Present and Future Approaches to Cortical Motor Prosthetics

Alex Storer

May 2, 2006

Introduction

As science and technology continue their endless progression forward, one important goal involves integrative rehabilitation of lost motor function,particularly loss of use of a limb due to spinal cord injury or amputation. Wheelchairs may be controlled with convoluted mechanisms such as breath control, while most artificial limbs are afforded almost no control whatsoever. Advances in computational power now allow the transcription of neural signals into clear control signals in real time, which can allow the direct manipulation of the aforementioned devices. Building a robust system to control a wheelchair or an artificial limb (or even a veridical limb through functional electric stimulation) comes itself with a multitude of challenges, at many levels of neurobiology, manufacturing, electrical engineering and computer science, nearly all of which are outside the scope of this paper. Rather, I will focus on the basis of system design, as opposed to the "implementation details". Even here there is a marked lack of consensus within the field as to the optimal approaches for system development, and in this paper I will offer a general set of principles that guide the success of a neural prosthetic device, explain and critique past attempts at prostheses and finally offer my own set of suggestions as to the best way to proceed.

Guiding Principles of Neural Prostheses

In spite of the lack of a unified approach to neural prosthetic devices, the principles which shape their development are extremely intuitive and nearly tautological. For an eloquent and thorough review of these factors there are a large number of book chapters and review articles, but the work of Paul Cheney in particular comes to mind, and sets the stage for much of the discussion herein [4]. Similarly, Stephen I. Helms Tillery's work is slightly more recent and focuses on the most pragmatic elements of neural prostheses in terms of basic design principles, and this work is likewise indebted to such concise summaries [25].

Any neural prosthesis that is developed for use in rehabilitative situations must be **robust**. That is to say that its behavior must be consistent over a long period of time, and across a wide-array of circumstances. To give an example of why this is a crucial property, consider the use of EEG-based brain-computer interfacing systems (for a review of such systems, see [29]). EEG is a non-invasive technology that uses small numbers of electrodes (e.g., 64) on the surface of the scalp to measure brain activity. While many EEG signals are themselves readily reproducible in a laboratory setting, even proponents of such systems are forced to accept that performing tasks such as having a conversation must be avoided to avoid disturbing system performance, as must any type of movement or electromagnetic noise (generated from, for example, the blinking of eyes). Thus, although EEG systems can easily record stereotyped signals for long periods of time, such systems are not robust outside of a laboratory setting. In terms of implantable neural arrays, many systems can last for months or perhaps even years, but very few of these systems can boast a consistent signal across various related tasks.

In addition to consistent functioning across varied situations for long periods of time, these systems must be, simply stated, **useful**. Many systems are developed primarily in the laboratory for research purposes, and require a strict set of circumstances to function appropriately. Consider, for example, systems in which an artificial limb is controlled neurally and its processing depends on timing with respect to the presentation of a light. Even if a system such as this one achieved 100% accuracy, which no systems have established, it would still suffer from the central issue of not being useful. In a natural environment, no signal will be presented by which temporal calibration may transpire, and systems must be designed with this natural interaction built into place. Even systems that ostensibly acknowledge this goal often do not provide a mechanism by which the underlying signals can be analyzed context-free and in real-time.

Finally, all systems should strive to be **minimally invasive**, even if they are useful and robust. While this certainly applies to biological concerns, there are sociological concerns as well - a successful system will not draw attention to itself. While this may seem to be a trivial concern for neuroscientists, patients are not interested in drawing attention to their disability if it can be at all avoided. The biological concerns, on the other hand, are more straightforward. The ideal system does not dot disturb brain function whenever possible, which functionally means the use of fewer electrodes in the brain. When electrodes are inserted into the brain, they can create a significant amount of damage, depending on the manufacturing of the electrode and skill of the surgeon. Many techniques rely on the brute force use of many electrodes to extract a signal, and indeed, these systems are almost certainly more successful in the aforementioned metrics, but until electrode technology increases quite dramatically, it will be important to minimize the number of necessary electrodes. On a related note, surgery should be performed as infrequently as possible, which is a consideration in terms of interfacing electrodes with the outside world. Including a spike-sorter under the scalp could improve the efficiency of processing, but it must be weighed against the potential for infection and need for replacement or repair. When electrodes are inserted into the brain, they cause damage throughout their path, which makes areas on the surface of the brain much more attractive to developers of neuroprosthetic systems [13]. Thus are eliminated many useful areas of the brain, such as the basal ganglia and cerebellum, which may cause a formidable challenge in meeting the goals stated above.

