

Modelling spinal organization of motor behaviors in the frog

Simon Giszter,
Dept. Brain and Cognitive Sciences,
MIT, 77 Massachusetts Ave,
Cambridge, MA 02139, USA
giszter@ai.mit.edu
giszter@18.88.0.69

Abstract

The goal of this paper is to present (1) results of a simulation of frog wiping behaviors based on recent data, (2) a brief review of experimental and theoretical approaches to understanding these protective frog behaviors and how these influence the simulations, and (3) outline some possible future directions. The wiping behaviors can be elicited from the spinal cord in isolation from the brain. All the basic decision making and computational apparatus for these behaviors thus resides low in the control hierarchy and can operate independently of supervision by the brain. It is shown that a Maes network (Maes, 1989) driving force field primitives can mimic many aspects of the frog wiping responses. This simulation draws on recent discoveries in the frog spinal cord. The simulation is able to deal with multiple stimuli, choose responses and show a type of exploratory response. This network structure has much to recommend it as a plausible framework for modelling the spinal cord. In particular the global variables available in the behavior network may have some correspondence to spinal observables such as sensitization, depression, habituation, and actions of neuromodulatory substances. The responses of the network are quite 'life-like' and fulfill the goal of providing an impetus and direction to experiment. It is unlikely that all the elements necessary to simulate spinal behaviors are present in the simulation. However experimental data that suggests possible additional network modifications is discussed. The possible relationships of reactive local control and descending systems in the spinal cord is discussed in the context of (1) the simulations presented here and (2) the basis field approach described in Mussa-Ivaldi and Giszter 1992. An understanding of the organization of wiping behaviors, their organization and relations to other motor control systems may be useful in the synthesis and design of artificial systems and lifeforms.

Introduction

Motor behaviors of animals might be broadly classified into those concerned with moving the organism about the environment and those concerned with manipulating objects in the environment at a location in space. Much recent work in the area of artificial life has focussed on the former. This paper looks at a type of external manipulation. Many animals survive well without limbs. Others manipulate objects in the environment through fluid vortices rather than by mechanical contact. Much manipulation by animals is actually performed 'on board' in the mouth or bowel (for example the stomatogastric systems of lobsters). This is particularly true of pelagic marine creatures. Many marine mammals have lost any grasping capabilities they may have possessed. Among the pelagic marine lifeforms the cephalopods possess sophisticated external manipulation. Manipulation is more generally associated with creatures dwelling on surfaces. For example, even some surface dwelling protozoa can construct a protective carapace with pseudopods. However the best examples of external mechanical manipulation, besides the cephalopods, are clearly among the arthropods and vertebrates.

Manipulation introduces several sets of sophisticated computational problems not directly associated with other behaviors. For jointed appendages these problems can be clearly defined, and this may be a computational advantage for animals with jointed limbs. Jointed appendages must be positioned in space so that the tool or effector (for example the hand) is in the appropriate relationship to the manipulated object with an appropriate mechanical impedance. This controlled configuration of the limb must occur under the constraints imposed by the relations of the sets of joints and their mobilities. This requires the solution of several non-linear problems which may often be ill-posed. A typical computational scheme goes as follows: Firstly, the coordinates of the manipulated object in sensor derived coordinates must be somehow transformed into a position in body centered coordinates. Next a limb configuration and trajectory bringing the effector or tool to the object must be chosen. Finally a pattern of actuation of muscles must be chosen to generate this motion under the sets of environmental constraints in effect. These are often ill-posed problems. That is to say, in most limbs there are excess degrees of freedom in both configuration and actuation. Thus the choice of an orientation of the limb must be made from a family of several sets of possible solutions. At the same time the impedance properties (or mechanical interactions) at the interface between the object and the effector must be controlled.

2: Frog spinal behaviors:

A class of manipulative behaviors which involve fewer hypothetical transformations than external manipulation using vision are the grooming and protective reflexes. In general these entail movements to the body surface and not to arbitrary points in external space. The task is thus essentially that of closing a kinematic chain. This may be a simpler task than general manipulation of external objects. At the same time many similar problems of configuration of the limb and trajectory control must be solved in these behaviors. Grooming behaviors have been examined in a variety of animals. Remarkably, many of the frog's grooming and protective reflexes can be elicited in the spinal cord alone.

