

MODELING THE MIRROR: GRASP LEARNING AND
ACTION RECOGNITION

by

Erhan Oztop

A Dissertation Presented to the
FACULTY OF THE GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA

In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(COMPUTER SCIENCE)

August 2002

Copyright 2002

Erhan Oztop

DEDICATION

To my Grandmother

ACKNOWLEDGEMENTS

The time I had at USC during the Ph.D. route was a very enriching period of my life. I had the opportunity to work in an exciting and stimulating research environment, led by Michael Arbib.

I would like to present my deepest gratitude to The Scientific and Technical Research Council of Turkey (TUBITAK) for providing me the scholarship that made it possible for me to attempt and complete the Ph.D. study presented in this thesis. The study would not have been possible if TUBITAK did not provide support for the very first semester and the final semester at USC.

I would like to thank Michael Arbib for guiding and educating me throughout the my years at USC. He has been an extraordinary advisor and mentor, whom I owe all the *brain theory* I learned. I would also like to present my gratitude to Michael Arbib for providing support via HFSP and USCBP grants. Without HFSP and USCBP support, this study would not be possible.

Stefan Schaal is a great mentor, who has been a constant support and source of inspiration with never-ending energy. Besides introducing me to Robotics, he and his colleague, Sethu Vijayakumar were very influential in maturing the concept of *machine learning* in my mind.

I also owe a great debt of gratitude to Nina Bradley for being a source positive energy and mentoring me especially in *infant motor development*. Without her, the thesis would be lacking a major component.

I am also full of gratitude to my Ph.D. qualification exam committee members Maja Mataric and Christoph von der Malsburg for their guidance and support.

I am greatly indebted to Giacomo Rizzolatti for enabling me to visit his lab and providing the opportunity to communicate with not only himself but also with Vittorio Gallese and Leonardo Fogassi who have provided invaluable insights about mirror neurons. In addition, I am very thankful to Massimo Matelli and Giuseppe Luppino, for the first hand information on the mirror neuron system connectivity, and to Luciano Fadiga for stimulating discussions. I would also like to thank to Christian Keysers and Evelyne (Kohler) for not only actively involving me in their recording sessions but also offering their sincere friendship.

I am very thankful to Hideo Sakata for giving me the opportunity to visit his lab in Tokyo and interact with many researchers including Murata-san with whom I had very stimulating discussions.

I am deeply thankful to Mitsuo Kawato, for giving me the opportunity to interact with various researchers in Japan by having me at ATR during the summer of 2001. My research experience at ATR was very rewarding; I greatly expanded my knowledge on motor control and motor learning. I would like to salute the staff at ATR for all their help. I also would like to present my thanks to the friends at ATR for welcoming me and making me feel at home.

I would like to present my appreciation and thanks to my mentors at Middle East Technical University in Turkey. I present my sincere thanks to my masters advisor Marifi Guler for introducing me to neural computation and to Volkan Atalay for introducing me to computer vision, and supporting my Ph.D. application. Especially, I would like to present my gratitude to Fatos Yarman Vural for her guidance and support during my Masters study and for preparing me for the Ph.D. work presented in

this dissertation. Other influential Computer Science professors to whom I am grateful for educating me are Nese Yalabik, Gokturk Ucoluk and Sinan Neftci.

I would like to present my gratitude to Tosun Terzioglu, Albert Ekip, Turgut Onder and Semih Koray who were professors of the Mathematics Department at the Middle East Technical University. They taught me how to think 'right' and exposed the beauty of mathematics to me.

Throughout these six years at USC, I had the pleasure to meet several valuable people who contributed to this dissertation. Firstly, I am very thankful to Aude Billard and Auke Jan Ijspeert for all their support and scientific interaction and feedback. They have a huge role in helping me get through the tough times during my Ph.D. work. In addition, I would like to thank Aude Billard for providing me the computing hardware and helping me have a nice working environment, which was very essential for the progress of my Ph.D. study. I am thankful to my great friend Sethu Vijayakumar for his support and stimulating discussions. Jun Nakanishi, Jun Mei Zu, Aaron D'Souza, Jan Peters and Kazunori Okada besides being of great support, were always there to discuss issues and helped me overcome various obstacles. I am deeply thankful to Shrija Kumari for offering not only her smile and friendship but also her energy to proofread my manuscript. I owe a lot to Jun Mei Zu: she has always been there as a great friend and has always offered her help and support when I needed it most. I am indebted to my ex-officemate and a very valuable friend Mihail Bota for his constant support and interactions for improving the thesis and providing me the psychological support to overcome many obstacles throughout my Ph.D. years. Finally, I would like to thank Juan Salvador Yahya Marmol for being a good friend and sharing my workload during various periods of my Ph.D. I count myself very lucky to have these great friends and colleagues whom once again I present my gratitude: Thank you guys!

Not a Hedco Neuroscience inhabitant but a very valuable friend, Lay Khim Ong was always there for offering her help both psychologically and physically (Hey Khim: thank you for your great editing!). I would like to thank other great friends who supported me (in spite of my negligence in keeping in touch with them). Kyle Reynolds, Awe Vinijkul, Aye Vinijkul Reynolds, Alper Sen, Ebru Dincel: please accept my sincere thanks and appreciation.

I owe deep gratitude to my wife Mika Satomi for her patience in dealing with me in difficult times. She was always there. Her contribution to this thesis is indispensable. I especially celebrate and thank her for the artistic talents and hard work that she generously offered me throughout the Ph.D. study.

I am greatly indebted to Paulina Tagle and Laura Lopez for their support and help over all these years. I also would like to thank Laura Lopez and Yolanda Solis for their kind friendship and support. My gratitude to Luigi Manna, who helped me with the hardware and software issues during the Ph.D. study.

I would like to thank Laurent Itti for his generous help for improving our lab environment and providing partial support for my research. Also, I would like to salute his lab members for their support and friendship. Florence Miao, in particular, had always offered her warm friendship during her internship at USC.

I would like to present my appreciation to the good things in life particularly, I would like to thank the ocean for comforting and rejuvenating me during difficult times.

Finally, I am deeply indebted to my family, to whom I owe much more than what can be expressed. This work would not be possible without the help of my parents. (Anne ve Baba, Evrim ve Nurdan: Hersey icin cok cok tesekkurler!)

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	x
1 INTRODUCTION.....	1
2 CHAPTER II: BIOLOGICAL BACKGROUND	4
2.1 Abbreviations	4
2.2 Premotor areas.....	4
2.2.1 Area F5.....	5
2.2.2 Area F4.....	10
2.2.3 Areas F2 and F7 (dorsolateral prefrontal cortex).....	11
2.2.4 Area F1 (the primary motor cortex).....	12
2.2.5 Areas F3 (SMA proper), F6 (pre-SMA).....	12
2.3 The superior temporal sulcus.....	12
2.4 Parietal Areas.....	13
2.4.1 The anterior intraparietal area (AIP)	13
2.4.2 The caudal intraparietal sulcus (c-IPS).....	15
2.4.3 Areas VIP, MIP and LIP	17
2.4.4 Areas 7a and 7b (PG and PF).....	18
2.5 Connectivity and other brain regions.....	19
2.6 Mirror neurons in humans.....	23
2.7 Summary	25
3 CHAPTER III: MIRROR NEURON SYSTEM MODEL.....	26
3.1 The mirror neuron system for grasping and FARS model.....	26
3.2 The hand-state hypothesis	29
3.2.1 Virtual fingers.....	29

3.2.2	The hand-state hypothesis	30
3.3	The MNS (mirror neuron system) model	31
3.3.1	Overall function.....	33
3.3.2	Schemas explained.....	33
3.4	Schema implementation.....	36
3.4.1	Grand schema 1: reach and grasp	36
3.4.2	Grand schema 2: visual analysis of hand state.....	38
3.4.3	Grand Schema 3: core mirror circuit.....	42
3.5	Simulation results	46
3.5.1	Non-explicit affordance coding experiments	46
3.5.2	Explicit affordance coding experiments.....	50
3.5.3	Justifying the visual analysis of hand state schema	54
3.6	Discussion and predictions.....	56
3.6.1	The hand state hypothesis	56
3.6.2	Neurophysiological predictions.....	57
4	CHAPTER IV: MULTILAYER SUPERVISED HEBBIAN LEARNING AND PROBABILITY CODING.....	60
4.1	Neural coding.....	60
4.2	Operation of the proposed network.....	61
4.3	Testing the proposed architecture	64
4.3.1	Deterministic environment.....	64
4.3.2	Stochastic environment	66
4.3.3	Combining multiple layers	68
4.4	Summary	69
5	CHAPTER V: INFANT GRASP LEARNING	71
5.1	Motivation.....	71
5.2	Infant reach and grasp.....	71
5.3	Neural maturation versus interactive learning.....	74
5.4	Infant Learning to Grasp Model (ILGM)	75
5.4.1	Layers of infant learning to grasp model.....	76
5.4.2	Functional description of ILGM layers	77
5.5	Joy of grasping.....	78

5.5.1	Mechanical grasp stability	78
5.5.2	Implementing the grasp stability	79
5.6	Learning approach direction with palm orienting behavior.....	80
5.6.1	Simulation results	80
5.6.2	Conclusions and predictions	82
5.7	Is infant palm orienting learned or innate? Learning the wrist orientation.....	82
5.7.1	Simulation results	83
5.7.2	Conclusions and predictions	84
5.8	Task constraints shape infant grasping.....	84
5.8.1	Simulation results	85
5.8.2	Conclusions and predictions	86
5.9	Affordance input matters.....	87
5.9.1	Simulation results	88
5.9.2	Comparison of ILGM with Lockman et al. (1984)	88
5.9.3	ILGM kinematics analysis (five months of age).....	88
5.9.4	ILGM kinematics analysis (nine months of age).....	89
5.10	Summary and conclusion	91
5.11	Discussion	92
6	CHAPTER VI: NEUROPHYSIOLOGICAL VIEW OF LEARNING TO GRASP.....	93
6.1	Grasp learning circuit and mirror neurons are complementary networks.....	93
6.2	Introduction to primate grasping	94
6.3	Neural correlates of infant reach and grasp	95
6.4	Primate grasp development hypotheses.....	97
6.4.1	Hypothesis I: two coexistent grasping circuits	97
6.4.2	Hypothesis II: single grasping circuit.....	98
6.5	Affordance-based learning to grasp model (LGM).....	98
6.5.1	Localizing learning to grasp model in primate cortex	99
6.5.2	What does cerebral cortex know about a grasp?	101
6.5.3	Simulation level description of LGM layers.....	102
6.5.4	Why LGM is relevant: good model versus bad model	102
6.6	Wrist orientation-learning revisited: neural level analysis	103

6.6.1	Neural level analysis.....	103
6.6.2	LGM represents a 'menu' of grasps in terms of neural activity	104
6.6.3	Predictions and discussion	106
6.7	Object axis selectivity: neural level analysis.....	106
6.7.1	Neural level analysis.....	107
6.7.2	Conclusions and neurophysiological predictions	108
6.8	Object size selectivity.....	109
6.8.1	Simulation results	109
6.8.2	Conclusions and neurophysiological predictions	111
6.9	Generalization: learning to plan based on object location.....	113
6.9.1	Simulation results	113
6.9.2	Summary and Conclusion	115
7	CHAPTER VII: BIOLOGICALLY REALISTIC F5 VISUAL SERVO CIRCUITS FOR GRASPING AND EMERGENCE OF MIRROR NEURONS.....	117
7.1	Motivation.....	117
7.2	The link between the mirror neuron system and grasp learning.....	118
7.2.1	Two Visual Control Hypotheses.....	119
7.2.2	Mirror neurons in feed-forward control (alternative I)	121
7.2.3	Mirror Neurons in feedback control (alternative II).....	122
7.2.4	The target of implementation.....	122
7.2.5	The visual servo task	123
7.2.6	The feed-forward model learning.....	123
7.3	Implementation: F5 manual visual control circuit for 2D arm	123
7.3.1	A leaky integrator model for F5 manual visual feedback circuit	124
7.3.2	Simulation: visual feedback control with leaky integrators.....	127
7.3.3	Simulation: feedback and lower motor centers.....	128
7.3.4	F5 Feed-forward visual control and mirror neurons	130
7.3.5	Simulation: trajectory planning and controller performance.....	135
7.4	Feed-forward unit activity and mirror neurons	137
8	CHAPTER VIII: CONCLUSIONS	142
8.1	Mirror neurons	142
8.2	Grasp learning: infant development.....	144
8.3	Grasp learning: monkey neurophysiology.....	144
8.4	Grasp learning: neural architecture.....	145

9	CHAPTER IX: FUTURE WORK	146
9.1	Simultaneous learning in MNS model and LGM.....	146
9.2	Going beyond grasping: learning the 'hand state'	146
9.3	Tactile feedback for grasping	146
9.4	Learning to extract the right affordance	147
9.5	More realistic models of the limb and the brain regions	147
9.6	Sensitivity analyses for the simulation parameters.....	147
10	REFERENCES	149
11	APPENDIX	164
11.1	Mirror neuron system model (MNS).....	164
11.1.1	Color segmentation.....	164
11.1.2	Reach and grasp schema precision grasp planning and execution.....	165
11.2	Learning to grasp models (ILGM and LGM)	165

LIST OF FIGURES

Figure 2.1 Lateral view of macaque brain showing the areas of agranular frontal cortex and posterior parietal cortex (adapted from Geyer et al. 2000). The naming conventions: frontal regions, Matelli et al.(1991); parietal regions, Pandya and Seltzer (1982).....	5
Figure 2.2 A canonical neuron response during grasping of various objects in the <i>dark</i> (left to right and top to bottom: plate, ring, cube, cylinder, cone and sphere. The rasters and histograms are aligned with object presentation. Small grey bars in each raster marks onset of key press, go signal, key release, onset of object pulling, release signal, and object release, respectively. The peaks in ring and sphere object cases correspond to the grasping of the object by the monkey (adapted from Murata et al. 1997a).....	6
Figure 2.3 The motor responses of the same neuron shown in Figure 2.2. The motor preference of the neuron is also carried over to the visual preference (compare the ring and sphere histograms of both figures) (adapted from Murata et al. 1997a)	7
Figure 2.4 Activity of a cell during action observation (left) and action execution (right). There is no activity in presentation of the object during both initial presentation and bringing the tray towards the monkey. The vertical line over the histogram indicates the hand-object contact onset. (from Gallese et al., 1996).	8
Figure 2.5 Visual response of a mirror neuron. A. Precision grasp B. power grasp C. mimicking of precision grasp. The vertical lines over the histograms indicate the hand-object contact onset. (adapted from Gallese et al., 1996).....	9
Figure 2.6 Example of a strictly congruent manipulating mirror neuron: A) The experimenter retrieved the food from a well in a tray. B) Same action, but performed by the monkey. C) The monkey grasped a small piece of food using a precision grip. The vertical lines over the histograms indicate the hand-object contact onset (adapted from Gallese et al., 1996).	9
Figure 2.7 The classification of area F5 neurons derived from published literature (Dipellegrino et al. 1992; Gallese 2002; Gallese et al. 1996; Murata et al. 1997a; Murata et al. 1997b; Rizzolatti et al. 1996a; Rizzolatti and Gallese 2001). All F5 neurons fire in response to some motor action. In addition, canonical neurons fire for object presentation while the mirror neurons fire for action observation. The majority of hand related F5 neurons are purely motor (Gallese 2002)(labelled as Motor Neurons in the figure).....	10

Figure 2.8 The macaque parieto-frontal projections from mesial parietal cortex, medial bank of the intraparietal sulcus and the surface of the superior parietal lobule (adapted from Rizzolatti et al. 1998). Note that the Brodmann's area 7m corresponds to Pandya and Seltzer's (1982) area PGm.....	11
Figure 2.9 The intraparietal sulcus opened to show the anatomical location of AIP in the macaque (adapted from Geyer et al. 2000)	13
Figure 2.10 An AIP visual-dominant neuron activity under three task conditions: Object manipulation in the light, object manipulation in the dark and object fixation in the light. The neuron is active during fixation and holding phase when the action is performed in light condition. However, during grasping in dark the neuron shows no activity. The fixation of the object alone without grasping also produces a discharge (adapted from Sakata et al. 1997a).....	14
Figure 2.11 Activity the same neuron in Figure 2.10 during fixation of different objects. The neuron show selectivity for horizontal plate (adapted from Sakata et al. 1997a).....	14
Figure 2.12 An AIP visual-dominant neuron's axis orientation tuning and object fixation response is shown. The neuron fires maximally during the fixation of a vertical bar or a cylinder. The tuning is demonstrated in the lower half of the figure (adapted from Sakata et al. 1999).....	15
Figure 2.13 Response of an axis-orientation-selective (AOS) neuron in the caudal part of the lateral bank of the intraparietal sulcus (c-IPS) to a luminous bar tilted 45° forward (<i>left</i>) or 45 backward (<i>right</i>) in the sagittal plane. The monkey views the bar with binocular vision. The line segment under the histograms mark the fixation start and the period of 1 second. (adapted from Sakata et al. 1999)	16
Figure 2.14 The response of the same neuron in Figure 2.13, for monocular vision conditions for the left and right eyes. (adapted from Sakata et al. 1999)	16
Figure 2.15 Orientation tuning of a surface-orientation selective (SOS) neuron. First row: Stimuli presented. Middle row: responses of the cell with binocular view. Last row: responses of the cell with monocular view (adapted from Sakata et al. 1997a).....	17
Figure 2.16 The reconstructed connectivity of area 7a. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001).....	20
Figure 2.17 The reconstructed connectivity of area 7b. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001).....	21
Figure 2.18 The reconstructed connectivity of area AIP. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001).....	22

Figure 3.1 Lateral view of the monkey cerebral cortex (IPS, STS and lunate sulcus opened). The visuomotor stream for hand action is indicated by arrows (adapted from Sakata et al., 1997).....	27
Figure 3.2 AIP extracts the affordances and F5 selects the appropriate grasp from the AIP 'menu'. Various biases are sent to F5 by Prefrontal Cortex (PFC) which relies on the recognition of the object by Inferotemporal Cortex (IT). The dorsal stream through AIP to F5 is replicated in the MNS model.....	28
Figure 3.3 Each of the 3 grasp types here is defined by specifying two "virtual fingers", VF1 and VF2, which are groups of fingers or a part of the hand such as the palm which are brought to bear on either side of an object to grasp it. The specification of the virtual fingers includes specification of the region on each virtual finger to be brought in contact with the object. A successful grasp involves the alignment of two "opposition axes": the <i>opposition axis in the hand</i> joining the virtual finger regions to be opposed to each other, and the <i>opposition axis in the object</i> joining the regions where the virtual fingers contact the object. (Iberall and Arbib 1990)	29
Figure 3.4 The components of hand state $F(t) = (d(t), v(t), a(t), o_1(t), o_2(t), o_3(t), o_4(t))$. Note that some of the components are purely hand configuration parameters (namely v, o_3, o_4, a) whereas others are parameters relating hand to the object.....	31
Figure 3.5 The MNS (Mirror Neuron System) model. (i) Top diagonal: a portion of the FARS model. Object features are processed by cIPS and AIP to extract grasp affordances, these are sent on to the canonical neurons of F5 that choose a particular grasp. (ii) Bottom right. Recognizing the location of the object provides parameters to the motor programming area F4 which computes the reach. The information about the reach and the grasp is taken by the motor cortex M1 to control the hand and the arm. (iii) New elements of the MNS model: Bottom left are two schemas, one to recognize the shape of the hand, and the other to recognize how that hand is moving. (iv) Just to the right of these is the schema for hand-object spatial relation analysis. It takes information about object features, the motion of the hand and the location of the object to infer the relation between hand and object. (v) The center two regions marked by the gray rectangle form the core mirror circuit. This complex associates the visually derived input (hand state) with the motor program input from region F5 canonical neurons during the learning process for the mirror neurons. The grand schemas introduced in section 3.2 are illustrated as the following. The "Core Mirror Circuit" schema is marked by the center grey box; The "Visual Analysis of Hand State" schema is outlined by solid lines just below it, and the "Reach and Grasp" schema is outlined by dashed lines. (Solid arrows: established connections; dashed arrows: postulated connections).....	32

