

Cognitive Based Neural Prosthetics*

R. A. Andersen and S. Musallam

*Division of Biology
California Institute of Technology
Pasadena, CA 91125*

J. W. Burdick and J. G. Cham

*Division of Engineering and Applied Science
California Institute of Technology
Pasadena, CA 91125*

jgcham@robotics.caltech.edu

Abstract – Intense activity in neural prosthetic research has recently demonstrated the possibility of robotic interfaces that respond directly to the nervous system. The question remains of how the flow of information between the patient and the prosthetic device should be designed to provide a safe, effective system that maximizes the patient's access to the outside world. Much recent work by other investigators has focused on using decoded neural signals as low-level commands to directly control the trajectory of screen cursors or robotic end-effectors. Here we review results that show that high-level, or cognitive, signals can be decoded from planned arm movements. These results, coupled with fundamental limitations in signal recording technology, motivate an approach in which cognitive neural signals play a larger role in the neural interface. This proposed paradigm predicates that neural signals should be used to *instruct* external devices, rather than control their detailed movement. This approach will reduce the effort required of the patient and will take advantage of established and on-going robotics research in intelligent systems and human-robot interfaces.

Index Terms – Neural prosthetics, brain-machine interfaces.

I. INTRODUCTION

Recent advances in neural prosthetics research have demonstrated the possibility of computer interfaces and robot arms that interact directly with the nervous system [1]-[6]. The development of neural prosthetic systems can one day allow persons with lost motor function due to spinal cord injury, stroke or neurodegenerative disorders to regain the ability to communicate and interact with their surroundings. Such systems will arguably represent the ultimate human-robot interface, with the robotic device literally becoming an extension of the patient's conscious sense of self.

Despite recent breakthroughs, many challenges still remain [7]. The development of neural prosthetics requires advances across many disciplines, including neuroscience, engineering, neurosurgery and neural informatics. An underlying challenge is to improve technology for acquiring neural signals and maintaining signal quality for long periods of time. Long-term tracking of neural signals is difficult due to the invasive nature of prevalent recording techniques. Another challenge is to understand how the brain encodes movements such as reaches, and how to use these signals for neural prosthetic applications. The question remains of how the flow of information between the patient and the prosthetic device should be designed to maximize the safety and effectiveness of neural prosthetic systems.

Given these challenges, in this paper we review and elaborate upon a framework for neural prosthetics that

emphasizes the use of *cognitive* signals in interacting with external devices [8][15]. Most recent successes in neural prosthetics have used signals primarily from the motor and pre-motor cortices (see Figure 1) to directly control the trajectory (position and velocity) of a screen cursor or robot arm. This approach, though proven successful, has its limitations. Using neural signals to directly control a robot arm for a wide variety of movements, manipulations and postures will likely place higher demands on the number and quality of signals that need to be recorded. We believe that neural prosthetics should acquire signals to *instruct* external devices, rather than to control every detail of their movement. Recent work by some of the authors, reviewed here, shows that it is possible to capture the intention, or goal, of reaching movements, and also the expected reward, or motivation, of the subject before the reach is made. Utilizing such cognitive variables may decrease the effort required of the user and reduce the informational burden placed on the neural recording interface.

Future neural prosthetic devices will likely require an integration of both higher-level (cognitive) and lower-level (motor) brain activity. This approach can take advantage of established and on-going robotics research in intelligent and supervisory systems and human-robot interfaces in order to develop new paradigms that can adapt to the needs and cognitive states of the user while maintaining safety.




Figure 1. Signals used for neural prosthetic applications. In patients with lost motor function, microelectrodes can be used to record the neural activity (spiking or field potentials) from brain areas that encode movements such as reaching and grasping to drive external devices such as robot arms.

* This work is partially supported by National Eye Institute, DARPA, ONR, the Boswell Foundation, NSF, the Christopher Reeve Paralysis Foundation, the Sloan-Swartz Center for Theoretical Neurobiology at Caltech, and the Human Frontier Science Program.