While many systems have varied strengths and weaknesses, all will ultimately be judged on these parameters, which will guide their use by people who need them. Developing these devices is indeed a challenging task, and a number of different approaches have been taken to best maximize the available resources. Herein I will present some developments in the field and propose new mechanisms for their improvement.

Present Research in Neural Prosthetics

Many distinct systems have been developed to take volitional neural signals and translate them into control signals for external devices, such as wheelchairs or prosthetic arms. It is, once again, important to seek out commonalities in these systems and discuss them in the context of these commonalities. Systems are made up of several basic components - neural recordings, a decoding algorithm and a task which the system itself performs.

Population Vectors, Issues Affecting Them and Related Computational Approaches

All such systems owe a great debt to the work of Georgopoulos, who in the 1980s pioneered a paradigm that would prove to be exceptionally popular in neuroprosthetic research [6],[7]. He used what is known as a center-out task, in which monkeys are trained to keep their hand centered on a 2-dimensional plane (e.g., a table while holding a "lightweight frictionless manipulandum". A light is then illuminated at one of eight equidistant and equally spaced locations, and the monkey moves the manipulandum to this light. Meanwhile, recordings were made in motor cortex (using a single microwire electrode) and Georgopoulos observed that the neural recording alone were sufficient to deduce the movement the monkey made. His decoding technique is known as the population vector and is a popular starting point for many other decoding algorithms. In this approach every single unit is fitted by a cosine tuning curve, such that at a given angle (recall that angles are clearly defined in a 2D plane), the cell is expected to maximally discharge. This discharge can be thought of as a vector along this two-dimensional plane, whose direction is given by this maximally tuned location (θ) and whose length is given by the cell-firing rate. Thus, the movement itself may be reconstructed through vector addition of all tuned cells.

The advantages to this approach are immediately obvious and quite enticing - within this approach it seems possible to ignore a good many of very difficult questions that are entrenched in the study of motor control. Consider, for example, the question of frames of reference. Cortical motor neurons may code in any number of different reference frames - extrapersonal space, joint space, or it could even depend on the task at hand, but this population vector approach seems to work relatively well without considering this approach. Other research (described in [4]) has focused on specific pathways from single neurons to muscles. Many modeling studies have considered this as well, and disagree on the mechanisms by which the muscles, spinal cord, cortex and subcortical areas interact [26], [3]. By keeping their mechanisms as simple as possible, Georgopolous and colleagues are able to skirt these issues to create a surprisingly effective way to use neural recordings to predict (and later control) arm movement.

Problems for Population Vectors

Of course, with any system so simple, there are myriad difficulties which hamper its performance. To begin with, a large number of neurons are needed to give an accurate degree of control to the system, which would force any system using this approach to be increasingly invasive. Additionally, the population vector approach requires that the system choose a frame of reference in which to operate. That is, although large number of cells may tune for movement direction in extrapersonal space, that is not necessarily true on an individual cell level. This could, in fact, be one reason for large numbers of cells to be necessary to control this system. While most research chooses an extrapersonal frame of reference, some research has suggested that joint space may be a better candidate [1]. Another issue is that cells must be uniformly distributed across preferred directions, which cannot be guaranteed aside from recording from as many cells as possible [7].

Another issue inherent in this approach to neural prosthetics is the use of a small number of "actual movement" directions in training. While the system can be trivially expanded to three dimensions by adding another term in the linear best-fit equation, it is unclear the role such a small number of training movements plays on the overall effectiveness of the system. One reason for this is that researchers often both train and test on the same small number of target locations. Additionally, lights are present which by default give a certain window in time for systems to move the effector in question. That is to say, this approach does not give any mechanism for a "start" or "begin" signal, and given the small number of neurons, could result in choreatic movements at all times. It is important to note that while this problem is certainly present in the center-out task, it is a general issue that must be addressed in the design of neural decoding.

A large number of very interesting results from the laboratory of John Kalaska have also cast doubt on this population vector decoding scheme. One interesting issue arises from the effect of varying the force at the end effector. This is, in fact, a crucial issue for many neuroprostheses, such as, for example, an artificial arm. If the application of force disturbs the population vector, then it will be impossible to move a force. While this may seem trivially clear, it is crucial to be able to manipulate objects in the world. No amputee would have brain surgery to obtain an ineffectual robotic arm. Much of Kalaska's research has focused on how variable application of forces can in fact change the direction of the population vector dramatically [18, 17]. This is in stark contrast to the expectations of other researchers that movements of greater force will simply scale the magnitude of the population vector [4]. Of course, not all cells are modulated by forces, and in particular, there is a rostral-caudal gradient of force modulation in motor cortex. Most researchers record from cells that are thought to be unaffected by force, but it is certainly a consideration to be taken into account when developing a neural prosthesis. Notably, it does not seem to be the case that any neuroprosthesis labs have tested using prosthetic devices with variable forces, and it is thus unclear as to what the overall effect may be from these studies.

