Modeling the CNS as the CPU for the Primate Motor Control System of Voluntary Movement: Control Schemes, State Estimation, and Input/Output Transformations

James Solberg

Final paper for NUIN 480

Course Director: Dr. James Houk

Due: June 16th, 2003

1. Preface

The enormous kinematic, information processing, and inter-neuronal complexities of the primate motor control system make it extremely difficult to model. This paper addresses some of those difficulties and then attempts to establish its own model. Parts of the model are described in detail while other parts are just briefly mentioned (mostly because space constraints). This model is schematically illustrated in figure 3, which was constructed in a separate Word file. If you are reading this paper electronically, then it would be good to keep figure 3 open in a separate window because this paper is built around that diagram and referred to quite often. This model was spliced together from the various literature discussed in this class. The specific articles used are listed below along with the week of class that the article was discussed (full citation can be found at the end of the paper).

Week 1 (Jim Houk):

- Kinematic Properties of Rapid Hand Movements in a Knob Turning Task (Novak, 2000)
- Features of Motor Performance that Drive Adaptation in Rapid Hand Movements (Novak, 2002)

Week 2 (Lee Miller):

- On the Relations Between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex (Georgopoulos, 1982)
- Coding of Finger and Wrist Movements (Georgopoulos, 1999)
- Muscle and Movement Representations in the Primary Motor Cortex (Kakei, 1999)
- Direction of Action is Represented in the Ventral Premotor Cortex. (Kakei, 2001)
- A Comparison of Movement Direction-Related Versus Load Direction-Related Activity in Primate Motor Cortex, Using a Two-Dimensional Reaching Task. (Kalaska, 1989)

Week 6 (Sandro Mussa-Ivaldi):

- Computational Nature of Human Adaptive Control During Learning of Reaching Movements in Force Fields. (Bhushan, 1999)
- Optimal Feedback Control as a Theory of Motor Coordination. (Todorov, 2002)

Week 9 (Sara Solla):

- Exploring the Neurophysiology of Decisions. (Leon, 1998)
- Motion perception: Seeing and deciding. (Shadlen, 1996)

Week 10 (Jim Houk):

- Model of Cortical-Basal Ganglionic Processing: Encoding the Serial Order of Sensory Events (Beiser, 1998) (week 10: Jim Houk)

2. Modeling a Primate Motor Control System

The International Space Station is one of the largest and most complex engineering endeavors ever attempted. In total, there are about 50,000 sensors that can generate observations every 10 seconds. This results in a massive telemetry stream. Clearly, human monitors can only watch a fraction of these signals and thus require computational aids to help maintain this complex system. Precision controlled thruster and robotic systems are used throughout the station to maintain proper operation. Feedback information from the sensors is fed into a grid of networked computers containing hundreds of millions transistors and performing 10s of MegaFLOPS (million floating-point operations per second). This enormous computing power is needed to compile the exorbitant amount of data being generated about the state of the space station. Then once there is an estimation of the state, action must be taken to obtain the desired state (or maintain current state for a regulatory controller). Every sensor, logic circuit, and actuator was individually built and then positioned in its niche in the space station with scrutinizing accuracy. Millions of lines of software code were written to ensure that every possible situation (state) could be accounted for. Hundreds of thousands of pages of blueprints and technical documentation are needed to explain how the entire station operates. Any engineer would tremble at the thought of such a monstrous project.

Now imagine a system of even greater complexity, having much more than 50,000 sensors being sampled much faster than every 10 seconds. This system has on the order of a trillion basic logic elements (analogous to transistors), and has many times more actuators than the space station. This is the system that neuroscientists look at every day. But what is even more intimidating than the numbers just presented is the fact that there is no 'owner's manual'; everything must be reverse-engineered from a system identification approach. And to make matters even worse, biological circuits and dynamics usually don't follow the laws of linear system theory that engineers have developed and perfected over the past few centuries.

Engineers have the luxury of being able to model their systems to almost arbitrary accuracy. In the space station case, given enough information about all the components and software used, the entire space station, in theory, could be completely modeled. Specifications could be extracted from the technical manuals and the dynamics could be constrained to follow some first principle PDEs (Newton's laws, Navier-Stokes, and Maxwell equations). It would be preposterous to think that a neuroscientist (or anyone else) could develop such a model of the primate central nervous systems (CNS) from first principles. The problem before the neuroscientist studying the control system is constructing a model that accurately emulates the input/output characteristics of the system being studied. Essentially, they must extract an abstraction of the physical system. A successful model of the CNS (or of its subsystems) will often demonstrate an emergent property of that system. The notion of *emergence* is an underappreciated concept that is essential for understanding complex systems. An emergent property arises out of a complex system when interactions between the individual agents of that system orchestrate a (complex) global behavior that is not intuitively deduced from the simple actions of those individual agents. Individual neurons are relatively simple agents when studied by themselves. If a million of those similar neurons are connected together, then the complexity of that system is too great to comprehend the dynamics from a cellular level.

The study of motor control is basically the study of sensorimotor transformations. For the effectors of the motor control system to move or to apply forces on objects in the world, it must process a variety of forms of sensory and motor data. This information is generally in different formats and can even refer to the same entities but in a different representation (frame of reference). By allowing motor and sensory data to be related, transformations between these various representations can close the sensorimotor loop. From a computational perspective the motor system can be considered as a system whose inputs are the motor commands generated from the controller within the CNS. In order to determine the behavior of the system in response to this input an additional set of variables, called state variables, must also be known. Taken together, the inputs and the state variables are sufficient to determine the future behavior of the system. It is unrealistic, however, to assume that the controller in the CNS has direct access to the state of the system that it is controlling; rather, we generally assume that the controller has access to a sensory feedback signal that is a function of the state. This signal is treated as the output of the abstract computational system.

