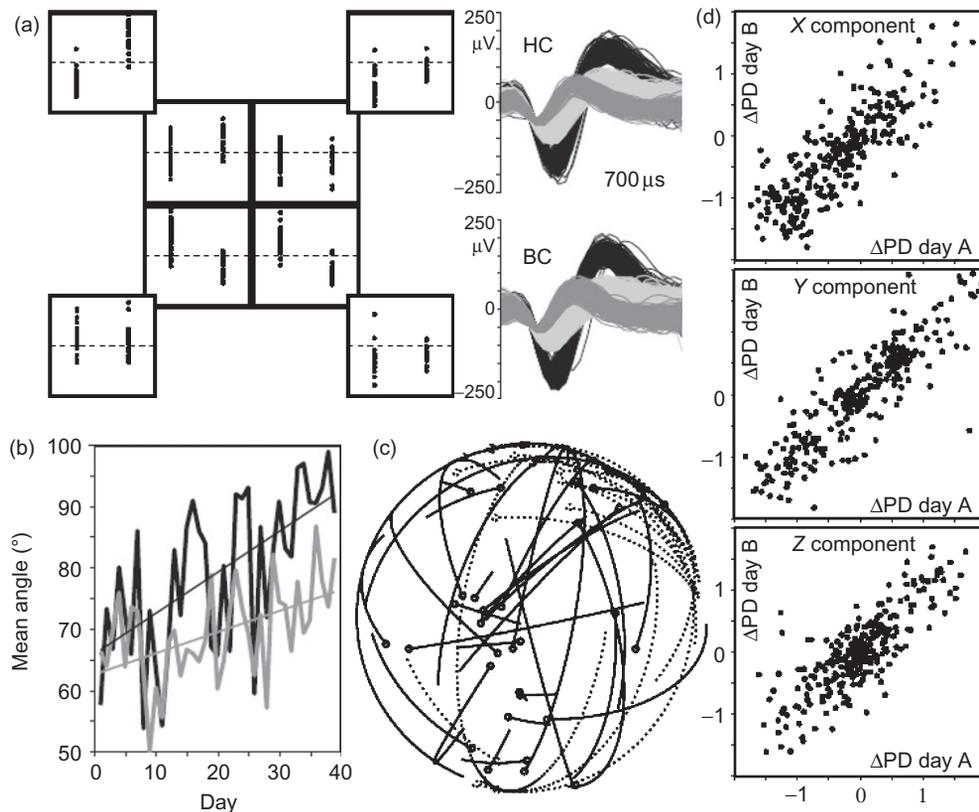


Fig. 1. Changes in cortical activity between hand-control (HC) and brain-control (BC) tasks in subject M. (a) Cell with a 107° change in tuning direction (the unit waveform is shown in black). Each dot is the mean firing rate during one movement. HC rates are in the right column and BC rates are in the left column of each square. The eight target directions (center four=distal; outer four=proximal). (b) Daily mean angles (thick lines) between HC and BC preferred directions for all cells significantly tuned during both tasks (black, contralateral; gray, ipsilateral to the arm moved during the HC task). (c) Lines connecting HC preferred directions with BC preferred directions (circle ends) projected onto a unit sphere (day 8, only cells significantly tuned in both tasks; black, contra; dotted, ipsi). (d) Change in the X, Y, and Z components of the preferred direction unit vectors between the HC and BC tasks plotted day-against-day for 8 random pairs of days (day 27 or later, only units that were significantly tuned in both tasks on both days; 35 ± 3 units per pair of days). With permission from Taylor et al. (2002).

with parameters from arm data (as in Taylor et al., 2002) or from a motor observation or imagery session (as in Hochberg et al., 2006; Wahnoun et al., 2006). Targets are then presented, one at a time in random order, and left on the screen until a movement time-out period elapses. Subjects are instructed to attempt to hit the target, and cells in the motor cortex tend to modulate as a function of target direction, indicating that an attempt is

being made. Depending on the initialization, however, the attempt may or may not be successful; with the random assignment initialization, the cursor hardly moves at all (Chase et al., 2009). Regardless of success or failure, after each target has been presented, firing rates are regressed against target direction to build an encoding model describing each cell's tuning to desired movement, and these tuning curves are used to

recompute decoding parameters. Then another round of targets is presented. It typically takes only a few minutes or so to converge to a stable set of decoding parameters. Once convergence is reached, the decoding parameters are fixed.