Another downside of the population vector method is that all it learns about a cell is its preferred direction - there is no temporal coding accounted for in the algorithm [15]. This approach is thusly problematic because there are distinct cell types with diverse properties.

For example, in Kalaska's earlier work (discussed in [3]), a number of canonical cell types are covered, including phasic, tonic and phasic-tonic cells. These cell types experience differential responses to loads, for example, and have very different postulated roles in motor circuitry. Specifically, phasic cells, which are active for short bursts of time, can be thought of as overcoming the inertia of beginning a movement. Cells such as these may give responses that are, in the eyes of the population vector algorithm, poorly tuned, but may in actuality be very well tuned for smaller windows of analysis. Clearly, excluding these cells from processing is a non-optimal approach to a neuroprosthetic device, both because more cells will need to be recorded from to obtain an appropriately large sample of neurons, but also because information is lost. Phasic cells, for example, could be used to indicate the onset of a movement, which is not a necessary issue in a laboratory setting, but a crucial concern in a real-world device.

Other work from Kalaska's laboratory has also shown the inflexibility of the population vector approach, as described above. In particular, a change as simple as a new arm posture can have a dramatic impact on the direction of the population vector [16]. This certainly casts doubt on the implicit notion in the population vector approach that movement is encoded in extrapersonal space. Rather, motor cortical cells may encode location in a different coordinate system altogether, and, frighteningly, there may not be a reason to assume a consistent coordinate frame across all tasks, all cells or all postures. From the standpoint of the population vector approach, this means that neural control of a robotic arm could be quite a disaster. If training took place over different postures, then no useful preferred direction could be learned at all, for example. Similarly, if trained on the same posture, shifting posture could result in unexpected movements for the system user.

Related Approaches

While these problems are certainly downfalls to the population vector approach, many other approaches are quite similar and suffer from the same problems discussed above. Consider, for example, the research done by Hatsopolous and colleagues in 2004 - more than 20 years after the initial research of Georgopoulos [8]. Even so, the task utilized a standard center-out procedure, which has the same problems as above. Similarly, although they do not use the specific population vector approach, the authors focus on a discrete-time filter as well as a backpropagation-based neural network to decode the neural information into a preferred direction. Once more, this uses no knowledge of different cell types, and assumes an extra-spatial reference frame in which movements occur. And finally, they test their data on the same tasks on which they gathered their data, providing no form of cross-validation.

Another mechanism that has received attention are statistical classifiers, which are discussed in several different publications [15, 13, 22]. While these classifiers are significantly more complex than the prior approaches, they likewise suffer from the same difficulties. Some statistical approaches, such as maximum likelihood estimation, rely on choosing a distribution for firing rates a priori, and generally assume that neural firing is independent over time. Thus, different cells may be fit to their respective preferred directions, but once more, these preferred directions will change with the addition of forces or changing posture.

The statistical encoding model developed by Shoham and colleagues is a particularly complex model that merits closer attention [22]. Whereas the past approaches ventured to find only the preferred direction, this particular model uses a Bayesian system to approximate all of the relevant kinematic parameters from the neural information.

$$\vec{x} = \{x, y, z, \dot{x}, \dot{y}, \dot{z}, \ddot{x}, \ddot{y}, \ddot{z}\} \quad (1)$$

$$P(N^i_{t-\tau}) = \frac{P(\vec{x}|N^i_{t-\tau} = 1)P(N^i_{t-\tau})}{P(\vec{x_t})} \quad (2)$$

In which the quantity N^i_t is the number of spikes in neuron i at time t. Empirical estimates for $P(\vec{x})$ can be obtained by simply tracking all movements, and estimates for $P(N^i_t|\vec{x})$ can be computed from cell firing histograms taken at each value of $\vec{x}$. And finally, the term $P(\vec{x_t})$ can be computed in terms of the neuron's instantaneous firing rate, which can be approximated at a given lat τ. What this system affords, then, is a way to estimate the probability of any given set of kinematic properties given the firing rate. This is an improvement to a population vector algorithm because it gives a complete set of locations and their derivatives as opposed to a simple direction. Phasic cells, could, for example, have a strong predictive power in acceleration but not in other kinematic properties. Moreover, this may be generalized across many locations because the complete probability distribution is known and is simply scaled by the firing rate of the cell. Thus, this mechanism is (thusfar) very, very similar to that of the population vector, in that it learns a tuning curve for a single neuron, albeit in more dimensions than simply velocity.