The purpose of the rest of this paper is to compile the papers and discussions of this class in a way that presents some coherent story about the dynamics of the primate motor control system. Using the space station as an analogy, this paper attempts to determine the software code of the controller. What scheme does the CNS use to control its motor system? How is the CNS able estimate its own state? What are the basic mechanisms of generating a command signal? This paper will address each of these questions.

3. Rational Decisions

A model is responsible for constructing a relationship between the inputs and the outputs of a system. In its simplest form the schematic of a model would look like figure 1:

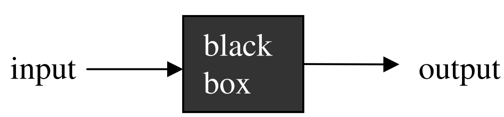

Figure 1. *schematic of simple model*

But before we can get into any of the details of the ‘black box’, we must first establish that it is even possible to model the primate CNS. For any model to be valid it must behave in a predictable manner. This means that for a given state of the system and a given set of inputs, the ‘black box’ must produce a consistent output. So if we are studying the CNS as a controller, then the CNS must be able exhibit logic that is repeatable. In other words, we must demonstrate rational (consistent) decisions.

In week 9 of the class Dr. Sara Solla presented a lecture that seemed to point to the conclusion that primates do indeed use rational decisions when controlling the motor system. Using the model from figure 1, the input is the sensory information, the output is the motor behavior, and the ‘black box’ is the logic of the neural circuitry that dictates how the sensory input is used to determine the motor output (sensorimotor transformation). The basic paradigm of the model for the CNS is that the brain must integrate the influx of new sensory information (input) with its prior relevant knowledge about the situation (black box), and then select the appropriate behavior (output). This nonreflexive linkage between afferent (from senses) and efferent (output command) data involves interpretation and behavioral selection. This is referred to as a decision process.

Intuitively we know that the CNS is capable of making decisions. All of our actions depend on that fact. The problem lies in uncovering the mechanism within the CNS that actually makes these decisions. Scientists have devised many experiments over the years that attempt to explain the mechanisms behind decision-making. The decision process can only be studied in alert animals that are able to perform sufficiently complex tasks. Also, there are several key elements that the task being performed must possess:

- The sensory input must require some sort of interpretation.
- There must be multiple options for the animal to choose from.
- The sensorimotor transformation (decision process) must happen within a time span of sufficient duration to permit a physiological dissection.

Monkeys performed such a task in an experiment where they had to make a decision on which direction to move their eyes for a given stimulus. In this paradigm (Shadlen, et al 1996) a monkey is trained to look forward at a fixation point on a screen as a field of moving dots is displayed. The dots individually move in random directions, but they have a net effect of moving to either the right or left. The fraction of dots that move in the net direction (right or left) is a parameter controlled in the experiment. The monkey reports the direction of coherent motion by making a saccadic eye movement to one of two visual targets, each corresponding to one of the possible directions of motion. Please refer to figure 2 on the next page for a detailed account of the experiment. Note that figure 2 (including the entire caption) was copied directly from the Shadlen paper (Shadlen, et al 1996). It is now possible to investigate the neural basis of simple visual decisions that link sensation to action in this type of experiment because of the extensive knowledge of the primate visual and oculomotor systems.

Neurons in the lateral intraparietal region (LIP) were recorded from in order to explore the link between sensation and action. These particular neurons fire in high levels during the initial