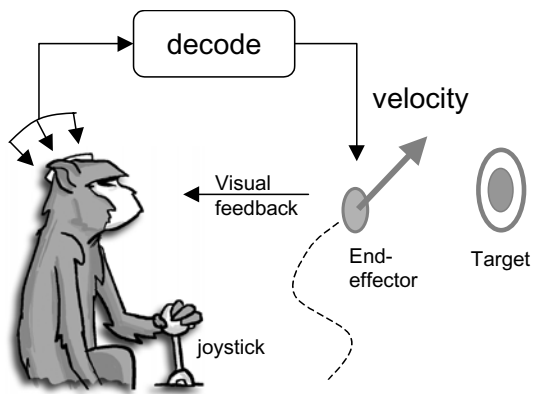


Figure 2. In current studies, non-human primates learned to control an end-effector using brain signals alone. Signals recorded during normal arm movement tasks were used to build a decode model that can predict the desired trajectory (primarily the instantaneous velocity of the end-effector).

II. RECENT ADVANCES IN NEURAL PROSTHETICS

Several investigators have demonstrated systems in which neural signals from non-human primates were used to control the motion of an end effector (a screen cursor or a robotic arm) [2][3][4]. In these studies, the activity of populations of neurons was recorded using arrays of microelectrodes implanted primarily in motor and pre-motor cortex (see Figure 1). The electrode arrays recorded the electrical or “spike” activity of individual or groups of neurons while the subjects performed arm movement tasks. This activity was then used to calibrate models, usually linear regression models, that relate the firing rate of the neurons to the observed movements. The subjects were then shown capable of moving the end-effector in “brain control” tasks in which current neural activity and the calibrated models were used to predict the end-effector trajectory without actual movement of the subject’s arm.

In addition, similar progress has been made in using EEG-based signals to derive neuroprosthetic commands from motor related areas [6]. Though EEG methods are less invasive, their poor spatial resolution may limit the amount of information that can be obtained with them.

These groups showed performance well above the chance level, and, more interestingly, reported that the performance of the subjects in the brain control task improved over time, giving evidence that the neural system learned to use the prosthetic device by adapting its output. In fact, [2] showed that the subjects learned to compensate for the dynamics of the robotic arm used. This is an encouraging result that allows a certain margin of error in the acquisition and decoding of neural signals, as the system can be relied upon to adapt to these errors.

While these studies have shown the possibility of neural-controlled prosthetic devices, and have sparked renewed interest in the field, the question arises of what other types of signals could be used in a neural interface. These studies used neuronal populations located primarily in motor and pre-motor cortex, which tend to encode the specific commands sent to our muscles for the control of limb movements. More important than the particular brain

region the signals are derived from, however, is the type of information that is decoded from them.

In the current approaches, the neural activity is used as a *control* signal to directly specify where the location of the end-effector needs to be at any instant in time. In many cases, the primary signal decoded from the neural activity was the velocity of the end-effector. Thus, the trajectories of the brain-controlled end-effector were the result of velocity control on the part of the subject, who had to continually command online corrections based only on visual feedback of the task (see Figure 2). Here the neural activity is used for low-level control commands, and the effectiveness of the neural prosthetic is limited to the subject’s ability to perform the closed-loop task.

It could be argued that motor cortex neurons should be used exclusively for interfacing to external devices, given their demonstrated ability to adapt to different tasks. In this approach, the motor cortex is used as a generic source of neural signals that can be adapted to communicate cognitive states and/or control movements of robot arms. However, there are at least two reasons not to depend solely on motor cortex areas. The first is that this dependence creates an informational bottleneck that will reduce the number of cognitive variables that can be read out at any one time. For example, a patient’s mood could be determined by asking him or her to move a cursor on a computer interface to answer sets of questions about their emotional status. However, this would preclude the patient from performing other tasks at the same time. It would be more efficient to decode this signal directly from an area that processes the mood of the subject.