Figure 3.6 (a) For purposes of simulation, we aggregate the schemas of the MNS (Mirror Neuron System) model of Figure 3.5 into three "grand schemas" for Visual Analysis of Hand State, Reach and Grasp, Core Mirror Circuit. (b) For detailed analysis of the Core Mirror Circuit, we dispense with simulation of the other two grand schemas and use other computational means to provide the three key inputs to this grand schema34

Figure 3.7 (Left) The final state of arm and hand achieved by the reach/grasp simulator in executing a power grasp on the object shown. (Right) The hand state trajectory read off from the simulated arm and hand during the movement whose end-state is shown at left. The hand state components are: $d(t)$, distance to target at time t ; $v(t)$, tangential velocity of the wrist; $a(t)$, Index and thumb finger tip aperture; $o1(t)$, cosine of the angle between the object axis and the (index finger tip – thumb tip) vector; $o2(t)$, cosine of the angle between the object axis and the (index finger knuckle – thumb tip) vector; $o3(t)$, The angle between the thumb and the palm plane; $o4(t)$, The angle between the thumb and the index finger37

Figure 3.8 Grasps generated by the simulator. (a) A precision grasp. (b) A power grasp. (c) A side grasp.....38

Figure 3.9 (a) Training the color expert, based on colored images of a hand whose joints are covered with distinctively colored patches. The trained network will be used in the subsequent phase for segmenting image. (b) A hand image (not from the training sample) is fed to the augmented segmentation program. The color decision during segmentation is done by consulting to the Color Expert. Note that a smoothing step (not shown) is performed before segmentation40

Figure 3.10 Illustration of the model matching system. Left: markers located by feature extraction schema. Middle and Right: initial and final stages of model matching. After matching is performed a number of parameters for the Hand configuration are extracted from the matched 3D model41

Figure 3.11 The scaling of an incomplete input to form the full spatial representation of the hand state As an example, only one component of the hand state, the aperture is shown. When the 66 percent of the action is completed, the pre-processing we apply effectively causes the network to receive the stretched hand state (the dotted curve) as input as a re-representation of the hand state information accessible to that time (represented by the solid curve; the dashed curve shows the remaining, unobserved part of the hand state)44

Figure 3.12 The solid curve shows the effective input that the network receives as the action progresses. At each simulation cycle the scaled curves are sampled (30 samples each) to form the spatial input for the network. Towards the end of the action the networks input gets closer to the final hand state.....45

- Figure 3.13 (a) A single grasp trajectory viewed from three different angles to clearly show its 3D pattern. The wrist trajectory during the grasp is marked by square traces, with the distance between any two consecutive trace marks traveled in equal time intervals. (b) Left: The input to the network. Each component of the hand state is labelled. (b) Right: How the network classifies the action as a power grasp: squares: power grasp output; triangles: precision grasp; circles: side grasp output. Note that the response for precision and side grasp is almost zero47
- Figure 3.14 Power and precision grasp resolution. The conventions used are as in the previous figure. (a) The curves for power and precision cross towards the end of the action showing the change of decision of the network. (b) The left shows the initial configuration and the right shows the final configuration of the hand48
- Figure 3.15: (Top) Strong precision grip mirror response for a reaching movement with a precision pinch. (Bottom) Spatial location perturbation experiment. The mirror response is greatly reduced when the grasp is not directed at a target object. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)48
- Figure 3.16 Altered kinematics experiment. Left: The simulator executes the grasp with bell-shaped velocity profile. Right: The simulator executes the same grasp with constant velocity. Top row shows the graphical representation of the grasps and the bottom row shows the corresponding output of the network. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)49
- Figure 3.17 Grasp and object axes mismatch experiment. Rightmost: the change of the object from cylinder to a plate (an object axis change of 90 degrees). Leftmost: the output of the network before the change (the network turns on the precision grip mirror neuron). Middle: the output of the network after the object change. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)50
- Figure 3.18 The plots show the level of mirror responses of the explicit affordance coding object for an observed precision pinch for four cases (tiny, small, medium, big objects). The filled circles indicate the precision activity while the empty squares indicate the power grasp related activity52
- Figure 3.19 The solid curve: the precision grasp output, for the non-explicit affordance case, directed to a tiny object. The dashed curve: the precision grasp output of the model to the explicit affordance case, for the same object52

Figure 3.20: Empty squares indicate the precision grasp related cell activity, while the filled squares represent the power grasp related cell activity. The grasps show the effect of changing the object affordance, while keeping a constant hand state trajectory. In each case, the hand-state trajectory provided to the network is appropriate to the medium-sized object, but the affordance input to the network encodes the size shown. In the case of the biggest object affordance, the effect is enough to overwhelm the hand state's precision bias.	53
Figure 3.21 The graph is drawn to show the decision switch time versus object size. The minimum is not at the boundary, that is, the network will detect a precision pinch quickest with a medium object size. Note that the graph does not include a point for "Biggest object" since there is no resolution point in this case (see the final panel of Figure 3.19)	54
Figure 3.22 The precision grasp action used to test our visual system is depicted by superimposed frames (not all the frames are shown).....	54
Figure 3.23 The video sequence used to test the visual system is shown together with the 3D hand matching result (over each frame). Again not all the frames are shown.....	55
Figure 3.24 The plot shows the output of the MNS model when driven by the visual recognition system while observing the action depicted in Figure 3.22. It must be emphasized that the training was performed using the synthetic data from the grasp simulator while testing is performed using the hand state extracted by the visual system only. Dashed line: Side grasp related activity; Solid line: Precision grasp related activity. Power grasp activity is not visible as it coincides with the time axis.....	56
Figure 4.1 The elevated circular region corresponds to the area defined by the equation $(x^2+y^2) < 0.25$. The environment returns +1 as the reward if the action falls into the circular region, otherwise -1 is returned.....	64
Figure 4.2 The adaptation of the firing potential of the stochastic units are shown as a series of evolving 3D maps. (left to right and top to bottom).....	65
Figure 4.3. The normalized histogram of the actions generated over 60000 samples. Note that the actions generated captured the environment's reward distribution (see Figure 4.1).	66
Figure 4.4 The stochastic environment's double peaked reward distribution (see text for the explanation).....	66
Figure 4.5 Some snapshots showing the phases of learning of the layer in the stochastic environment where the reward distribution has two peaks (see Figure 4.4)......	67
Figure 4.6. The normalized histogram of 60000 data points (actions) generated by the trained layer in the stochastic environment depicted in Figure 4.4.	68

Figure 5.1 Infant grip configurations can be divided in two categories: power and precision grips. Infants tend to switch from power grips to precision grips as they grow (adapted from Butterworth et al. 1997).....	73
Figure 5.2 The structure of the Infant Learning to Grasp Model. The individual layers are trained based on somatosensory feedback	76
Figure 5.3 Hand Position layer specifies the approach direction of the hand towards the object. The representation is allocentric (centred on the object). Geometrically the space around the object can be uniquely specified with the vector (azimuth, elevation, radius). The Hand Position layer generates the vector by a local population vector computation. The locus of the local neighbourhood is determined by the probability distribution represented in the firing potential of Hand Position layer neurons (see Chapter 4, for details).....	77
Figure 5.4: The grasp stability we used in the simulations is illustrated for a hypothetical precision pinch grip (note that this is a simplified, the actual hand used in the simulations has five fingers).....	79
Figure 5.5 The trained model's Hand Position layer is shown as a 3D plot. One dimension is summed to reduce the 4D map to a 3D map. Intuitively the map says: 'when the object is above the shoulder and in front grasp it from the bottom'	81
Figure 5.6: The output of the trained model's target position layer is shown as a 3D plot. One dimension is summed to reduce the 4D map to a 3D map. The object is on the left side of the (right handed) arm. Intuitively, the map says 'when the object is on the left side grasp it from the right side of the object'	81
Figure 5.7 The learning evolution of the distribution of the Hand Position layer is shown as a 3D plot. Note that the 1000 neurons shown represent the probability distribution of approach directions. Initially, the layer is not trained and responds in a random fashion to the given input. As the learning progresses, the neurons gain specificity for this object location.....	82
Figure 5.8 ILGM planned and performed a power grasp after learning. Note the supination (and to a lesser extent extension) of the wrist required to grasp the object from the bottom side	83
Figure 5.9 Two learned precision grips (left: three fingered; right four fingered) are shown. Note that the wrist configuration for each case. ILGM learned to combine the wrist location with the correct wrist rotations to secure the object.....	83
Figure 5.10 ILGM was able to generate two fingered precision grips. However these were less than the three or four finger grips	84

Figure 5.11 The <i>cube on the table</i> simulation set up. ILGM interacts with the object with the physical constraint that it has to avoid collision with the table	85
Figure 5.12 ILGM learned a ‘menu’ of precision grips with the common property that the wrist was placed well above the object. The orientation of the hand and the contact points on the object showed some variability. Two example precision grips are shown in the figure.....	85
Figure 5.13. ILGM acquired thumb opposing index finger precision grips	86
Figure 5.14 The three cylinder orientations and grasp attempts by the poor vision condition.	87
Figure 5.15 The orientation match of the hand and the cylinder is illustrated. Dashed line with diamonds: 5 months old infants; Solid line with diamonds: 9 months old infants; Dashed line with circles: ILGM with no affordance; Solid line with circles: ILGM with affordance (infant data from Lockman et al. (1984)). Right panel illustrates the object orientation used for the simulation and for the infants in this comparison	88
Figure 5.16 The hand orientation and cylinder orientation difference curves for individual trials. The columns from left to right correspond to horizontal, diagonal and vertical orientations. The upper row <i>flat</i> class of error curves, lower row <i>non-flat</i> class for error curves (see text for explanation)	89
Figure 5.17 The hand orientation and cylinder orientation difference curves while ILGM was executing four types of grasp in the full-vision condition. Left two figures are two typical error curves for the horizontal cylinder. Note that the two horizontal case error patterns reflect the two possible grasps: from the bottom and from the top. The third and fourth are typical error curves for the diagonal and vertical cylinders respectively	90
Figure 5.18 The grasps performed after ILGM learned the association between the wrist rotations and the object affordance (orientation)	91
Figure 6.1: The overall MNS model. The grey background rectangle shows the focus of this chapter. In addition to the areas shown, area F2 will be posited as being involved in grasp planning.	94

Figure 6.2 Left: precision grasp (pad opposition); Middle: Power grasp (palm opposition); Right: Side grasp (side opposition). Each of the 3 grasp types here is defined by specifying two 'virtual fingers', VF1 and VF2, which are groups of fingers or a part of the hand such as the palm which are brought to bear on either side of an object to grasp it. The specification of the virtual fingers includes specification of the region on each virtual finger to be brought in contact with the object. A successful grasp involves the alignment of two "opposition axes": the *opposition axis in the hand* joining the virtual finger regions to be opposed to each other, and the *opposition axis in the object* joining the regions where the virtual fingers contact the object (adapted from Iberall and Arbib 1990).....95

Figure 6.3 The two possible organization of learning to grasp circuit are shown. According to Hypothesis I, two grasping circuits exist; the phylogenetically older one located in area F1 (hatched background) and the newer one in the premotor cortex (solid background). According to Hypothesis II, F1 is involved in only executing the premotor cortex instructed movements. LGM is based on the latter hypothesis. The details of LGM are shown in Figure 6.4. Note that we introduced area F2 for complementing the MNS structures. The visual input to area F2 originates from MIP (not shown) and V6a98

Figure 6.4 The Learning to Grasp Model. F5 is implicated in all grasp related parameters. Dashed connections indicate the direct corticospinal projections of premotor areas. Area F5 works with area F2 and F4 to transform visual affordances signalled by parietal areas into a grasp plan. The grasp plan is then, relayed to primary motor cortex (F1) and spinal cord for execution. The tactile feedback of the action is integrated in the first somatosensory cortex (SI), which mediates the adaptation of the parietal-premotor and inter-premotor connections100

Figure 6.5 The top-left shows the Hand Position layer output summed over the radius (approach direction is encoded in spherical coordinates) as a 3D plot. The top-centre panel shows the sample generated from the Hand Position distribution. Bottom-left shows the Wrist Rotation layer output summed over the heading axis as a 3D plot. The bottom-centre panel shows the parameters picked from the Wrist Rotation layer distribution. Note that Wrist Rotation layer distribution depends on (i.e. represents a conditional distribution) the sample picked from the Hand Position layer. The rightmost panel shows the executed grasp104

Figure 6.6 Using the same LGM used for Figure 6.5, another grasp plan is generated (left four panels). The resulting grasp is shown on the right. By comparing the grasp plan shown on the left four panels with of Figure 6.5's grasp plan we see how the selection of a different approach direction (see the centre-top panels of both figures) changed the Wrist Orientation distribution.....105

Figure 6.7:Two very different grasp generation from the same LGM. Upper panel: Grasping with maximum wrist extension with some pronation. Lower panel: Grasping with maximum wrist supination and small wrist extension. Note that the Wrist Layer probability map is the same since the approach direction was chosen the same (the small dots in the right most panels).	105
Figure 6.8 The grasps performed after LGM learned the association of hand rotations with the object orientation input (full vision condition). Note that the left panel shows a bottom side grasp. All of the shown grasp configurations satisfied grasp stability criterion	107
Figure 6.9 In the poor-vision case, the hand rotation neurons in LGM show the same response for horizontal (left panel), diagonal (centre panel) and vertical (right panel) object presentations because of the lack of axis orientation input	107
Figure 6.10 When LGM has access to axis orientation information the Hand Rotation neurons represent different plans in response to horizontal (left panel), diagonal (centre panel) and vertical (right panel) object presentations	108
Figure 6.11 The small object presentation produced two peaks of activity in the Hand Position layer corresponding to the probability distribution of approach directions. The right panel shows the executed grasp when the data generation was localized in the area pointed by the leftmost arrow.....	109
Figure 6.12: A large cube was grasped by securing the object between the thumb and the other fingers (right panel). The Hand Position layer activity is shown on the left panel. The neuron with largest activity is marked with an arrow.....	110
Figure 6.13 The largest object presentation and grasping. The Hand Position reflects a single reach direction as indicated with an arrow	110
Figure 6.14:The Hand Position layer activity is superimposed to demonstrate that the maximum activity loci are separated for each object indicating <i>selectivity for object size</i>	111
Figure 6.15 The trained Learning to Grasp Model executed grasps to objects located at nine different locations in the workspace. The grasp locations were not used in the training. All of the grasps shown were stable.....	114
Figure 6.16 The same model used in generating Figure 6.15 was used to generate a different set of grasps. Again all the grasps were stable.	114

Figure 6.17 The internal mechanisms of representing and generating multiple grasp plans are shown. Solid arrows (except object encoding) denote learned connections while empty arrows indicate data generation. The flow of operation starts with the presentation of the object (the bottom centre) and follows the arrows. At the top-centre, the data generation can yield multiple approach directions. The two possible approach directions are shown creating two streams (left column and right column), each of which yields different grasp execution (bottom pictures of left and right column).....	115
Figure 7.1 The MNS model repartitioned to show the focus of this chapter. The grey background marks area of interest.....	118
Figure 7.2 One alternative visual control structure for manipulation is shown within the MNS framework. The mirror neurons generate feed-forward commands.....	119
Figure 7.3 Another alternative visual control structure for manipulation is shown within the MNS framework (compare with Figure 7.2). The mirror neurons generate feedback commands.....	120
Figure 7.4 The feedback and feed-forward control view of the F5 grasping circuit, alternative I: F5mirror neurons learn to generate feed-forward command. The desired state is assumed to be available and is converted to a correction motor command by F5motor-only units using stochastic gradient descent. F5canonical neurons gate the feed-forward and feedback pairs.....	121
Figure 7.5. The feedback and feed-forward control view of the F5 grasp circuit, alternative II: F5mirror neurons learn to compute the error. The error is then converted to a correction motor command by F5motor-only units. F5canonical generates the feed-forward command signal.....	122
Figure 7.6 The schema level view of the feedback controller. The visual processing encapsulates the process of extracting an error based on the vision of the hand and the object. Lower Motor Centers encapsulates the functionality involved in transforming the motor signal into actual commands sent to muscles.....	125
Figure 7.7 The leaky integrator implementation of the feedback circuit that solves the inverse kinematics problem for precision grasping. See text for the explanation	126
Figure 7.8 Three grasping tasks executed by the feedback circuit proposed shown on the upper half of the figure. The change of arm/hand configuration during the execution is illustrated by snapshots of the arm/hand. Each hand figure is accompanied (lower half) by the error plot. The grasp execution is stopped (success) when the sum of finger distances to their target was less than 2mms.	128
Figure 7.9.The Visual feedback circuit generating desired trajectories for the 'lower level motor centre' (implemented as a PD controller).....	129

Figure 7.10 The slow (2 seconds) (lower half) and fast (0.5 seconds) (upper half) performance of the 'visual feedback servo' + 'PD controller' system is shown. The right hand side graphs show the tracking error (of the wrist) versus time. In the slow case, the object can be grasped but in the fast case, it is missed.....	130
Figure 7.11 The F5 mirror neurons viewed as the memory based feed-forward controller. The arrows below the sheet of neurons indicate outputs while the arrows coming above the sheet indicate inputs.....	133
Figure 7.12 The arm configurations that were acquired during 6 object grasping actions are shown. Each of the superimposed configurations is represented by a unit in the feed-forward layer	134
Figure 7.13 The trajectory generation with feedback and feed-forward control is illustrated for comparison with Figure 7.8 (feedback-only system). In the lower panel the error graphs are plotted as error versus iteration. The error is the sum of squared distances of the fingertips to their targets. The rightmost object was not grasped in the training (a novel object/location). Thus the system could not make use of the feed-forward signal, approximately after iteration 25 and switched to feedback only mode, resulting in slower positioning of the fingers on the target locations	135
Figure 7.14 The feedback, feed-forward and lower level motor servo and the dynamic arm was simulated all together. Upper half: The grasp lasted 0.5 seconds. Lower half: the grasp lasted 2 seconds. The fast movement error reduced with a factor of 6 while the slow movement reduced with a factor of 10 in terms (compare with Figure 7.10).....	136
Figure 7.15 The top row demonstrates two trajectory-planning examples for grasping without obstacle. The bottom row demonstrates how new trajectories are formed by introduction of an obstacle as a local inhibition on the feed-forward controller units	137
Figure 7.16 The feed-forward unit activations for four grasp observations shown as unit versus time. Each graph consists of 157 neurons acquired during the learning phase (the rows). The columns represent the time.....	138
Figure 7.17 A mirror neuron recorded during a grasp observation. On the left the raster plots; on the right the histogram. The recording data shown spans 2 seconds. In addition, the hand start to move approximately at time = 1 second indicated by the vertical bar at the centre of the raster panel (Rizzolatti and Gallese 2001)	139
Figure 7.18 Top row: Real mirror neuron recording during a precision grasp. Bottom row: One of the feed-forward controller unit's responses to vision of a grasping action. In the left panels, each raster row corresponds to a trial (Poisson spike generation for the model). The right panels show the histograms. The rasters aligned according to the contact of the hand with the object	140

Figure 7.19. The similarity of a real neuron and model unit is demonstrated. Left two panels real mirror neuron rasters and histogram. Right two panels are the model generated rasters and histogram. A slow increasing activity is observed in both cases140

Figure 7.20 Left: a sharp mirror neuron activity, which could only be replicated with our simulator by reducing the receptive field. Right: Similar response profile obtained from one of the feed-forward module units.....141

Figure 7.21 The population activity of feed-forward units with smaller receptive fields. The feed-forward unit we used to match the real mirror firing profile in Figure 7.20 is marked with an ellipse141

1 INTRODUCTION

Imagine two friends chatting over a coffee table. At the instant when one of them wants to get a sip of coffee, he effortlessly reaches and grasps his cup, shaping his hand according to the properties of the cup. If it were a mug, he would reach and grab the handle so as to counteract gravity; if it were a paper cup, probably he would grasp the cup with his whole hand covering the surface of the cup unless the coffee is too hot, forcing him to grab the cup from its outer rim. When he starts reaching for the cup, his friend easily understands that he wants to drink coffee even before his hand contacts the mug. Probably the way he reaches and shapes his hand conveys the information that he is not aiming at other objects on the table. He could possibly even infer that his coffee is hot before he grabs it. The situation for the two is reciprocal; they switch roles of being observer and actor as they sip their coffee.

