Neurobiological and neurorobotic approaches to a control architecture for a humanoid motor system.

Simon F. Giszter[1], Karen A. Moxon[2], Ilya A. Rybak[3], John K. Chapin[4]

1 Department of Neurobiology and Anatomy, MCPHU school of Medicine, 3200 Henry Ave, Philadelphia, PA 19129, USA
simon@swampthing.neurobio.mcphu.edu

2,3 School of Biomedical Engineering, Science and Health Systems, Drexel University, Philadelphia, PA 19104, USA
karen.moxon@drexel.edu
rybak@drexel.edu

4 Dept Physiology, Health Science Center, SUNY Downstate, Brooklyn, NY, USA
john_chapin@netmail.hscbklyn.edu

Abstract. In the mammalian nervous system the sensorimotor interface between cortex and spinal cord plays a key role in novel skill formation and motor learning. We seek to understand the principles of motor learning at this interface using a multidisciplinary approach. We believe this approach will prove relevant to the development of biomimetic control architectures for humanoid robots. Learning at this interface requires an understanding of the spinal output structures. These may guide, constrain and ultimately must form the basis of the algorithms needed for adaptive motor learning. These spinal structures interact with descending cortical control to produce accurate limb trajectories, and novel motor behavior. Our research integrates three approaches: (1) defining the properties of a basis set of specific temporal dynamics (force-field primitives) organized at the spinal level, (2) simulating how spinal systems may be organized into a hierarchical neural circuitry with a modular organization using detailed computational models of the spinal neural circuitry, and (3) experimenting with a variety of neuroprosthetic interfaces to investigate cortical coding, learning and execution of sensorimotor tasks. Combining these three approaches will create a framework for testing a set of neurobiologically inspired learning models.

1 Introduction

The architecture of the mammalian central nervous system (CNS) that produces accurate limb trajectories, multi-limb locomotion, novel motor tasks and the motor learning algorithms implemented by this architecture can be used as a model system for the development of humanoid robots. To investigate this architecutre our approach has been to insert interfaces that interact with the biological system at different levels of organization. Our goal is to systematically replace, augment or co-opt components of the biological system with biologically inspired artificial components that interact with the original biological system in real-time. We aim to (1) monitor the response changes of the local neurons (neural plasticity) and the behavioral adaptations of animals interacting with these interfaces and (2) characterize the output behavior and capabilities of the interfaces. The results of these studies, outlined below, suggest a learning system which is partly hierachical but also possesses much local autonomy, and layered competences.

1.1 The degrees of freedom problem and motor 'bootstraps'

A biological strategy for the control of movement that designers of biomorphic robots have considered using is progressive development of skills in mammals (e.g. Drescher). This progressive learning of skill immediately introduces the problem of too many degrees of freedom – to allow simple searches for adequate strategies and control. The degrees of freedom problem first discussed by Bernstein presents difficulty to the animal, robot or designer of a neurorobotic or neuroprosthetic interface. How many degrees of freedom can be easily handled by a learning system? How long will learning take? Can catastrophic failure be avoided? This problem encompasses selection of effectors, linkage configurations, actuators and feedback and measurement use in execution. The high numbers of degrees of freedom available to biological systems and especially human systems can make the search space appear to rapidly approach astronomical size.

In light of this issue it has been suggested that organisms require a 'boot-strap" system so that new born organisms can immediately interact with their surroundings. For many organisms this is essential to survive.. It would seem reasonable that even with a strongly nurtured infancy, there would be selective pressure to ensure that some type of method would evolve to foster rapid searches of potential successful movement strategies. Spinal circuits or primitives that we have explored (e.g. Giszter et al. 1993, Kargo and Giszter 2000) could provide a basic competence or set of 'seed skills', installed by evolution, from which efficient locomotory and manipulation skills can be rapidly elaborated after birth. If and how this occurs is an important issue that is relevant to biomorphic robot design, to neuroprosthetic interface design and to neuroscience.

The first half of this paper presents a review of neurorobotic approaches that we anticipate may allow us to examine the solutions to these issues in cortical representations and at the corticospinal interface. The second half of this paper focuses on a conceptual understanding of evaluating how learning strategies might be elaborated at this interface based on our current biological understanding of spinal motor primitives.

1.2 Neurorobotics as an exploratory tool

Data(Chapin et al. 1999) suggest that designing hybrid, brain-machine organisms will provide further insights into the functioning of truly intelligent, adaptive systems which will aid in the development of humanoid robotics.

We first give a brief introduction to the design principles and implementation of a direct brain to machine interface through a neurorobotic device and then provide background data and isues with a theoretical framework for development and learning defined by motor primitives. The challenge to the machine learning and roboticist community is to discover the advantages and disadvantages of different possible learning methods and architectures. The challenge to our experimental programs is to discover which of these several methods the biological system actually employs and use these to design a neuroprosthetic interfaces well suited to rapid motor learning.

2 Biomimetic neurorobotics

The biomimetic approach to the development of humanoid robotics is founded on the premise that if the design mimics the human system sufficiently well, an intelligent, autonomous artificial system will be the result. One problem with this approach is that we still know very little about how mammalian (and human) biological systems that produce intelligent autonomous behavior. Mammalian brains all possess several key features, which may be essential for learned intelligent behaviors. These include a distributed, redundant system that contains stable processing systems at several, hierarchically organized levels. At each level the task is represented with different degrees of precision and flexibility. For example, at the spinal level, basic competence of locomotor patterning and support is organized and available after separation from the brain. At successively 'higher' levels of the brain, the control and organization of locomotion is further developed. Lesion studies show that this locomotion organization progresses from unbalanced locomotion in the decerebrated spinal cord (as observed on a treadmill), through obstinate progression with good balance (in the decorticate cat), to voluntary locomotion when cortical systems are engaged. At each level of organization there exists an expertise or local autonomy such that, as information flows through the CNS, it can be acted upon at the local level or passed on two successive levels after some local processing. This local autonomy results in gating or filtering sensory input and an ultimate reduction in information flow to the highest levels of the system. This preserves their computational resources for the input from selected high priority activities. The layered hierarchy provides a mechanism for delegating some details of execution to lower levels and some of the data connected with these details may often be hidden from upper levels. For example, we are mostly completely unaware of the extensive nerve information traffic into the brain in some spinal tracts such as the dorsal and ventral spinocerebellar tracts. In the hierarchy decisions may be executed simultaneously on several levels. The cat whose leg encounters an obstacle in pursuit of a mouse or in flight from a dog may not have to concern itself with the correction of that leg's individual motion but only with the overall balance in anticipation of the pounce or jump to follow. The integration of these layers of controls into an adaptive seamless motion control is a central problem of neurobiological motor control. These observations were part of the inspiration of the subsumption program of Brooks' group, and the various "nouvelle AI" approaches (e.g. see Meyer and Wilson 1991).

3 Design and implementation of a neurorobotic control

In the mammalian CNS several brain areas cooperate in motion control and planning: the cerebral cortex, the cerebellum and the basal ganglia. All must ultimately exert their controls through direct or indirect actions on the low level spinal cord and brainstem.

Our approach to resolving the function and integration of these areas is to create interfaces at various levels of the system which we use to define the levels and function of organization. At each interface we extract information or deliver control information or subsume the circuit to our experimental goal. We (Chapin et al. 1999) have developed interfaces with motor centers in the brain to directly examine cortical motor learning using a neurorobotic interface in rats. This interface effectively provides the brain with a completely novel output channel and ultimately a novel sensorimotor interface. The expectations of the animal using a motor interface, and the properties of an interface that can successfully be exploited will help us to understand the representations, integration and functions of these brain areas with output primitives and their developmental processes.

3.1 Electrodes for neurorobotic control

Recent technological developments in the design of brain-machine interfaces provide the possibility of producing a neuroprosthetic device capable of restoring motor functions in paralyzed patients by creating a direct interface with the brain. These devices would extract "motor commands" from the brain through use of neuronal population recordings. These signals would be processed and the information ultimately employed to control robotic devices, or even to use functional electrical stimulation (FES) to control the patients' own muscles directly or through surviving spinal cord below a lesion.