As far as weaknesses are concerned, this approach also has many. For example, $\vec{x}$ is still taken to be hand position, although this could be replaced with other coordinate systems without loss of generality. Another is a lack of simplicity in combining probability distributions of $\vec{x}$. In the population vector approach, each neuron is modeled as a vector with a given angle and amplitude, whereas here there is an order of magnitude more information present. Thus, probabilistic methods must be undertaken to generate a single response $\vec{x}$, but such methods are guided both by the connection between movements and cell firing, but also between past and future movements. That is to say, movement parameters change in a fairly stereotyped fashion, and are modeled by an auto-regressive process. Further details are provided in [22], but this complex approach still suffers from many of the same problems as the population vector approach. The issues of posture-dependent modulation of firing discussed in [16] will likely not be an issue in an approach such as this one. The rationale for this is that all of the information regarding posture is included in $\vec{x}$, and the firing rates are given with respect to all position and movement properties. For this to be true, however, once must choose the position and movement properties wisely - if only the end effector is recorded, as is previously discussed, then no posture information will be given (as the hand will be located in an identical location), and cells will not have consistent firing properties for the same $\vec{x}$. If, on the other hand, the movement vector consists of the positions of other relevant body parts as well (e.g., elbow), then it would be simple to extend this model to such a system. Of course, each additional parameter adds time in predicting movement, but the ability to include posture and end effector location simultaneously is certainly present in the mathematical formalism described herein.

Regarding the force considerations discussed earlier, however, this particular model will perhaps not be as successful. Responding to an external force need not have any effect whatsoever on $\vec{x}$ as discussed earlier, but it will have an effect on the firing rate of many cells. Within the framework of this paper, it may be possible to augment, once more, $\vec{x}$ with an external force, but this is not nearly so elegant a solution as it was in the case of posture. External forces are most likely not well modeled as an autoregressive process. Consider the process by which one picks up or drops an object - here the force is binary, a step function, which will not be modeled well by a stochastic process. Moreover, consider the generation of $\vec{x}_{t+\tau}$, the next state. Here, the force and acceleration terms are clearly connected, and may interact with each other in complex ways that may not be well covered by a sampling of the probability distribution.

Real-time Approaches to Continuous Control

Many systems have been developed to control extraneural factors in real-time, that is to say, continuous control. Examples of this include moving a robotic arm or moving a cursor on a screen, as opposed to making a selection from a menu, which would be classified as discrete control. Several of these systems have been quite successful, notably those developed in the laboratories of Schwartz, Nicolelis and Donoghue [23, 24, 14, 9, 12, 19]. While these three labs take similar approaches to neural prostheses, they also exhibited marked differences which distinguish them from each other. Let us begin by discussing the mechanisms of neural recordings.

Donoghue and Schwartz both make a point of recording from areas that are primarily associated with motor cortex, in monkeys, traditionally, and more recently, in humans [28]. This is in line with the tradition championed by Georgopoulos in the early 1980s, correlating motor cortical activity with movement. Donoghue reports that electrodes are implanted directly in M1, while Schwartz uses both Premotor (PMv) and Motor cortices [19, 14]. Nicolelis, on the other hand, uses a radically different approach, which is to record from a wide array of areas on the brain [9]. In particular, these recordings have originated in M1, PMd, SMA, S1 and Posterior Parietal cortices, bilaterally, which requires an order of magnitude more surgery, and thus may be less useful for use outside the laboratory on that criterion alone. Moreover, cells in non-primary motor areas may be active for a wide array of other tasks as well, which could generate errant actions from a neuroprosthetic device.

As far as control is concerned, Schwartz focuses primarily on control of a robotic arm in three dimensions. This affords a large number of very realistic and testable issues, for example, the effect of force and posture, and it provides another dimension of complexity to control. Such complexities push the boundaries of what neurons can control, and also afford an important interface between robotics and neural signals. In particular, a system of such complexity demonstrates the need to optimize wherever possible, and is a more appropriate model of systems that will eventually come to fruition in reality. In contrast, both Nicolelis and Donoghue have focused on 2D control of either manipulanda or cursors. While these systems are much more simple, they have the potential to be much more reliable and robust. It is, however, important to develop systems which are both well within the boundaries of what is simple and straightforward as well as systems that are not.