When considered individually each of them is engaged in two tasks -*grasping* and *observing*. The former is a *goal directed movement* while the latter is a *perceptual task*. The grasping task requires the integration of information from a variety of sources (MacKenzie and Iberall 1994). The context and visual analysis of an object determines the general grasp plan. Proprioceptive (during reach), tactile and kinesthetic information (after contact) are used to guide the hand and arm to complete the grasp. Thus, the task of grasping involves, at the least, the determination of the following information:

1. How to conform the hand to the object. Based on object properties such as the shape, orientation and size the questions of 'which side of the object the hand should approach'; 'which fingers should be engaged', and 'what should the appropriate wrist rotation be' must be answered.
2. How to control the limbs. According to the physical properties of the environment and the biomechanics of the limbs, a control mechanism must engage muscles to transport the hand and shape it to match the specifications given by (1).

The former is the problem of *grasp planning*; the latter is the problem of *grasp execution*, which involves dynamics, that is, the adjustment of forces that the muscles must exert to achieve the specified plan. Many other information sources are integrated to refine the reach and grasp plan. For example, obstacle avoidance, speed, and accuracy requirements affect the reach component while the force requirements to secure a heavy slippery object or to handle a delicate object affect the grasp component.

The perceptual task, in essence, does not involve determining movement related parameters, as no movement has to be made. Nevertheless, the observer recognizes the action even before it is complete. Thus, the observer has to analyze the motion of the hand and its relevance to the target to determine whether the hand approach and preshape would yield a grasping behavior. It is interesting to note that there is some similarity in action recognition and action generation in terms of the underlying computations.

In fact, if one could compare the activity of observer's *brain* while he is engaged in grasping versus while he is observing his peer grasping, one would see that the *motor related* regions of the observer's brain was active in both observation and execution. Thus, one could conclude that the observer's brain *mirrored* the action of his peer by establishing equivalence between the observed action and his grasp plan.

The mechanism I caricaturized above is the focus of this thesis. The execution-observation matching system as introduced above *does* exist in monkey. There is strong evidence that human brain is also endowed with similar mechanism.

To be precise, this thesis investigates the cortical mechanisms involved in:

1. Translation of a visual description of an object into an appropriate grasp plan that is, *learning to make motor plans* that yield grasps that are appropriate for the target objects
2. Mapping of observed grasp actions into internal motor representations
3. Developmental processes shaping neural circuits to provide the functionality (1)
4. Developmental processes of (2), that is *learning to recognize* observed grasp actions based on self-executed grasps

The thesis analyzes (human and monkey) behavior and monkey neurophysiology from a developmental perspective, and constantly asks the questions: what is the underlying factor that give rise to such mechanisms? How does it shape the basic schemas of newborns into a functional form? The aim of the thesis is to give answers to these questions via computational models that learn and adapt starting from minimal bootstrapping behaviors. The models and the hypotheses developed in the thesis are based on:

1. Human infant motor development studies
2. Human behavioral and neuroimaging studies
3. Monkey neurophysiology and neuroanatomy studies

The thesis also presents significant predictions that can be experimentally tested with the hope that experimentalists will be stimulated to conduct the experiments suggested or design new experiments to test the model predictions and further uncover details of the cortical mechanisms of action recognition (mirror neurons), visuomotor learning and motor planning. The results of these experiments then would feedback into the modeling presented here, leading to validation (or rejection) and refinements of the models developed.

The organization of the thesis is as follows:

Chapter II presents the basic biological background with an emphasis on the brain areas that will be the focus of modeling. ‘Mirror Neurons’ (Dipellegrino et al. 1992; Gallese et al. 1996; Rizzolatti et al. 1996a) of the monkey premotor cortex are introduced in this chapter. The research on locating mirror neurons in human is also reviewed in Chapter II.

Chapter III develops the *Mirror Neuron System (MNS)* model based on the hypothesis that self-observation of grasping movements mediates the adaptation of parietal-premotor and premotor-premotor connections. Using simulation results, the chapter presents predictions on the *timing* of mirror neuron responses (and others) and suggests neurophysiological experiments for testing the predictions. In addition, Chapter III introduces the grand schemas of *Visual Analysis* and *Reach and Grasp*. The simulated 3D arm/hand of the *Reach and Grasp* schema is used in other chapters to graphically display the learned grasp actions.

Chapter IV develops a reinforcement learning based neural architecture that is capable of representing multiple values of a variable in terms of its probability distribution. The probability distributions are shaped through interaction with the environment so as to reflect the reward distribution in the environment. Chapter IV shows how layers can be connected to build multilayer reinforcement networks that are capable of representing conditional probability distributions. The architecture present in Chapter IV is used by Chapter V and Chapter VI.

Chapter V develops the Infant Learning to Grasp Model (ILGM) based on infant literature. ILGM is a schema level behavioral model that reproduces many infant behaviors and produces testable

hypotheses through simulation results. The notable property is that ILGM starts with a very basic repertoire of action mimicking neonates and show how a range of grasping categories can emerge via explorative interaction with the environment.

Chapter VI introduces the neurophysiological Learning to Grasp Model (LGM) as a neural level instantiation of ILGM constrained by monkey neurophysiology and neuroanatomy. LGM replicates existing premotor cortex findings such as the object selectivity of F5 canonical neurons and yields testable predictions about the grasp learning circuit in monkeys. This chapter also, poses serious questions to neurophysiologists on the assessment used to relate neuron firing to behavior. In particular, Chapter VI argues, by simulation results and examples from literature, that behavior-neural activity *correlation* is not an appropriate measure for investigating brain mechanisms of movement planning. The chapter suggests an experimental methodology for deciphering the neural substrates of movement planning suitable for finding the neuron populations that encode the true movement *control variables (parameters)*.

Chapter VII asks the question, 'why do the mirror neurons exist?' The hypothesis of the chapter (introduced in Chapter III) is that mirror neurons evolved initially to provide visual feedback for manual manipulative actions, and later became effective in recognizing actions of others. Chapter VII presents a simplified model of grasping (planar arm/hand) with increasing degrees of complexity in its control. First, a biologically realistic stochastic gradient following visual feedback circuit that can perform precision grips is developed. Then a memory based neural feed-forward circuit is introduced to augment the visual feedback servo circuit. After assessing the behavioral properties of the integrated visual servo circuit, the chapter analyzes the activities of individual memory units in the feed-forward circuit. The activities are converted into spike patterns using Poisson model of neural firing for a direct comparison with mirror neuron data. The results indicate that some of the feed-forward units' firing patterns are very similar to mirror neurons', in spite the fact that the feed-forward activity is characterized by a grasp error map. This supports the hypothesis that mirror neuron system initially evolved to serve as a visual servo circuit for manual manipulation.

Chapter VIII summarizes the key results and presents the predictions with high impact potential.

Chapter IX concludes the thesis by pointing out some open questions and recommending future research directions.

2 CHAPTER II: BIOLOGICAL BACKGROUND

In this chapter, we review the literature on the brain areas involved in the mirror neuron system functioning and the grasp related visuomotor transformation. We present major findings from neurophysiological and neuroanatomical (connectivity) literature to form a background for modeling chapters. The main regions of interest are intraparietal sulcus and premotor areas. A supplementary review of the superior temporal sulcus and other parietal areas is included. The posterior parietal cortex is involved in sensory-motor transformations, combining various sensory inputs and computing representations that are used by the motor system to generate movements. In turn, premotor cortex is involved in generating motor plans based on parietal representations of object affordances. The superior temporal sulcus performs visual analysis of motion and form including biological motion providing visual input to the parietal areas.

2.1 Abbreviations

7a, parietal area 7a; area PG	IT, inferotemporal cortex
7b, parietal area 7b; area PF	AIP, anterior intraparietal area; part of area 7
7ab, PFG, parietal area 7ab, area PFG	cIPS, caudal intraparietal sulcus; part of area 7
7ip, parietal area 7ip	LIP, lateral intraparietal area, part of area 7
7m, mesial part of area 7; area PGm	MIP, medial intraparietal area; part of area 5/7
F1, primary motor cortex; M1	VIP, ventral intraparietal area, part of area 5/7
F2, caudal part of dorsal premotor area; PMdc	MST, medial superior temporal area; part of STS
F3, supplementary motor area; SMA	MT, middle temporal area; part of STS
F4, caudal part of ventral premotor area; PMvc	PE, parietal area PE; part of area 5
F5, rostral part of ventral premotor area; PMvr	PEa, parietal area PEa; part of area 5
F6, pre-supplementary motor area; pre-SMA	PEc, parietal area PEc; part of area 5
F7, rostral part of dorsal premotor area; PMdr	PO, parieto-occipital cortex
SPL, superior parietal lobule	V6A, visual area 6A, area 19
IPL, inferior parietal lobule	S1, first somatosensory cortex, SI
STS, superior temporal sulcus	S2, second somatosensory cortex, SII
IPS, intraparietal sulcus	

2.2 Premotor areas

The macaque inferior premotor cortex is located ventral from the spur of the arcuate sulcus (see Figure 2.1) and considered to be involved in reaching and grasping movements (Rizzolatti et al. 1988). This region has been further partitioned into two sub-regions: F5, the rostral region, located along the arcuate and F4, the caudal part (see Figure 2.1).

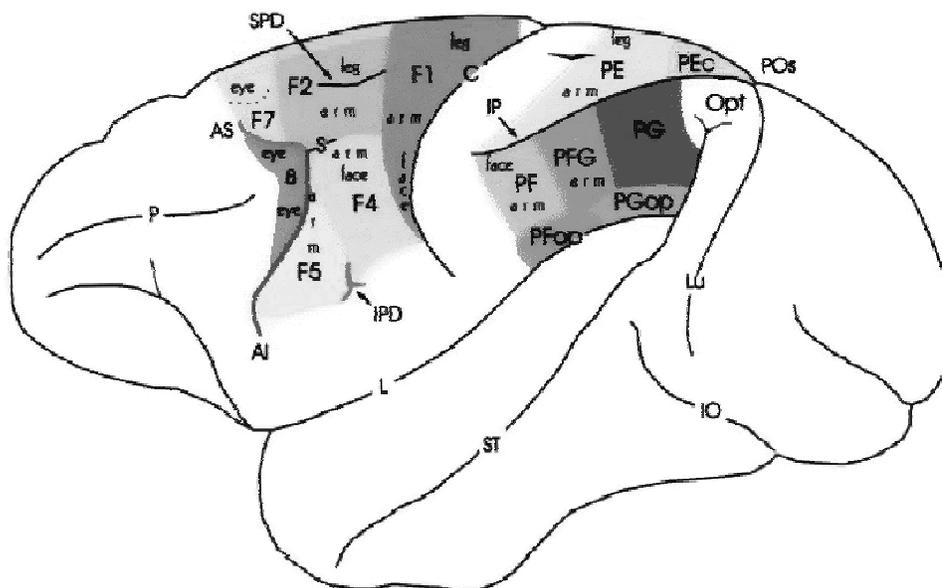


Figure 2.1 Lateral view of macaque brain showing the areas of agranular frontal cortex and posterior parietal cortex (adapted from Geyer et al. 2000). The naming conventions: frontal regions, Matelli et al.(1991); parietal regions, Pandya and Seltzer (1982)

The neurons in F4 appear to be primarily involved in the control of proximal movements (Gentilucci et al. 1988), whereas the neurons of F5 are involved in distal control (Rizzolatti et al. 1988).

2.2.1 Area F5

Area F5 is one of the various agranular frontal areas of particular interest due to its complex function (Matelli 1986). In the monkey, this area lies immediately caudal to the inferior arm of the arcuate sulcus. Stimulation and recording experiments showed that F5 is concerned with both hand and mouth movements. Hand movements are represented mostly in its dorsal part while mouth movements tend to be ventrally (Rizzolatti et al. 1988). Little is known about the functional properties of mouth neurons, however hand neurons were extensively studied.

2.2.1.1 Motor properties

Hand neurons discharge during specific goal-related movements such as grasping, tearing, manipulating and holding (Rizzolatti et al. 1988). Many of them are specific for a particular type of hand movement (Rizzolatti et al. 1988). In addition, some F5 neurons become active at the presentation of three-dimensional objects, in the absence of any overt movement, similar to AIP neurons that become active when the monkey fixates on a presented object. In many cases these visually triggered discharge requires a congruence of the presented object to the grip coded by the neuron (Murata et al. 1997a).

Rizzolatti et al. (1988) found that most F5 neuron firings correlated with specific goal related distal motor acts rather than with single movements made by the animal¹. Using the motor acts as the classification criterion, they subdivided the neurons into different classes such as *grasping-with-the-hand-and-the-mouth*, *grasping-with-the-hand* and *holding* neurons. The discharge of many F5 neurons depended on the way in which the hand was shaped during the motor act. For example the three main type of neurons found by Rizzolatti et al. (1988) were *precision grip*, *finger prehension* and *whole hand prehension* neurons. Furthermore, almost all of the neurons would discharge when the action was performed with

¹ Some proximal type neurons have been identified as well.

either hand. In addition, Rizzolatti et al. (1988) reported that 20% of the recorded neurons had visual response properties and they required motivationally meaningful visual stimuli to be triggered. Furthermore, they observed that, in the case of distal neurons, there was a relationship between the type of prehension coded by the cells and the size of the stimulus (presented object) effective in triggering the neurons. However, note that the purely motor related neurons constitute the majority of F5 neurons (Gallese 2002).

2.2.1.2 Visual properties: canonical neurons

Murata et al. (1997a) studied the properties of object related activity of F5 neurons. The result of their study indicates that some F5 neurons encode object shapes in motor terms. That is, every time an object is presented, its visual features are automatically translated into an internal motor representation. The translation takes place whether a motor response is required or not. Therefore, these neurons are not intention related.

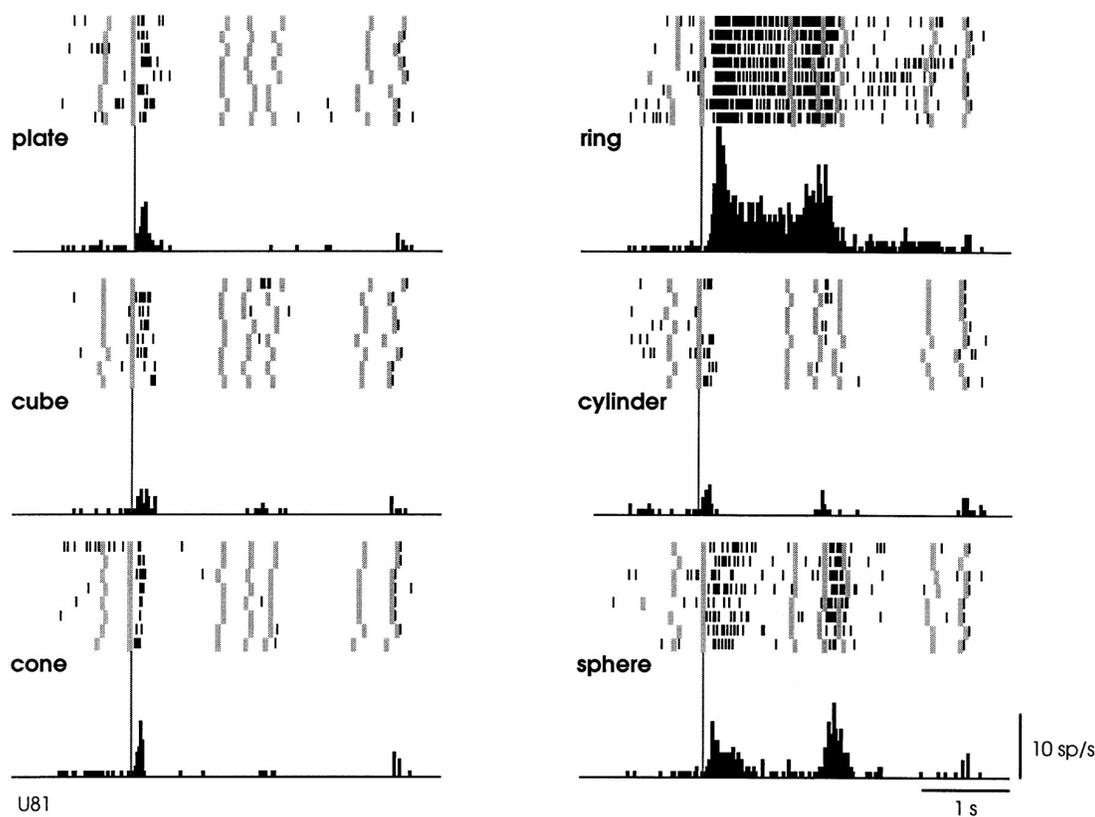


Figure 2.2 A canonical neuron response during grasping of various objects in the *dark* (left to right and top to bottom: plate, ring, cube, cylinder, cone and sphere). The rasters and histograms are aligned with object presentation. Small grey bars in each raster marks onset of key press, go signal, key release, onset of object pulling, release signal, and object release, respectively. The peaks in ring and sphere object cases correspond to the grasping of the object by the monkey (adapted from Murata et al. 1997a)

The similarity of the AIP and F5 visual neuron responses suggests that they may be part of a visuomotor transformation circuit. This view is supported by the reciprocal connections between F5 and AIP (Sakata et al. 1997a). Figure 2.2 shows a canonical neuron's response during motor execution. To test whether the motor related activity was due to the vision of the object, the trial was performed in the dark. The neuron was primarily responsive for ring grasping and a lesser extend the sphere grasping.

The same neuron's response in the object fixation, without any subsequent grasp requirement, is shown in Figure 2.3. It is important to note that the motor preference of the neuron is reflected in the visual fixation condition as well.