Not only does this iterative procedure result in tuning curves that are substantially different from those observed during arm movement tasks (Taylor et al., 2002) or passive observation (Wahnoun et al., 2006), it also results in better control. Wahnoun and colleagues characterized various parameters of control during the iterative procedure (Fig. 2). In less than 5 min, they found substantial straightening of the trajectories, and the average time to hit the targets decreased by about 20%. Further, their improvements in

control were accompanied by a decrease in the average firing rates of the cells. This might indicate that the iterative procedure resulted in control that required less effort, an important consideration for a device which might need to be used by patients for long periods of time.

Although motor imagery approaches have been successfully applied to prosthetic control, as with arm movement approaches they still suffer from one major failing: they do not account for types of neuronal modulation that might not be included in the tested range of movements. For example, if some of the recorded neurons happen to encode neck movement, and neck movements are not included in the set of instructed movements or imagined movements,

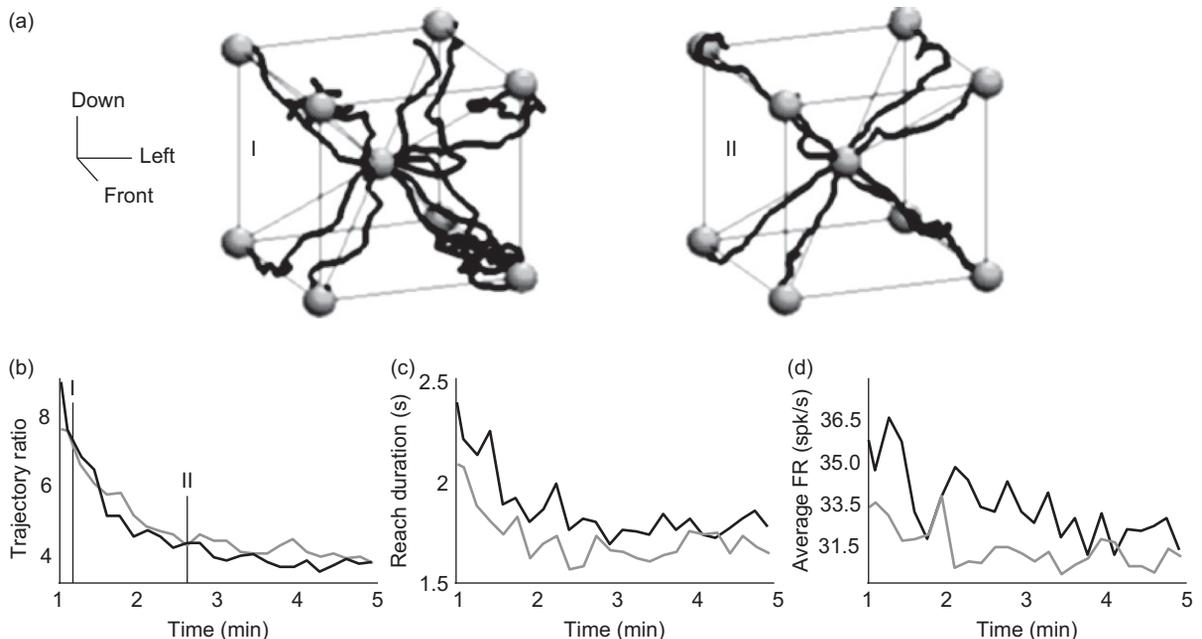


Fig. 2. Change in performance as parameters for the control algorithm are tuned. (a) illustrates individual cortically controlled cursor movements to each of the eight targets at two times following the beginning of tuning: (I) at the end of the visual following task; (II) after 60 s of additional data accumulation. Figures are drawn as shown on the screen prior to reflection in the mirror, leading to a left-handed coordinate system. (b, c, and d) show general characteristics of the movements and neural activity. (b) Ratio of the summed path length to the direct distance from the center to the target. Smaller values indicate straighter, more direct movements. (c) Time in seconds taken to reach the target. (d) Average firing rate of the entire neural ensemble as a function of time. With permission from Wahnoun et al. (2006).

this potential source of control signal might be overlooked. This could in turn limit device performance, as this extra control signal could have been co-opted to extend the number of controllable degrees of freedom of the effector. If we are to achieve control of a device that can replace the functionality of a limb, we will need to utilize every available control signal.