The neural decoding algorithms used by these laboratories are also quite distinct. Nicolelis and Schwartz both use a population vector approach to assign preferred directions to neurons, which are then scaled by their overall activity and summed to form a trajectory. Schwartz, however, discusses in detail his alterations to the decoding algorithm, which include a sign-reversal at the peak of the cosine-tuning function, for example. In addition, cells are retuned on a daily basis to achieve a maximal fit, even using an open-loop system (in which there is no visual feedback of the decoder's success). In contrast, Nicolelis uses a closed-loop system to allow the neurons being recorded to learn the appropriate mapping to best control the device at hand. He makes the very interesting point that it may be more important to allow the system to adjust itself rather than working slavishly to fine tune the relevant details. Also interesting, however, is the relative simplicity of his task, and the large number of recorded units used in his tasks (up to 177 units, Schwartz's lab used as few as 30 units). Donoghue, on the other hand, uses a statistical processing approach described in great detail in [22]. None of these approaches speculate as to the effect of changing posture or manipulating force constraints, although the power of the neural circuitry to adapt may make the entire discussion of such concerns a moot point.

Real-time Approaches to Discrete Control

An alternative to continuous control that many find appealing is the notion of discrete control, that is, control of signals with a small number of outputs. While this may seem like a useless category of neural prostheses, they may in fact perform as well as the above systems when evaluated with the criteria put forth in the introduction. That is to say, **robust**, **useful** and **minimally-invasive**. Because systems such as these need control fewer outputs, it is reasonable to a cleverly designed system to work even after signals have degraded noticeably and across many diverse situations. As far as whether they are useful, many systems can be parceled into discrete segments, such as the control of a motorized wheelchair. Precise control of a few cardinal directions could afford a fair amount of control to such a system. Similarly, menu-driven systems are common across rehabilitative technology, and they are well-suited to discrete control, as opposed to moving a shakily controlled cursor to a selection. And finally, because these systems are not as dependent on complex signals for input, fewer electrodes could be used as input for such a system. This could prove less invasive, as fewer cells would be damaged in electrode insertion, and more expensive, specialized electrodes could be used. A small set of expensive but less invasive and more robust neurotrophic electrodes could be used in place of large microelectrode arrays, and still produce a sufficient signal [11].

One interesting system uses recordings from motor cortex to build a SVM to classify between the choices "left" and "right" [10]. A support vector machine algorithm uses a set of training points to build a decision surface that is maximally distant from the training points. This algorithm runs quickly once it has trained, and although quite complicated, is readily available and implementable. The study itself utilized small numbers of implanted electrodes - 8 specifically, which is an order of magnitude smaller than electrodes used in, for example, the work of Nicolelis. Moreover, training to nearly 70% correct took place in a single session, and there is evidence that performance could approach perfection with a very small amount of training, on the order of a week, perhaps. Even more notably, these studies took place in rats - humans would likely train more quickly and perform better, even in more complicated tasks.

While the goals of discrete control systems are not inspiring, it is clear how they could be useful for patients with extreme disabilities. One large issue, however, revolves around the alignment of a timed window in which to perform the relevant processing. The authors do not provide a robust mechanism for establishing the beginning of a trial, and it is not obvious how trials could be initiated by a user. Systems such as these must additionally be weighed against completely non-invasive systems, such as EEG-based Brain-Computer Interfaces that allow similar capabilities, such as menu selection and even cursor movement, depending solely on the processing algorithm.

Other studies have focused on the role of the premotor cortex in providing discrete signals. In particular, in a center-out task, it appears that premotor cortex is more efficient at predicting the final location of a manipulandum than is motor cortex [8]. This, as the authors point out, provides a possible way to obtain both continuous and discrete signals simultaneously from the motor cortex. Thus, this could afford the best of both paradigms if a clever algorithm were used to either use these signals independently, or perhaps to combine them. Another recent study has examined other properties of premotor cortex that could prove useful in the development of control signals [2]. This study examines the variability of neurons as a function of motor preparedness. While their analysis transpires primarily independently of the realm of neural prosthetics, their discussions provide hope for a signal that could augment motor signals in a very meaningful way. One problem to which I have alluded is the high probability of choreatic movements without some form of temporal gating (although labs do not report this, it seems unlikely that control of an arm by 30 neurons is ever perfectly still given tonic firing rates), that is to say, no description of movement which I have described herein has placed the end points for motion and imposed stillness otherwise. The decline in variance occurs before the onset of the movement itself, and, interestingly, larger variances correspond to longer reaction times. Regardless of any underlying biological interpretation, it seems that a signal such as premotor cortical variance could be computed and used as a gate to control the output from a neuroprosthetic device. While many of the details of such a gating system may need to be explored in more detail before being attempted in a laboratory, signals that allow signaling of the beginning and ending of a movement should be utilized in motor neuroprostheses.