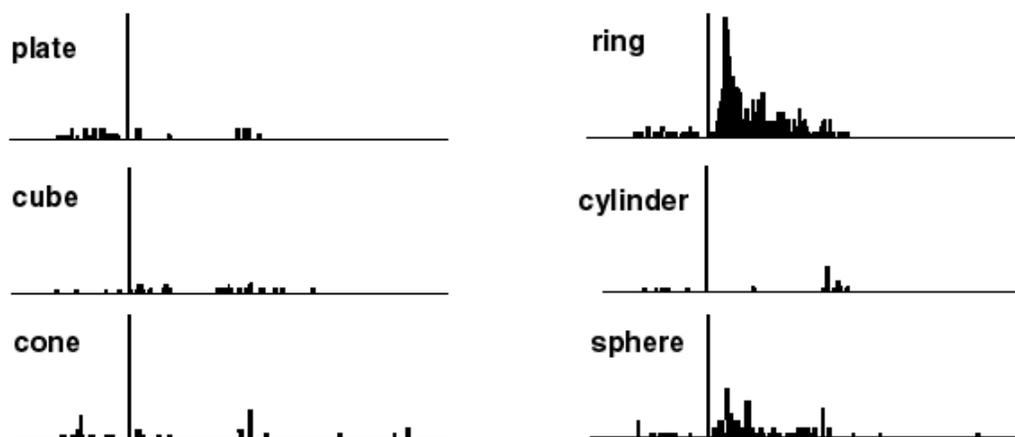


Figure 2.3 The motor responses of the same neuron shown in Figure 2.2. The motor preference of the neuron is also carried over to the visual preference (compare the ring and sphere histograms of both figures) (adapted from Murata et al. 1997a)

2.2.1.3 Visual properties: mirror neurons

Recording studies of the rostral part of inferior area 6 (area F5) region showed that some of the visual neurons were responsive to action observation (Gallese et al. 1996; Rizzolatti et al. 1996a; Dipellegrino et al. 1992). The cells with action observation property have been located on the convexity of the bank of arcuate sulcus. Like other F5 neurons, mirror neurons were active when the monkey performs a particular class of actions. However, in addition the mirror neurons became active when the monkey observes the experimenter or another monkey performing an action (Gallese et al. 1996; Rizzolatti et al. 1996a; Dipellegrino et al. 1992). In most of the mirror neurons, there was a clear relation between the coded observed and executed action. The actions studied so far include grasping, manipulating and placing. The congruence between the observed and executed action varied. For some of the mirror neurons, the congruence was quite loose; for others, the general action (e.g. grasping) and the way the action was executed (e.g. power grasp) had to match in order to activate to neuron (Gallese et al. 1996; Rizzolatti et al. 1996a). An important observation was that mirror neurons required an interaction between the experimenter and the object; the sight of the experimenter or the object alone was not enough to trigger mirror activity. (Gallese et al. 1996; Rizzolatti et al. 1996a) All the neurons were studied by examining their discharge while the experimenter performed a series of motor actions in front of the monkey. These actions were related to food grasping and manipulation and other objects grasping and manipulation. In order to verify whether the recorded neuron coded specifically hand-object interactions a series of actions such as mimicking grasping without any object, prehension actions with tools, mimicking grasp with spatially separated object were performed. All experimenter's actions were repeated on different positions (e.g. left,-right, far-close). Of the 532 recorded neurons, 92 of them showed mirror property (i.e. they discharged both when the monkey made active movements and when it observed specific meaningful actions performed by the experimenter) (Gallese et al. 1996).

The two important aspects of the mirror neurons are (1) they are robust, they don't habituate and (2) the distance of the experimenter to the monkey does not affect the response intensity of the cell.

Most of the neurons are active during observation of a single action: for example in the study of Gallese et al. (1996). 51/92 of the cells preferred only single action; 38/92 of the cells preferred two or three actions; 3/92 of the cells were active for both hand or mouth grasps. The motor properties of these neurons were indistinguishable from those of other F5 neurons. They had preference for certain actions: 60/92 cells responded when the animal performed only a grasping action. 9/92 cells fired when the animal grasped with his mouth. 11/92 of cells fired for both hand and mouth grasps (Gallese et al. 1996). The remaining 14 neurons had the distribution: tearing (2), bringing to the mouth (2), manipulating (8). The light and dark conditions were employed for 14 cells to test whether the motor property was a result of self-hand vision. All the tested neurons confirmed that, the discharge was not due to self-vision (Gallese et al. 1996).

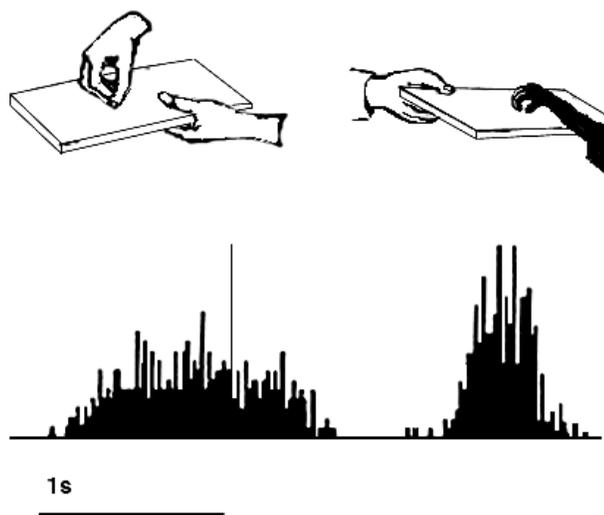


Figure 2.4 Activity of a cell during action observation (left) and action execution (right). There is no activity in presentation of the object during both initial presentation and bringing the tray towards the monkey. The vertical line over the histogram indicates the hand-object contact onset. (from Gallese et al., 1996).

Figure 2.4 shows the dual response property of mirror neurons. The recorded neuron in the figure was silent during the presentation of the object, but started firing when the experimenter picked up the object. The neuron, interestingly, did not fire during the time the tray was moved towards the monkey and finally it started firing again when the monkey picked up the object. Note that during the period when the tray was moved towards to monkey it could predict that he would grasp the object (Gallese et al. 1996)

Figure 2.5 shows the specificity of a grasp related mirror neuron where the experimenter performed (A) a precision grip, (B) a whole hand prehension, and (C) mimicked a precision grip. The notable property of this neuron is that miming the action was not effective in activating the neuron.

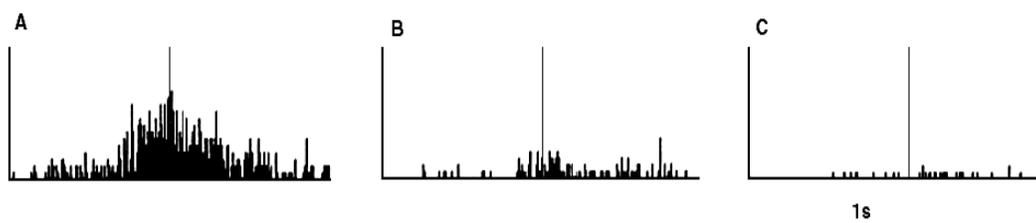


Figure 2.5 Visual response of a mirror neuron. A. Precision grasp B. power grasp C. mimicking of precision grasp. The vertical lines over the histograms indicate the hand-object contact onset. (adapted from Gallese et al., 1996)

In most mirror neurons, there is a relationship between the visual action they respond and the motor action they code. The mirror neurons studied by Gallese et al. (1996) were divided into three, according to their visuomotor congruence: *strictly congruent*, *broadly congruent* and *non-congruent*. A neuron was labeled as strictly congruent when the effective observed and executed actions match both in terms of general action type (e.g. grasp) and in terms of how the action was executed (e.g. power grasp). Figure 2.6 shows a strictly congruent neuron.

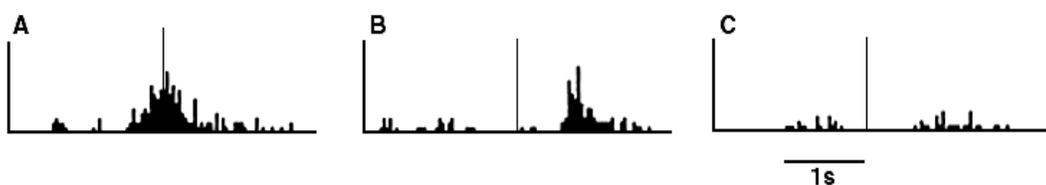


Figure 2.6 Example of a strictly congruent manipulating mirror neuron: A) The experimenter retrieved the food from a well in a tray. B) Same action, but performed by the monkey. C) The monkey grasped a small piece of food using a precision grip. The vertical lines over the histograms indicate the hand-object contact onset (adapted from Gallese et al., 1996).

The number of strictly congruent neurons found was 29/92. The number of broadly congruent neurons was 56/92 (Gallese et al. 1996). In the case of broadly congruent neurons, there was a link between the executed action and the observed preferred action. These neurons were further sub-classified according to their motor strictness: If a broadly congruent neuron fired only for one motor act (e.g. grasp) with only a single hand configuration (e.g. precision) then it would be of the first type. On the other hand, if the neuron fired for one motor act but the way the action was performed did not affect the firing then it would be of the second type. The third and last type of broadly congruent neurons appear to be activated by the goal of the observed action (Gallese et al. 1996). Finally, the neurons with no apparent congruence were labeled as non-congruent (7/92). Figure 2.7 shows the classification of F5 neurons including the mirror neuron types discussed.

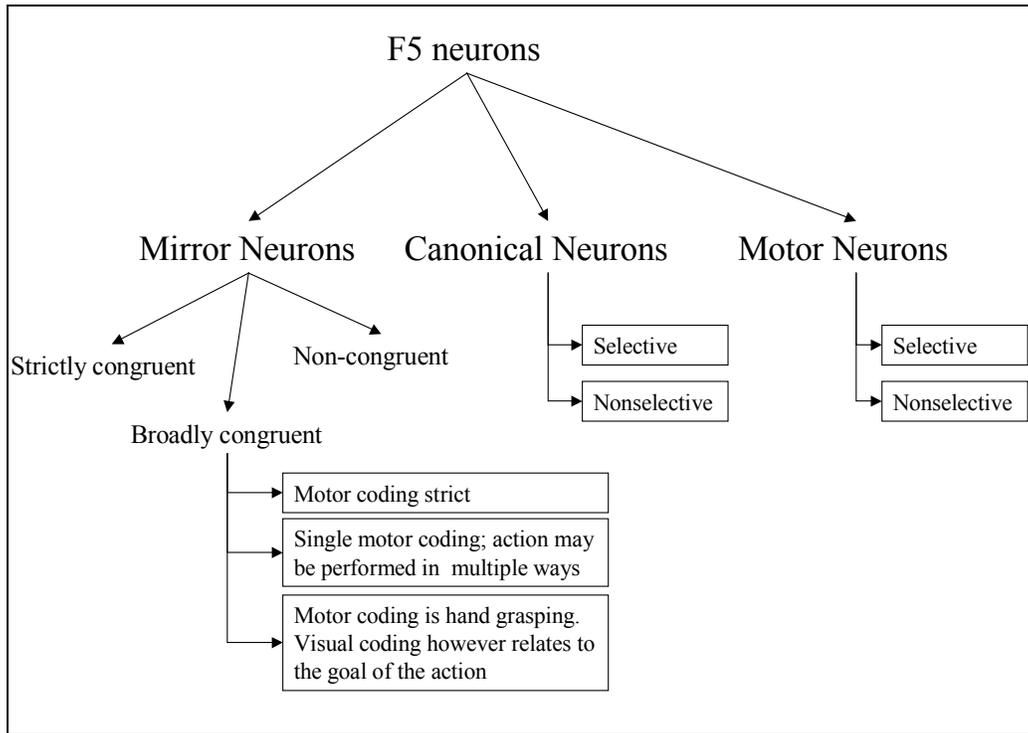


Figure 2.7 The classification of area F5 neurons derived from published literature (Dipellegrino et al. 1992; Gallese 2002; Gallese et al. 1996; Murata et al. 1997a; Murata et al. 1997b; Rizzolatti et al. 1996a; Rizzolatti and Gallese 2001). All F5 neurons fire in response to some motor action. In addition, canonical neurons fire for object presentation while the mirror neurons fire for action observation. The majority of hand related F5 neurons are purely motor (Gallese 2002)(labelled as Motor Neurons in the figure)

Fogassi et al. (1998) found that area F5 was not the only area that had mirror neurons. The rostral part of the inferior parietal lobule of the macaque monkey (area 7b or PF) also has neurons with similar mirror properties. Although some neurons with strict congruence of the executed and observed action have been found, the majority of the neurons studied had limited congruence (similarity) or no congruence at all (Fogassi et al. 1998). 8/43 PF mirror neurons were strictly congruent; 9/43 had low level of congruence (a similarity); and the majority (26/43) were non-congruent (Fogassi et al. 1998). The main cortical input to area F5 comes from the inferior parietal lobe, and in particular areas AIP and PF (Matelli 1986). The similar properties of F5 canonical neurons with AIP neurons, and F5 mirror neurons with PF neurons, suggests that these three areas work together for visuomotor transformation and action recognition.

2.2.2 Area F4

Area F4 (see Figure 2.1) is connected with area F3 and to a lesser extent, to area F6 (Geyer et al. 2000). Area F4 projects to primary motor cortex (F1). The main parietal input to area F4 comes from VIP (Geyer et al. 2000). In area F4 the space is coded in body-parts-centred coordinate frame (e.g. centred on the hand) (Fogassi et al. 1996). When the body-part-moves the coordinate system follows, but when the gaze moves the coordinate frame stays anchored on the body-part (Fogassi et al. 1996). Many of F4 neurons fire during reaching movements of the proximal arm but not the movements of the distal arm. The neurons usually have somatosensory receptive fields that match the movement direction of the limb (Gentilucci et al. 1988). It is suggested that VIP-F4 circuit transforms object locations into motor

plans to reach towards them as area F4 sends descending projections to the brain stem and spinal cord (Rizzolatti et al. 1998).

2.2.3 Areas F2 and F7 (dorsolateral prefrontal cortex)

Area F2 (see Figure 2.1) neurons can be grouped into three different classes: (1) signal related neurons, (2) set-related neurons, and (3) movement-related neurons. (see Geyer et al. 2000 for a review). Signal related neurons are activated right after visual instruction stimuli and have phasic response. Set-related neurons show sustained activity after the instruction stimulus during the delay period. Movement related neurons start firing after the trigger signal. Area F2 receives somatosensory input from areas PEip and PEc, and visual input from areas MIP and V6A. Rizzolatti et al. (1998) suggested that F2 can use the MIP and V6A inputs in controlling arm position during the transport of the hand to spatial targets.

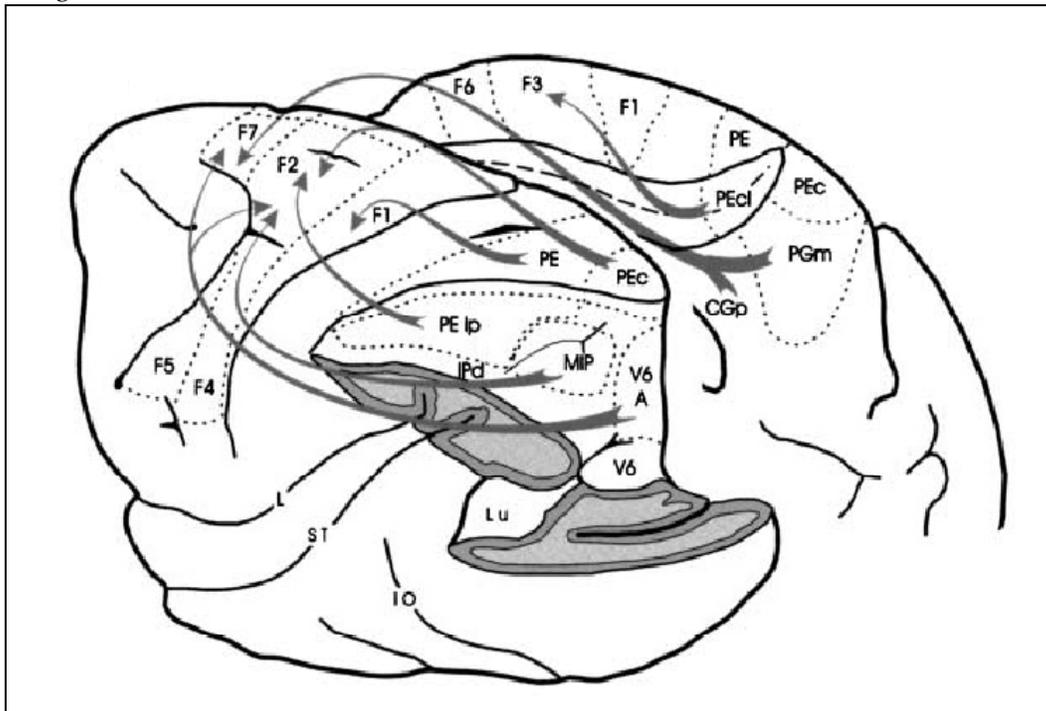


Figure 2.8 The macaque parieto-frontal projections from mesial parietal cortex, medial bank of the intraparietal sulcus and the surface of the superior parietal lobule (adapted from Rizzolatti et al. 1998). Note that the Brodmann's area 7m corresponds to Pandya and Seltzer's (1982) area PGm

Area F7 receives inputs from area 7m (Ferraina et al. 1997a; Ferraina et al. 1997b) (see

Figure 2.8). The neurons in area F7 fire in response to arm movements (Caminiti et al. 1991; Crammond and Kalaska 1996) or visual stimuli (Shen and Alexander 1997b). However, in contrast to area F2, area F7 visual response does not depend on a pending movement (di Pellegrino and Wise 1991). It appears that the 7m-F7 circuit is important for conditional movement selection (Geyer et al. 2000). The other projection to area F7 is from LIP (Lewis and Van Essen 2000), where saccade related target memory activity is represented. The neuronal activity in LIP area can be modulated by attention and eye position (see Colby and Duhamel 1996 for a review of LIP neuron responses). Thus, LIP-F7 circuit may be important for complex saccade control (Geyer et al. 2000).

2.2.4 Area F1 (the primary motor cortex)

Area F1 (see Figure 2.1) is organized somatotopically, where the body parts that require finer movements are represented over a larger cortical surface than the body parts that require less precision. Each neuron may contribute to multiple spinal neuron pools. The motor parameters that are encoded by F1 neurons are usually a combination of the following physical parameters: force, rate of change of force, joint position or the velocity of the movement (Pandya and Seltzer 1982). However, it is possible to get meaningful physical parameters using a *population* of F1 neurons. Georgopoulos et al. (Georgopoulos et al. 1982) trained monkeys to perform radial outward reaches to a target light. Recording over a population of primary cortex neurons they showed that each neuron fired maximally for a direction (preferred direction), and fired less and less as the direction deviated from the preferred direction. Given a population, the weighted sum of the preferred direction vectors, the population vector, predicted the monkeys reaching direction.

The subcortical input to F1 is relayed by thalamus (see Matelli et al. 1989 for the distinct nuclei projecting to F1). The corticocortical inputs to hand area of F1 comes primarily from supplementary motor area (SMA) and to a lesser extent from the lateral premotor cortex. The other inputs are from area 1, 2 and 5 (Ghosh et al. 1987). Approximately half of the corticospinal projections are formed by area F1 neurons (Dum and Strick 1991).

2.2.5 Areas F3 (SMA proper), F6 (pre-SMA)

Area F3 is somatotopically organized where arm and leg representations run as two oblique dorsorostral-to-ventrocaudal directions (see Figure 2.1). In addition, area F3 has an orofacial representation, while area F6 has only an arm representation (Luppino et al. 1991).

Areas F3 and F6 have different patterns of thalamic input indicating that they are part of different motor loops with different functions (Luppino et al. 1991). Cortical input to area F3 originate mainly from areas F2, F4, F5, F6 and F7, and the primary and secondary somatosensory cortices and the posterior parietal areas PE and Peci, and the cingulate and the primary motor cortex. On the other hand, area F6 is mainly connected with areas F5 and F7, followed by the prefrontal and cingulate cortex, F2, F3 and F4, and to lesser extent with the posterior areas PG, PFG and superior temporal sulcus (Geyer et al. 2000).