A more general approach

To take advantage of all of the potential control signals that can influence the neural activity, it is necessary to make no *a priori* assumptions about what those control signals might actually be. Instead of assuming that the neurons will respond to a particular type of motor imagery, it is possible to use operant conditioning to allow the subject to discover the underlying control signal by trial and error. In the late 1960s and early 1970s, work by Dr. Eberhard Fetz and colleagues demonstrated a technique by which one could train monkeys to gain volitional control over the firing rates of individual neurons in primary motor cortex (Fetz, 1969; Fetz and Baker, 1973). By providing feedback only about the firing rate of a target neuron, the subjects quickly (typically, within 10 min or so) learned to ramp the firing rate of that neuron up and down to achieve reward. These firing rate increases and decreases were often, though not always, associated with overt movements. Even when the firing rate of the neuron was consistently related to muscle activity, however, the activity could typically be dissociated by properly conditioning the feedback signal (Fetz and Finocchio, 1971, 1975). In a follow on to this work, they were even able to train monkeys to use single neurons to control a muscle stimulation device (Moritz et al., 2008).

In a sense, Fetz's technique could be considered as a way of identifying the volitional control signals that affect single neurons in motor cortex, without resorting to any kind of explicit motor task. In fact, this ability is not unique to motor cortex; there are now numerous examples of

subjects using operant feedback to gain volitional control over single neurons in a number of cortical areas (for review, see Fetz, 2007). Can this technique be extended to populations of neurons? We have recorded from groups of neurons while using operant conditioning to train monkeys to gain control of the firing rates of single cells. Invariably, we find that even when only one neuron is being used to provide feedback, there are typically large correlations in the firing rates of many of the other simultaneously recorded cells (Fig. 3). Clearly, these correlations indicate that the volitional control signal used to drive the target neuron also influences the firing rates of other neurons. In essence, the operant-conditioning task has uncovered a controllable *pattern* of neural activity. We have found that by performing the single-neuron operant-conditioning task with multiple target neurons, a variety of controllable patterns can be uncovered.

Of course, the patterns of activity uncovered in these operant-conditioning tasks may themselves represent collections of correlated volitional control signals. Neurons in motor cortex tend to be tuned to multiple parameters (for review, see Scott, 2003), and so may be best driven by multiple control signals. Further, control signals used to drive two different neurons may be partially overlapping. Using dimensionality reduction techniques, like principal component analysis or independent component analysis, it may be possible to resolve the effect of individual volitional drivers on the neural population response. Recently, a new technique was developed that combined dimensionality reduction with temporal smoothing to infer the presence of latent driving signals from populations of neurons recorded during an arm movement task (Yu et al., 2009). Called Gaussian process factor analysis, the method allows one to disentangle the multiple driving signals inducing correlations across the population; a description of the technique is given in Fig. 4. These factor analysis methods have already been shown to account better for correlated noise that might otherwise degrade prosthetic decode performance (Santhanam et al., 2009).

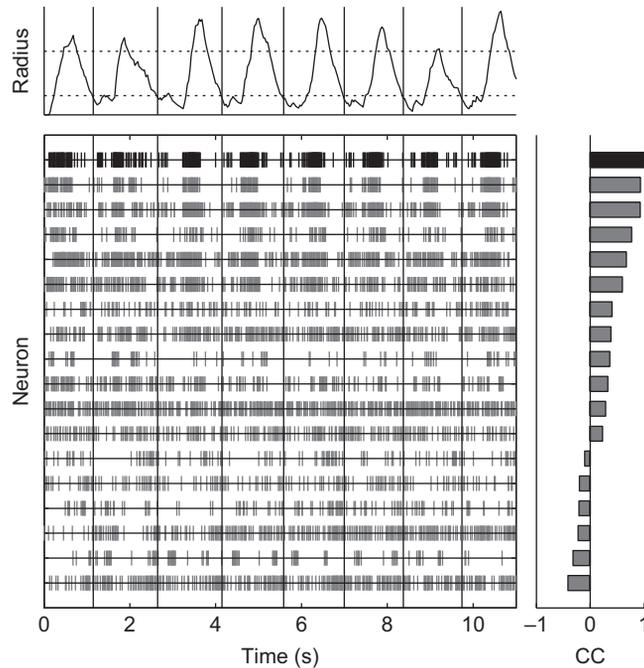


Fig. 3. Correlations in neural activity observed during an operant-conditioning task. *Top*: radius of a ring controlled by the firing rate of one neuron. The subject had to oscillate the firing rate to hit target ring sizes located at the dashed lines. Vertical lines denote successful completions. *Bottom, left*: rasters of 18 simultaneously recorded neurons during this task. The black raster is the neuron being used to control the ring. *Bottom, right*: correlation coefficient of each neuron with the neuron controlling the ring.