Frog wiping behaviors can be simply elicited in the spinalized frog using a small piece of paper soaked in acid or a local electrical stimulus. The frog removes the irritant with a graceful coordinated motion of the hindlimb involving some 7 degrees of freedom. To do this it solves the aforementioned ill-posed problems.

In different areas of skin several qualitatively different reach strategies are necessary for removal of the stimulus. In this case the frog chooses strategy adaptively. Some of these skin areas overlap, so that one of several possible strategies must be chosen. The choice in these region is based on the overall configuration of stimuli on the body and the history of responses executed prior to the current response. These areas in which a configuration can be chosen from one of several disjoint sets of joint angles have been called transition zones.

A second adaptation which occurs for multiple stimuli on the skin is blending of responses. In blending responses the frog adapts the kinematics of two different movements elicited in different areas of skin in order to reach both stimuli in a single movement.

Finally, within a strategy the frog can make small adjustments of configuration, presumably in order to optimize the relationship of the wiping effector to the irritant. These adjustments are continuous across an area of skin.

The circuitry that is used to perform these sets of movements is fully contained in the spinal cord. This has made the spinal frog an attractive preparation with which to examine mechanisms of limb positioning and control. Understanding the neural architecture which supports this set of behaviors may help in the design of artificial creatures and mechanisms.

3: The force field decomposition and movement primitives:

Recently we have shown that many postural phases and trajectory fragments of spinal behaviors can be simply expressed as force fields (Mussa-Ivaldi, Giszter and Bizzi, 1990, Bizzi, Mussa-Ivaldi and Giszter 1991, Giszter, Mussa-Ivaldi and Bizzi, 1991a,b, 1993, Giszter 1992). A force field is a continuous function relating a position to an associated force. A force field thus summarizes the force that will be generated in the limb at any point in the reachable space. The fields measured were generally obtained from non-redundant arrangements of the limbs. To measure such force fields a force transducer was attached to a spinal frog's leg close to the ankle. The limb was positioned at a location in space. The skin was then stimulated. The result of this was an attempt by the frog to move of the leg. The forces generated were recorded by the sensor. The leg was then positioned at a new location. The procedure was applied repeatedly. A set of samples of a 2 or 3 dimensional force field was thus collected.

The force fields associated with three different tasks executed by a hind-limb of the spinal frog have been examined in detail. The force fields associated with these behaviors had fixed structures. These fields were hypothesized to represent movement primitives, by which we mean the smallest fragments of meaningful behaviors. We have also found that even spinal circuitry microstimulation produces fields that are remarkably similar to the fields underlying behavior that are generated by skin stimulation. In summary, the main finding described here was that invariant force field patterns underlay each of these behaviors.


The spinally generated force fields were considered to include both position and force control strategies. The spinal cord thus appears to support both equilibrium point control strategies specifying absolute location and other force field strategies specifying relative movements or contact forces. This suggests that there may be no explicit trajectory planning in the spinal cord. The path emerges from the interaction of limb dynamics with a force field. This force field results from both environmental constraints and from the neurally generated force field. It has

now been shown that in several instances these fields may exhibit vector summation when activated in parallel (Bizzi et al, 1991).

It seems that reflex spinal behaviors of some kinematic complexity might be constructed using simple fixed force field strategies. These force field or movement primitives might be thought of as permanently instantiated schemas for motor behavior (see Arbib and Cobas, 1991)

4: Modelling force field responses:

If one hypothesizes that force field primitives can be taken as behavioral movement primitives then the critical issue for the nervous system is to choose sequences and combinations of these primitives which are both relevant to the situation and successful in satisfying the behavioral goals. The combination of force fields by vector summation appears to be an interesting mechanism to generate adaptive responses, but some summations will prevent attaining the behavioral goal of either force field behavior in the 'average response'. It thus appears that some method of choosing allowable combinations of force fields is needed. The task of the driving network for these primitives is then to select, sequence, and activate sets of primitives while preserving their behavioral relevance. Examining such models of the generation of wiping behaviors can fulfill several purposes: (1) the model may suggest new experiments (2) the experimental data can form a critical test of model function (3) models constructed during this theory/data dialog which are not accurate models of the biological system may nonetheless be useful to artificial system development. One iteration of the data/model interaction will be discussed below.