2.3 The superior temporal sulcus

In the macaques's brain, posterior parietal cortex and the cortex of caudal superior temporal sulcus (STS) have been subdivided into numerous areas mainly involved in spatial analysis of the visual environment and in the control of spatially oriented behaviour (Maioli et al. 1998). The cortex of superior temporal sulcus (STS) contains neurons that are selective for biological motion observation such as limb movements and full body motion. Perrett et al. (1990b) reported STS neurons that were responsive to goal directed hand motion (Perret et al. 1990b; Perret et al. 1990a). PET studies showed that STS in human shows strong activation during biologically meaningful visual stimuli (Bonda et al. 1996) including goal-directed hand actions. In monkeys, some of the STS neurons that are triggered by biologically meaningful stimuli have two notable properties. Firstly, these neurons show responses to goal directed hand motion in a translation/scale/rotation invariant way (Perret et al. 1990b; Perret et al. 1990a). Secondly, these neurons do not require a pictorially realistic stimulus; they respond to point light stimuli (Perret et al. 1990b; Perret et al. 1990a) where the stimulus is just the movement of a small number of points. Bonda et al. (1996) also used this kind of stimulus – 3 lights for the arm and 2 for each finger – when they scanned the subjects during action observation.

2.4 Parietal Areas

Based on cytoarchitectonic and connectional criteria the inferior parietal lobule (Brodmann's area 7) includes areas 7a, 7b and 7ip (Cavada and Goldman-Rakic 1989). Area 7 reaches its highest development in primates (Cavada and Goldman-Rakic 1989). Damage to this area can cause impairments in spatial perception, neglect of sensory stimuli contralateral to the damage side, defects in visually guided reaching and oculomotor control (Ratcliff 1991; Stein 1991). Cavada & Goldman-Rakic (1989) divides area 7 in sub-areas of 7m, 7a, 7b, and 7ip. Area 7m is located on the medial surface of the hemisphere. This corresponds to Pandya and Seltzer's (1982) area PGm. Areas 7a, 7b lie on the convexity of the posterior parietal lobule (Cavada and Goldman-Rakic 1989). These regions correspond to Pandya and Seltzer's (1982) PG and PF respectively. Pandya and Seltzer's (1982) also distinguish the subdivisions of PGop and PFop in the lateral opercular part of PG and PF and area Opt in caudal PG. Area 7ip is situated in the posterior bank of intraparietal sulcus and referred as POa by Pandya and Seltzer (1982). In addition, the posterior half of 7ip corresponds to functionally defined areas VIP (Maunsell & Van Essen, 1983) and LIP (Andersen et al., 1985). Figure 2.9 shows the intraparietal sulcus (opened) and neighbouring parietal regions using Pandya and Seltzer (1982) nomenclature.

2.4.1 The anterior intraparietal area (AIP)

The anterior part of the lateral bank of the intraparietal sulcus (area AIP) (see Figure 2.9) is involved in extracting visual properties of objects relevant for grasping (Sakata et al. 1997a; Sakata et al. 1998; Sakata et al. 1995; Murata et al. 1996).

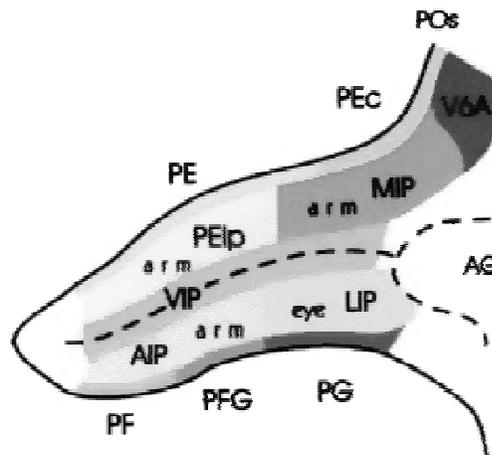


Figure 2.9 The intraparietal sulcus opened to show the anatomical location of AIP in the macaque (adapted from Geyer et al. 2000)

Neurons in area AIP are active either in relation to the grasping behavior alone or in relation to the vision of objects (Sakata et al. 1998; Sakata et al. 1997b; Taira et al. 1990). Some of the latter type are active exclusively for visual fixation. In one study, 21% of cells studied responded to simply fixating an object (visual-related), others (37%) were active only when a movement is being made to manipulate the object (motor-related) (Taira et al. 1990). However, many cells (37%) fell somewhere between these two extremes (visual-dominant) (Taira et al. 1990). Figure 2.10 shows the response of a visual-dominant neuron during different experimental conditions.

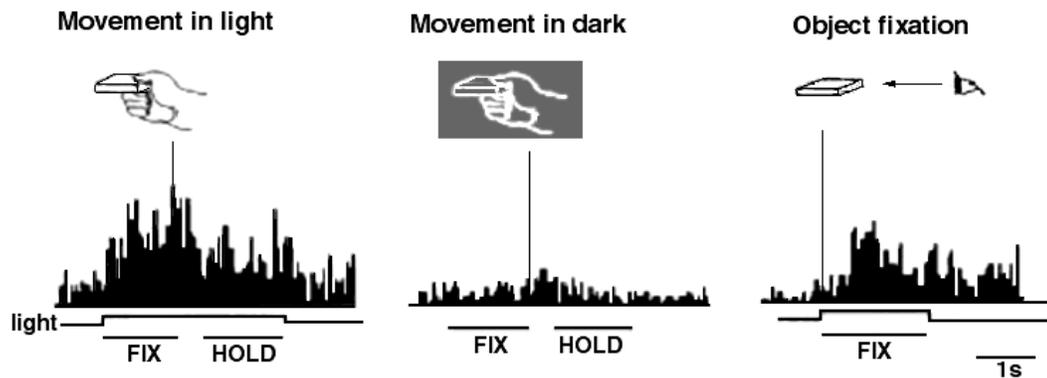


Figure 2.10 An AIP visual-dominant neuron activity under three task conditions: Object manipulation in the light, object manipulation in the dark and object fixation in the light. The neuron is active during fixation and holding phase when the action is performed in light condition. However, during grasping in dark the neuron shows no activity. The fixation of the object alone without grasping also produces a discharge (adapted from Sakata et al. 1997a)

The neuron shown in Figure 2.10 is active during fixation and holding phase when the action is performed in light condition. However, in grasping-in-dark condition the neuron shows no activity. The fixation of the object alone without grasping also produces a discharge, however the activity is less than the grasping-in-light condition.

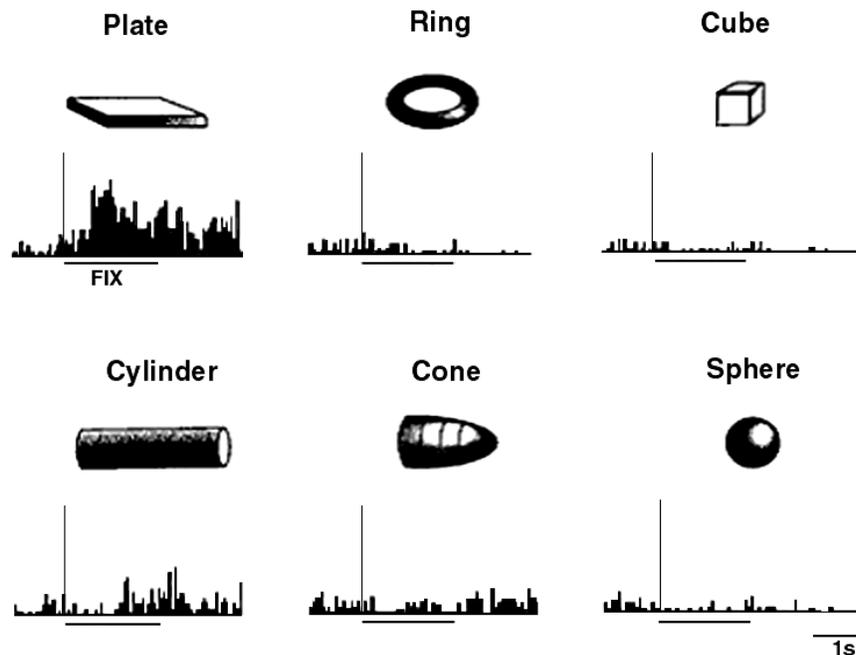


Figure 2.11 Activity the same neuron in Figure 2.10 during fixation of different objects. The neuron show selectivity for horizontal plate (adapted from Sakata et al. 1997a)

In addition, some of these neurons show object specificity (object-type visual-dominant neurons) which responds to the sight of complex objects such as a knob-in-groove and a plate-in-groove (Sakata

et al. 1997a). Figure 2.11 shows response profile of the same neuron in Figure 2.10 for different objects during fixation. The neuron has a strong preference for the plate shaped object.

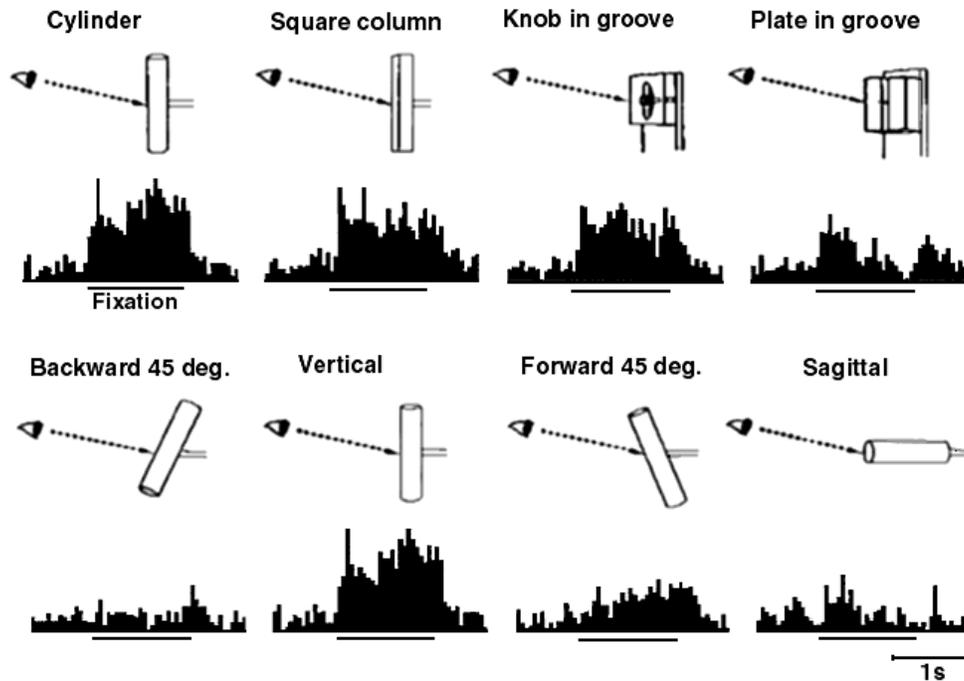


Figure 2.12 An AIP visual-dominant neuron's axis orientation tuning and object fixation response is shown. The neuron fires maximally during the fixation of a vertical bar or a cylinder. The tuning is demonstrated in the lower half of the figure (adapted from Sakata et al. 1999)

Furthermore some object-type visual-dominant neurons, show tuning according to the orientation of the longitudinal axis or the surface orientation of flat objects (Sakata et al. 1999; Murata et al. 2000). An example of an object-type visual-dominant neuron that showed tuning for the axis orientation regardless of the shape is presented in Figure 2.12. Top row shows the strong response of the neuron to a vertical cylinder, a square column, and a vertical knob-in-groove in the fixation condition. The bottom row of Figure 2.12 demonstrates the tuning of the neuron for different axis orientations.

The muscimol-induced lesions of area AIP lead to a significant deficit in monkey's ability to grasp objects (Sakata et al. 1997a; Gallese et al. 1994). The grasping movements become clumsy and uncoordinated, and as a result, the monkey is unable to shape his hand and orient his wrist appropriately for objects that are presented. However, the monkey can still execute the basic sequence of the task employed (Sakata et al. 1997a; Gallese et al. 1994).

2.4.2 The caudal intraparietal sulcus (c-IPS)

The lateral bank of the intraparietal sulcus (c-IPS) is involved in three-dimensional analysis of objects (Sakata et al. 1997a; Sakata et al. 1999). Some of these binocular visual neurons are selective for the orientation of the axis of the objects (AOS neurons) and some are selective for the surface orientation of the objects (SOS neurons) (Sakata et al. 1997a; Sakata et al. 1999). AOS neurons prefer long and thin objects as visual stimuli and are tuned to the three-dimensional axis orientation of the objects in space. Figure 2.13 shows the response of an AOS neuron when the object is viewed in binocular viewing

condition. Figure 2.14 shows the same neuron's response when the visual information is limited to the left or right eye indicating that binocular cues are important for driving the AOS neuron shown. SOS neurons prefer broad and flat objects as visual stimuli (Sakata et al. 1999). Complementary to AOS neurons; they are tuned to the surface orientation of objects in three-dimensional space (see Figure 2.15).

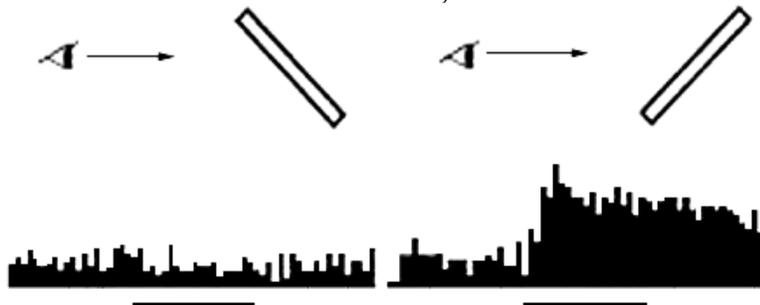
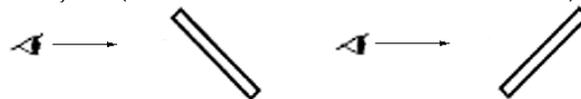


Figure 2.13 Response of an axis-orientation-selective (AOS) neuron in the caudal part of the lateral bank of the intraparietal sulcus (c-IPS) to a luminous bar tilted 45° forward (*left*) or 45 backward (*right*) in the sagittal plane. The monkey views the bar with binocular vision. The line segment under the histograms mark the fixation start and the period of 1 second. (adapted from Sakata et al. 1999)

It is suggested that c-IPS is a higher center for stereopsis, which integrates various binocular disparity signals received from the V3 complex and other prestriate areas to represent the neural code for geometric features of objects (Sakata et al. 1997a; Sakata et al. 1997b).



B. Monocular (Left eye)



C. Monocular (Right eye)



Figure 2.14 The response of the same neuron in Figure 2.13, for monocular vision conditions for the left and right eyes. (adapted from Sakata et al. 1999)

Sakata et al. (1997a) suggested that c-IPS could send projections to AIP and thus, contribute to the visual adjustment of the shape of the hand grip and/or hand orientation for manipulation and grasping. Figure 2.15 shows a SOS neuron that is selective for a surface that is 135 degrees tilted around the sagittal axis

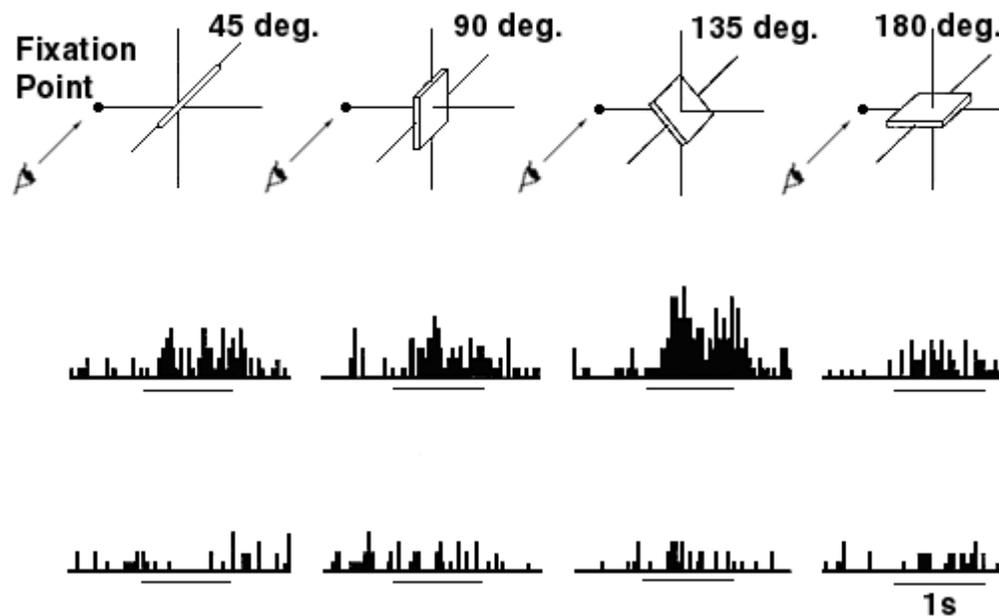


Figure 2.15 Orientation tuning of a surface-orientation selective (SOS) neuron. First row: Stimuli presented. Middle row: responses of the cell with binocular view. Last row: responses of the cell with monocular view (adapted from Sakata et al. 1997a)

2.4.3 Areas VIP, MIP and LIP

The intraparietal regions VIP, MIP and LIP (see Figure 2.9) encode the space around the animal with multiple reference frames for different movement purposes (Colby and Goldberg 1999). A crude separation is that VIP is involved in ultra-near space (less than 5cm from the face) (Colby et al. 1993b), MIP with stimuli within reaching distance (Colby and Duhamel 1991) and LIP with far visual stimuli (Colby and Goldberg 1999).

LIP coding has been implicated as attentional (Gottlieb et al. 1998), decision-related (Shadlen and Newsome 2001; Shadlen and Newsome 1996), visual target memory related (Gnadt and Andersen 1988) and motor intention related (Snyder et al. 2000; Snyder et al. 1997). Colby and Goldberg (1999) suggested a unifying functional role for LIP that it encodes the representation of salient spatial locations (with attentional tuning). They noted the distinctive property of neurons in LIP that their firing was not tied to any particular modality and the representation was limited to attended objects and their locations.

Neurons in LIP have retinotopic receptive fields, where they carry visual, memory, and saccade-related signals that describe stimuli in terms of the distance and direction of the stimulus or saccade location relative to the center of gaze (Colby and Goldberg 1999). VIP neurons represent visual locations using a continuum of eye centred to a head centred spatial reference frames (Bremmer et al. 1999; Duhamel et al. 1997). Eskandar and Assad (1999) found neurons with reaching-related activity encoding stimulus features, such as location and direction of stimulus motion. In addition, MIP neurons maintained the memory of a reach target during the delay period of a memory-guided reach task or when the target is obscured (Eskandar and Assad 1999; Snyder et al. 1997). When the hand direction and the visual target direction were disassociated through a well designed set up², it was found that

² The monkeys were trained to use a joystick to guide a spot of light to a circular target. While the animal fixated, two spots within two circular targets appeared, oriented along the preferred/null axis of the cell. The monkey used

MIP neuron activity correlated more with the hand direction than the object location. The opposite was true for LIP neurons (Eskandar and Assad 1999) .

2.4.4 Areas 7a and 7b (PG and PF)

The experimental findings indicate that area 7a, together with other inferior parietal lobule sectors, is involved in spatial coding. Researchers suggested various types of spatial encoding for area 7a. Stein (1991) suggested that area 7a represented extra-personal space. Andersen et al. (1999) suggested that area 7a represents targets in a world-centered coordinate frame. It has been shown that area 7a neurons are involved in the analysis of motion evoked during locomotion or by the manipulation of objects by the hands (Siegel and Read 1997). The different interpretation of area 7a responses can be due to either the non-homogenous functional distributions of neurons or due to the experimental setup differences (see the reviews: Andersen et al. 1997; Wise et al. 1997).