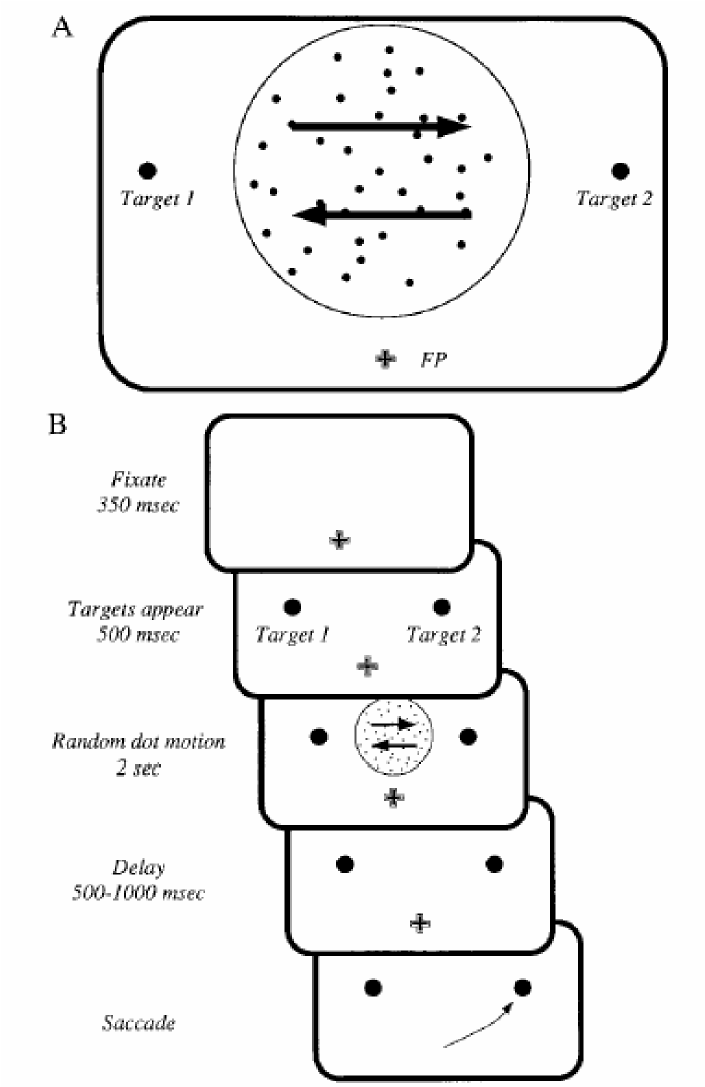

Figure 2 The psychophysical task. Two rhesus monkeys performed a single-interval, two-alternative, forced choice discrimination of motion direction. (*A*) The monkey judged the direction of motion of a dynamic random dot stimulus that appeared within an aperture 4–8° in diameter. In this example, the monkey made a saccadic eye movement to target 1 (T1) if leftward motion was detected; conversely, the monkey made a saccade to target 2 (T2) if rightward motion was detected. Each experiment included several stimulus conditions—two directions of motion for several nonzero coherences, plus the zero coherence condition, which does not contain a coherent direction of motion. All stimulus conditions were presented in random order until a specified number of repetitions was acquired for each condition (typically 15). The experiment was designed so that T1 fell within the movement field of the LIP neuron; T2 and the motion stimulus were placed outside the neuron's movement field. (*B*) The sequence of events in a discrimination trial; see text for details. Throughout the trial, the monkey maintained its gaze within a 1–2° window centered on the fixation point (FP). Failure to do so resulted in abortion of the trial and a brief time-out period. Eye movements were measured continuously at high resolution by the scleral search coil technique (19), enabling us to enforce fixation requirements and detect the monkey's choices. The monkey received a liquid reward for each correct choice.

states of planning a saccade. This demonstrates that they are likely a part of the decision process. Once a neuron was found that was active during the delay period of a remembered saccade task, one of the targets (called T1) was defined in that neuron's movement field. T2 was placed well outside the movement field (often in the opposite hemifield).

The results show that neural firing patterns in the LIP carry signals that can predict the motor action (saccade) that the monkey will make. These signals typically appear early in the trial while the dots are being displayed and are sustained during the delay period after the stimulus has disappeared. The authors thus conclude that the evolution of predictive signals in LIP comprises a neural correlate of decision formation within the CNS.

Is it possible that these neurons in the LIP fired as a result of the sensory input or the motor command output? A sensory account of the predictive activity can be ruled out quickly by examining the responses at 0% coherence (random noise). In these trials the monkeys are not given any relevant sensory data that would allow them to make the 'correct' decision (because there is, in fact, no correct decision). Yet, the firing rate of the neuron in the LIP can still accurately predict which way the monkey will move its eyes. The authors ruled out the firing of the neuron as a result of a motor command by showing that the firing increases as the coherence increases. A motor signal commanding a saccade should only depend on the metrics of the planned movement, not on the strength of the sensory signal that evoked the decision to move.

The evidence just presented points to a conclusion that the primate CNS controller does make rational decisions that are based on its prior knowledge. Too often we observe humans making decisions that seem irrational. But, it is not the fundamentals of the decision process (neural circuit) that are at fault. Rather, we perceive their decision to be irrational because their *prior knowledge* is different or foreign than our own. Referring back to figure 1, their 'black box' parameters are different than ours, not the decision mechanisms. Now that the model building foundation is laid we can proceed to constructing a more complete model of the primate motor control system.

4. The Controller

The stage is now set for looking at the workings of the primate muscular-skeletal controller. The basic purpose of the controller is to send the appropriate efferent signals to the actuators. In the case of the space station, the controller is the software that determines what commands to send to its actuators given the state of the system and the current data from the sensors. In this case the lines of code are explicitly written out. These lines of code represent the exact (explicit) control law. Neuroscientists have to extract a control law from the input/output characteristics of the CNS. To gain a little perspective of where the controller fits refer to the block diagram of figure 3 (If you are reading an electronic copy of this paper, then the block diagram is a separate Word file). The "Control Law" block is on the far left of the diagram and labeled "A". The control law block (block A) sends out its control signal that gets transformed from block B (if feedforward is used for the particular model), then from block C, then noise is added, then through block D, and the output of block D is the actual movement.

The A-C-D transformation (see figure 3) just described is considered a feedforward process because the control law does account for sensory feedback. A paper by Todorov and Jordan (Todorov and Jordan, 2002) opens up by implying that the body of literature on human control schemes uses this feedforward model to explain how CNS controls for a given objective. This model says that the CNS has a detailed plan of the exact motion (exactly which muscles to activate, for how long, and when) that needs to be executed to accomplish a given task objective. Essentially, the processes of planning and executing are distinct and serial. The only use of feedback in this situation is through the controller adaptation process illustrated by the H-I loop (blue loop) in figure 3 (The dynamics of this process are much slower than the rest of the system.).

Todorov and Jordan are not satisfied with this model. They proposed that the human CNS uses an optimal feedback control scheme for motor coordination. R.E. Kalman developed much of the theory behind optimal control while developing estimation techniques of noisy data for the space program. A Kalman filter is simply an optimal recursive data processing algorithm. It is a set of mathematical equations that provides an efficient computational solution of the least-squares method. The filter is very powerful in several aspects: it supports estimation of past, present, and even future states, and it can do so even when the precise nature of the modeled system is unknown. As its name implies, optimal control - which usually takes the form of a feedback control law - is the best possible control algorithm given the desired task and the constraints of the system. Whenever the task allows redundant solutions, movement duration exceeds the shortest sensorimotor delay, and either the internal state of the plant is uncertain or the consequences of the control signals are uncertain, optimal performance is achieved by a feedback control law that resolves redundancy moment-by-moment, using all available information to choose the best action under the circumstances. By making decisions regarding movement details at the last possible moment, this control scheme is able to take advantage of better alternatives (as they arise) once the actual trajectory has deviated from the desired (expected) path. This model makes no distinction between trajectory planning and execution.