Tapping into the neural system at multiple levels requires an appropriate neural-machine interface. Many investigators have attempted to tap into the sensory motor system at various levels to achieve control of limb systems. These methods include direct spinal stimulation (see Giszter et al. 2000), nerve cuff electrodes and intra-cortical electrodes (see Moxon and Chapin 2000). Several of these systems have developed into complete clinical prosthetics and are successful examples of neural interfaces that replace part of the neural system to restore function (often limited function). The most obvious missing piece in each of these devices is somatosensory feedback. Most are guided by visual feedback and sometimes auditory cues.

We have been developing a cortically controlled neural interface to examine the role of interacting motor systems in the control of movement. There are roughly four major subassemblies of a cortically controlled neural interface (1) the electrodes subassembly, (2) signal conditioning subassembly, (3) signal processing subassembly, and (4) transmitter subassemblies. While a wide variety of each of these subassemblies have been developed, a complete cortically controlled neural interface for human use that is capable of interacting with and replacing biological subassemblies requires major technological advances. We will attempt to point out some of the possible solutions.

From a technical standpoint there are several issues that must be addressed when using a brain-machine interface. This includes minimizing the damage from invasive microelectrodes, which will require re-engineering the response of neural tissue to the introduction of the electrode. Recent studies of the response of neural tissue to stab wounds have shown the responses to be stereotypic. In the future molecular biologists advances may allow the brain's response to the electrode to be re-engineered so as to produce only a limited response to the initial insult and to stabilize the brain's interface with the electrode.

Another issue for the development of a direct brain-machine interface is obtaining large numbers of direct contact points (or recording sites) with the CNS with minimal tissue disruption. Several ingenious methods have been developed to optimize the number of neurons recorded wile minimizing tissue damage. One, developed by the University of Michigan Center for Sensor technology is a silicon based electrode (Wise et al. 1970). By using boron doped silicon to define a small (approximately 15 micron thick) electrode and ribbon cable (Hetke, 1994), upto 16 recording sites can be etched onto the silicon using standard photolithography processes. The silicon ribbon is flexible and deforms with the movement of the brain ensuring minimal damage. A second method was developed by the Univeristy of Utah (Jones, 1992) and is an array of electrodes etched out of a block of silicon such that the pad holding all of the electrodes floats on the surface of the brain. While each of these anchoring methods holds out promise, none have shown consistent recordings of single neurons in chronic preparation. One reason may the semiconductor properties of silicon. When a silicon electrode is embedded into the brain, the constant background activity and corrosive nature of the brain microenvironment reduces the signal-to-noise making it difficult to record single unit activity

over long periods of time. We have developed a ceramic electrode using 99.9% alumina, similar to the silicon electrode that has a high density of recording sites but, because of the insulating properties of pure alumina, they can continuously record single unit activity for at least 3 months.

Critical components of a neuroprosthetic control device for interfacing with neural systems must also include an integrated circuit to amplify and filter the neural signals.   (we should either remove this paragraph or add another sentence or two). While commercial available system exist for processing the raw neural signals in a laboratory setting, ultra-miniature circuits for use with a neuroprosthetic system are still being developed.3.2 Implementation of a Neurorobotic system

The Chapin laboratory has recently provided the first demonstration that large scale multineuron recordings in experimental animals can be used to control robotic devices with sufficient accuracy to obtain a water reward from a dropper (Chapin et al., 1999). This provides a novel means to examine how voluntary representations and computations in the motor cortex are organized. If the cortex is disconnected from its normal motor effectors, but is then given the opportunity to directly control an external robot, how does the brain use this novel environment and interface? Does the cortex seek out regularities in the controlled movement and reinvent its own primitives? Alternatively, does cortex expect certain primitives that it must use in order to begin to learn? The issues of biological layering of controls, hierarchy and execution can be addressed in new ways in this paradigm. If understood sufficiently well this research direction will ultimately provide direct neural interfaces into novel devices or humanoid robots or robot limbs.

3.2.1 Implementing a neural interface in an awake behaving rat:

As described in detail in Chapin et al. 1999, we have recently demonstrated the feasibility of using simultaneously recorded neuronal populations in the brain to control a robot arm.

Fig. 1. Experimental protocol. (a) 'Lever-movement/robot-arm' mode: rats were trained to press a lever (b) for a water reward; displacement was electronically translated to pro-portionally move a robot arm (c) from rest position through a slot in barrier (d) to a water dropper (e). The robot arm/water drop moved passively to the rest position (to the rat). (f) 'Neuronal-population-function/robot-arm' mode: Rats were chronically implanted with multi-electrode recording arrays in the MI cortex and VL thalamus, yielding simultaneous recordings of up to 46 discriminated single neurons. (g) Superimposed waveforms of 24 such neurons. (h) Sample spike trains of two neurons (N1, N2) over 2.0 s. (i) Neuronal-population (NP) function extracting the first principal component of a 32-neuron population. (j) Switch to determine input source (lever movement or NP function) for controlling robot-arm position. In experiments, rats typically began moving the lever. The input was then switched to the NP function, allowing the animal to obtain water through direct neural control of the robot arm.
(reproduced from Chapin et al. 1999)
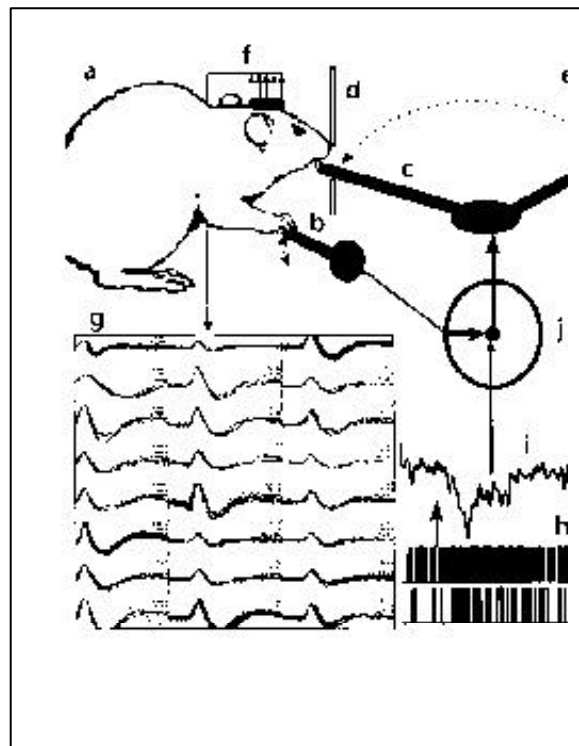
As shown in Fig. 1, rats were trained to use their forelimbs to position a spring loaded lever to obtain a water reward. This apparatus was designed to translate the position of the lever ? the proportional control of the position of a robot arm. The robot arm was configured to move radially from a rest position near the animal's mouth to the position of a water dropper. In the initial phase of training, rats obtained their water by pressing down the lever to position the robot arm under the water dropper. A small water drop was captured in a small cup at the tip of the robot arm. The rat then released the lever, allowing the robot arm to return the water to the rat's mouth. Neuronal population recordings obtained during this "lever movement controls robot arm mode" (lever drives robot arm) were then used to investigate the neural encoding of the forelimb movements, and then

to derive neuronal population functions (NPFs) capable of predicting movement in realtime. A multi-channel electronic device was constructed to calculate, in realtime, the inner product of such an NPF (a 32 number weighted matrix) times the time-integrated spiking activity of 32 simultaneously recorded neurons. The output was a single analog voltage signal that could be used in proportional control of the position of the robot arm. A switch was used to arbitrarily change the control the robot arm from the lever (necessarily driven by real forelimb movement) to this brain-derived NPF signal.

Rats were chronically implanted with multi-electrode recording arrays in the MI cortex (16 electrodes) and VL thalamus (8 electrodes) for a total of 24 recording channels and, with four single neuron discriminators per channel, a theoretical possibility of recordings of 96 single neurons. Initially the rats were trained to press down a spring loaded lever for a water reward, retrieved from a dropper to the mouth by a "robot arm" (lever drives robot arm mode). The robot arm position was unconstrained and continuously variable and could thus be moved beyond the water dropper range. Positioning was actuated by a radial motor driven by an analog voltage output from an angular transducer on the manipulandum. A switch was then used to determine whether the robot arm position could be controlled from the lever or the NPF integrator (brain activity (NPF) drives robot arm). In experiments, rats typically began working in the "lever drives robot arm" mode. Without warning control of the robot arm was then switched to the "brain activity (NPF) drives robot arm" mode, allowing testing of whether the animal could routinely obtain its water through direct neural control of the robot arm.