It has been found that reach-related activity in area 7a signaled specific phases of the motor performance (MacKay 1992). Further, it has been suggested that it could be used by the frontal lobe to facilitate upcoming elements of a motor sequence, including terminal corrections (MacKay 1992). Motter et al. (1987) identified visually sensitive and insensitive neurons in area 7a (Motter and Mountcastle 1981; Motter et al. 1987). The Neurons insensitive to visual stimuli comprised the fixation, oculomotor, and projection-manipulation classes, which were suggested to be involved in initiatives toward action (Motter and Mountcastle 1981). Most of the visually sensitive neurons were activated from large and bilateral response areas that excluded the foveal region. The visually sensitive neurons were responsive to stimulus movement and direction over a wide range of velocities. The movement vectors pointed either inward toward the center or outward toward the perimeter of the visual field. For bilaterally activated neurons, the vectors pointed in opposite directions in the two half-fields (opponent vector organization). Motter and Mountcastle (1981) suggested that the neurons could signal motion in the immediate surround.

Constantinidis and Steinmetz (1996) showed that a population of neurons in area 7a was active during the delay period of a spatial memory task that did not require a motor response directed toward the stimulus. Thus, it is suggested that the activity could represent a short-term memory trace for the spatial location of the stimuli (Constantinidis and Steinmetz 1996). In accordance with the spatial memory hypotheses, Maunsell (1995) indicated that the object location coding in area 7a was capable of representing visual stimuli without ever falling into the corresponding receptive field.

Another functional aspect of area 7a, the attentional tuning was studied by Constantinidis and Steinmetz (2001). Their results indicate that area 7a neurons represent the location of the stimulus that attracts the animal's attention and can provide the spatial information required for directing attention to a salient stimulus in a complex scene (Constantinidis and Steinmetz 2001).

According to our view the fundamental and unifying property of area 7a neurons, is that they can potentially be used to monitor the relation of body parts with respect to objects once they are fixated. A population of neurons that detect the motion of visual stimuli inwards to (or outwards from) the fixation point can encode the kinematics aspects (e.g. proximity) of a movement to satisfy a goal such as reaching or grasping. There is evidence that when humans perform reaching movements, they fixate to target objects or obstacles to plan reach actions (Johansson et al. 2001), which can be thought of

the joystick to move one of the spots toward the opposite target. On 'visible' trials, the moving spot remained visible throughout its trajectory, and the opposite spot disappeared at the start of movement. On 'occluded' trials, both spots disappeared without moving as soon as the animal moved the joystick, and the spot being guided then reappeared near the target, as if it had been moving smoothly behind an occluder.

registering the relevant locations in area 7a as a *saliency* map. This proposal is supported by the fact that the removal of areas 7a, 7ab and LIP caused marked inaccuracy in reaching in the light to visual targets but had no effect on reaching in the dark (Rushworth et al. 1997). In contrast, the removal of areas 5, 7b and MIP caused misreaching in the dark, but had little effect on reaching in the light. Therefore, Rushworth et al. (1997) suggested that the two divisions of the parietal cortex organize limb movements in distinct spatial coordinate systems: area 7a/7ab/LIP are essential for spatial coordination of visual motor transformations whereas areas 5/7b/MIP is essential for the spatial coordination of arm movements in relation to proprioceptive and efference copy information.

Other parietal areas that can be involved in hand-object relation signals are area 7m (Ferraina et al. 1997a; Ferraina et al. 1997b), and area V6a and area PEc (Caminiti et al. 1999; Battaglia-Mayer et al. 2000; Ferraina et al. 2001; Marconi et al. 2001).

Conventionally, area 7b is considered to be a somatosensory area (Andersen et al. 1990). Robinson and Burton (1980b) studied the somatic response properties of neurons from area SII and area 7b. One-half of the recorded 7b neurons responded only to somatic stimulation. Many neurons in the lateral parts of area 7b were vigorously activated by tactile stimulation. In spite the majority of somatic responses, some visual responses from area 7b were noted. The visual responses of 7b neurons were not studied in detail either because it was not the focus of interest (as in Robinson and Burton 1980a) or due to the complex response properties. In fact, it is possible to find considerable unimodal visual 7b neurons as well as the neurons that respond *only* to visual stimulation (Dong et al. 1994). The visual responses of 7b neurons can be based on the signals carried by the small projections from the visual cortical areas (Andersen et al. 1990).

Fogassi et al. (1998) studied some of area 7b neurons' visual properties. They found that the activity of some neurons were triggered by the observation of various hand actions performed by the experimenter. The neurons had motor properties similar to mirror neurons of area F5 (see section 2.2.1.3). The congruence between the action performed by the monkey and the observed action was usually low. The connection of area F5 with area 7b (Fogassi et al. 1998) indicates an intimate relation between 7b and F5 mirror neurons. Currently there are no detailed data on 7b mirror activity. However, unpublished results (Fogassi 1999) indicate that in addition to those neurons that have similar properties as F5 mirror neurons there exist mirror-like neurons that fire for simple arm/hand movement observations (in contrast to complete action observations).

2.5 Connectivity and other brain regions

According to Cavada and Goldman-Rakic (1989) 7m, 7a, 7ip are extensively connected with a number of visual areas located on the medial surface of the hemisphere and in the depths of parieto-occipital and intraparietal sulci. Areas 7m, 7a, 7ip, and to a much lesser extend 7b, are reciprocally connected with the visual temporal cortex, principally with the cortex of the superior temporal sulcus (STS) (Cavada and Goldmanrakic 1989). Although the density of 7b connections with the visual motion cortex of STS is largely surpassed by the extensive connections of 7b with somatosensory areas the interconnections of 7b with the visual regions are established through anterior 7ip, and the transitional cortex 7ab between 7a and 7b (Cavada and Goldmanrakic 1989).

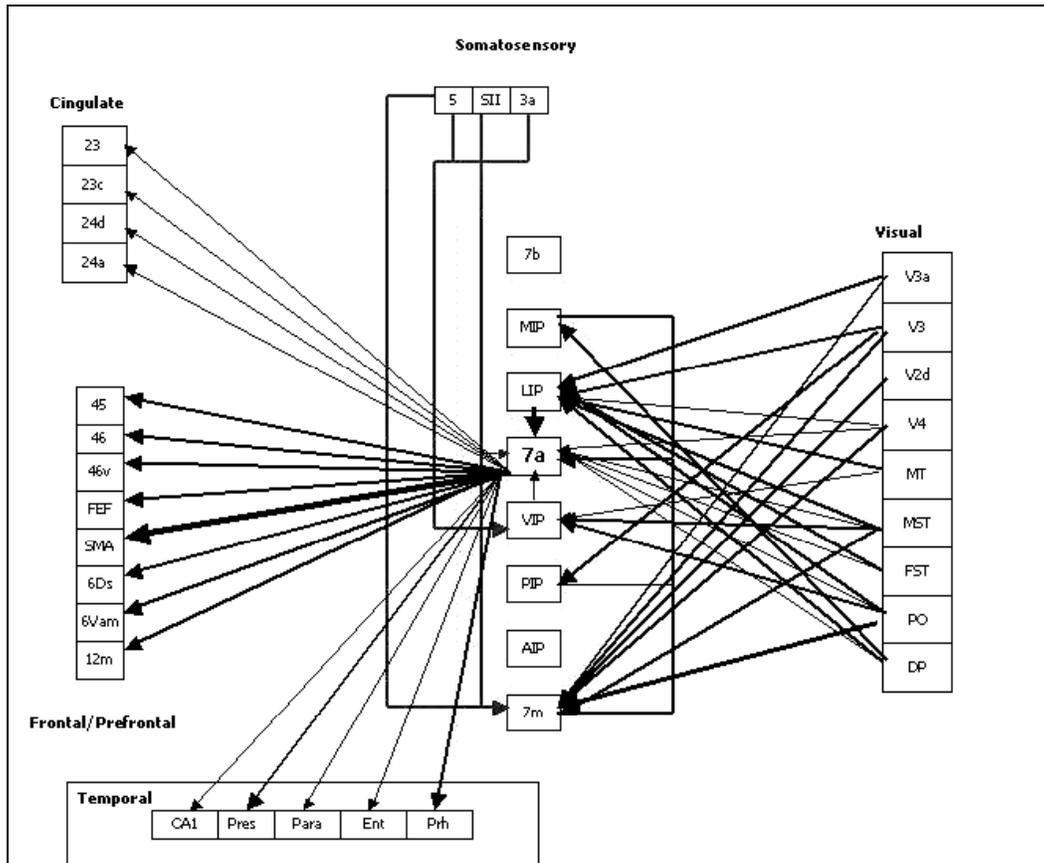


Figure 2.16 The reconstructed connectivity of area 7a. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001)

Findings from the same study also confirm that AIP is connected with area 7b. Area 7ip is unique among posterior parietal areas in its direct and indirect connections with the IT cortex (Cavada and Goldmanrakis 1989) and may form one of the object information channel to area AIP (Sakata et al. 1997b). Figure 2.16 and Figure 2.17 shows the reconstructed connectivity of areas 7a and 7b; while Figure 2.18 shows the reconstructed connectivity of AIP (Bota 2001).

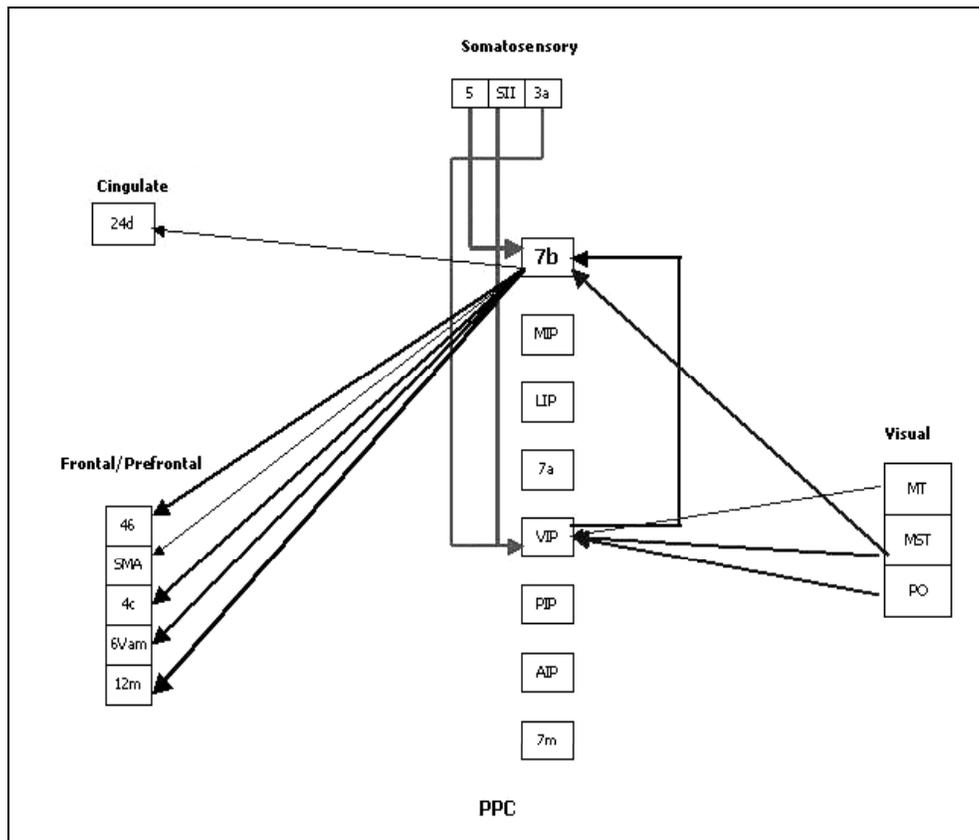


Figure 2.17 The reconstructed connectivity of area 7b. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001)

Andersen et al. (1990) suggests two types of processing for area 7a, each one following a different path. First path originates from visual area V4, which is believed to have an important role in pattern and color processing, and reaching to area 7a. Second path is the motion processing input originating from the middle temporal area (MT) and relayed via medial superior temporal area (MST) or LIP (Andersen et al. 1990). MT lies on the posterior bank of the superior temporal sulcus, while MST lies on the anterior bank of the same sulcus (Kandel et al. 2000; Maioli et al. 1998). MT projects to MST and to other areas in the parietal cortex concerned with visuospatial function. The preprocessed visual input from V1 is further elaborated in MT, where the firing pattern of neurons reflect the speed and direction of motion of visual targets (Kandel et al. 2000). Barnes and Pandya (1992) report that area 7a (PG-Opt) is reciprocally connected to STS and suggest that the visuospatial analysis that is associated with posterior intraparietal lobule could be amplified in the multimodal regions of STS. Therefore, the neurons of multimodal areas of the STS could be involved in analyzing the position of the body in relation to the environment (Barnes and Pandya 1992).

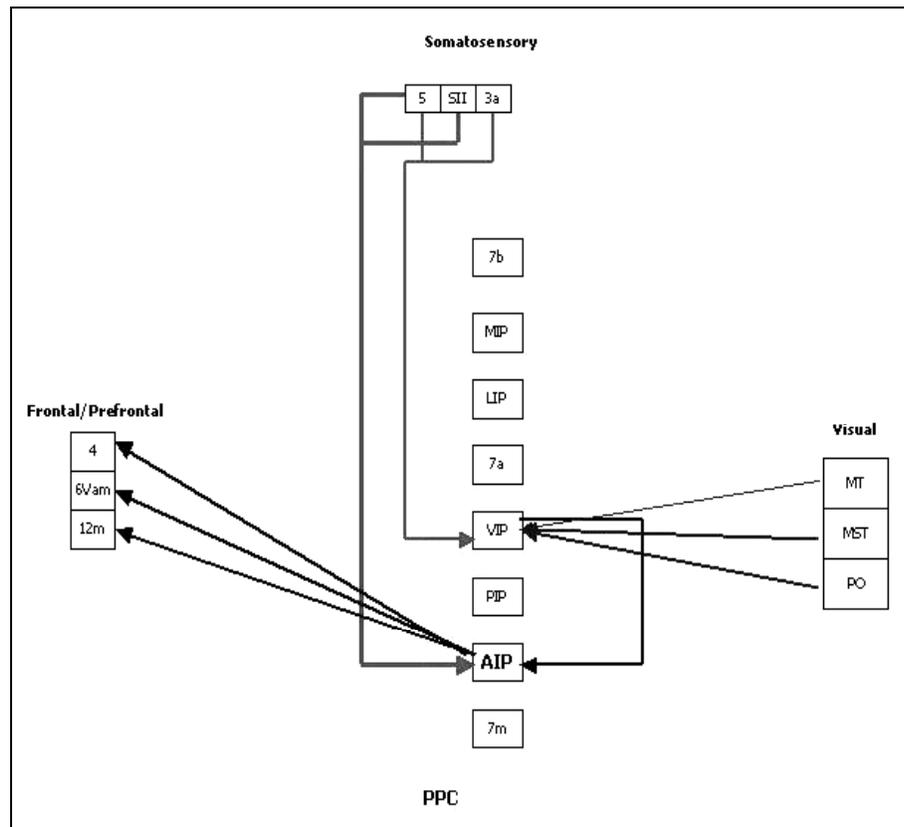


Figure 2.18 The reconstructed connectivity of area AIP. The thickness of the arrows represent the strength of the connection. (adapted from Bota 2001)

AIP receives input from other areas of the posterior parietal cortex such as 7b (Neal et al. 1990). In addition, this region has very significant recurrent cortico-cortical projections with area F5 of the inferior premotor cortex (Matelli et al. 1994; Sakata et al. 1997b). Figure 3.1 illustrates the visuomotor stream for hand action as well other related structures. Also see Figure 2.18 for the reconstructed connectivity diagram (Bota 2001) for AIP.

The anterior cingulate cortex is somatotopically organized and has direct connections with the motor and premotor cortices, suggesting that among 7 areas 7b has preferential access to motor centers (Cavada and Goldman-Rakic 1989). Area 7b is distinguished from other areas with its prominent connections with somatosensory related areas including S1, S2, the vestibular cortex, area 5 and the granular insular cortex. The only subdivision of area 7 that is connected to primary sensory cortex (S1) is area 7b (Cavada and Goldman-Rakic 1989). The heaviest connection of area 7b is with S2 it is likely that all body representation in S2 is connected to 7b (Cavada and Goldman-Rakic 1989). The connection to granular insular cortex is wide spread which contains high proportion of somatic-sensitive neurons. The connections of area 5 with 7b are topographic: the region of area 5 buried in the anterior bank of IPS, which is involved in forelimb mechanisms, is the source of strongest projection from area 5 to area 7b (Cavada and Goldman-Rakic 1989).

We conclude our discussion of anatomical connections by summarizing the connectivity of functionally defined intraparietal regions. Area 7a receives input from LIP (Andersen et al. 1990; Lewis and Van Essen 2000), MIP (Boussaoud et al. 1990; Lewis and Van Essen 2000; Bota 2001) and VIP

(reviewed in Maunsell 1995; Lewis and Van Essen 2000). Interested readers can find more details about these connections at the *NeuroHomology Database Website*³ (Bota 2001 and citations therein). The premotor projections of these intraparietal areas include the regions F2 and F4 (Luppino et al. 1999) as reviewed by Geyer et al.(2000).

2.6 Mirror neurons in humans

There is an unsettled debate about mirror neurons' function. It is suggested that mirror neurons may form the basis of understanding (Fadiga et al. 2000; Umiltà et al. 2001), and imitation (Arbib 2001; Rizzolatti and Arbib 1999) and even language in human (Rizzolatti and Arbib 1998). Thus, research for mirror neuron existence in human became necessary to support the idea that mirror neuron involvement in cognitive tasks.

Grafton et al. (1996) using positron emission tomography (PET), scanned subjects under three conditions, one of them being the control condition (object viewing). The other two were observing grasping actions of common objects and imagining themselves doing the same grasp actions. Grafton et al. (1996) used only precision grasps. The imagined-minus-control and observation-minus-control results were compared. The activation pattern was different. In their analysis, they categorized the activations into lateral activations and medial/dorsal activations. The lateral activation is relevant for our discussion⁴. In the observation condition, the activity locations were left rostral superior temporal sulcus (STS), left inferior frontal cortex (area 45), and the left rostral inferior parietal cortex (area 40). In addition, there was some activation found in the rostral part of the left intraparietal sulcus. However, the imagined grasping activated the left inferior frontal (Broca's area or area 44) and middle frontal cortex, left caudal inferior parietal cortex (area 40)⁵. Based on these findings, Grafton et al. (1996) suggested that the areas active during grasping observation might form a circuit for recognition of hand-object interactions, whereas the areas active during imagined grasping might be a human homologue of the action observation and execution matching system found in monkeys (mirror neurons). Their conclusion was that humans, as in monkeys, had a similar cortical circuit that was involved in representing observed grasping. Unfortunately, Grafton et al. (1996) did not include the self-execution condition in the experimental setup. Therefore, it cannot be concluded that the areas activated in this study have the dual property of the mirror neurons (the activation during self-action and observation of the same action performed by the demonstrator). In addition, note the discrepancy that the human homologue of the monkey F5, namely Broca's area (Rizzolatti and Arbib 1998), was not activated during grasp observation but only during imagined grasping.

In another study Grafton et al. (1997) used positron emission tomography (PET) imaging to test whether the observation of tools activates premotor areas without any overt motor demand⁶. Tool observation strongly activated the left dorsal premotor cortex. Silent tool-use naming activated Broca's area, the left dorsal premotor cortex (more than the observation case), the left supplementary motor area and the left ventral premotor cortex. These data indicate that, in human, F5 canonical type of neurons may exist in the left ventral premotor cortex, which can be triggered by object observation.