The second reason is that the normal functional architecture of motor cortex is for generating commands for movement trajectories. While it may be possible for the motor cortex to be treated like an undifferentiated neural network and trained to perform any task, it has been shown that neural networks trained to do a large number of different tasks tend to do each one poorly compared to being trained to perform a small number of tasks [9].

Can higher-level, or cognitive, signals be used to interface to an assistive neural prosthetic device? In the next section, we review current work that shows that cognitive signals such as reach goals and reach motivation can be used in a neural prosthetic system.

III. COGNITIVE NEURAL SIGNALS

In theory, cognitive control signals appropriate for reaching tasks could be derived from many higher cortical areas related to sensory-motor integration in the parietal and frontal lobes. Here we focus on the posterior parietal reach region (PRR) and the dorsal premotor cortex (PMd). Again, the primary distinction is not the specific brain area where the signals are obtained, but rather the type of information that is being decoded, and how that information is used. Similar approaches to the one presented can be used for interpreting cognitive signals from other brain areas. It is likely that some areas will yield better results than others depending on the cognitive signals to be decoded and the parts of the brain that are damaged.

PRR in non-human primates lies within a broader area of cortex called the posterior parietal cortex (PPC) [10][11] (see Figure 1). The PPC is located functionally at a transition between sensory and motor areas and is involved in sensory-motor integration, that is, it helps transform sensory inputs into plans for action. PRR is known to be primarily active when a subject is preparing and executing a movement [11][13]. However, the region receives direct visual projections and vision is perhaps its primary sensory input.

Neural activity in PRR has been found to encode the targets for a reach in visual coordinates relative to the current direction of gaze (also called retinal or eye-centered coordinates) [13]. In other words, this area contains information about where in the field of view the subject is planning to reach. This coding of planning information in visual coordinates underscores the cognitive nature of the signal within PRR. It is coding the desired goal of a movement, rather than the intrinsic limb variables required to reach to the target. The human homologue of PRR has recently been identified in fMRI experiments [14].

Recent experiments with monkeys have demonstrated that reach goals decoded from PRR can be used to drive a neural prosthetic computer interface to position a cursor on a screen [15]. These experiments are described in more detail in the following section.

A. Decoding the goal of a reach

In the experiments, arrays of electrodes were placed in the medial intraparietal area (MIP), a portion of PRR, area 5 (also in the posterior parietal cortex), and the PMd. These electrodes record the electrical activity of neurons in the vicinity of the electrode tips. Each experimental session began with the monkeys performing a series of reaches to touch different locations on a computer screen (Figure 3a). As shown in the figure, the reaching task consisted of four phases. First, the monkey is instructed to fixate on the center cue. Next, a target location is presented for a brief period of time. The target then disappears followed by a delay period. Finally, the monkey is given a “go” signal instructing him to reach to the location where the target was.

The neural activity recorded during the delay period, in which the monkey plans the reach movement, was used to build a database that relates the firing rate of the neurons to the target location. For example, as illustrated in Figure 3a, the database would store which location in the visual field a particular neuron exhibited higher firing rates when movements were planned to that location.

After enough trials were performed to build an acceptable database, the monkey was switched to the “brain control” task. Here, the monkeys were instructed with a briefly flashed cue to plan to reach to different locations but without making a reach movement (Figure 3b). The activity during the delay period was then compared to that in the database and, using a Bayesian decode algorithm, the location where the monkeys were planning the reach was predicted. If the predicted reach direction corresponded to the cued location, then the animals received a drop of fluid reward and visual feedback was provided by re-illuminating the cued location. If the animals moved their arm, the trial

was cancelled and no reward was given. This approach was necessary because we cannot simply say to the monkeys “think about reaching to the target without actually reaching to it.”