The aim of Chapin et al. was to determine whether the animal could utilize its own brain activity, as electronically transformed in realtime from on-line neuron activity recordings, to control the robot arm with sufficient accuracy to obtain water. We hypothesized that a peak in neuronal activity which preceded voluntary manipulations by the rat coded for forelimb movement as the rat pressed the lever down. We tested the hypothesis that this neural code could be used to directly achieve the same positioning of the robot arm that the forelimb controlled lever produced. It was not immediately obvious if the animal could behaviorally utilize this brief, pre-movement brain activity signal as a direct surrogate for the relatively long duration limb movement with which it was normally associated. To test this notion, an NPF was generated using neurons selected for their strong expression of the pre- motion peak, and then "optimized" using PCA. This NPF was converted into an analog voltage for controlling the robot arm in realtime. Through a switch, one could suddenly change the control of the robot arm from the "lever drives robot arm" mode to a "brain activity (NPF) drives robot arm mode". This allowed investigation of whether the NPF signal could substitute for the lever motion as the operant behavior necessary to obtain the water reward. To maintain the normal association between the lever motion and the robot arm, the animals were allowed approximately 5 minutes at the beginning of each experimental day to work in the lever drives robot arm mode. This ensured that the animal would continue to initiate conditioned movements when the mode was suddenly switched to the brain activity (NPF) drives robot arm mode. These movements provided a direct measure of the timing of the animal's motor-behavioral "intention", as defined by the conditioning task.

The success of the NPF in maintaining the conditioned behavior was found to be dependent on the strength of its pre-movement activity relative to its background variability (Moore et al., 1966), and this was most importantly related to the number and selectivity of the neurons whose activity was integrated to generate the NPF. In one rat for example, the threshold for the NPF to move the robot arm to the water dropper was 10.25 SDs above the mean, but among the 32 neurons used to encode it the mean pre-movement peak activity was 3.2 SDs above the mean. Thus, the NPF resolved the magnitude and timing of limb movement much more clearly than any of its constituent single neurons. Across the 4 successful animals, the pre-movement peaks were a mean 9.2 SDs above the mean (range: 7.64-10.25), but only 3.6 for the individual neurons. In the 2 unsuccessful rats, the yield of task related neurons was relatively low (16 and 19), yielding relatively weaker pre-motion peaks. Thus pre-movement peaks of 6-8 SDs above the mean, and a NP of 20-25 task related neurons were required here to maintain the operantly conditioned behavior and successfully control the robot arm.

Of particular interest was a gradual reduction of forepaw movement associated with the NPF signals used to control the robot in the "brain activity (NPF) drives robot arm" mode of the task. This reduction in forepaw movement indicated to the authors that some of the rats were successfully dissociating the NPF from the normal physical effectors direct role in the task in the lever drives robot arm condition and building new task representations.

3.2.3 Understanding the structure and integration of motor learning mechanisms at the cortical and spinal level

Despite the feasibility demonstration recounted in the preceding sections, certain critical scientific advances must be made to successfully integrate the approaches described: In particular, a "neurocontroller" must be developed to translate brain-derived motor information into multi-dimensional multi degree of freedom

movements. The problem is to find a computationally efficient approach to emulate the spinal cord's ability to "decode" descending signals from supraspinal centers and use them to implement movement strategies. This neurocontroller might well ultimately be isomorphic with the lower level controls needed in a humanoid robot.

Successful designs with the neuro-robotic paradigm will be able to ultimately control posture and locomotion as well as limb movement. Such designs may translate readily into autonomous or semiautonomous architectures, in a humanoid design, if they emulate biological design principles. In such designs one must recognize the important computational roles played by each level of the interacting and cooperating layers of control. The biological system demonstrates how different computational tasks may be shared in a collection of cooperating supraspinal motor structures, plus the low level execution system of the spinal cord. Each level not only processes descending signals from the higher brain, but may also generate its own motor patterns and primitives. Learning how to integrate devices with this architecture will provide us one type of blueprint and design strategy for a viable humanoid.

3.2.3 Next Generation Devices

We have been successful in our lab in recording from arrays of electrodes chronically implanted at different levels of the brain and recording command signals that code for movement. We have successfully used the neural command signal to control an external robotic device [Chapin et al. 1999]. This method can be used to test our hypothesis about the brain-spinal cord interface. The neural signals, recorded from the brain's motor systems can be combined into a command signal and introduced into the spinal cord to activate local spinal circuits. Successful implementation of this would allow us to test our hypothesis about how neural signals, recorded across multiple levels of the brain are combined for the control of movement by the spinal cord. We are currently exploring two methods to introduce the brain signal into the spinal cord. The first method uses our ceramic based multi-site electrodes for direct stimulation of the spinal cord. However, this method has not been tested in chronic freely moving preparations. A second method, currently being explored in the Giszter Laboratory is to use light pulses delivered through implanted fiberoptic light guides to uncage or activate caged neurotransmitters. We expect the fiber optic cables to flex within the spinal cord and be effective in a chronic freely moving animal prep. In addition, interfacing directly with the neurotransmitter system should allow us more fine-grained temporal and spatial activiation of the spinal circuits. Lastly, this system will allow for the selective use of inhibition and modulatory neurotransmitters, by uncaging inhibitory neurotransmitters. Inhibition and modulation state are difficult to control with electrical microstimulation.

Figure 2 : ideal neuroprosthetic interface
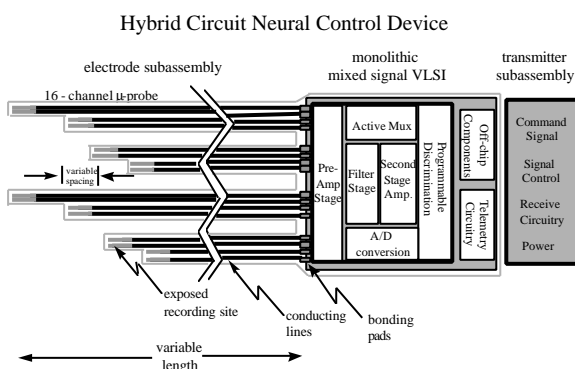
These initial studies provided a clear feasibility for developing a brain-machine interface. However, there are several obvious improvements that must be implemented before a multi-degree of freedom, multi-taksing adaptive interface can be realized. The most obvious is increasing the number of single neurons recorded from, maintaining these signals for long periods of time (i.e. years) and producing ultra-miniature, on-board signal processing. We expect our ceramic electrodes, with several advantages over micorwires and silicon electrode to provide this interface.

We are currently designing and testing a variety of neurorobotic interfaces that are based on the computational principles of the corticospinal interface and modular output structures in order to put the neural

command signal back into the spinal cord, rather than a completely external device. This framework will allow us to test various hypotheses about the computational processes that are carried out at this interface. Understanding this interface adequately may also enable future novel man-machine interfaces that are rapidly learned and intuitively managed by users. The expectations of a user of a direct interface about the structure, behavior and control principles of an effector system may be constrained by evolution. These expectations certainly develop to match the biological principles of control. Embedding these principles in humanoid robot parts and subsystems may be critical to good teleoperation or imitation.

4 Force-field motor primitives : the cortical and spinal bases?

4.1 Autonomous Systems At the Spinal Level

Neurobiologists have attacked the problem of the spinal cord systems in a variety of ways and an emerging scheme is becoming accepted. First, the spinal system (or ganglion chain of invertebrate species) can as a whole exhibit an ability to generate many patterns of neural activity required for motion in a feedforward fashion. The collections of neural circuits in spinal cord that produce these patterns of neural activity are referred to by neurobiologists as "central pattern generators". These "pattern generators" are operationally defined. The experimenter records neural activity elicited from a spinal cord severed from the brain, after the body has been paralyzed with for example curare or flaxidil. The spinal cord in this condition exhibits activity that would, if the muscles were not paralyzed, generate particular patterns of behavior. Of course feedback systems play a role in modulating, phaselocking and controlling the execution of the pattern and the resulting motion in the unparalyzed animal. By manipulating sensory feedback systems in paralyzed animals some of these controls can be examined. Taken together the experiments show there is a dynamical pattern organizing system in the low level spinal cord.