Premotor cortex, need not, however, encode solely discrete signals. In fact, some researchers have postulated that there is a differential representation of perceived movement and actual movement, with the former encoded in premotor cortex [14]. In this study, the authors taught monkeys to trace ovals in 3D-space, enabling them to observe their progress via a cursor updated in real-time. Meanwhile, recordings were made from both M1 and PMv, and preferred directions were calculated using an approach similar to the population vector approach. Following this, a motor illusion was introduced such that the resistance of the cursor was increased in the horizontal direction, the result of which being that monkeys appeared to be tracing ovals, but were in fact tracing circles. Upon observation of these population vectors, the premotor vector clearly showed activity of the observed shape, the oval, while the motor cortical vector showed the circle. In terms of motor neuroprostheses, this gives a control signal that could be more adaptable for interfacing with a robotic arm. In particular, if population vectors within premotor cortex are in fact closely tied to visual feedback, one would expect them to be less modulated by force, a property that is important in neuroprosthetic control signals, as has been discussed in detail earlier.

Miscellaneous Critiques

It is important to emphasize the pragmatic nature of motor cortical prostheses - one must take into account not only the functioning of biological structures in healthy animals, but more importantly how these structures are utilized in patients who are likely to use motor prosthetic devices. Motor cortex, for example, is known to be an extremely plastic region, and horizontal connections enable a great deal of variability across maps due to any number of even slight behavioral changes, let alone traumatic injury [12]. Different injuries, however, result in many different effects in motor representations in cortex, as well as distinct pathologies. Amputees, for example, generally maintain a primary motor representation of their amputated limb even years after the injury has transpired [27]. Interestingly, such subjects also experience phantom limb sensations that can be very frustrating and even painful. On the other side of the spectrum, there are subjects with spinal cord injuries. These people likewise cannot use certain parts of their body, but the damage is generally much more widespread, depending on the location of the injury itself along the spinal cord. Studies have been inconclusive as to the results of motor cortical mappings in patients with these injuries. Using fMRI and asking subjects to attempt to move limbs that are out of their control, a

BOLD signal can be acquired even in the absence of muscle activity. Some studies (e.g., [21]) have demonstrated that subjects with spinal cord injuries still have BOLD responses in the expected area of motor activation, while others (e.g., [27]) have lost nearly all such organized signals in motor cortex.

Any system being developed for implant in the motor cortex of disabled individuals must keep these issues in mind. As such, it may indeed be the case that innate plasticity in the motor cortex is able to adjust to systems that are presented in a closed-loop fashion, regardless of the present activity mappings that may be observed. Of course, it is unclear as to the full extent of adaptability in motor cortex, indeed, in many cortical areas, because so few experiments may be carried out. Even algorithms that have been painstakingly developed over the course of years may be misrepresented in terms of quality due to the exceedingly large number of variable situations. In any system, there are various electrode types, surgeons of different skills, spike sorters and, of course, subjects. Significantly altering any one of these could have a dramatic effect on system performance.

Additionally, many systems operate primarily open-loop or closed-loop, which present additional difficulties in comparison. There is no reason to assume that because a decoding algorithm is a closer fit to the underlying motor properties it will be more easily learned. In fact, there have been very, very few comparisons of the rate of learning of various decoding algorithms and their final accuracy, a point which is probably the most crucial to the success of the system on the whole. Moreover, subjects using such systems are only able to train for small amounts of time per day, which puts a damper on the ability of motor cortical neurons to learn an optimal control scheme for an external device. In fact, the problem is most likely compounded by the inability of subjects to use a given system all day, every day, which is a rather unnatural circumstance in the realm of learning to control a limb.