Thus, the reach goals were decoded from activity present when the monkeys were planning the reach movements, but otherwise were sitting motionless in the dark and were not making eye movements. Figure 4a shows results from PRR recordings. The figure shows the cumulative accuracy of the brain control system for four target locations as the experimental session progresses. As shown in Figure 4b for a recording session made from PMd, only a small number of cells were required for successful performance of the task, with performance increasing with the number of neurons. Figure 4c shows neural activity in both the reaching and brain control tasks from a recording site in PRR, demonstrating that the cognitive signals in the brain control task were free of any sensory or motor related activity.

In addition, the animals showed considerable learning in the brain control task, as evidenced by a significant increase in their performance over the course of one to two months [15]. This behaviour is consistent with a number of studies of cortical plasticity [16], and the time scale for learning is similar to that seen in motor cortex for trajectory decoding in previous studies [2][4]. In this case, the improvement in performance was found to be related to an




Figure 3. Normal reaching and brain control tasks used in decoding reach goal and expected reward cognitive signals. Adapted from [19].

increase in the amount of information encoded by the neurons in the brain control task [15], which was calculated using a mutual information measure. This measure quantifies the degree to which a neuron's firing rate encodes a particular direction. In essence, the neurons improved their tuning to increase performance in the task. Plastic behaviour such as this will be important in enabling patients to optimize neural prosthetic systems with training.

B. Decoding expectation

It has also been shown possible to decode neural activity that foretells a subject's expectation of a reward for performing a task. Signals related to reward prediction are

found in a number of brain areas [17]. In area LIP of the PPC, which is involved in planning and executing eye movements, it was found that cells code the expected value of rewards [18]. Using an eye motion task, it was found that the neurons increased their activity when the animal expected a larger reward or the instructed eye movement was more likely to be in the cells' preferred locations. Similar effects have recently been found for PRR neurons for amount of reward in both the reaching and brain control task previously described [15]. PRR cells are also more active and better encode reach goals when the animal is cued to expect a higher probability of reward at the end of a successful trial. Remarkably, PRR cells also encode reward preference, with higher activity seen when the monkey expects delivery of a preferred citrus juice reward rather than water.

This expectation signal tells us something about the cognitive state of the subject, possibly indicating the subject's motivation or interest in the task's outcome. Moreover, this expectation of reward could be read out simultaneously with the intended reach goal using offline analysis of the brain control trials [15], thus showing that multiple cognitive variables can be read from the brain at the same time.

IV. NEUROPROSTHETIC CONTROL SYSTEMS BASED ON INTELLIGENT DEVICES AND SUPERVISORY CONTROL

A neural prosthetic system must be flexible and capable of adapting to the needs and cognitive states of the patient. It is not clear how the sole use of motor-based signals can provide enough information for such subtle and complex interaction. Given that it is possible to obtain cognitive signals, we have proposed a framework for the control of neural prosthetics that takes advantage of the high-level nature of these signals. In this approach, cognitive signals like the intended goal of a reach or the motivation or expected reward of a reach are used to *instruct* the neural prosthetic system, rather than control it directly, as illustrated in Figure 5.

In this case, interaction between the patient and the mechanical device is monitored by an intelligent supervisory system [20]. This system monitors the cognitive state of the patient, and combines it with knowledge of the workspace and patient information such as gaze direction to assess the situation and calculate the most appropriate course of action. The cognitive neural signals can operate much like "body language" by providing, on-line and in parallel with readouts of other cognitive variables, the preferences, mood, and motivational level of the patient. Implants in emotional centers could also provide real time readouts of the patient's emotional states, such as alarm, urgency or displeasure. This augmentation of the information channels derived from the patient is particularly important for locked-in patients that cannot move or speak.