Simulations were performed to explore the validity of such a theory. A simple simulation of an optimal control law required the movement (in a 2-D plane) from an initial (random) position to a position where the sum of the two coordinates equal 2. The other objective of the control law was to minimize the energy used to get there:

> *Given the state variables x_i, choose the control signals u_i that minimize the expected cost $E_\varepsilon(x_1^{final} + x_2^{final} - X^*)^2 + r(u_1^2 + u_2^2)$ where the stochastic dynamics are $x_i^{final} = ax_i + u_i(1 + \sigma\varepsilon_i)$; $i \in \{1,2\}$, and ε_i are independent random variables with mean 0 and variance 1. In other words, the (redundant) task is to make the sum $x_1 + x_2$ of the two state variables equal to the target value X^*, with minimal effort*

The results of the simulation are shown in figure 4 (taken directly from the Todorov and Jordan paper).

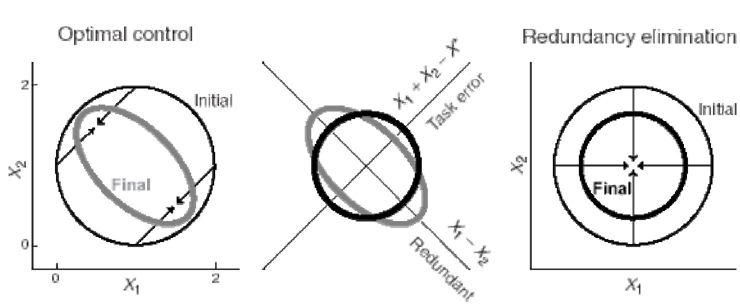

Redundancy exploitation. The system described in the text (with $X^* = 2$, $a = \sigma = 0.8$) was initialized 20,000 times from a circular two-dimensional Gaussian with mean (1, 1) and variance 1. The control signals given by the two control laws were applied, the system dynamics simulated, and the covariance of the final state measured. The plots show one standard deviation ellipses for the initial and final state distributions, for the optimal (left) and desired-state (right) control laws. The arrows correspond to the effects of the control signals at four different initial states (scaled by 0.9 for clarity).

Figure 4. Simulation using optimal control and redundancy elimination schemes

Essentially, the simulation attempts to minimize the function: $x_1 + x_2 - 2$ while keeping the control signal minimal. Trials were initialized randomly with the standard deviation shown by the thin black line (figure 4). The optimal control simulation (figure 4, left) takes advantage of redundancy by allowing the solution to end anywhere along a line where $x_1+x_2=2$. In the redundancy elimination simulation (figure 4, right) they used the average behavior of the optimal control solution from the previous simulation as the desired trajectory: $x_1=x_2=1$. On average this trajectory has to travel further (larger control signal) to get to its solution (cost function is usually greater). So, in the case of the optimal controller, by opening up the redundant dimension to the solution space, more solution exploration is possible (more chances to find a better solution). In this particular simulation there is only one constraint in 2-D, so now the solution (without noise) is a line instead of a point. The authors use their *minimal intervention principle* to explain the variance in the dimensions that do not affect the net outcome: Deviations from the average trajectory are corrected only when they interfere with task performance. In contrast, the redundancy elimination (trajectory planned) model deals with redundancy by constraining excess degrees-of-freedom, even if those constraints don't help the performance. Simply put, every actuator has an explicit plan.

The paper contained some more simulations that endorsed their optimal control theory. The paper even had a few experiments that demonstrated that the simulated trajectories of the optimal controller looked pretty close to the actual trajectory that human subjects followed. While the methods in the paper appeared to be convincing, the conclusions drawn from these simulations and experiments cannot be clearly deduced from the evidence presented. The paper was able to establish that the optimal controller performed better than the trajectory planning model. It then inferred that we use an optimal control scheme instead of a trajectory planning scheme. Implicit in this argument is the premise that the human motor control system uses the best possible scheme. Is this a valid premise?

The paper provided a few experiments on human subjects where the trajectory followed by the subject in a goal-oriented task seemed to resemble the trajectory of a simulated optimal controller (for the same task). Should this lead one to believe that we use an optimal control scheme for motor control? In order for the argument to be convincing the similarities must be demonstrated across a large domain of tasks, motions, and frequencies. Two models (e.g. simulation and actual data) can be proven to be identical only if their input/output (e.g. input: task; output: trajectory) relations are identical for all possible inputs. Not only did they use a very small range of inputs (only 2 distinct tasks), but the input/output relations had some significant discrepancies.

Their simulation used rigorous mathematics to define their controller. For a modern PC the number of matrix multiplications, additions, and inversions needed in each sample for Kalman filtering isn't too overwhelming. But, to imply that the human CNS performs such calculations in real-time is simply absurd. While the paper does not explicitly state this conjecture, they offer no alternative explanation, let alone speculate on any neural mechanism that could perform such processing.

Of course this paper isn't the first to suggest that the human CNS motor controller can be modeled with complex mathematics. One widely accepted property of human movement is the notion of minimum jerk. Depending on the nature of the afferent signals returning to the controller, the CNS would have to compute several derivatives to determine the jerk of the movement, and then the parameters of the controller would have to adapt in order to develop a control scheme that does minimize the jerk (H-I learning loop in figure 3). Suggesting that the human CNS computes multiple derivatives of its afferent input may be far-fetched, but it is a widely accepted notion that an *emergent* property of the human motor control system does minimize the jerk.