³ <http://brancusi.usc.edu/scripts/webmerger.exe?/database/homologies-main.html>

⁴ The rostral part of the left supplementary motor area (SMA-proper), and the right dorsal premotor cortex were also activated in grasp observation

⁵ The left rostral SMA-proper and left dorsal premotor cortex were also activated in imagined grasping

⁶ Silent naming of the presented tools was also studied

Iacoboni et al (1999) used functional magnetic resonance imaging (fMRI) to study the brain regions involved in imitation. Their paradigm had three observation conditions and three observation-execution conditions. In the observation-execution conditions, imitative and non-imitative behavior of simple finger movements was compared. In the imitative condition, participants had to execute the observed finger movement. In the two non-imitative conditions, participants had to execute the same movement in response to spatial or symbolic cues. The imitation task, when contrasted to non-imitative tasks, activated three areas: the left frontal operculum (Broca's area or area 44), the right anterior parietal region, and the right parietal operculum. The Broca's area and right anterior parietal region was also active during observation conditions. Iacoboni et al (1999) argued that Broca's area was activated due to the action-observation as Broca's area is the human homologue of monkey area F5 (Rizzolatti and Arbib 1998).

However, the data is not conclusive since, the Broca's area was active for all observation cases, not only the action observation. Krams et al (1998) in a similar study found that the Broca's region was more active during action preparation compared to action preparation-and-execution conditions. In both conditions the visual stimuli presented was the same and consisted of a hand drawing with a mark on a finger indicating the action to be prepared for. Krams et al. (1998) argued that Broca's area was involved in action suppression (see Krams et al. 1998 for a detailed discussion). However, in both studies, the actions were intransitive; they did not involve an object to be manipulated. In contrast, the majority of mirror neurons require an object and the action together; the miming of the action is not effective (Gallese et al. 1996).

In one study the motor cortex was stimulated using transcranial magnetic stimulation technique while the subjects (1) observed an experimenter grasping 3D-objects, (2) looked at the same 3D-objects, (3) observed an experimenter tracing geometrical figures in the air with his arm and (4) detected the dimming of a light (Fadiga et al. 1995). During the conditions of (1)-(4) the motor evoked potentials were recorded from the hand muscles. Fadiga et al. (1995) found that motor evoked potentials increased when the subjects observed movements. The motor evoked potential patterns reflected the pattern of muscle activity recorded when the subjects executed the observed actions. Therefore Fadiga et al. (1995) concluded that in humans there is an action observation and execution matching system, which is similar to monkey action recognition system (mirror neurons). This study showed that the effect of executing and observing the same action performed by others is similar. However, the localization of action observation and execution matching system was not possible with motor evoked potential recordings.

Hari et al. (1998) using a different technique (magnetoencephalogram) showed that the observation of object manipulation activated the primary motor cortex. Hari et al. (1998) recorded neuromagnetic oscillatory activity of the human precentral cortex while subjects were (1) idle, (2) manipulating a small object, and (3) observing another individual performing the same task. The left and right median nerves were stimulated alternately (inter stimulus interval, 1.5s) at intensities exceeding motor threshold, and the poststimulus rebound of the rolandic 15-to 25-Hz activity was quantified (Hari et al. 1998). The rebound was diminished during action observation as it did in action execution case (the observation suppression magnitude was 31-46% of the suppression during object manipulation). Hari et al. (1998) concluded that the human primary motor cortex was activated during observation as well as execution of the motor tasks since the 15-to 25-Hz activity mainly originates from the precentral motor cortex.

Nishitani and Hari (2000) showed that the inferior frontal area was active during both execution and observation of hand actions which confirmed the existence of a mirror system in human. In contrast to several PET studies (e.g. Grafton et al. 1996; Decety et al. 1997; Rizzolatti et al. 1996b), Broca's area

was active during action observation while area 45 was not active. Therefore, the study of Nishitani and Hari (2000) not only shows that the human brain is endowed with a mirror neuron system but also supports the hypothesis that Broca's area is the locus of action observation and execution matching system, which is consistent with the homology between Broca's area and area F5 (Rizzolatti and Arbib 1998).

Buccino et al. (2001) used fMRI to localize action recognition circuitry in humans for actions performed with different effectors. The subjects were presented with transitive and intransitive actions performed with mouth, hand and foot. Observation of both object- and non-object-related actions determined a somatotopically organized activation of premotor cortex. In addition, Buccino et al. (2001) found that during the observation of *object-related* actions, an activation -also somatotopically organized- was present in the posterior parietal lobe (Buccino et al. 2001). Buccino et al. (2001) argued that when individuals observe object-related actions, an internal replica of the motor act and the result of an object-related analysis are automatically generated in the ventral premotor cortex and the parietal lobe respectively. This result suggests that the observation and execution matching system is not constrained to hand actions but could be a general strategy used in the primate brain for interacting with the environment.

2.7 Summary

The posterior parietal cortex is involved in sensory-motor transformations, combining various sensory inputs and computing representations that are used by the motor system to generate movements. In particular AIP extract object features relevant for grasping. Other parietal areas such as VIP, MIP and LIP are involved in spatial aspects of object representations. These areas project to motor and premotor cortices enabling specific movement planning. Area F5 is involved in grasp planning while F4 is involved in reaching movement planning. The visual areas in the superior temporal sulcus perform visual analysis of form and motion including biological stimuli and provide parietal networks with motion related and, for some sectors, highly processed visual input. Chapter 3 will factor the connectivity specified in this chapter when developing Mirror Neuron System (MNS) model. The neurophysiology of area F5 will guide the modelling presented in this thesis throughout. We will implicate AIP and c-IPS as coding the object affordances serving as inputs to MNS and Learning to Grasp Models (LGM) of Chapters 5 and 6. We implicate target location schema to be represented in areas MIP/VIP/LIP, without specifying the neural region level assignment. Area 7a will combine the hand and object related visual inputs into an internal representation on which area 7b and F5 can be adapted to form mirror neurons.

3 CHAPTER III: MIRROR NEURON SYSTEM MODEL

Mirror neurons within a monkey's premotor area F5 fire not only when the monkey performs a certain class of actions but also when the monkey observes another monkey (or the experimenter) perform a similar action. It has thus been argued that these neurons are crucial for understanding of actions by others. This chapter offers the 'hand-state' hypothesis as a new explanation of the evolution of this capability: the basic functionality of the F5 mirror system is to elaborate the appropriate feedback – what we call the *hand state* – for opposition-space based control of manual grasping of an object. Given this functionality, the social role of the F5 mirror system in understanding the actions of others may be seen as an exaptation gained by generalizing from self-hand to other's-hand. In other words, mirror neurons first evolved to augment the 'canonical' F5 neurons by providing visual feedback on 'hand state', relating the shape of the hand to the shape of the object.

First, we introduce the MNS (Mirror Neuron System) model of F5 and related brain regions in terms of basic schemas. Then we aggregate them into three 'grand schemas' – Visual Analysis of Hand State, Reach and Grasp, and the Core Mirror Circuit – for each of which we present a useful implementation. The MNS model shows how the mirror system can *learn* to recognize actions already in the repertoire of the F5 canonical neurons. The chapter, in particular, shows how the connectivity pattern of mirror neuron circuitry can be established through training, and that the resultant network can exhibit a range of novel, physiologically interesting, behaviors during the process of action recognition.

3.1 The mirror neuron system for grasping and FARS model

The macaque inferior premotor cortex has been identified as being involved in reaching and grasping movements (Rizzolatti et al. 1988). This region has been further partitioned into two sub-regions: F5, the rostral region, located along the arcuate and F4, the caudal part (see Figure 3.1). The neurons in F4 appear to be primarily involved in the control of proximal movements (Gentilucci et al. 1988), whereas the neurons of F5 are involved in distal control (Rizzolatti et al. 1988). Rizzolatti et al. (1996a; Gallese et al. 1996). discovered a subset of F5 hand cells, which they called *mirror neurons* (Gallese et al. 1996; Rizzolatti et al. 1996a). Like other F5 neurons, mirror neurons are active when the monkey performs a particular class of actions, such as grasping, manipulating and placing. However, in addition, the mirror neurons become active when the monkey observes the experimenter or another monkey performing an action. The term F5 *canonical neurons* is used to distinguish the F5 hand cells which do *not* possess the mirror property but are instead responsive to visual input concerning a suitably graspable object. The canonical neurons are indistinguishable from the mirror neurons with respect to their firing during self-action. However they are different in their visual properties – they respond to object presentation not action observation per se (Murata et al. 1997a).

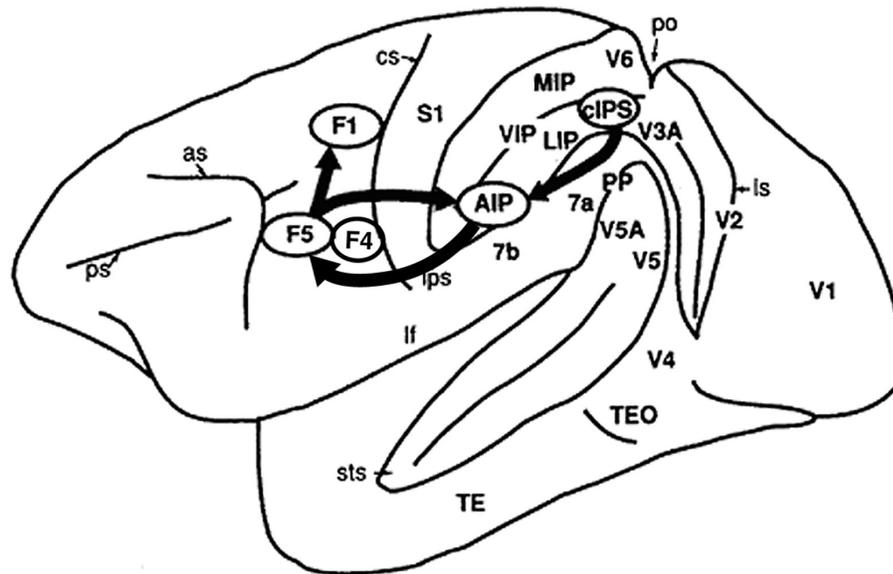


Figure 3.1 Lateral view of the monkey cerebral cortex (IPS, STS and lunate sulcus opened). The visuomotor stream for hand action is indicated by arrows (adapted from Sakata et al., 1997)

Most mirror neurons exhibit a clear relation between the observed and executed actions for which they are active. The congruence between the observed and executed action varies. For some of the mirror neurons, the congruence is quite loose; for others, not only must the general action (e.g., grasping) match but also the way the action is executed (e.g., power grasp) must match as well. To be triggered, the mirror neurons require an interaction between the hand motion and the object. The vision of the hand motion or the object alone does not trigger mirror activity (Gallese et al. 1996).

It has thus been argued that the importance of mirror neurons is that they provide a neural representation that is common to execution and observation of grasping actions and thus that these neurons are crucial to the social interactions of monkeys, providing the basis for *understanding* of actions by others through their linkage of action and perception (Rizzolatti and Fadiga 1998). Below, we offer the Hand-State Hypothesis, suggesting that this important role is an exaptation of a more primitive role, namely that of providing feedback for visually-guided grasping movements. By *exaptation* we mean the exploitation of an adaptation of a system to serve a different purpose (in this case for social understanding) than it initially developed for (in this case, visual control of grasping). We will then develop the MNS (Mirror Neuron System) model and show that the system can exploit its ability to relate self-hand movements to objects to recognize the manual actions being performed by others, thus yielding the mirror property. We also conduct a number of simulation experiments with the model and show that these yield novel predictions, suggesting new neurophysiological experiments to further probe the monkey mirror system. However, before introducing the Hand-State Hypothesis and the MNS model, we first outline the FARS model of the circuitry that includes the F5 canonical neurons and provides the conceptual basis for the MNS model.

Studies of the anterior intraparietal sulcus (AIP; Figure 3.1) revealed cells that were activated by the sight of objects for manipulation. In addition, this region has very significant recurrent cortico-cortical projections with area F5 (Matelli 1984; Sakata et al. 1997a). In their computational model for

primate control of grasping (the FARS – Fagg-Arbib-Rizzolatti-Sakata – model), Fagg and Arbib (1998) analyzed these findings of Sakata and Rizzolatti to show how F5 and AIP may act as part of a visuo-motor transformation circuit, which carries the brain from sight of an object to the execution of a particular grasp. In FARS model, the findings of Sakata (on AIP) and Rizzolatti (on F5) were interpreted as showing that AIP represents the grasps *afforded* by the object while F5 selects and drives the execution of the grasp (Fagg and Arbib 1998). The term *affordance* (adapted from Gibson 1966) refers to parameters for motor interaction that are signaled by sensory cues without invocation of high-level object recognition processes. The model also suggests how F5 may use task information and other constraints encoded in prefrontal cortex (PFC) to resolve the action opportunities provided by multiple affordances. Here we emphasize the essential components of the model (Figure 3.2) that will ground the version of the MNS model presented below. We focus on the linkage between viewing an affordance of an object and the generation of a single grasp.

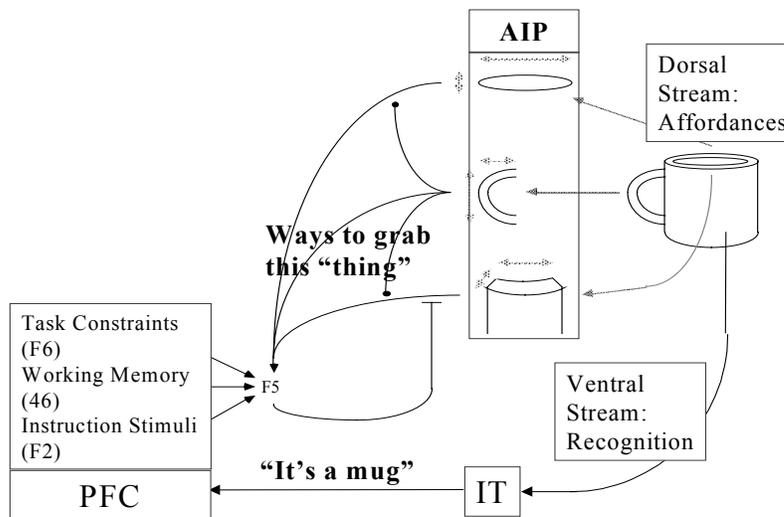


Figure 3.2 AIP extracts the affordances and F5 selects the appropriate grasp from the AIP 'menu'. Various biases are sent to F5 by Prefrontal Cortex (PFC) which relies on the recognition of the object by Inferotemporal Cortex (IT). The dorsal stream through AIP to F5 is replicated in the MNS model

(1) The dorsal visual stream (parietal cortex) extracts parametric information about the object being attended. It does not "know" what the object is; it can only see the object as a set of possible affordances. The ventral stream (from primary visual cortex to inferotemporal cortex, IT), by contrast, recognizes what the object is and passes this information to prefrontal cortex (PFC) which can then, on the basis of the current goals of the organism and the recognition of the nature of the object, bias F5 to choose the affordance appropriate to the task at hand.

(2) AIP is hypothesized as playing a dual role in the seeing/reaching/grasping process, not only computing affordances exhibited by the object but also, as one of these affordances is selected and execution of the grasp begins, serving as an active memory of the one selected affordance and updating this memory to correspond to the grasp that is actually executed.

(3) F5 is hypothesized as first being responsible for integrating task constraints with the set of grasps that are afforded by the attended object in order to select a single grasp. After selection of a single grasp, F5 unfolds this represented grasp in time to govern the role of primary motor cortex (F1) in its execution.

(4) In addition, the FARS model represents the way in which F5 may accept signals from areas F6 (pre-SMA), 46 (dorsolateral prefrontal cortex), and F2 (dorsal premotor cortex) to respond to task constraints, working memory, and instruction stimuli, respectively, and how these in turn may be influenced by object recognition processes in IT (see Fagg and Arbib 1988 for more details), but these aspects of the FARS model are included in MNS model.

3.2 The hand-state hypothesis

The key notion of the MNS model is that the brain augments the mechanisms modeled by the FARS model, for recognizing the grasping-affordances of an object (AIP) and transforming these into a program of action, by mechanisms which can recognize an action in terms of the hand state which makes explicit the relation between the unfolding trajectory of a hand and the affordances of an object. Our radical departure from all prior studies of the mirror system is to hypothesize that this system evolved in the first place to provide feedback for visually-directed grasping, with the social role of the mirror system being an exaptation as the hand state mechanisms become applied to the hands of others as well as to the hand of the animal itself.

3.2.1 Virtual fingers

As background for the Hand-State Hypothesis, we first present a conceptual analysis of grasping. Iberall and Arbib (1990) introduced the theory of *virtual fingers* and *opposition space*. The term *virtual finger* is used to describe the physical entity (one or more fingers, the palm of the hand, etc.) that is used in applying force and thus includes specification of the region to be brought in contact with the object (what we might call the 'virtual fingertip'). Figure 3.3 shows three types of opposition: those for the precision grip, power grasp, and side opposition. Each of the grasp types is defined by specifying two virtual fingers, VF1 and VF2, and the regions on VF1 and VF2 which are to be brought into contact with the object to grasp it. Note that the "virtual fingertip" for VF1 in palm opposition is the surface of the palm, while that for VF2 in side opposition is the side of the index finger.

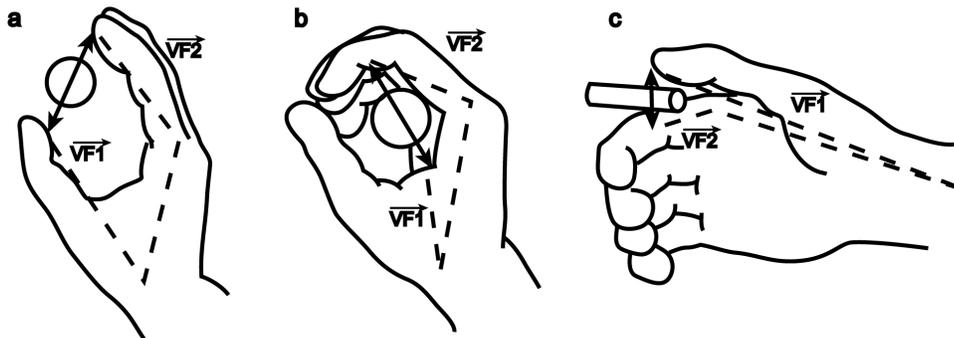


Figure 3.3 Each of the 3 grasp types here is defined by specifying two "virtual fingers", VF1 and VF2, which are groups of fingers or a part of the hand such as the palm which are brought to bear on either side of an object to grasp it. The specification of the virtual fingers includes specification of the region on each virtual finger to be brought in contact with the object. A successful grasp involves the alignment of two "opposition axes": the *opposition axis in the hand* joining the virtual finger regions to be opposed to each other, and the *opposition axis in the object* joining the regions where the virtual fingers contact the object. (Iberall and Arbib 1990)

The grasp defines two 'opposition axes': the *opposition axis in the hand* joining the virtual finger regions to be opposed to each other, and the *opposition axis in the object* joining the regions where the virtual fingers contact the object. Visual perception provides *affordances* (different ways to grasp the object); once an affordance is selected, an appropriate opposition axis in the object can be determined.

The task of motor control is to preshape the hand to form an opposition axis appropriate to the chosen affordance, and to so move the arm as to transport the hand to bring the hand and object axes into alignment. During the last stage of transport, the virtual fingers move down the opposition axis (the 'enclose' phase) to grasp the object just as the hand reaches the appropriate position.