Given this information, combined with readings of the patient's intended reaching goals from regions such as PRR, the intelligent system can then compute the appropriate trajectory and posture of the robotic arm. For example, given the Cartesian coordinates of an intended object for grasping and knowledge of the environment, a

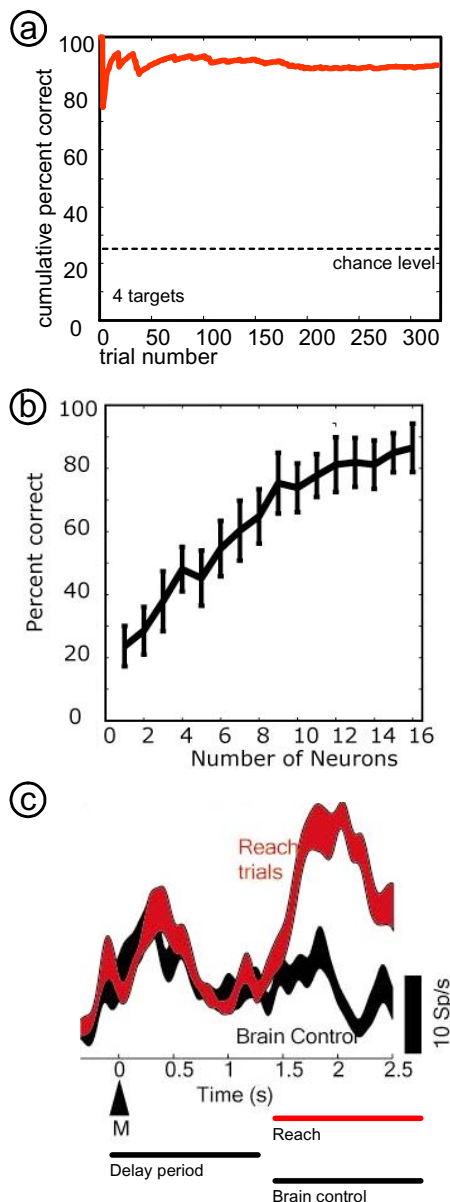


Figure 4. Results of brain control tasks in which cognitive signals (reach goals) were used to drive a cursor to target positions on a screen. a) Task performance as a function of time for a four-target brain control task. b) Performance as a function of the number of neurons used in the decode model for an eight-target task. c) Comparison of neural activity during normal reaching and brain control tasks (b and c adapted from [15]).