Another interesting point to ponder is the role of modeling literature in the development of cortical motor neuroprostheses. While many, many mathematically rigorous models have been developed to both account for psychophysical and physiological tuning curves, it remains unclear how best to apply this knowledge to the field of the neural prosthesis. To begin with, few if any models of motor cortical function address the questions that have been developed in neuroprosthesis research, that is, primarily, how is it possible to generate a robust arm movement signal from a limited sample of imperfectly recorded neurons. Many models have been developed to convert a set of ideal signals into the canonical waveforms that have been observed in recordings of motor cortex. Consider, for example, models such as [3] and [5] that produce excellent fits to psychophysical data. One input to these models comes in the form of a TPV, or “target position vector”, which the article cites as originating in parietal regions. Such signals, however, are not so easy to come by, as is evidenced by further research in this area [20]. In fact, if a signal such as a “target position vector” were easily measurable, then it would form an excellent basis for a motor prosthetic device, as a target vector could be easily converted to the appropriate coordinate space, would (ideally) have no effects from forces and were it robust, would be the only necessary signal from the brain. Of course, herein lies the distinction between feasible engineering solutions and neuroscience modeling. In the former, you are forced to work with what you can easily come by, whereas the latter allows the flexibility to postulate signals that likely exist, although perhaps not in the form that is proposed.

Although models may not provide explicit equations that can be used in modeling prostheses, they can indeed help guide the factors necessary to develop a successful device. In particular, models strive for perfection and simplicity in small parameter spaces that provide a maximal output. A succinct model of trajectory generation could take as input a minimal amount of signal from the brain and generate a trajectory. Given, such models have not been developed to directly interface with recordable signals, but this interface seems to be a natural step in the coevolution of engineering and neuroscience.

Discussion

Having examined a large number of studies describing cortical function, processing algorithms and the overall functioning of neuroprosthetic systems, some observations may be made as to fruitful directions of future research and system design. Perhaps the prevalent view amongst researchers in motor prostheses is that the most important discovery that is yet to be made concerns the use of advanced electrodes. Current electrode technology makes chronic implants for more than a few months quite difficult, and there has thusfar been no elegant solution to this problem. While this would indeed constitute a large advance in the entire field of neural prostheses, there are still areas in which systems may be optimized for functionality. The most important of these optimizations is to develop systems that are built strongly enough to exist outside of a laboratory setting. A neuroprosthetic arm could likely be used with a great deal of accuracy if only monkeys were able to practice with the arm constantly.

Of all the models and decoding algorithms that have been developed, it is quite clear that the most flexible, adaptable and successful of them all is the cortical structure itself.

Systems should be developed exclusively for closed-loop use, such that their performance can be rapidly optimized. Doing so, it would be possible to determine which algorithms and neural codes are most flexible in terms of the ability of the subject to control them. This is, after all, crucially important, especially when considering that many motor cortical cells may be recruited for other purposes in the wake of a serious injury. In addition, no systems posit an elegant interaction between different cortical areas and indeed, simple population vector approaches are quite successful in many circumstances. It is not unreasonable, however, to use signals from disparate parts of the brain to modulate activity from primary motor cortex. Variance changes in premotor cortex, for example, can give clues as to optimal control strategies for motor cortical ensembles. Other variants that seem trivially obvious but have not been considered include varying the frame of reference in which a neuron is fit, which would likely yield higher proportions of tuned neurons, and thus more accurate systems.

Clearly, as algorithms for decoding become more advanced and electrode technology becomes more safe and commonplace, cortically controlled motor prostheses will become increasingly viable systems, whose abilities may exceed our wildest dreams, allowing rehabilitation for scores of injured people across the world.

References

[1] Robert Ajemian, Daniel Bullock, and Stephen Grossberg. Kinematic coordinates in which motor cortical cells encode movement direction. *J Neurophysiol*, 84, 2000.

[2] Mark M. Churchland an Byron M. Yu, Stephen I. Ryu, Gopal Santhanam, and Krishna V. Shenoy. Neural variability in premotor cortex provides a signature of motor preparation. *Journal of Neuroscience*, 26(14), 2006.

[3] Daniel Bullock, Paul Cisek, and Stephen Grossberg. Cortical networks for control of voluntary arm movements under variable force conditions. *Cerebral Cortex*, 8, 1998.

[4] Paul D. Cheney, Jennifer Hill-Karrer, Abderraouf Belhaj-Saif, Brian J. McKiernan, Michael C. Park, and Joanne K. Marcario. Cortical motor areas and their properties: Implications for neuroprosthetics. *Progress in Brain Research*, 128(12), 2000.

[5] Stephen Grossberg Daniel Bullock and Frank H. Guenther. A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, 5(4), 1993.

[6] Apostolos P. Georgopoulos, John F. Kalaska, Roberto Caminiti, and Joe T. Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2(11), 1982.