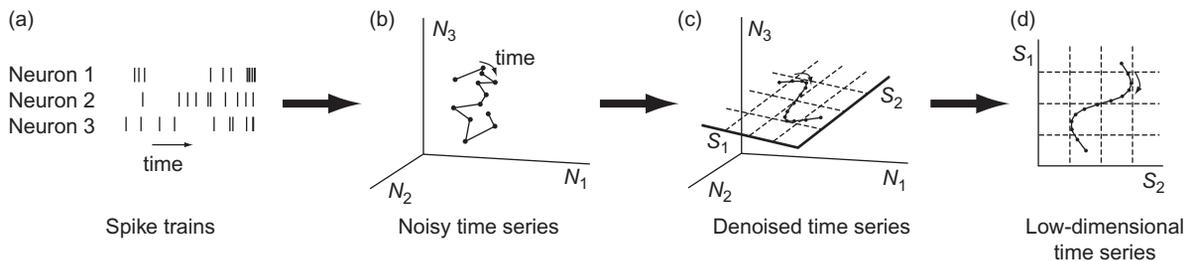


Fig. 4. Extracting a neural trajectory from multiple spike trains. For clarity, the activity of only three neurons are considered in this illustration. (a) Spike trains recorded simultaneously from three neurons. (b) The time evolution of the recorded neural activity plotted in a three-dimensional space, where each axis measures the instantaneous firing rate of a neuron (e.g., N_1 refers to neuron 1). (c) The neural trajectory (a “denoised” version of the trajectory in (b) is shown to lie within a two-dimensional space with coordinates S_1 and S_2 . (d) The neural trajectory can be directly visualized in the low-dimensional space and be referred to using its low-dimensional coordinates (S_1 , S_2). With permission from Yu et al. (2009).

One of the real benefits of these dimensionality reduction methods is that they can recover a description of the latent driving signals using only the correlations they induce in the population, without reference to any kind of external, movement-related parameters. In essence, these techniques make no assumptions about what these driving signals might actually encode. This is especially useful since it has been shown that, in motor cortex, firing rate predictions based only on the firing rates of simultaneously recorded neurons tend to outperform predictions based on external parameters, when recording from 50 or more cells (Stevenson and Kording, 2011). We have begun to apply these dimensionality reduction techniques in the operant-conditioning framework. After performing conditioning on several single neurons and observing the correlations in the recorded population, we apply the dimensionality reduction methods to recover canonical patterns of population activity. We then perform the same ring-control task described in Fig. 3, but instead of reinforcing the firing rate of a single neuron, we condition on the overlap of the population activity with a particular target pattern. Eventually, we hope to use these techniques to identify putative control signals for a prosthetic device.

By marrying operant-conditioning experimental procedures with nonparametric, correlation-based techniques for source identification, it is possible to develop a fully nonparametric description of the volitional control signals that influence a neural population. These signals have the potential to extend the capability of prosthetic devices to enable more complicated grasping and hand-shaping movements, even if we do not understand how these movements are represented during natural behavior. At the very least, they ought to be useful in identifying volitionally controllable signals inherent in the population that might not otherwise be found. They might also serve to facilitate learning. Current methods of prosthetic decoding try to identify patterns neural activity that covary together while the subjects are attempting to learn how

to control the various degrees-of-freedom of the prosthetic device. With these new techniques, it may be possible to separate this into two sequential processes: first identify the patterns of activity that naturally covary, *then* apply them to decoding. For example, after the patterns are identified, they could be applied to decoding by actuating each degree-of-freedom according to the projection of the neural activity along each pattern. The subject must then solve the easier problem of associating each controllable pattern with a particular device function.

It will be interesting to see how BCI control based on volitional signals identified through operant conditioning will compare with BCI control based on standard, parametric encoding models. Ultimately, we hope that synergizing these approaches will allow us both to develop better prosthetic devices, and gain insight into the cognitive mechanisms of volition.

References

- Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., et al. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, *1*, E42.
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. A. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, *2*, 664–670.
- Chase, S. M., Schwartz, A. B., & Kass, R. E. (2009). Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural Networks*, *22*, 1203–1213.
- Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science*, *163*, 955–958.
- Fetz, E. E. (2007). Volitional control of neural activity: Implications for brain-computer interfaces. *Journal of Physiology*, *579*, 571–579.
- Fetz, E. E., & Baker, M. A. (1973). Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *Journal of Neurophysiology*, *36*, 179–204.
- Fetz, E. E., & Finocchio, D. V. (1971). Operant conditioning of specific patterns of neural and muscular activity. *Science*, *174*, 431–435.