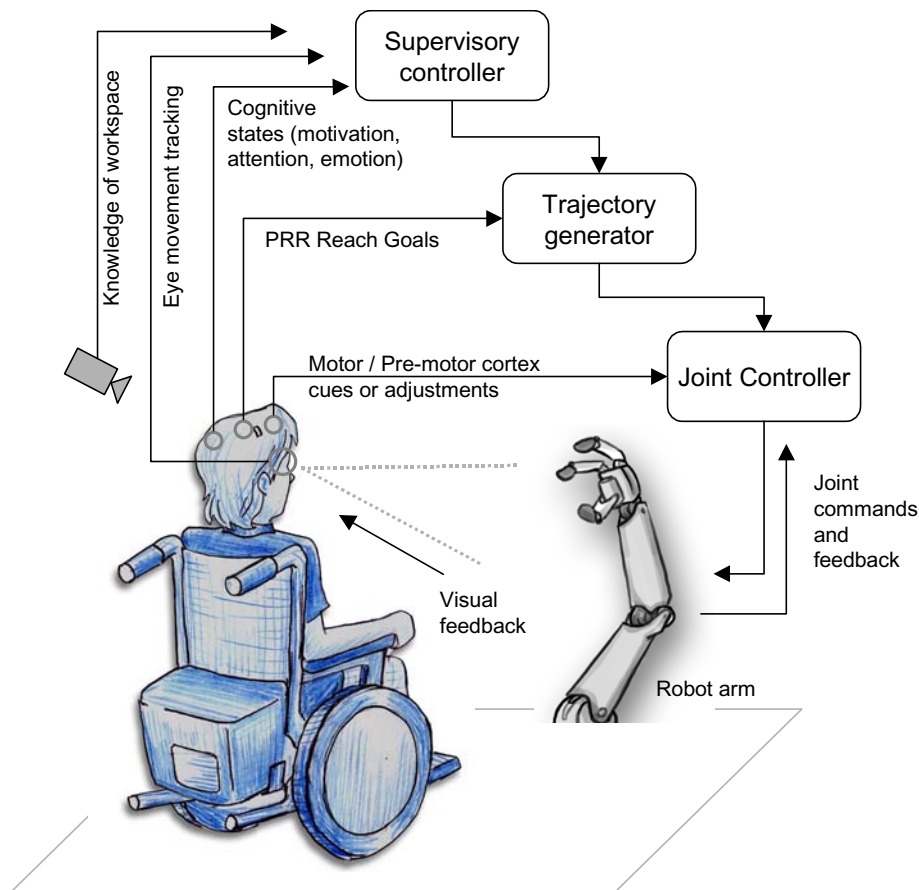


Figure 5. The proposed framework emphasizes parallel decoding of high-level cognitive variables (reach goals, motivation) and low-level motor variables under the supervision of an intelligent system, which manages the interaction between the patient and the robot arm.

robotic motion planner [21] can determine the detailed joint trajectories that will transport a prosthetic hand to the desired location in the safest, most efficient way.

This trajectory can then be sent to the robot arm's low-level joint controller for execution. Sensors embedded in the mechanical arm ensure that it follows the commanded trajectories, thereby replacing the function of proprioceptive feedback that is often lost in paralysis. Other sensors can allow the artificial arm and gripper to avoid sudden obstacles and regulate the interaction forces with its surroundings, including grasping forces, thereby replacing somatosensory feedback.

If available, motor signals can augment low level plans to help "guide" the end-effector to the appropriate action by providing cues or corrections to the trajectory planned by the supervisory controller. Thus, future applications are likely to involve recordings from many areas to read out a substantial number of cognitive and motor variables. Results from cognitive-based and motor-based approaches will likely be combined in single prosthetic systems to capitalize on the benefits of both.

IV. CONCLUSIONS AND FUTURE WORK

We have reviewed a cognitive based framework for the control of neural prosthetic systems, based on evidence that

such cognitive signals can be directly read from the nervous system. In this approach, neural signals are used to *instruct* an intelligent supervisory system, rather than directly *control* an external device such as a robot arm. The proposed supervisory system in turn manages the interaction between the user and the external device. This approach has many potential benefits for both the neuroprosthetic user and for the implementation of the prosthetic system.

The proposed framework would reduce the effort required of the user in executing tasks such as reaching, since they are not continually involved in the task of controlling the position and trajectory of the robot arm. What we are proposing is to use the intrinsic organization of the nervous system to provide multiple hierarchical channels of communication and control. Task execution will more closely resemble normal function, in which low-level control of limb movements often occurs without conscious attention. Similarly, the hierarchical nature of supervisory control should allow patients to learn much more quickly how to command a new device.

For the systems engineer, this approach has the benefit that to adapt the neural interface to different electromechanical devices (e.g. different types of robotic arms or communication devices), only the lowest level of the control hierarchy need be re-engineered for the specific mechanical device. In addition, the approach may possibly

reduce the number of signals needed to be extracted from the nervous system, since lower numbers of variables may be needed to specify cognitive states such as goals and intentions in comparison to those needed to control a wide set of movements of an articulated robot arm. This in turn can alleviate the informational burden placed on the neural recording device, making them more practical.

Finally, this approach can take advantage of the long history and on-going research in robotics that focuses on intelligent systems, autonomous navigation and path planning, and human-robot interaction. In turn, this application will present exciting new challenges for these areas of robotics research.

Future neural prosthetic devices will likely require an integration of both higher-level (cognitive) and lower-level (motor) brain activity. By using activity from several different parts of the brain and decoding a number of cognitive variables, a neural prosthetic can provide a patient with maximum access to the outside world.