The Todorov and Jordan model of optimal control relies on sensory feedback for the state estimation that is needed to determine the optimal corrective movement that will accomplish the task. This is illustrated as the A-C-D-E-F-G loop in figure 3. For a motor control system feedback is an extremely powerful tool. But, if used improperly, it can cause catastrophic instabilities. Let's suppose that once the control signal, u, is computed, block C (figure 3) induces a 50 msec transmission delay in the signal. The signal then recruits muscles that move the skeleton (block D). Afferent signals are encoded (block E), and block F induces another 50 msec to the signal. Upon arriving at the state estimator (block G), let's assume a small overhead computation time of only 25 msec. So, without even considering the phase lags induced by block D or the additional computation that would probably actually be needed by blocks G and A, there is pure delay of 125 msec (whether or not is an accurate estimate it will be used in this example). Let's say that a particular task to be performed requires a movement path that contains frequency components up to 4 Hz (a bandwidth of 4 Hz). By the time the 4 Hz component makes the A-C-D-E-F-G loop, it will have a 125 msec delay (just from blocks C, D, F, and G), which corresponds to a 180^{o} phase lag. The system is now unstable. In reality, time delays can be longer, and phase lags (from muscular-skeletal mechanics) are finite and significant. This analysis would lead one to believe that any quick motion would cause the controller the go unstable and exhibit uncontrollable oscillations.

From a theoretical point of view a controller can become more stable (larger range of stability) by taking advantage of a feedforward compensator (figure 3, block B) and a forward dynamics model (block J). A feedforward transformation is able to induce a phase lead that could cancel an anticipated phase lag. A forward dynamics transformation is able to predict the dynamics of the muscles from the state of the system and a copy of controller output. This means the CNS can predict the state variables and the errors without having to wait for the sensorimotor delay. The stability of such a system will strongly rely on the accuracy of the forward model and the ability to cancel the response from the remote system.

A paper by Bhushan and Shadmehr (Bhushan and Shadmehr, 1999) investigated a model of the human motor controller that utilized feedforward and a forward dynamics computation. The hand of a human subject was coupled to a robotic arm during some simple reaching tasks. The robotic arm was able to apply an arbitrary force field that could change the dynamics of the reaching task. After a series of experiments, they found that a control scheme that utilized an adaptive feedforward model (A-B-C-D-E-F-G and H-I-B loops in figure 3) was pretty similar to the data observed in their subjects, but lacked some essential characteristics. It seemed that the controller was using descending commands to predict the dynamics of the plant. But, after constructing a controller architecture that incorporated an adaptive forward model (A-C-D-E-F-G, A-J-G, and H-I-J loops in figure 3), they found that this model didn't exactly match the data either. So, a third model was constructed that used the feedforward transformation in conjunction with the forward model computation (A-B-C-D-E-F-G, A-J-G, H-I-B, and H-I-J loops), and they found that the resulting dynamics were remarkably similar to that observed in the experimental data.

5. Command Generation

We now shift our focus from the controller dynamics to the generation of the command signal that leaves the CNS and is directed to individual motor units (control signal, **u**; the signal entering block C in figure 3). We know that the CNS generates some sort of signal that orchestrates a pattern of muscle movements. But, where does this motor command originate from? How does the CNS encode the command? And what is the nature of such a signal? In week 2 Dr. Lee Miller discussed the details of how neurons in the motor cortex control the movement of distal muscles.

One of the pivotal papers on this specific topic came from the Georgopoulos lab in 1982 (Georgopoulos, et al, 1982). In their study they recorded the activity of single cells in the motor cortex while monkeys made arm movements in eight directions (at 45^{o} intervals) in a two-dimensional apparatus. All of the movements started from the same point and were the same amplitude. Of the 606 cells that related to proximal arm movements 323 were active in that task and were studied in detail. The frequency of discharge for 241 of the 323 cells modulated in a predictable pattern that varied with the direction of movement. Each of the 241 cells had a preferred direction in which they fired most frequently, and their firing rate gradually decreased as the movement direction went farther away from their preferred direction. The study also revealed that a cell's discharge changed with several (not just one) movement directions. This means that the individual cells did not just encode movement directions in a simple, one-to-one fashion. This suggests that movements in a particular direction are not subserved by motor cortical cells uniquely related to that movement. But rather, the orderly, sinusoidal variation in the tuning curve suggest that cells with overlapping tuning curves might cooperate to generate a movement trajectory in a desired direction.

Sixteen years later Georgopoulos published another paper that investigated neural codings on distal movements, but this time the study focused on finger and wrist movements (Georgopoulos, et al, 1998). Previous studies had observed that control of each individual finger or wrist movement was achieved not by activation of a specific territory of motor cortex but rather by the net activity of overlapping neuronal populations distributed throughout the hand region of the motor cortex. In their study they apply a population vector analysis to the problem of coding individual finger and wrist movements to determent whether single motor cortical neurons are tuned in an abstract instructed finger movement space, and to determine whether the movement performed is specified by the neuronal population activity. They were able to conclude that the firing of neurons in the motor cortex that are related to individual finger movements do contain information that represents the spatial geometry of the hand embedded within the flexion and extension movement domain. These results are consistent with their previous work where they found that the direction of arm movements is encoded (in a complex fashion) in the motor cortex.

Georgopoulos et al. were able to show that neuronal firing patterns in the motor cortex represent the kinematics of distal movements. While it is convenient to observe kinematics of movements, kinematics alone cannot predict what the muscles are being commanded to do. Lifting a volleyball and lifting a bowling ball could be done with very similar kinematics, but the intensity of the commands to the individual muscles would be significantly different.

Kalaska et al (Kalaska et al 1989) looked at direction-related as well as load direction-related activity in the primate motor cortex during a two-dimensional reaching task. As one might imagine, the responses of single cells to various combinations of movements and load directions is very complex. But, averaging data from the sample population revealed that any combination of movement direction and load direction could be described reasonably well by firing patterns of neurons in the motor cortex as the summation of movement-related firing without any load and the extra neuronal firing caused purely by the load (as measured prior to movement). Basically, the firing pattern is the linear superposition of the firing pattern from just the movement and the tonic activity from just the load.