Stimulus:

$$S_i = \{ x_i, y_i, A_i \}$$

where x, y are position and A amplitude.

Conditions:

$$c_j = \sum_i g_j(S_i)$$

Maes Network:

$$A_{i+1} = M * A_i + B * C_i$$

where A is the vector of node activations and C the vector of conditions.

The entries of Matrices M and B are determined from the boolean state of the conditions ($c_j < 1$ or ≥ 1), and the state of each node (executable/nonexecutable).

Behavior nodes drive k force field primitives f_1 to f_k . Each node is associated with one or more primitives.

The force field strength:

$$p_k = \sum (\text{associated executable activations of nodes})$$

Limb field:

$$F_i = \sum_k p_k f_k$$

Figure 1:

Overall network: Left: Diagrammatic representation. Skin stimuli set conditions, stimulus values are summed after modulation by spatial gradients. Conditions activate and control the operation of behavior nodes in the behavior network. The behavior nodes which are both executable and non-excluded activate the collection of force field primitives. Force field primitives are summed vectorially to determine the limb kinematics (3). Right: A brief summary of operations occurring at each level of the network. The principal non-linear operations are contained in the Maes network and exclusion system.

The model structure that will be described here consists of three parts. The first is a simulation of areas of skin and associated sensory systems. These areas have associated overlapping sensory fields which transform skin stimulation into the activation of a layer of nodes which act as the conditions or propositions for input to an augmented Maes network. This portion of the simulation allows the stimulation of a path of 'skin' to activate condition nodes which would lead to several different possible behaviors. See Figure 1. The choice of behaviors is resolved by a Maes network, which forms the second part of the simulation.

The choice of a Maes network for the core of the simulation was based on several criteria. Firstly the Maes network in its original form chooses a course or sequence of actions based on the global situation while using local interactions. Thus adaptive behavior is an emergent property of the interaction of network and environment. Secondly, there is potentially a simple relationship between the network structure, its global parameters and neural entities such as pattern generators and modulators in the spinal cord. Thirdly, the one to one correspondence of nodes and actions seemed suitable for the activation of force field primitives.

The behavior network design used here is built upon the organization originally proposed by Maes 1989 (and see Maes 1991a,b). It is worthwhile summarizing the core operation of the basic network first, which is rather elegant.

The overall flow of activation propagated by the links leads to increasing activation of those behavior nodes that are currently executable and also increasing activation of those nodes that will most likely shortly become executable as a result of environmental variations or current actions. Chains of nodes are therefore readied in anticipation of action. Thus threads and trees of actions and potentially useful actions emerge from the interactions of the network with the current conditions.

The augmentations added to the Maes network in the simulations here are described in detail in Giszter 1993. They consist of adding (1) extended durations for behaviors, (2) concurrent operation of behaviors, (3) intensity variation of behaviors, (4) an exclusion layer to the basic network and (5) allowing that the network links to vary continuously rather than boolean switching. The first three are essentially changes in the network output. The last two alter the structure and operation of the network.

The exclusion system controls access to the effectors. This mechanism prevents conflicting behaviors while allowing activation buildup to proceed unaffected. Concurrent execution is often desirable. Concurrency can be disastrous for some combinations of behaviors however. In the frog two rostral and caudal wipe types cannot be combined and still retain any functional value. The exclusion network in this simulation is very simple. It excludes behaviors which conflict with a currently executing behavior. Exclusion occurs on a first come first served basis. There is no pre-empting although this can be achieved very simply. The exclusion network design avoids starvation and deadlock problems as a result of handcrafting and because it is a particularly intelligent implementation. However the principal goal of the implementation used was simply to ensure that access to the effector system was separated from the process of activation buildup from which 'planning' emerges. The reason for choosing explicit exclusion over the addition of novel 'virtual' conditions not anchored to the external environment is (1) that this would destroy the biological analogy of conditions to afferents, and more importantly (2) that such 'virtual' conditions would alter the network operation substantially for behaviors that were mutually exclusive. Thus for example adding a 'virtual' condition to a flexion behavior could have several results. Flexion normally follows extension. Thus it is a natural candidate for receipt of activation from extension but clearly it should also be excluded by extension, as their goals conflict. One