- Fetz, E. E., & Finocchio, D. V. (1975). Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. *Experimental Brain Research*, *23*, 217–240.
- Fraser, G. W., Chase, S. M., Whitford, A., & Schwartz, A. B. (2009). Control of a brain-computer interface without spike sorting. *Journal of Neural Engineering*, *6*, 055004.
- Ganguly, K., & Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*, *7*, e1000153.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*, *2*, 1527–1537.
- Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, *8*, 2928–2937.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.
- Hendrix, C. M., Mason, C. R., & Ebner, T. J. (2009). Signaling of grasp dimension and grasp force in dorsal premotor cortex and primary motor cortex neurons during reach to grasp in the monkey. *Journal of Neurophysiology*, *102*, 132–145.
- Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., et al. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, *442*, 164–171.
- Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., & Schwartz, A. B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 19486–19491.
- Kennedy, P. R., Bakay, R. A., Moore, M. M., Adams, K., & Goldwithe, J. (2000). Direct control of a computer from the human central nervous system. *IEEE Transactions on Rehabilitation Engineering*, *8*, 198–202.
- Kettner, R. E., Schwartz, A. B., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. III. Positional gradients and population coding of movement direction from various movement origins. *Journal of Neuroscience*, *8*, 2938–2947.
- Lemon, R. N. (1993). The GL Brown Prize Lecture. Cortical control of the primate hand. *Experimental Physiology*, *78*, 263–301.
- Miller, K. J., Schalk, G., Fetz, E. E., den Nijs, M., Ojemann, J. G., & Rao, R. P. (2010). Cortical activity during motor execution, motor imagery, and imagery-based online feedback. *Proceedings of the National Academy of Sciences of the United States of America*, *107*, 4430–4435.
- Moritz, C. T., Perlmutter, S. I., & Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, *456*, 639–642.
- Mulliken, G. H., Musallam, S., & Andersen, R. A. (2008). Decoding trajectories from posterior parietal cortex ensembles. *The Journal of Neuroscience*, *28*, 12913–12926.
- Musallam, S., Corneil, B. D., Greger, B., Scherberger, H., & Andersen, R. A. (2004). Cognitive control signals for neural prosthetics. *Science*, *305*, 258–262.
- Pohlmeyer, E. A., Oby, E. R., Perreault, E. J., Solla, S. A., Kilgore, K. L., Kirsch, R. F., et al. (2009). Toward the restoration of hand use to a paralyzed monkey: Brain-controlled functional electrical stimulation of forearm muscles. *PLoS ONE*, *4*, e5924.
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, *442*, 195–198.
- Santhanam, G., Yu, B. M., Gilja, V., Ryu, S. I., Afshar, A., Sahani, M., et al. (2009). Factor-analysis methods for higher-performance neural prostheses. *Journal of Neurophysiology*, *102*, 1315–1330.
- Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience*, *8*, 2913–2927.
- Scott, S. H. (2003). The role of primary motor cortex in goal-directed movements: Insights from neurophysiological studies on non-human primates. *Current Opinion in Neurobiology*, *13*, 671–677.
- Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., & Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, *416*, 141–142.
- Stevenson, I. H., & Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience*, *14*, 139–142.
- Suminski, A. J., Tkach, D. C., Fagg, A. H., & Hatsopoulos, N. G. (2010). Incorporating feedback from multiple sensory modalities enhances brain-machine interface control. *The Journal of Neuroscience*, *30*, 16777–16787.
- Taylor, D. M., Tillery, S. I., & Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, *296*, 1829–1832.
- Tkach, D., Reimer, J., & Hatsopoulos, N. G. (2007). Congruent activity during action and action observation in motor cortex. *The Journal of Neuroscience*, *27*, 13241–13250.
- Vargas-Irwin, C. E., Shakhnarovich, G., Yadollahpour, P., Mislow, J. M., Black, M. J., & Donoghue, J. P. (2010). Decoding complete reach and grasp actions from local primary motor cortex populations. *The Journal of Neuroscience*, *30*, 9659–9669.

- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, *453*, 1098–1101.
- Wahnoun, R., He, J., & Helms Tillery, S. I. (2006). Selection and parameterization of cortical neurons for neuroprosthetic control. *Journal of Neural Engineering*, *3*, 162–171.
- Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., et al. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, *408*, 361–365.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*, 614–635.