More recently, experiments in unparalyzed spinalized animals (rat and frog) have shown that there are a collection of separate modular circuits which are likely to be recruited by different pattern generators in different behaviors and perhaps also used by descending systems. These modules or primitives offer features which may considerably impact on the construction and organization of motion at spinal and higher levels. They have also inspired some robotic designs and they show clear correspondences to algorithms developed independently. Primitives may form a core design feature in biological motor competences which may assist in designing humanoid systems. We will review these primitives in more detail.

4.2 Spinal motor primitives

The neural code to control novel voluntary motor acts in man and other mammals are believed to be developed in the cortex, which uses the spinal cord for execution. It is now clear the interface between cortex and spinal cord is not a simple relay like a telephone exchange. Rather, many functions are embedded in the spinal cord and the spinal cord itself can execute complex, graceful and well planned movements. Indeed some aspects of motor learning are believed to be directly embedded in the spinal cord. The spinal cord is now thought to contain a modular organization of motor primitives.

Flexible adaptable motor patterns must presumably be elaborated from spinal central pattern generators (or CPGs, to be distinguished from "Control Plan Generators"). CPGs are circuits in spinal cord able to organize complex adaptable behavior even when the rest of the brain is removed (see Grillner and Wallen 1985, and Kiehn et al. 1997). The perspective we favor views the CPG as modular: the tasks of rhythm generation, pattern assembly and output adjustment are separated into modular components. The output of the CPG is believed by many to comprise motor primitives (Bizzi et al.1991, and Giszter et al. 1993). The motor primitives are seen as separate from the timing or rhythm generation system in this perspective. It is speculated that motor primitives may be reused in different tasks (Bizzi et al. 1994, Mussa-Ivaldi and Giszter 1992, Mussa-Ivaldi et al. 1995, Field and Stein 1997a,b, and Berkinblitt et al. 1989). The precise composition of these motor output elements has been a point of discussion.

Motor primitives are defined by us as fixed groups of muscles, their motoneuron pools and their afferent regulatory circuits, which together act as a unit to produce a fixed motor output.

At the simpliest level, a motor primitive could define a single muscle. On the one hand, this would allow complete flexibility to produce any motor output by combining activity across multiple muscles. However, this definition ids not useful as a computational framework and ,since the combinations of muscle activity are constrained by the joints and limbs, a more useful definition of motor primitives is a unit that acts together to represent a biomechanical force-field. A force-field, in this context, is a function that maps the forces generated in the limb, to the limb's configuration.   Ideally, these force-fields would combine to synthesize behaviors.

In 1991, based on microstimulation, we suggested that spinal circuits are in part organized into modules that support and control such force patterns in the limb. Remarkably, only a few types of force-fields were observed in the spinal cord. It was discovered in dual stimulations that these force-fields could be combined, in a simple linear fashion, by vector summation. It therefore seemed that a collection of spinal force-field primitives could form a basis for the construction of arbitrary force-fields or movements by recruiting them in combinations. It was hypothesized that such recruitment might be performed either by spinal pattern generators or by descending pathways. Subsequent studies by our laboratory (Giszter) and that of Bizzi have tested the organization and control of spinal primitives and have demonstrated that reflex behaviors are likely to be built and adjusted using these primitives (Mussa-Ivaldi et al. 1995, Tresch et al. 1999, and Kargo and Giszter, 1999a,b). Further, Tresch has shown that closely comparable force-field primitives exist in mammals (Tresch 1997, and Tresch and Bizzi 1999), and can be elicited from lumbar spinal cord in chronically transected adult rats. Modular muscle groups have also been observed in man, in cycling (Ting et al. 1998) and finger controls (Valero-Cuevas et al. 1999).

Ten years after the original publications the idea of the force-field primitive in frog and rat has survived several sets of critical tests and corroborating experiments. We have shown (Kargo and Giszter 2000) that in the precisely targeted multi-phase wiping reflex, both aimed trajectory formation and rapid on-line corrections to circumvent obstacles can be described as summation of primitives. Further, and perhaps most significant of all, in this and other studies we have found that each and every force-field primitive (and its underlying motor activity) which contributed to the behavior or to trajectory adjustments exhibited a common activation dynamics. Taken together, our data strongly support a modular control framework for these behaviors based on force-field primitives. Moreover, the data suggest that the force-field primitive forms a very particular kind of building block for movement, and this leads naturally to several testable hypotheses and learning schemes, which may be of wide reaching importance.

4.3 Force-field primitives as building blocks for motor learning

A final aspect of biology that designers of biomorphic robots have been sensitive to is development of skills in mammals (Drescher, Cog etc).   It has been suggested that organisms require a 'boot-strap" system so that new born organisms can immediately interact with their surroundings. Spinal circuits or primitives that we have explored could provide a basic competence, installed by evolution, from which efficient locomotory skill can be rapidly eleaborated after birth.
Force-field Primitives were defined as active force-fields that exhibited invariant force vector directions and magnitude balances over time, i.e. primitives had conserved structures through time (Giszter et al. 1993). Force-field primitives are the effect of ensemble feedforward and feedback circuits, which together act as a unit. These circuits generate a biomechanical multi-joint force-field.

We will follow the definition used in Kargo and Giszter 2000a: a force-field primitive is observed as a structurally invariant force-field over time. In its most general form, a force-field primitive would be a vector function of both position and velocity, modulated by a time varying scalar. We use r here as a configuration vector. The optimal coordinate system to express primitives (joint based, end-point, polar, cartesian) is not addressed here directly.

A primitive can thus be expressed in the form

$$F(r, \dot{r}, t) = a(t) f(r, \dot{r}) \tag{1}$$

F(r,r',t) is the observed field derived from a primitive which can be expressed as a scaling a(t) of a fixed field structure φ(r,r'), a function only of position and velocity.

In the experiments described above isometric testing was used (in isometric testing the limb is held rigidly clamped and not allowed to move) and in this case these descriptions can be written without velocity terms. Under such isometric conditions, at a single limb position, the force vectors generated by a primitive will

increase and decrease in magnitude with activation-deactivation dynamics of the primitive. At the same time, force magnitude ratios among the sampled limb positions remain constant, i.e. a primitive's time evolution can be expressed in the form a(t) • f(r). This is the description most often seen in neurobiological papers on primitives. Further this is the framework in which linear superposition was tested (Mussa-Ivaldi et al. 1994).

Using elements of this form Mussa-Ivaldi (1992) showed that arbitrary smooth force-fields could be approximated using combinations of appropriately chosen basis fields. In particular, he showed that an arbitrary smooth field F(r) could be generated by appropriate linear combinations of non-linear conservative and rotational basis fields :

$$F(r) = \sum_i A_i \boldsymbol{f}_i(r) + \sum_i B_i \boldsymbol{y}_i(r)$$

(2)

Where the field F(r) is constructed from conservative primitives $\Phi_I(r)$and rotational or circulating primitives $\Psi_i(r)$ modulate by the individual scaling parameters A_i and B_i

This theoretical framework was particularly interesting for two reasons. First and foremost these studies raised the interesting possibility that motion control by the brain and by spinal reflexes might synthesize the range of dynamic force patterns needed for real limb behaviors by recruiting, scaling, and combining the force-field primitives. Second, the theory showed arbitrary smooth fields could only be created by inclusion of non-conservative fields in the basis set (i.e. by inclusion of the second summation term $\Psi_i(r)$ in equation (2). Non-conservative fields could be used as energy sources. However, non-conservative fields might present a serious challenge to stability of the limb.