could use a virtual condition in order to suppress the flexion. This would not allow flexion to ever become executable during periods of extension. However flexion must become executable in order to promote activation of any third behavior. For example suppose the next behavior were a locomotory push: its activation would be independent of the excluded flexion since the flexion would remain non executable so long as it was excluded. The implementation here uses a separate exclusion system that allows such interactions. Thus more stable behavior chains and limit cycles are anticipated in this network design. The exclusions proposed here were implemented as a simple message passing between the behavior nodes. This implementation thus preserves the encapsulation of behaviors in nodes.

Virtual conditions in the type of network described here would allow hierarchical structuring and subsumption of networks of this type into larger structures by descending control of the executable / non-executable switch. This mechanism might be useful in the broader context of descending control of spinal motor systems (see below).

The Maes network augmented in this way was used to drive force field summation which formed the third and final element of the simulation. The limbs overall force field was used to calculate motions of a three link limb. This limb was given the ability to remove all or part of irritant stimuli 'applied' to the model by moving the last link (the foot) across them in a defined direction. The effectiveness of the 'wipe' was proportional to the velocity of the motion. The overall simulation structure is summarized in figure 1.

5: Simulation results:

Although the network and limb were set up to deal with three-dimensional data the principle results of this simulation are well expressed in two dimensions. The simulation shows a very 'lifelike' set of response patterns. The main results demonstrated are (1) successful stimulus removal (2) switching between multiple stimuli (3) choice of the strongest irritant on first wipe (4) blending of responses for multiple stimuli that preserves function (5) trial to trial kinematic variability of response (6) apparent searching and exploration to reach 'difficult' stimuli. This latter exploration comprised a variety of activation patterns which caused the limb to cover more workspace and use trajectories not expected in the initial hand coded design.

Stimulus removal for a single stimulus:

The limb was moved close to a posture specified by a placing field and then swept the area of the stimulus. For weak stimuli a small and ineffectual 'intention' like movement could precede a full fledged response. Similar twitching and abbreviated response patterns for weak stimuli are observed in real frogs. After completion of the irritant removal the simulation recovered the base posture and became quiescent for weak and moderate stimuli. For strong responses a short period of twitching and brief extensions were produced as the activation of the network died down. This type of behavior is also seen in real spinal frogs and was called 'irradiation' by Sherrington. Figure 2 shows sample kinematics of the model.

Stimulus removal for multiple stimuli:

1: blending: When multiple stimuli could be reached in a single trial using non exclusive fields these were utilized by the network. The network could also operate on a single stimulus initially and ignore the second. This depended on stimulus strength and prior network state.

SAMPLE WIPING KINEMATICS (2D)




Figure 2:

Sample behaviors: Skin areas 1, 2, 3, and 4 were used. Several different types of wipes can be elicited in different skin regions. Two different forelimb wipes are shown for comparison in the upper panels. In the second upper panel the system executes a forelimb wipe using summation of a flexion field to sweep back across the forelimb. In the third upper panel the system sweeps the stimulus using an extension. In the lower panels two wipes to areas 3 and 4 are shown. Finally in the bottom right a blended wipe is shown which sweeps both areas 3 and 4. These last panels also demonstrate exploration and variety in the network. Areas are swept with differing methods and this variety can presumably be fine tuned by a learning scheme.

2: switching: When multiple stimuli could only be reached in different ways switching occurred. This was not surprising since the exclusion mechanism explicitly coded into the network would be expected to generate this switching behavior. Figure 3 shows switching and variable trajectories for strong stimuli.