ACKNOWLEDGMENT

We would like to thank the members of the Andersen lab at Caltech, especially C. Buneo, B. Pesaran, B. Corneil and B. Greger. We also thank E. Branchaud and Z. Nenadic.

REFERENCES

- [1] Wessberg, J., et al., Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 2000. 408(6810): p. 361-365.
- [2] Carmena, J.M., M.A. Lebedev, R.E. Crist, J.E. O'Doherty, D.M. Santucci, D.F. Dimitrov, P.G. Patil, C.S. Henriquez, M.A.L. Nicolelis, Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS*, 2003. 1: p. 193-208.
- [3] Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R., Donoghue, J.P., Instant neural control of a movement signal. *Nature*, 2002. 416: p. 141-142.
- [4] Taylor, D.M., Tillery, S.I.H., Schwartz, A.B., Direct cortical control of 3D neuroprosthetic devices. *Science*, 2002. 296: p. 1829-1832.
- [5] Kennedy, P.R. and R.A. Bakay, Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport*, 1998. 9(8): p. 1707-11.
- [6] Wolpaw JR, B.N., McFarland DJ, Pfurtscheller G, Vaughan TM, Brain-computer interfaces for communication and control. *Clin Neurophysiol*, 2002. 113: p. 767-791.
- [7] Donoghue, J. P., Connecting cortex to machines: recent advances in brain interfaces. *Nature Neurosci*, 2002 vol. 5, p. 1085-1088.
- [8] Andersen, R.A., Burdick, J.W., Musallam, S., Pesaran, B., and Cham, J.G. Cognitive Neural Prosthetics. *Trends in Cognitive Science* (In press).
- [9] Kosslyn, S.M., C.F. Chabris, C.J. Marsolek, and O. Koenig, Categorical versus coordinate spatial relations: computational analyses and computer simulations. *Journal of Experimental Psychology*, 1992. 18: p. 562-577.
- [10] Andersen, R.A. and C.A. Buneo, Intentional maps in posterior parietal cortex. *Ann. Review Neuroscience*, 2002. 25: p. 189-220.
- [11] Snyder, L.H., Batista, A.P., Andersen, R.A., Coding of intention in the posterior parietal cortex. *Nature*, 1997. 386(6621): p. 167-170.
- [12] Sakata, H., et al., Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey. *Cereb Cortex*, 1995. 5(5): p. 429-38.
- [13] Batista, A.P., et al., Reach plans in eye-centered coordinates. *Science*, 1999. 285: p. 257-260.
- [14] Connolly, J.D., R.A. Andersen and M.A. Goodale, fMRI evidence for a 'parietal reach region' in the human brain. *Exp Brain Res*, 2003. 153: p. 140-145.
- [15] Musallam, S., B.D. Corneil, B. Greger, H. Scherberger, and R.A. Andersen, Cognitive control signals for neural prosthetics. *Science*, 2004. 305: p. 258-262.

- [16] Buonomano, D.V., M.M. Merzenich, Cortical plasticity: from Synapses to Maps. *Ann. Review Neuroscience*, 1998. 21: p. 149-186.
- [17] Schultz, W., Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Current Opinion in Neurobiology*, 2004. 14: p. 139-147.
- [18] Platt, M.L. and P.W. Glimcher, Neural correlates of decision variables in parietal cortex. *Nature*, 1999. 400(6741): p. 233-8.
- [19] Buneo, C. A., Jarvis, M. R., Batista, A. P. and Andersen, R. A. Properties of spike train spectra in two parietal reach areas. *Exp. Brain Res* (2003) 153:134-139.
- [20] Sheridan, T.B., *Telerobotics, Automation, and Human Supervisory Control*. 1992, Cambridge, MA: MIT Press.
- [21] Murray, R.M., S. Li, and S.S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. 1994, Boca Raton, FL: CRC Press.