The repertoire of force patterns that could be constructed from primitives embedded in spinal cord all appeared to be solely conservative. Was this a failing or a design feature? Colgate and Hogan have shown that the class of conservative and passive systems could be coupled to other passive systems without danger of instability. The set of isometric fields which could be constructed by combinations of frog's primitives belong to this class. More recently the stability of classes of non-linear systems created from linear combinations of stable non-linear dynamical systems has also been demonstrated by Lohmiller and Slotine (1998). These results suggested it was feasible to design a simple principled motion control based around dynamic force-field primitives (Mussa-Ivaldi and Giszter, 1992). Primitives could be recruited and combined by sets of planning, patterning and control systems. Further, this design might offer specific stability guarantees during interactions with different environments. The laboratories of Mussa-Ivaldi (1997,1999), Bizzi and and Shadmehr among others, have examined biological motor learning and developed algorithms and descriptions of motor adaptation using this framework (see Shadmehr and Mussa-Ivaldi, 1994). The task for the motor system that is learned is the tracking of an arbitrary (usually straight) trajectory in different mechanical and visual environments. More complex tasks such as ball catching, ball bouncing or pole balancing could also be used.

Recently, we (Kargo and Giszter 2000a, and Giszter and Kargo2000) observed in our experiments that not only field structure appeared conserved in spinal behaviors. The temporal dynamics of activation of primitives also seemed conserved in different spinal behaviors and corrections. Our data from the frog currently support the idea that all primitives used in a (spinal reflex) behavior can be described as having similar temporal dynamics, i.e. are driven by a common function a(t). Similar results can be found in human motion (Burdet and Milner, 1998, Doreinger and Hogan 1998, Krebs et al. 1999). We currently believe that during motion control in the spinal frog any primitive θ_I can be expressed in the form:

$$\boldsymbol{q}_i(r, \&, t) = A_i.a(t)\, \boldsymbol{f}_i(r, \&)$$

(3)

Where a(t) is similar in all primitives, (resembling a fixed half cycle oscillation, or a fixed impulse response, such as a cosine packet) and A_i is a scaling factor. This structure constrains the method of construction of a time varying force-field F(r,r,t). The free parameters for generation of a behavior are the selection of the scalings A_i and the phasings $\tau_?$ of the component primitives. In this formulation the motion control field F(r,ṙ,t) is constructed as:

$$F(r, \dot{r}, t) = \sum_i A_i \, a(t + t_i) \, f_i(r, \dot{r})$$

<div align="right">(4)</div>

4.4 Data supporting the decomposition of behavior into primitives:

It is worth reviewing the data that supports the decomposition in equation 4 above. Giszter and Kargo have now amassed several sets of experiments supprting this framework. The strong predictions of the framework are (1) that new primitives added in to a behavior should exhibit dynamics a(t), (2) that primitives deleted from a behavior should cause loss of a patterns of the same temporal form a(t), and (3) that feedback modulation of primitives by natural or artificial stimuli should cause independent alterations of amplitude (A_i) or phase (τ_{Ii}) but not alter a(t). Each of these is observed in controlled experiments. In collision with obstacles frogs make trajectory corrections my adding in an amplitude modulated and externally triggered primitive with identical temporal pattern at different phases of the motion (Kargo and Giszter 2000a). In immobilized wipes, deletions of phases occur (Giszter and Kargo 2000 and Kargo 2000). Each deleted phase has similar temporal dynamics. In figure 3 we compare data from deletion of a hip extension primitive reported previously in Giszter and Kargo 2000 and data from deletion of knee extension reported in Kargo 2000. Kargo 2000 also describes effects of muscle and tendon vibration modifying proprioceptive feedback control signals. These manipulations alter the phase and amplitude of primitives while preserving temporal activation patterns.Moving beyond purely spinal motor patterns D'Avella and Bizzi have demonstrated decomposition of more complex brain driven (vestibular system) behaviors into force-field elements.

Use of the ideas of primitives to describe data in human motor learning has also been highly successful. In motor adaptation both Shadmehr and Thoroughman, and Matsuoka and Bizzi have shown varying degrees of generalization which can be described as adaptation of individual primitives or groups of primitives. Tresch and Bizzi (1999) in rat and in cat Lemay and Grill (see Giszter, Grill, Lemay, Mushawhar and Prochazka 2000) have shown force-field primitives exist in the mammalian spinal cord and are not a quirk of amphibian evolution. Primitives may also bear on imitation and perception of motion (Mataric, 200, Jenkins et al. 2000, and Billard and Mataric 2000). Perhaps the most interesting issue in this area as a whole is the relation among the various primitives in these several lines of research.

The framework expressed in equation 4 might be adequate for reflex behaviors, but seems limited. In motor skill development and learning it seems likely there should be the capacity to fully explore the degrees of freedom of the motor apparatus. Each primitive reduces the potential degrees of freedom implicit in the n muscles comprising the primitive to a single field. Similarly, the activation of the primitive is constrained to a fixed pattern. It is likely that very small time increments τ_i are not possible. The value of the primitive may be to place the motor learning system within an initial useful domain of the degrees of freedom, enabling initial adaptive survival skills to execute.

We speculate that from this initial constrained and adapted framework motor skill learning elaborates more flexible controls. Our perspective on the spinal primitives described above in lower vertebrates and in rats and cats is thus that these form a bootstrap system for motor development and motor learning. The idea is that a collection of primitives defined by evolution 'seed' subsequent motor development.
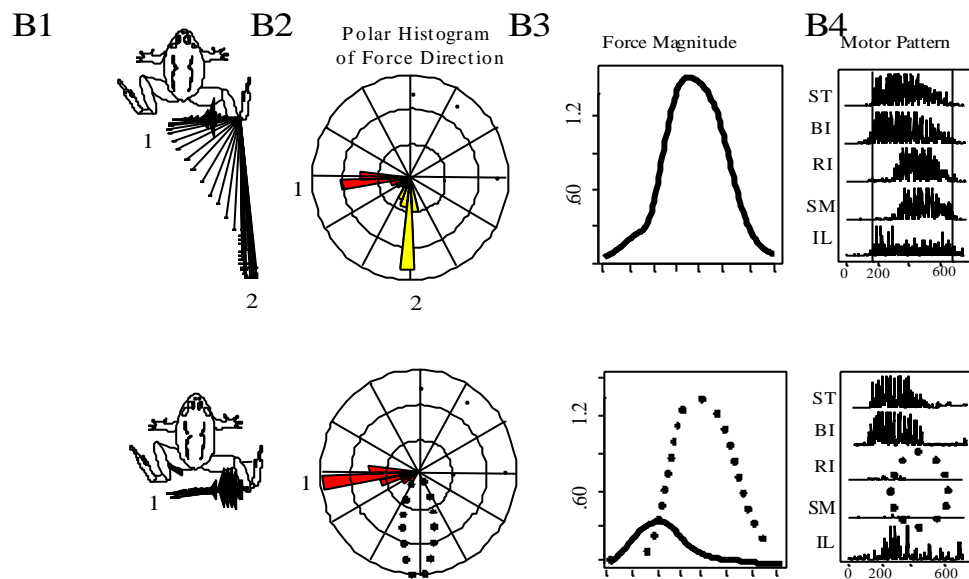
Hindlimb to Hindlimb Wiping

Construction from force field primitives -   deletion data

A: Hindlimb to Hindlimb Wiping : knee extensor deletion

A1   A2   A3   A4

Polar Histogram   Force Magnitude   Motor Pattern
of Force Direction
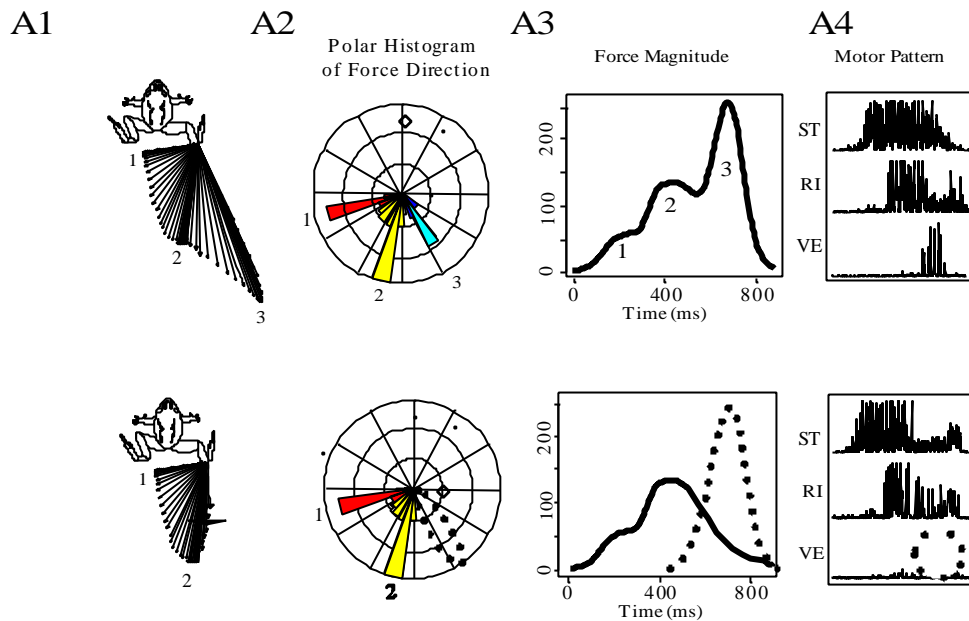
B: Hindlimb to Hindlimb Wiping : hip extensor deletion