'Exploration':

Difficult to reach regions were regions for which very precise and strong activation of single force fields or combinations were needed to bring the limb to the stimulus. These difficult to reach regions were due to very approximate hand coding of force field positions by the author. Interesting behavior was generated in the simulation in these regions. The persistent flow of activation into the network from the almost unreachable stimulus coupled with a relatively sparse set of conditions and exclusions allowed a collection of increasingly variable responses in the vicinity of the stimulus to be generated. Whether to call this exploration or flailing is a moot point. This is true for both the simulation and for the real spinal frog which also exhibits similar increasingly violent responses. However the net effect of these violent behaviors was usually stimulus removal. Presumably, a learning mechanism added to a network of this type could utilize the system behaving in this way. The learning system could learn the appropriate activation patterns for these 'hard' stimuli, and tune the behavior network to cope with these stimuli more gracefully. This would form an interesting extension of this framework. Figure 3 shows a variety of explorations generated with a 'difficult' stimulus.




Figure 3:

Switching behavior: Stimuli are in areas 2 and 3 (stimuli are the diamonds). These cannot be reached together. Switching occurs in panels 1,2 and 3. between areas 2 and 3. In panels 5 and 6 a motion is generated which blends the kinematics of motions to the two areas in a way which is physically impossible for a real frog due to kinematic constraints. These last two panels demonstrate types of explorations possible in the network.

Understanding the network:

The behavior of networks of this form is very difficult to understand. The number of links which must be defined, learned or evolved increases rapidly to outstrip machine resources and learning algorithms. (See Maes 1991). To try to grasp some of the operation of the simulations here I devised a graphical device to display the networks. This gives some intuition as to the network operation. The method is as follows : The network in a given environment can be represented as a graph. This graph has a linear realization as some geometric form in a multidimensional space. To represent this graph approximately in a manner which is easily grasped the graph of the network can always be imbedded in R^3 . Clearly the graph topology must be distorted in order to do this. To attempt an optimal (least distorted) representation of the graph distances and topology in this constrained low dimensional representation a physical model of the graph was used. Nodes were represented as three dimensional objects in three space. The size of a node was proportional to the activation. Each node was assumed to repel other nodes according to an inverse square law. The activation of the node was used to parameterize this repulsion. The nodes were attached by a non linear spring for each link in the network. Each spring was of stiffness proportional to the activation flow through the link. This mechanical system was solved numerically for one of its equilibrium configurations and the system displayed. The distance of nodes was then a summary of the amount of activation flow between them. Threads of actions linked by activation exchange in the network were close in this virtual space while those not exchanging activation were repelled from the active clusters. In this way the dynamics of exchange of activation are expressed by the dynamics of movement of nodes in this




Figure 4:

'Information Theatre' of a Maes Net. Display of a four node network: The size of each octahedron represents the node activation. Position represents an equilibrium between mutual repulsion of the nodes and the elastic network representing the activation flow in the links (see text). Thus in frame 5 the top two nodes are exchanging energy rapidly while the nodes below are relatively isolated with very little energy flow among them. After execution completes node energy is zeroed. Hence some of the octahedra become miniscule in the transition from frames 7 to 8.

virtual space. Conditions can also be added to this 'information theatre' representation of the network. An example of the movements of nodes is shown in figure 4. This leads to a better grasp of the network. The approach also suggests experiments with learning systems based on distance measures in such low dimensional representations.

6: Relation of reactive and autonomous approaches to other frameworks:

Current work supports the idea that primitives in the spinal-cord (as described above) either specify position or force controls, but not trajectory details. Evidence from some experiments and the simulations presented here supports the idea that this description extends to wiping behaviors. These mechanisms involving local decision making by spinal circuits both constrain and simplify some aspects of descending control.

Other frameworks for use of force fields have been examined and are very powerful when provided with an explicit planning system (see Mussa-Ivaldi and Giszter 1992). The basis field approach described by Mussa-Ivaldi and Giszter allows generation of arbitrarily complex control fields by summation of a small set of primitives using linear summation mechanisms. This framework is ideally suited for implementing a descending control of spinal cord function in which all intelligent computation is performed elsewhere. The approach relies on vector summation of primitives in arbitrary combinations.