These studies just presented gave some great insight into where the command signals come from and how they are encoded in the primary motor cortex. It has been shown that certain neurons in the CNS are responsible for certain movements. These neurons send out motor commands (signal going into block C in figure 3) that eventually recruit the individual motor units. But what is the nature of these commands? Before answering that question, let's revisit the path of a command signal leaving the CNS (refer to figure 3). Recall that once the motor command leaves the CNS, there is a significant delay as it passes through blocks C and D. The command then actuates the muscle, and there is an additional delay through blocks E and F before afferent information returns to the CNS. While information is traversing through C-D-E-F, the CNS has no direct confirmation that its command has executed the desired motion. In simple, rapid finger and hand motion, this delay time is often longer than the actual movement time. For example, if the subject attempts to quickly move a low inertia knob only 30 degrees, they would execute a discrete and rapid motor command that would get the knob close to the target. But, because of the finite-time feedforward and feedback circuit, there may be some time before any secondary adjustment could be made.

In the first week of the class Dr. Jim Houk discussed some of the kinematic properties of rapid hand movements in a knob turning task (Novak et al. 2000). The study employed a novel approach to decomposing and characterizing the discrete movements during the knob turning task. They then performed an analysis of the properties of those pure movements without overlapping submovements. The paper demonstrated that subjects attempted to eliminate the need for corrective secondary movements by making their primary movements more accurate with practice. But, the amount of variability inherent in the rapid primary movements caused them to make corrective secondary movements. As the movements become faster, the signal-to-noise ratio increases. So, rapid movements have too much noise (thus, too much movement variability) to accomplish the goal in a single command. Information has to traverse through C-D-E-F (alternatively, visual feedback could bypass blocks E and F via a parallel visual feedback loop, which is not shown in figure 3) before the CNS can make its next iterative movement. The slower loops of H-I-B and H-I-A modify blocks A and B so that only one movement is needed to accomplish the goal. This is accomplished via practice.

Another paper by Novak, Miller, and Houk focused on the adaptation of these movements with an imposed perturbation. A motor attached to the knob applied a destabilizing negative viscosity perturbation to the subjects. Upon initial application of these perturbations, subjects required large secondary movements to accomplish the goal. But as subjects practiced with the negative viscosity, they were able to reduce the amplitude of the primary movement, and therefore needed to make fewer and smaller corrections. With enough practice the subjects' movements became as accurate as before the perturbation was applied. These observations supported their hypothesis that subjects adapt their movements such that they make more accurate primary movements. The increase in smoothness and the return to near normal kinematics is just a secondary effect of minimizing the need for a secondary movement.

6. State Estimation

In practical engineering systems, often the most difficult aspect of controlling complex systems is estimating the state of that system. Once the state of the system is known, control signals are often trivial solutions of a control law. In the example of the International Space Station, all the sensors, most of the computer code, and most of the computer hardware are dedicated to estimating the state of the ISS. A simple single-input-single-output (SISO) system may use a simple feedback signal to estimate its state, and then it renders an appropriate control signal. For SISO systems the state estimation/command generation process typically takes place in a single step and is often a closed-form control law. But, as the number of inputs and outputs increases, the complexity of the system rapidly increases. With the enormous amount of degrees-of-freedom (DOFs) of the muscular-skeletal system, it would take computational power comparable to that used on the ISS in order to have real-time state estimation. Yet, to our knowledge, the human CNS has no circuitry for explicitly solving PDEs.

In the previous section (see Command Generation) it was shown that the primary motor cortex (M1) is responsible for explicitly commanding the individual muscles during voluntary movement. But, in goal-directed movement tasks are defined in an external 3-D space (often with a 4th dimension of time). So, the CNS must translate the location of the target, specified in an external coordinate frame, into a set of muscle activation patterns, specified in an intrinsic coordinate from. If M1 is responsible for distributing the final descending commands to the motor units, then it would make sense that the individual neuronal firing patterns within M1 would correspond to specific muscle activity. But, we have seen from Georgopoulos that there are neurons in M1 that represent a movement (in a preferred) direction. A paper by Kakei et al. (Kakei et al. 1999) addresses the question: "What aspects of movement are represented in the primary motor cortex: relatively low-level parameters like muscle force, or more abstract parameters like handpath?" If I have learned anything from neuroscience, it is that simple questions do not have simple answers. This question is no different. As one would expect, the Kakei paper found that a significant portion of the neurons studied (28 out of 88) did have firing patterns that correlated with individual muscles. But, unexpectedly, an even larger portion of the neurons (44 out of 88) displayed changes in firing rates that were related to the direction of the wrist movement in space.

To the neural coding of our brain, 3-D space and time are abstract concepts. But yet we have no trouble orchestrating movements with strict time and space constraints. It is obvious that there must be hierarchical processing between M1 and other cortical regions that are able to transform information between various (and often extremely abstract) neural representations. For example, in another paper by the Kakei group (Kakei et al. 2001) they found that neurons in the ventral premotor area (PMv) seemed to encode the direction of movement in space, which was an independent representation of what was recorded in M1. From their data they concluded that the intracortical processing between PMv and M1 might be responsible for the transformation between an external reference frame (outside 3-D space and time) and a more abstract internal reference frame (neuronal circuit specific).