B1   B2   B3   B4

Polar Histogram   Force Magnitude   Motor Pattern
of Force Direction

Figure 3 Examples of deletion of primitives redrawn from Giszter and Kargo 2000 and Kargo 2000:
A Deletion of the third knee extensor phase: A1 frog and forces recorded at ankle in relation to frog over time at a single limb posture (limb held immobile during force recording). A2 polar histogram of distribution of forces. Note three peaks reduced to two following deletion. A3 force magnitude over time for entire pattern (above)and after deletion (below). Shown dotted is the lost phase's magnitude. This was obtained by vector subtraction of deletion from full pattern. A4 electromyogram recordings of critical muscles: note deletion of vastus externus (VE) burst. B Deletion of the second hip extensor phase in a second frog in which knee extension was already deleted: B1 frog and forces recorded at ankle in relation to frog over time. A2 polar histogram of distribution of forces. Note two peaks reduced to one following deletion. A3 force magnitude over time for entire pattern (above)and after deletion (below). Shown dotted is lost phase's magnitude. This was obtained by vector subtraction of deletion from full pattern. A4 electromyogram recordings of critical muscles: note deletion of RI (rectus internus) and SM (semimembranosus) burst. ST: semitendinosus, BI biceps (illiofibularis), IL iliopsoas, RI rectus internus, SM semimembranosus.

The natural question arising from this idea, for which there is currently no answer, is whether the primitives that comprise such a bootstrap survive or dissolve in the subsequent motor learning and plastic changes. This is an important issue from the point of view of design of a neuroprosthetic interface: is a bootstrap still available in spinal cord or expected by cortex? Does each adult cortex expect different "individualized" primitives as a result of ontogenetic modifications? Does a neuroprosthesis need a similar bootstrap set of primitives? How should the neuroprosthetic device recruit and be allowed to alter primitives?

Below, we outline several biologically plausible methods for motor learning using primitives as we currently understand them. Our long range approach is to interact with the neural circuits at different levels of the system to distinguish among these hypotheses and develop theories of how, given a collection of 'seed' primitives, subsequent motor learning occurs. Means of interaction with the neural system to be used include (1) simultaneous recording from neurons across the motor systems of the brains and EMG combined with microstimulation during behavioral tasks, (2) computational models of spinal systems that are driven by neural activity recorded from the brain and (3) biomorphic robotic devices that are controlled by neural activity in the brain. Our basic hypothesis is that in mammals and especially humans, these seed primitives are altered during learning by interactions at the brain-spinal cord interface. If this can be shown to be true, the question then remains, how are these basic seed primitives altered. Below we outline several ways the brain may modify the seed primitives during learning to create the flexible adult structure. Each is likely to have different strengths and weaknesses as a method for learning limb movements. Tasks to be learned by each, would include standard trajectory learning, and also gradual expansion of the ranges and independence of the degrees of freedom in the limb that are used in different combinations. By exploring several different schemes for motor learning using the interaction methods outlined above we will not only get closer to understanding how the mammalian neural system develops, we will develop a structure for possible methods of motor learning in humanoid robotics and with neuroprosthetic interfaces to these.

5 Model schemes for corticospinal motor learning using primitives based on the biological system

The concept of the force-field primitive as an element of motor behavior can be extended beyond spinal cord. To us it seems likely that motor skill learning and recovery from injury, or mastery of novel neurorobotic interfaces all require that the CNS devise an adequate strategy of creation, control and combination of cortical and/or spinal motor primitives to perform the task.

In addition, the stability of a force-field structure is due, at least in part to a maintained balance of activation of the n component muscles and their generated end-point forces M_i. Each muscle contributes torques which sum to drive motion or determine end-point force. A primitive measured in isometric conditions $\phi(r,t)$ can be expressed as

$$\boldsymbol{f}(r,t) = A(t)\boldsymbol{f'}(r) = \sum_{i=1}^{n} a_i(t)M_i(r)$$

(5)

Each $M(r)$ represents the force-field due to an individual muscle within the primitive (see Giszter et al. 1993, Mussa-Ivaldi et al. 1994 and Kargo 2000). Any alteration of the ratios of $a_i(t)$ at each time t will alter the field structure $\phi(r,t)$ in such a way that a simple scalar function for scaling (of the form A(t)) does not hold. Thus alteration of ratios of $a_i(t)$ by either descending or by feedback controls will violate the stability of the isometric field structure $\phi'(r)$. This stability of isometric force-field structure is one that we observed consistently in spinal animals. However, clearly the most flexible use of the motor apparatus is to control each $M(r)$ independently by varying $a_i(t)$ freely.

The modular scheme of elements in spinal cord is shown in the diagram in figure 3. How descending or feedback controls could violate equation 5 if used in motor learning is discussed below.
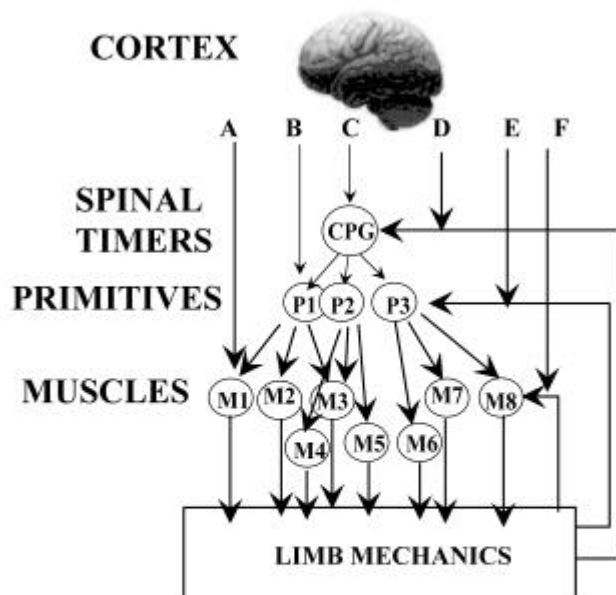
Figure 4. Sites of potential descending control of the CPG and spinal primitives. (see text) Feedback must be distributed in a constant ratio to muscles in a primitive if the description in equation 1-4 are to hold. Descending controls could act on spinal motor structures A muscles B primitives or C pattern generators or influence these indirectly by midifying how sensory information is used be these (D-F). Each has different implications for spinal modularity or control through a neural prosthesis.

In figure 4, at the spinal level a timing system recruits primitives of fixed intrinsic dynamics based on task and sensory conditions. This in turn recruits primitives which recruit muscles and feedback pathways in fixed ratios to one another. Motor learning can be expressed as locating a dynamic time evolution of the total force-field in the limb to complete the task. The cortex learning task is therefore to choose among the many free parameters A through E in order to create this force-field time evolution. The cortex could in principle act on individual muscle's recruitment (A), on primitives recruitment (B), on the operation of the timer or CPG (C) or on feedback effects on each of these (D through F). Sensory information from the plant or limb is used locally in the primitive (E,F), as well as in control of timing and recruitment of primitives (D). Descending controls are known to have the capability of modifying each of the connections shown, as indicated by the descending lines. Differential alteration of the balance of components of primitives may alter either their field structure or dynamics. In contrast, alteration of other interactions (B,C,D,E) will leave primitives unaltered.