Since descending controls and autonomously organized local spinal behaviors must be seamlessly integrated into an adaptive whole, the rules of interaction between descending and local control may be complex. If the same primitives can be used in local behaviors and in subsumed into a descending control scheme, the local behaviors must react appropriately to the context provided by the descending control and to the loss of those modules and actions that are coopted by the descending pathways. Similarly the local behaviors must not prevent access to primitives needed by descending control and must gracefully surrender these so as not to jeopardize the functioning of the system. The scheme presented here, based on the Maes




Figure 5:

The organization of spinal cord visualized in A is a strictly top down approach. The cord is a passive output device. In B the cord is seen as having two functions, one as a transmission system for parameterized basis fields and the second as an autonomous local controller into which various 'code' can be downloaded. This modulation of the reflex systems can set up the cord in support of the basis field generated behavior. Thus different reflex systems might operate in ice-skating as compared to walking, while other parallel processes are controlled by basis field summation. The local behaviors might be suppressed, tuned, recruited or left unchanged after the fashion of subsumption architectures.

approach may allow the integration of local and descending control in a very flexible fashion. I would suggest that setting up the decision making and summation rules for the local reactive spinal circuitry in preparation for, and in support of descending controls is an important aspect of normal motor control. Learning how to do this delegation may be a critical aspect of learning generally, see Figure 5.

7: Relations of the simulation to experimental data:

The predictions of this network are quite good. However some data recently collected as a result of these simulations suggests it is still insufficient to capture some aspects of spinal organization. In particular, rhythmically repeated force responses to brief but intense stimuli can be observed which have quite precise phasing of responses (see figure 6). This precision of timing independent of stimulus strength and duration suggests some clocking of decisions and actions, or data synchronization process in the spinal cord. The model framework presented here might need such a structure or set of structures to replicate the biology. However, it must be noted that the data concerning such cyclic responses are difficult to interpret unambiguously. Some views would regard the rhythmic response as pathological. Modelling the networks presented here coupled to a data synchronizing mechanism or oscillator / clock may give insight into the advantages or disadvantages of such an arrangement in the spinal cord and what experiments are critical.




Figure 6:

Physiological behavior that is difficult to account for simply in the present scheme. A single stimulus of 300 to 600ms can generate a chain of force responses extending several seconds. The measured traces are F_x , F_y , F_z force components. The responses appear to be precisely timed despite their amplitude variations. This suggests some clocking or precise reverberation of spinal circuitry. Total duration of each trial is 2 s. Stimulus onset occurred at 200ms into the traces.

Discussion

A simulation of frog wiping and protective reflexes based on a behavioral decomposition into primitives has been presented. This simulation draws on recent discoveries in the frog spinal cord and has suggested several experiments. The simulation is able to deal with multiple stimuli, choose responses and show a type of exploratory response. As discussed elsewhere this network structure has much to recommend it as a plausible framework for modelling the spinal cord. In particular the global variables available in the behavior network may have some correspondence to spinal observables such as sensitization, depression, habituation, and actions of neuromodulatory substances. The responses of the network are quite 'life-like' and fulfill the goal of providing an impetus and direction to experiment. It is unlikely that all the elements necessary to simulate spinal behaviors are present in the simulation but it encompasses a breadth of response patterns even with a relatively simple structure.

References:

- Arbib M Cobas A (1991) Schemas for prey catching in frog and toad. In: *From Animals to Animats: Proc. First Intl. Conf on Simulation of Adaptive Behavior*, ed Meyer J-A and Wilson SW., pp. 142-151.
- Arkin RC (1991) Integrating behavioral, perceptual and world knowledge in reactive navigation. In: *Designing Autonomous Agents* ed. Maes P. pp. 105-122 MIT/Elsevier.
- Berkinblitt MB, Feldman AG, Fukson OI (1989) Wiping reflex in the frog: Movement patterns, receptive fields, and blends. In: *Visuomotor coordination: amphibians, comparisons, models, and robots* (Ewert J-P, Arbib MA, eds), pp. 615-630. New York: Plenum.
- Bizzi E, Mussa-Ivaldi FA, Giszter SF (1991) Computations underlying the execution of movement: a novel biological perspective. *Science* 253: 287-291.
- Brooks R (1991) Elephants dont play chess. In: *Designing Autonomous Agents* ed. Maes P. pp.3-16 MIT/Elsevier.
- Chepelyugina MF (1947) (Ph.D. Thesis) Moscow: University of Moscow.
- Fukson OI, Berkinblitt MB, Feldman AG (1980) The spinal frog takes into account the scheme of its body during the wiping reflex. *Science* 209: 1261-1263.
- Giszter SF (1993) Behavior networks and force fields for simulating spinal redlex behaviors of the frog. *Second Intl. Conf. on the Simulation of Adaptive Behavior*, ed Meyer J-A, Roitblat HL, and Wilson SW, MIT Press, Cambridge MA.
- Giszter SF, McIntyre J, Bizzi E (1989) Kinematic strategies and sensorimotor transformations in the wiping

movements of frogs. *J Neurophysiol* 62: 750-767.

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1992a) Movement primitives in the frog spinal cord. *Neural systems, II* (Eeckman FH, ed), in press. Boston, MA: Kluwer.

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1992b) The organization of limb motor space in the spinal cord. in *Control of arm movement in space*. ed Caminiti, R., Johnson, P.B. and Burnod, Y., *Exp. Brain Res. Series* 22, Springer-Verlag, Berlin.

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1993) Convergent force fields organized in the frog spinal cord. *J. Neurosci.* 13(2): 467-491.

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1991) Equilibrium point mechanisms in the spinal frog. In: *Visual structures and integrated functions* (Arbib MA Ewert JP, eds), pp. 223-237. New York: Springer-Verlag.

Johnson MB 1990 Build-A-Dude. Master's Thesis MIT Media-Lab.

Khatib O (1986) Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. Robotics Res.* 5:1

Maes P (1989) The dynamics of action selection. *Proc IJCAI -89 Conf Detroit*

Maes P (1991a) Situated agents can have goals. In: *Designing Autonomous Agents* ed. Maes P. pp.49-70 MIT/Elsevier.

Maes P (1991b) A bottom up mechanism for behavior selection in an artificial creature. In: *From Animals to Animats: Proc. First Intl. Conf on Simulation of Adaptive Behavior*. ed Meyer J-A and Wilson SW., pp. 238-246

Minsky M (1986) *Society of Mind*, Simon and Schuster

Mussa-Ivaldi FA, Giszter SF (1992) Vector Field approximation: a computational paradigm for motor control and learning. *Biol. Cybern.* 67: 491-500.

Mussa-Ivaldi FA, Giszter SF, Bizzi E (1990) Motor-space coding in the central nervous system. In: *Cold Spring Harbor symposium on quantitative biology* 55: 827-835. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

Mussa-Ivaldi FA and Hogan N (1991) Integrable solutions of kinematic redundancy via impedance control. *Intl J. Robotics Res.* 10:481-491.

Ostry DJ, Feldman AG, Flanagan JR (1991) Kinematics and control of frog hindlimb movements. *J Neurophysiol* 65: 547-562.

Payton DW (1991) Internalized plans: a representation for action resources. In: *Designing Autonomous Agents* ed. Maes P. pp.89-103 MIT/Elsevier.

Schotland JL, Lee WA, Rymer WZ (1989) Wiping and flexion withdrawal reflexes display different EMG patterns prior to movement onset in the spinalized frog. *Exp Brain Res* 78: 649-653.

Sherrington C (1961) *The integrative action of the nervous system*. New Haven, CT: Yale University Press.

von Holst E, (1973) *The behavioral physiology of animals and man. Selected papers of Eric von Holst*. Coral Gables, Florida: University of Miami Press.

Acknowledgements:

Conversations with Pattie Maes and Michael Johnson were of great help in the design and implementation of these simulations. The simulation work presented was conducted independently and the interpretation of spinal organization presented should not be construed to necessarily reflect the views of all my colleagues in the Bizzi laboratory.