The prefrontal cortex (PF) is one of these areas of the brain that is responsible for storing abstract representations. It has been suggested that the PF is used in planning complex behavior and is crucial in analyzing serial events and in using the results to control behavior. How is the CNS able to encode serial events into spatial neural representations in the PF? A paper by Beiser and Houk (Beiser and Houk 1998) set out to answer this question. They present a neural network model of a cortical-basal ganglionic transformation of sequential sensory input into spatial patterns in area 46 of the prefrontal cortex. The process provides a means to encode serial events into spatial representations without requiring adaptive training mechanisms of any kind. In the model two types of units are modeled in the PF: recurrent (R) units and event (E) units. Since the R units are in a bistable, positive-feedback loop, past information is retained, plus new information can be integrated in. Because R units provide information about past events (state space representation) and can influence future states, temporally spaced sensory inputs can be linked. This model constructs PF patterns that provide an unambiguous coding of the input sequence. But, the model does not attempt to explain how the information can get from the neural representation in the CF to meaningful motor commands via M1. This *decoding* problem is another complex topic in the field of motor control that I will leave for another day.

7. Concluding Remarks

In week 6 of the course Dr. Sandro Mussa-Ivaldi raised a point that stuck with me throughout the remainder of the course. He began by describing the semantics and structure of the Roman numeral system in great detail. The reason for this was to make you realize how inefficient the system was and how it didn't have an exact structure the way modern Arabic numerals do. Then he made the statement that despite their great advances in politics, philosophy, and science, the Roman culture failed to produce any advanced mathematics. The moral of the story was that you have to have the right tools for the job. The inherent structure of the Roman numeral system prohibited the Romans from developing advanced mathematics because it was extremely difficult to construct formulas that governed the semantics and relationships of the numbering system. Arabic numbers did not invent higher mathematics; higher mathematics is just easier to represent in Arabic numbers. In this paper I have demonstrated the extreme complexities of the primate motor control system by comparing it to engineering systems. I have described the CNS using engineering methodologies that were specifically designed to analyze engineering systems. Just as relatively simple mathematics seems very complicated in Roman numerals, maybe the reason CNS dynamics seem so complicated is because we do not use the proper tools to analyze it.

In the optimal control model described earlier (Todorov 2002), the CNS state estimator is described as behaving like a Kalman filter. Even if it was true that the CNS is capable of estimating state variables that correspond to Kalman's recursive least-squares method, does this mean the CNS explicitly estimates its own state by using the exact same methodology as an engineer programming a digital computer would use? The CNS plays by different rules than digital computers or conventional mathematics. Perhaps by using the semantics of neural circuitry and logic, there is a relatively simple way to process data that results in a transformation similar to a Kalman filter.

In a different section of this paper I referred to a theorem that suggests that the CNS controller attempts to minimize the jerk of the movement. I suggested that it would be far-fetched to think that the human CNS could take multiple derivatives of many afferent signals in real-time. I was basing this on the fact it takes a digital computer many floating-point operations to compute just one derivative of one signal at a single point in time. So, by comparing the CNS to a digital computer, the CNS appears to be a computational juggernaut. But, if I instead compare the CNS to an *analog* computer, the processing power of the CNS becomes much less significant. Before the transistor became the common switching device in computers, the digital computer was rivaled by its older brother: the analog computer. For our case, what made the analog computer interesting is that it was capable of having multiple parallel circuits for processing the derivative of a signal (multiple differentiator circuits). So, a properly configured analog computer could easily process a large number of derivatives of multiple channels (in parallel) in real time. Perhaps the CNS has dedicated circuitry for estimating derivatives.

The term 'central processing unit' (CPU) is generally reserved for describing the processing circuitry of a digital computer, thus, the title of this paper seems to implicate that the CNS should be modeled and evaluated using the same standards as a digital computer. I have just shown the fallacy of such a comparison. To simulate all the computations done by the CNS (in real time) with a digital computer it would take a system many times larger than any supercomputer ever constructed. Computer engineers and scientists have been working on this problem since computers were invented. It is obvious now that the CNS uses a different methodology to store and process information than a digital computer. Similarly, the non-linear dynamics of the motor control system should probably not be described (approximated) by traditional linear system theory. In fact an entirely new method of representing the workings of the CNS may need to be developed if we want to truly understand the mechanisms of the human central nervous system.