According to the definition used here, a force-field primitive is observed as a structurally invariant force-field over time. In its most general form, a force-field primitive need not represent an assembly of muscles and

feedback circuits. The formulation can equally represent a single muscle's contractile force and torque delivery (i.e. n=1 in equation 5 above), which would be a function of both position and velocity. Ultimately the collection of muscles coalesced into primitives by the spinal cord should also be available as independent control elements. We presume this occurs after sufficient motor learning starting from the spinal primitives. Further this level of 'fractionation' may rarely by required in a motor task and may be avoided where possible.

5.1 Approximation of a dynamic field by adding independent muscle controls from cortex

5.1a Framework
A method of learning by recruiting individual muscles is the most flexible framework, fully utilizing all degrees of freedom in actuation of the limb, especially when combined with control of feedback (F). However it presents ill-posed problems to the cortical systems. Interestingly only man and the higher primates show much direct cortical control of muscles, and then most frequently in distal joints (e.g. fingers). This method is actually a special case of 5.2 below in which primitives are individual muscles and will be discussed in more detail in that section.

5.1b Neurorobotic implications
For our goal of designing a neuroprosthesis, rapid learning of control actions in a high dimensional actuation space by cortex would make the interface design very straightforward, potentially requiring the bare minimum of structure in the interface.

5.2 Approximation of a dynamic field by modulating recruitment of primitives directly using descending pathways

5.2a Framework
This method of approximation of time varying force-field patterns is relatively well understood. Combining force-field primitives to synthesize fields has been examined in some detail, (e.g. Mussa-Ivaldi 1992, 1997). The major variance with earlier work we see in the recent experimental data is the presence of fixed primitive dynamics and duration. This constrains approximations to the form in equation 4, modulating A_i for each primitive. However it is not clear if descending controls are able to modulate or time dilate the activation patterns of primitives. This would violate equation 4 and this is discussed below in section 5.4. The mechanism of field synthesis from unaltered primitives discussed here could be implemented either (1) by descending controls that completely 'take charge' by recruiting primitives directly and perhaps suppressing other spinal systems, or (2) by uniformly modulating sensory feedback gains within a primitive recruited by other means.Biologically, both types of mechanism are likely.

5.2b Neuroprosthetic implications
If learning mechanisms at the cortical level depend on sculpting spinal force-field primitives into driving fields for limb motion, a neuroprosthesis may have to provide a collection of similar elements and controls in order to expect rapid learning. This is our current 'best guess' at the structure the interface of a neurorobotic controller should take with controls described in 5.3 and 5.4 made available in parallel.

5.3 Approximation of a dynamic field by altering timing and rhythm generators that recruit primitives

5.3a Framework
This method of approximation amounts to fitting dynamic field structures by setting both A_i and τ_i in equation 4. Descending controls from the cortex are known to act on spinal timing systems. When this timing modulation

is combined with independent amplitude modulation this may provide an extremely flexible approximation technique. Schaal and Sternad have begun to explore this framework for human limb motor patterns.

5.3b Neuroprosthetic implications

Locating and extracting separate timing and amplitude signals for primitives from cortical populations is still an experimental problem. Whether the same motor cortical cells modulate both amplitude and timing is unclear, and whether some populations of cortical cells are more specialized is not clear. Combined with 5.2 this mechanism would provide a framework in which all behaviors were built from the bootstrap primitives following equation 4 above. Without further augmentation these alone would constrain the tasks available severely. However a stable and adaptive interface adequate for many kinds of tasks might be possible using this simple structure in a prosthetic or teleoperated humanoid.

5.4 Approximation of a dynamic field by altering temporal dynamics of primitives: dilations of activation

5.4a Framework

This approach supposes the cortex can use primitives in a manner akin to wave packets and wavelets, time dilating them to generate different shapes in time (i.e. altering individual $a_i(t)$ in equation 4), presumably in a continuous rather than dyadic manner. It seems likely the cortex can do this, either by a direct control of spinal dynamics, or by splitting distributed circuits involved in primitives and then phase-shifting subpopulations that are drawn from the total assembly of circuitry associated with primitives as discussed in 5.6 below. The basic force-field patterns of individual primitives ($\varphi(r)$ in equation 4) remain unaltered. Rather their individual timing is more flexibly adjusted as the original constraints on $a_i(t)$ are relaxed.

5.4b Neuroprosthetic implications

Designing an interface that adaptively splits primitives or dilates them while providing stable reflex or local functions seems to the authors a serious design problem unless the neuroprosthetic or humanoid control is a highly distributed and parallel. At what point splitting populations should occur and how to assign the population of neural cells recorded at the interface to the population of robot controls is a problem to which we currently have no solution but one which may need to be solved.

5.5 Approximation of a dynamic field by altering feedback patterns within primitives.

5.5a Framework

Altering the feedback patterns within a primitive will alter force-field structure (see equation 5: $a_i(t)M_i(r)$ will be differentially affected). This method of approximation may provide a way of generating differing sets of primitives in different contexts. Effectively, different approximation bases could be loaded by this mechanism. For example the cortex could modify primitives and spinal circuits like a local task expert by modulating the pattern, sign and gain of feedback. Many physiologists might argue that this is clearly the case, and that the notion of fixed primitives is a result of looking at amphibians or spinalized mammals and simpler tasks. However the degrees of freedom problem cannot be forgotten in this context.

An interesting special case of varying sensory feedback patterns is the case in which local feedback gain is modulated differently in response to configuration but between muscle ratios of feedback gains are not. The 'muscle synergy' is preserved though the force-field primitive is not since the spatial modulation of the synergy amplitude is changed. This has specific implications. In this case although field structure $\phi'(r)$ in equation 5 alters, at any point the two vector fields direction vectors do not. As an example, an isotropic attractor field could have its potential basin altered from a isotropic conical basin to an isotropic parabolic basin by this special type of feedback modulation. This feedback modulation preserves the field equilibrium or attractor, if it exists. However, it obviously guarantees nothing for field combinations.

5.5 b Neuroprosthetic implicationsWe do not know if this method of adjustment occurs as a result of descending controls. It sems not unlikely. However it may present serious difficulty to the designer of the neurorobotic interface. Locating signals of sufficient bandwidth in the recorded brain areas to allow precise modulation of basis sets and to allow learning may be very difficult with current neuroengineering technologies.

5.6 Splitting or creating new primitives

5.6a Framework

In spinal cord it is clear that primitives are probably composed of populations of distributed and redundant circuits. In the biological system it is therefore possible that splitting these populations or recombining subpopulations in novel ways is a means to construct novel field structures, or achieve time dilations. Under descending control novel muscle balances and force-field structures could be created as discussed in 5.1. By modulating and redirecting sensory feedback novel force-field structures could be created as discussed in 5.5. Splitting and phase shifting a population of neurons associated with a primitive might be one means of generating time dilation of primitives under descending control as described in 5.4.

5.6 b Neuroprosthetic implicationsUnderstanding spinal plasticity and flexibility sufficiently well to emulate it in a neural prosthetic design so the prosthetic and nervous system co-adapt to a seamless whole is an interesting but currently unattainable goal. However it may represent the key to design of both prosthetics and biomimetic and humanoid robots that truly emulate biological skills.

6 Conclusions

Primitives and population coding phenomena offer exciting prospects for understanding biological learning and designing new type of neuroprostheses and interfaces. They may relate closely to the structure and devlopment of internal models (see reviews of Mussa-Ivaldi 1999 and Kawato 1999). Similar ideas have been applied in machine learning frameworks (Mussa-Ivaldi 1992, Cannon and Slotine 1995, Mataric et al. 1998, Sanner and Slotine 1995 and Williamson 1996). However serious gaps in both our neurobiological and theoretical knowledge exist. Such gaps may be addressed by careful analyses of these frameworks for approximation. In parallel, experiments using neuroprostheses, modular interfaces and modular neurorobotic effector devices may help determine both what the cortex expects of its output and what the cortex can learn. The notion of a gradually accreting and changing set of primitives as a means to use and explore a very high degree of freedom actuation space is attractive, but beyond the findings in spinal cord is still a largely unproven but testable idea.

Acknowledgements

References

Berkinblit, MB, Feldman, AG, and Fukson, OI (1986) Adaptability of innate motor patterns and motor control mechanisms. Behavioral and.Brain.Sciences 9: 585-638.