[7] Apostolos P. Georgopoulos, Ronald E. Kettner, and Andrew B. Schwartz. 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(8), 1988.

[8] Nicholas Hatsopoulos, Jignesh Joshi, and John G. O'Leary. Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *Journal of Neurophysiology*, 92, 2004.

[9] Mikhail A. Lebedev, Jose M. Carmena, Joseph E. O'Doherty, Miriam Zacksenhouse, Craig S. Henriquez, Jose C. Principe, and Miguel A. L. Nicolelis. Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *Journal of Neuroscience*, 25(19), 2005.

[10] Byron P. Olson, Jennie Si, Jing Hu, and Jiping He. Closed-loop cortical control of direction using support vector machines. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(1), 2005.

[11] Kennedy P.R. and Bakay R.A.E. Restoration of neural output from a paralyzed patient by a direct brain connection. *NeuroReport*, 9(8), 1998.

[12] Jerome N. Sanes and John P. Donoghue. Plasticity and primary motor cortex. *Annu. Rev. Neurosci.*, 23, 2000.

[13] Andrew B. Schwartz. Cortical neural prosthetics. *Annual Reviews in Neuroscience*, 27, 2004.

[14] Andrew B. Schwartz, Daniel W. Moran, and G. Anthony Reina. Differential representation of perception and action in the frontal cortex. *Science*, 303, 2004.

[15] Andrew B. Schwartz, Dawn M. Taylor, and Stephen I. Helms Tillery. Extraction algorithms for cortical control of arm prosthetics. *Current Opinion in Neurobiology*, 11, 2001.

[16] Stephen H. Scott and John F. Kalaska. Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology*, 73(6), 1995.

[17] Lauren E. Sergio, Catherine Hamel-Paquet, and John F. Kalaska. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *Journal of Neurophysiology*, 94, 2005.

[18] Lauren E. Sergio and John F. Kalaska. Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *J. Neurophysiol.*, 80, 1998.

[19] Mijail D. Serruya, Nicholas G. Hatsopoulos, Liam Paninski, Matthew R. Fellows, and John P. Donoghue. Instant neural control of a movement signal. *Nature*, 416, 2002.

[20] Krishna V. Shenoy, Daniella Meeker, Shiyan Cao, Sohaib A. Kurshi, Bijan Pesaran, Christopher A. Buneo, Aaron P. Batista, Partha P. Mitra, Joel W. Burdick, and Richard A. Andersen. Neural prosthetic control signals from plan activity. *NeuroReport*, 14(4), 2003.

[21] Shy Shoham, Eric Halgren, Edwin M. Maynard, and Richard A. Normann. Motor-cortical activity in tetraplegics. *Nature*, 413, 2001.

[22] Shy Shoham, Liam M. Paninski, Matthew R. Fellows, Nicholas G. Hatsopoulos, John P. Donoghue, and Richard A. Normann. Statistical encoding model for a primary motor cortical brain-machine interface. *IEEE Transactions on Biomedical Engineering*, 52(7), 2005.

[23] Dawn M. Taylor, Stephen I. Helms Tillery, and Andrew B. Schwartz. Direct cortical control of 3d neuroprosthetic devices. *Science*, 296(7), 2002.

[24] S.I. Helms Tillery, D.M. Taylor, and A.B. Schwartz. Training in cortical control of neuroprosthetic devices improves signal extraction from small neuronal ensembles. *Reviews in the Neurosciences*, 14, 2003.

[25] Stephen I. Helms Tillery and Dawn M. Taylor. Signal acquisition and analysis for cortical control of neuroprosthetics. *Current Opinion in Neurobiology*, 14, 2004.

[26] Emanuel Todorov. Direct cortical control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience*, 3(4), 2000.

[27] Jessica A. Turner, Jae S. Lee, Osvaldo Martinez, Adam L. Medlin, Steven L. Schandler, and Michael J. Cohen. Somatotopy of the motor cortex after long-term spinal cord injury or amputation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 9(2), 2001.

[28] Wikipedia. Braingate — wikipedia, the free encyclopedia, 2006. [Online; accessed 1-May-2006].

[29] Jonathan R. Wolpaw, Niels Birbaumer, William J. Heetderks, Dennis J. McFarland, P. Hunter Peckham, Gerwin Schalk, Emanuel Donchin, Louis A. Quatrano, Charles J. Robinson, and Theresa M. Vaughan. Brain-computer interface technology: A review of the first international meeting. *IEEE Transactions on Rehabilitation Engineering*, 8(2), 2000.