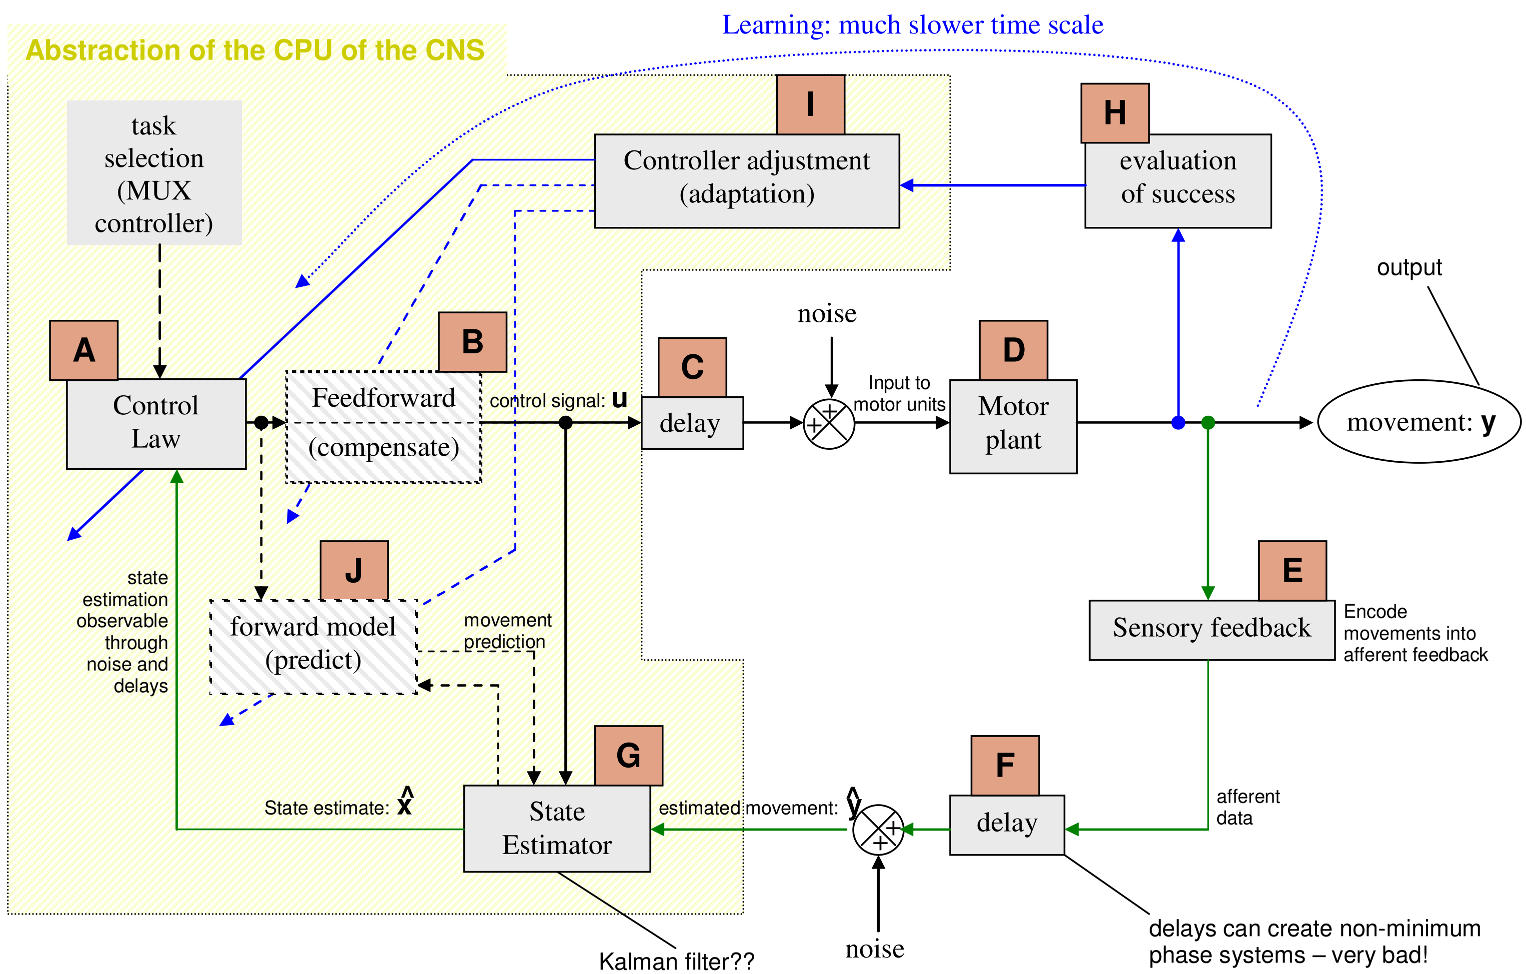

Figure 3. A model of the primate motor control system (for voluntary movement)

8. References

Beiser, D. G., Houk, J. C. Model of Cortical-Basal Ganglionic Processing: Encoding the Serial Order of Sensory Events. *American Physiological Society* 3168-3186, 1998 (week 10: Jim Houk)

Bhushan, N., Shadmehr, R. Computational Nature of Human Adaptive Control During Learning of Reaching Movements in Force Fields. *Biol. Cybern. 81.* 39-60, 1999. (week 6: Sandro Mussa-Ivaldi)

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. On the Relations Between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex. *The Journal of Neuroscience vol 2 no. 11.* 1527-1537, November 1982 (week 2: Lee Miller)

Georgopoulos, A. P., Pellizzer, G., Poliakov, A. V., Schieber, M. H. Neural Coding of Finger and Wrist Movements. *Journal of Computational Neuroscience 6.* 279-288, 1999. (week 2: Lee Miller)

Kakei, S., Hoffman, D. S., Strick, P. L. Muscle and Movement Representations in the Primary Motor Cortex. *Science vol 285*. 2136-2139, 24 September 1999 (Week 2: Lee Miller)

Kakei, S., Hoffman, D. S., Strick, P. L. Direction of Action is Represented in the Ventral Premotor Cortex. *Nature Neuroscience vol 4 no 10.* 1020-1025, October 2001. (Week 2: Lee Miller)

Kalaska, J. F., Cohen, D. A., Hyde, M. L., Prud'homme, M. A Comparison of Movement Direction-Related Versus Load Direction-Related Activity in Primate Motor Cortex, Using a Two-Dimensional Reaching Task. *The Journal of Neuroscience.* 2080-2102, June 1989 (week 2: Lee Miller)

Leon, M. I., Shadlen, M. N. Exploring the Neurophysiology of Decisions. *Neuron 21*, 669-672, October, 1998. (week 9: Sara Solla)

Novak, K. E., Miller, L. E., Houk, J. C. Kinematic Properties of Rapid Hand Movements in a Knob Turning Task. *Exp Brain Res, 132*. 419-433, 2000 (week 1: Jim Houk)

Novak, K. E., Miller, L. E., Houk, J. C. Features of Motor Performance that Drive Adaptation in Rapid Hand Movements. *Exp Brain Res, 148*. 388-400, 2002 (week 1: Jim Houk)

Shadlen, M. N., Newsome, W. T. Motion perception: Seeing and deciding. *Proc. Natl. Acad. Sci. USA. 93,* 628-633, January 1996 (week 9: Sara Solla)

Todorov, E. and Jordan, M. I. Optimal Feedback Control as a Theory of Motor Coordination. *Nature Neuroscience vol 5 no 11.* November 2002. (week 6: Sandro Mussa-Ivaldi).