Billard A and Mataric, MJ, (2000) "Learning human arm movements by imitation: Evaluation of a biologically inspired connectionist architecture", Proceedings, First IEEE-RAS International Conference on Humanoid Robotics (Humanoids-2000), MIT, Cambridge, MA, Sep 7-8, 2000.

Bizzi E, Giszter SF, Loeb E, Mussa-Ivaldi FA, Saltiel P (1995) Modular organization of motor behavior in the frog's spinal cord. Trends in Neuroscience 18(10):442-6.

Bizzi E, Mussa-Ivaldi FA, Giszter SF (1991) Computations underlying the execution of movement:a novel biological perspective. Science 253: 287-291.

Burdet E. Milner TE (1998) Quantization of human motions and learning of accurate movements. Biological Cybernetics 78(4):307-18.

Cannon, M and Slotine JJE (1995) Space frequency localized basis networks for nonlinear estimation and control. Neurocomputing 9(3). Chapin JK 1998 in Neuronal population recording (ed. Nicolelis MA) 193-228 CRC Press Boca Raton.

Chapin, JK, Moxon, KA, Markowitz, RS, Nicolelis, MAL (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex, *Nature Neurosci*ence, 2(7), 664-670.

d'Avella A. Bizzi E (1998) Low dimensionality of supraspinally induced force-fields. Proceedings of National Academy of Science 95(13):7711-4.

Dietz V Colombo G Jensen DM Baumgartner L. 1995 Locomotor capacity of spinal cord in paraplegic patients. Annals of Neurology 37:574-582

Doeringer JA and Hogan N (1998) Intermittancy in preplanned elbow movements persist in the absence of visual feedback. Journal of Neurophysiology 80(4):1787-99.

Drescher GL (1991) Made-up minds. MIT Press, Cambridge MA

Giszter, SF, Mussa-Ivaldi, FA, and Bizzi, E (1993) Convergent force-field organized in the frog's spinal cord. Journal of Neuroscience 13: 467-491.

Giszter SF, Grill W, Lemay M, Mushahwr V and Prochazka A (2000) Intraspinal microstimulation: techniques, perspectives and prospects for FES. In : Neural Prostheses. Ed KA Moxon and JK Chapin, CRC Press (in press)

Giszter SF, Kargo WJ and Davies MR (2000) Augmenting postural primitives in spinal cord: Dynamic force-field structures used in trajectory generation. In Biomechanics and Neural Control of Movement, eds. Winter J and Crago PJ. Springer Verlag.

Giszter, S.F., Mussa-Ivaldi F.A., and Bizzi E. (1993) Convergent force-fields organized in the frog's spinal cord. Journal of Neuroscience 13:467-491.

Giszter SF and Kargo WJ (2000) Conserved temporal dynamics and vector superposition of primitives in frog wiping reflexes during spontaneous extensor deletions. Neurocomputing. 32-33: 775-783.

Giszter SF Moxon KA Rybak I and Chapin JK (2000) A neurobiological perspective on humanoid robot design IEEE Intelligent Systems, Jul 2000, pp 64-69

Grillner S, Wallen P (1985) Central pattern generators for locomotion, with special reference to vertebrates. Ann Rev Neurosci 8: 233-261.

Hetke, J.F., Lund, J.L., Najafi, K., Wise, K.D., Anderson, D.J., Silicon ribbon cables for chronically implantable microelectrode arrays, IEEE Trans. Biomed. Eng. 41, 314, 1994.

Jenkins, OC, Mataric, MJ, and Weber, S "Primitive-Based Movement Classification for Humanoid Imitation", Proceedings, First IEEE-RAS International Conference on Humanoid Robotics (Humanoids-2000), MIT, Cambridge, MA, Sep 7-8, 2000.

Jones, K.E., Campbell, P.K., Normann, R.A., A glass/silicon composite intracortical electrode array, Ann. Biomed. Eng., 20, 423, 1992.

Kargo, WJ and Giszter, SF (2000a) Rapid correction of aimed movements by summation of force-field primitives. J. Neurosci. 20(1):409-426

Kargo, WJ and Giszter. SF (2000b). Afferent roles in hidlimb wipe-reflex trajectories: free-limb kinematics and motor patterns. *Journal of Neurophysiology* 83:1480-1501.

Kawato M 1999 Internal models for motor control and trajectory planning Current Opinion in Neurobiology 9:718-727.

Kiehn, O., Hounsgaard, J., and Sillar, K (1997) Basic building blocks of vertebrate spinal central pattern generators. In Neurons, Networks, and Motor Behavior (P.Stein, S.Grillner, A.Selverston, D.Stuart eds.), pp. 47-60. Boston, MA: MIT Press.

Krebs HI, Aisen ML, Volpe BT, Hogan N (1999) Quantization of continuous arm movements in humans with brain injury. Proceedings of National Academy of Science 96(8):4645-9.

Lohmiller W and Slotine JJ 1998 on contraction analysis for non-linear systems Automatica 34:683-696

Loeb G 2000

Maes P 1990 Situated agents can have goals Robotics and Autonomous Systems 6:49-70

Massaquoi, SG and Slotine, JE. (1996) The intermediate cerebellum may function as a wave-variable processor. Neuroscience Letters 215: 60-4.

Mataric MJ, Williamson MM, Demiris J and Molan A (1998) Behavior-based Primitives for Articulated Control. Proceedings Simulation of Adaptive Behavior 1998, Zurich, Switzerland, MIT Press.

Mataric, MJ "Getting Humanoids to Move and Imitate", IEEE Intelligent Systems, Jul 2000, 18-24.

Matsuoka Y Soc Neurosci Abstr 1999

Moxon, K.A., Multichannel elctrode design: Considerations for different applications, in *Methods for Neural Ensemble Recordings*, Nicolelis, M.A.L. Ed., CRC Press, New York, 1999, 25.

Mussa-Ivaldi 1999 Modular features of motor control and learning. Current Opinion in Neurobiology 9:713-717

Mussa-Ivaldi FA (1997) Nonlinear force-fields: a distributed system of control primitives for representing and learning movements. *Proceedings of the IEEE International Symposium on Computational Intelligence Robotics and Automation, Computer Society Press* 84-90

Mussa-Ivaldi FA (1992) From basis functions to basis fields: vector field approximation from sparse data. Biol Cybern 67(6):479-89.

Mussa-Ivaldi FA, Giszter SF (1992) Vector field approximation: a computational paradigm for motor control and learning. Biological Cybernetics 67(6): 491-500.

Mussa-Ivaldi FA Giszter SF and Bizzi E 1994 Linear superposition of primitives in motor control. Proceedings and National Academy of Sciences . 91:7534-7538

Rybak, I. A., Giszter. S., Wieland, S. J., and Chapin, J. K. (1999). Computational modeling of the spinal cord neural circuitry. Soc. Neurosci. Abstracts, 759.4, 1907.

Sanner R and Slotine JJE. (1995) Stable robotic learning using 'neural networks.' Neural Computation. 7(4).

Shadmehr R and Mussa-Ivaldi FA 1994 Adaptive representation of dynamics during learning of a motor task Journal of Neuroscience 14:3208-3224.

Sherrington, C.S. (1910) Flexion-reflex of the limb, crossed extension reflex and reflex stepping and standing. Journal of Physiology. 40(1910): 28-121.

Thoroughman KA and Shadmehr R Soc Neurosci Abstr 1999

Tresch, MC and E Bizzi. (1999) Responses from spinal microstimulation in the chronically spinalized rat and their relationship to spinal systems activated by low threshold cutaneous stimulation. Experimental Brain Research 129:401-416.

Vicario, DS and Ghez, C (1984) The control of rapid limb movement in the cat. IV. Updating of ongoing isometric responses. Exp Br Res 55: 134-44.

Williamson MM (1996) Postural Primitives: Interactive Behavior for a Humanoid Robot Arm. Proceedings of Simulation of Adaptive Behavior 1996, Cape Cod, MA.

Wise, K.D., Angell, J.B., Starr, A., An integrated-circuit approach to extracellular microelectrodes, IEEE Trans. Biomed., Eng., 17(3), 1970.