3.2.2 The hand-state hypothesis

We assert as a general principle of motor control that if a motor plant is used for a task, then a feedback system will evolve to better control its performance in the face of perturbations. We thus ask, as a sequel to the work of Iberall and Arbib (1990), what information would be needed by a feedback controller to control grasping in the manner described in the previous section. Modeling of this feedback control is presented in Chapter 7, using a simplified hand/arm. In this chapter, our aim is to show how the availability of such feedback signals in the primate cortex for self-action for manual grasping can provide the action recognition capabilities which characterize the mirror system. Specifically, we offer the following hypothesis.

The hand-state hypothesis: The basic functionality of the F5 mirror system is to elaborate the appropriate feedback – what we call the *hand state* – for opposition-space based control of manual grasping of an object. Given this functionality, the social role of the F5 mirror system in understanding the actions of others may be seen as an exaptation gained by generalizing from self-hand to other's-hand.

The key to the MNS model, then, is the notion of *hand state* as encompassing data required to determine whether the motion and preshape of a moving hand may be extrapolated to culminate in a grasp appropriate to one of the affordances of the observed object. Basically a mirror neuron must fire if the preshaping of the hand conforms to the grasp type with which the neuron is associated; and the extrapolation of *hand state* yields a time at which the hand is grasping the object along an axis for which that affordance is appropriate.

Our current representation of hand state defines a 7-dimensional trajectory

$$F(t) = (d(t), v(t), a(t), o_1(t), o_2(t), o_3(t), o_4(t))$$

with the following components (see Figure 3.4):

Three components are hand configuration parameters:

$a(t)$: Index finger-tip and thumb-tip aperture

$o_3(t), o_4(t)$: The two angles defining how close the thumb is to the hand as measured relative to the side of the hand and to the inner surface of the palm

The remaining four parameters relate the hand to the object. o_1 and o_2 components represent the orientation of different components of the hand relative to the opposition axis for the chosen affordance in the object whereas d and v represents the kinematics properties of the hand with reference to the target location.

$o_1(t)$: The cosine of the angle between the object axis and the (*index finger tip – thumb tip*) vector

$o_2(t)$: The cosine of the angle between the object axis and the (*index finger knuckle – thumb tip*) vector

$d(t)$: distance to target at time t

$v(t)$: tangential velocity of the wrist

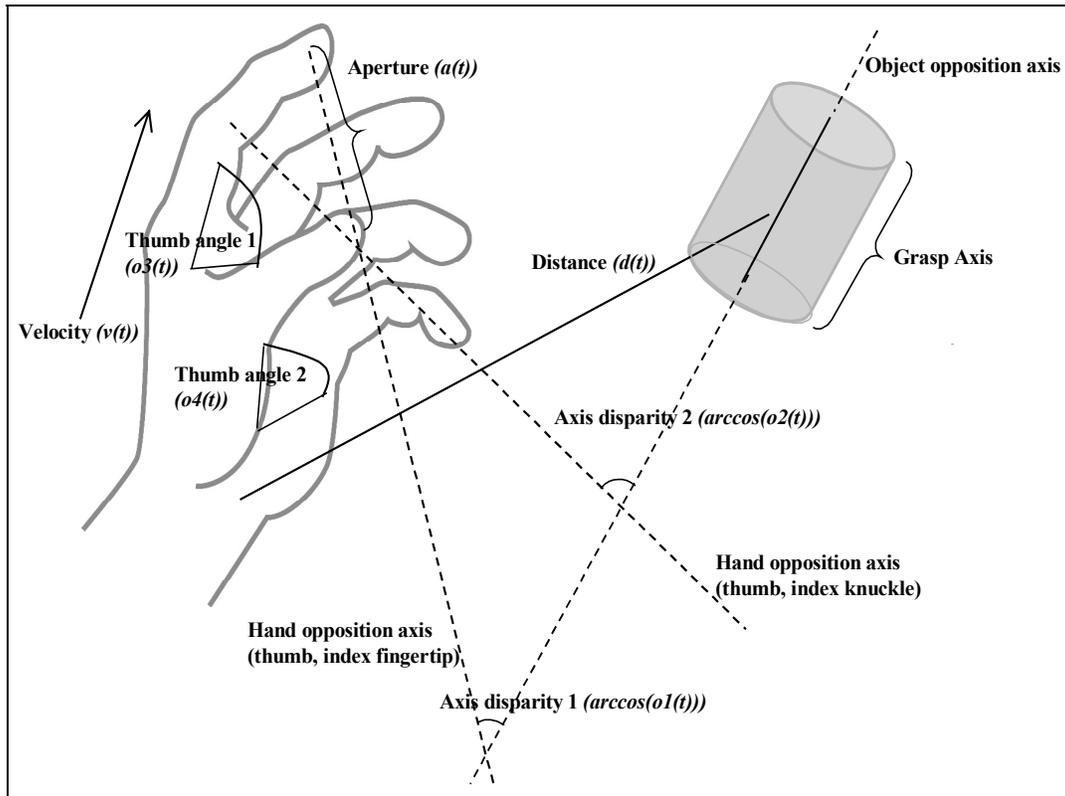


Figure 3.4 The components of hand state $F(t) = (d(t), v(t), a(t), o_1(t), o_2(t), o_3(t), o_4(t))$. Note that some of the components are purely hand configuration parameters (namely v, o_3, o_4, a) whereas others are parameters relating hand to the object

In considering the last four variables, note that only one or two of them will be relevant in generating a specific type of grasp, but they all must be available to monitor a wide range of possible grasps. We have chosen a set of variables of clear utility in monitoring the successful progress of grasping an object, but do not claim that these and only these variables are represented in the brain. Indeed, the brain's actual representation will be a distributed neural code, which we predict will correlate with such variables, but will not be decomposable into a coordinate-by-coordinate encoding. However, we believe that the explicit definition of hand state offered here will provide a firm foundation for the design of new experiments in kinesiology and neurophysiology.

The crucial point is that the availability of the hand state to provide feedback for visually-directed grasping makes action recognition possible. Notice that we have carefully defined the hand state in terms of relationships between hand and object (though the form of the definition must be subject to future research). This has the benefit that it will work just as well for measuring how the monkey's own hand is moving to grasp an object as for observing how well another monkey's hand is moving to grasp the object. This, we claim, is what allows self-observation by the monkey to train a system that can be used for observing the actions of others and recognizing just what those actions are.

3.3 The MNS (mirror neuron system) model

We now present a high level view of the MNS (Mirror Neuron System) model in terms of the set of interacting *schemas* (functional units: Arbib 1981; Arbib et al. 1998) shown in Figure 3.5, which define the

MNS (Mirror Neuron System) model of F5 and related brain regions. The connectivity shown in Figure 3.5 is constrained by the existing neurophysiology and neuroanatomy of the monkey brain (reviewed in Chapter 2). We have already introduced areas AIP and area F5, dividing the F5 grasp-related neurons into (i) F5 *mirror neurons* which are, when fully developed, active during certain self-movements of grasping by the monkey and during the observation of a similar grasp executed by others, and (ii) F5 *canonical neurons*, namely those active during self-movement and object vision but not for recognition of the action of others. Other brain regions also play an important role in mirror neuron system functioning in the macaque's brain for which the readers are referred to Chapter 2.

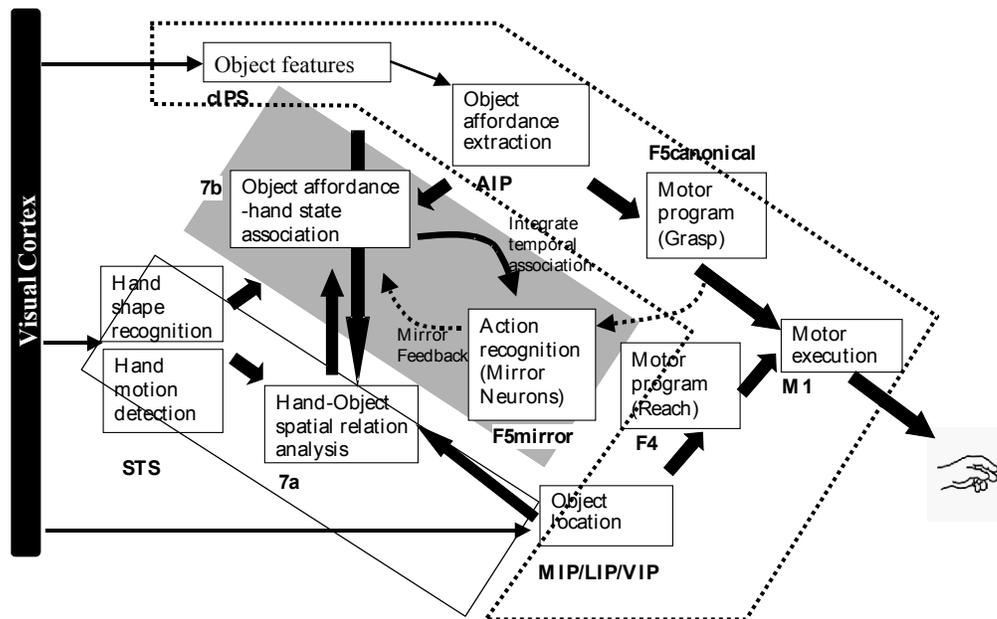


Figure 3.5 The MNS (Mirror Neuron System) model. (i) Top diagonal: a portion of the FARS model. Object features are processed by cIPS and AIP to extract grasp affordances, these are sent on to the canonical neurons of F5 that choose a particular grasp. (ii) Bottom right. Recognizing the location of the object provides parameters to the motor programming area F4 which computes the reach. The information about the reach and the grasp is taken by the motor cortex M1 to control the hand and the arm. (iii) New elements of the MNS model: Bottom left are two schemas, one to recognize the shape of the hand, and the other to recognize how that hand is moving. (iv) Just to the right of these is the schema for hand-object spatial relation analysis. It takes information about object features, the motion of the hand and the location of the object to infer the relation between hand and object. (v) The center two regions marked by the gray rectangle form the core mirror circuit. This complex associates the visually derived input (hand state) with the motor program input from region F5canonical neurons during the learning process for the mirror neurons. The grand schemas introduced in section 3.2 are illustrated as the following. The “Core Mirror Circuit” schema is marked by the center grey box; The “Visual Analysis of Hand State” schema is outlined by solid lines just below it, and the “Reach and Grasp” schema is outlined by dashed lines. (Solid arrows: established connections; dashed arrows: postulated connections)

The subsystem of the MNS model responsible for the visuo-motor transformation of objects into affordances and grasp configurations, linking AIP and F5 canonical neurons, corresponds to a key subsystem of the FARS model reviewed above. Our task is to complement the visual pathway via AIP by pathways directed toward F5 mirror neurons which allow the monkey to observe arm-hand trajectories and match them to the affordances and location of a potential target object. We will then

show how the mirror system may *learn* to recognize actions already in the repertoire of the F5 canonical neurons. In short, we will provide a mechanism whereby the actions of others are ‘recognized’ based on the circuitry involved in performing such actions. The Methods section provides the details of the implemented schemas and the Results section confronts the overall model with virtual experiments and produces testable predictions.

3.3.1 Overall function

In general, the visual input to the monkey represents a complex scene. However, we here sidestep much of this complexity (including attentional mechanisms) by assuming that the brain extracts two salient sub-scenes, a stationary object and in some cases a (possibly) moving hand. The overall system operates in two modes:

(i) Prehension: In this mode, the view of the stationary object is analyzed to extract affordances; then under prefrontal influence F5 may choose one of these to act upon, commanding the motor apparatus to perform the appropriate reach and grasp based on parameters supplied by the parietal cortex. The FARS model captures the linkage of F5 and AIP with PFC, prefrontal cortex (Figure 3.2). In the MNS model, we incorporate the F5 and AIP components from FARS (top diagonal of schemas in Figure 3.5), but omit IT and PFC from the present analysis.

(ii) Action recognition: In this mode, the view of the stationary object is again analyzed to extract affordances, but now the initial trajectory and preshape of an observed moving hand must be extrapolated to determine whether the current motion of the hand can be expected to culminate in a grasp of the object appropriate to one of its affordances.

We do not prespecify all the details of the MNS schemas. Instead, we offer a learning model which, given a grasp that is already in the motor repertoire of the F5 canonical neurons, can yield a set of F5 mirror neurons trained to be active during such grasps as a result of *self-observation* of the monkey's own hand grasping the target object. (How such grasps may be acquired in the first place is a topic of current research.) Consistent with the hand-state hypothesis, the result will be a system whose mirror neurons can respond to *similar actions observed being performed by others*. The current implementation of the MNS model exploits learning in artificial neural nets.

The heart of the learning model is provided by the *Object affordance-hand state association* schema and the *Action recognition (mirror neurons)* schema. These form the *core mirror (learning) circuit*, marked by the gray slanted rectangle in Figure 3.5, which mediates the development of mirror neurons via learning. The simulation results of this article will focus on this part of the model. Section 3.4.3.1 presents in detail the neural network structure of the core circuit. As we note further in the Discussion section, this leaves open many problems for further research, including the development of a basic action repertoire by F5 canonical neurons through trial-and-error in infancy and the expansion and refinement of this repertoire throughout life.

3.3.2 Schemas explained

As shown in the caption of Figure 3.5, we encapsulate the schemas shown there into the three ‘grand schemas’ of Figure 3.6(a). These guide our implementation of MNS. Our earlier review of the neuroscience literature in Chapter 2 justifies our initial hypotheses, made explicit in Figure 3.5, as to where these finer-grain schemas are realized in the monkey brain. However, after we explain these finer-grain schemas, we will then turn to our present simulation of the three grand schemas which is based on overall functionality. Nonetheless, the neural structure of Core Mirror Circuit yields interesting predictions for further neurophysiological experimentation.

3.3.2.1 Grand schema 1: reach and grasp

Object features schema: The output of this schema provides a coarse coding of geometrical features of the observed object. It thus provides suitable input to AIP and other regions/schemas.

Object affordance extraction schema: This schema transforms its input, the coarse coding of geometrical features of the observed object provided by the *Object features* schema, into a coarse coding for each affordance of the observed object.

Motor program (grasp) schema: We identify this schema with the canonical F5 neurons, as in the FARS model. Input is provided by AIP's coarse coding of affordances for the observed object. We assume that the output of the schema encodes a generic motor program for the AIP-coded affordances. This output serves as the learning signal to the *Action-recognition* (Mirror neurons) schema and drives the hand control functions of the *Motor execution* schema.

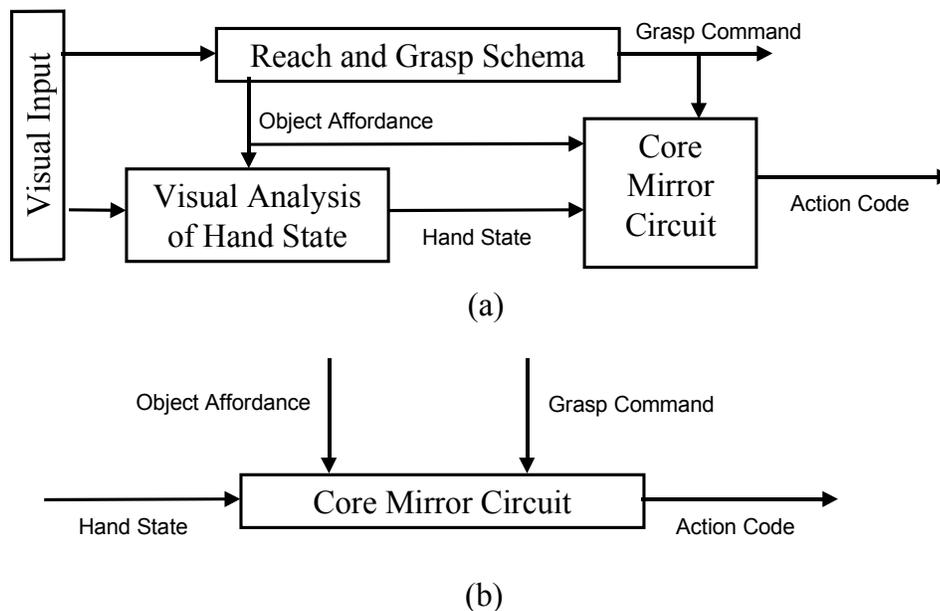


Figure 3.6 (a) For purposes of simulation, we aggregate the schemas of the MNS (Mirror Neuron System) model of Figure 3.5 into three "grand schemas" for Visual Analysis of Hand State, Reach and Grasp, Core Mirror Circuit. (b) For detailed analysis of the Core Mirror Circuit, we dispense with simulation of the other two grand schemas and use other computational means to provide the three key inputs to this grand schema

Object location schema: The output of this schema provides, in some body-centered coordinate frame, the location of the center of the opposition axis for the chosen affordance of the observed object.

Motor program (reach) schema: The input is the position coded by the *Object location* schema, while the output is the motor command required to transport the arm to bring the hand to the indicated location. This drives the arm control functions of the *Motor execution* schema.

The **motor execution schema** determines the course of movements via activity in primary motor cortex M1 and "lower" regions.

We next review the schemas which (in addition to the previously presented *Object features* and *Object affordance extraction* schemas) implement the visual system of the model:

3.3.2.2 Grand Schema 2: Visual Analysis of Hand State

The **hand shape recognition schema** takes as input a view of a hand, and its output is a specification of the hand shape, which thus forms some of the components of the hand state. In the

current implementation these are $a(t)$, $o_3(t)$ and $o_4(t)$. Note also that we implicitly assume that the schema includes a validity check to verify that the scene does contain a hand.

The **hand motion detection schema** takes as input a *sequence* of views of a hand and returns as output the wrist velocity, supplying the $v(t)$ component of the hand state.

The **hand-object spatial relation analysis schema** receives object-related signals from the *Object features* schema, as well as input from the *Object Location*, *Hand shape recognition* and *Hand motion detection* schemas. Its output is a set of vectors relating the current hand preshape to a selected affordance of the object. The schema computes such parameters as the distance of the object to the hand, and the disparity between the opposition axes of the object and the hand. Thus the hand state components $o_1(t)$, $o_2(t)$, and $d(t)$ are supplied by this schema. The *Hand-Object spatial relation analysis* schema is needed because, for almost all mirror neurons in the monkey, a hand mimicking a matching grasp would fail to elicit the mirror neuron's activity unless the hand's trajectory were taking it toward an object with a grasp that matches one of the affordances of the object. The output of this visual analysis is relayed to the *Object affordance-hand state association* schema which drives the F5 mirror neurons whose output is a signal expressing confidence that the observed trajectory will extrapolate to match the observed target object using the grasp encoded by that mirror neuron.

3.3.2.3 Grand Schema 3: Core Mirror Circuit

The **action recognition schema** – which is meant to correspond to the mirror neurons of area F5 – receives two inputs in our model. One is the motor program selected by the *Motor program* schema; the other comes from the *Object affordance-hand state association* schema. This schema works in two modes: learning and recognition. When a self-executed grasp is taking place the schema is in learning mode and the association between the observed hand-state (*Object affordance-hand state association* schema) and the motor program (*Motor program* schema) is learned. While in recognition mode, the motor program input is not active and the schema acts as a recognition circuit. If satisfactory learning (in terms of generalization and the range of actions learned) has taken place via self-observation then the schema will respond correctly while observing other's grasp actions.

The **object affordance-hand state association schema** combines all the hand related information as well as the object information available. Thus the inputs to the schema are from *Hand shape recognition* (components $a(t)$, $o_3(t)$, $o_4(t)$), *Hand motion detection* (component $v(t)$), *Hand-Object spatial relation analysis* ($o_1(t)$, $o_2(t)$, $d(t)$) and from *Object affordance extraction* schemas. As will be explained below, the schema needs a learning signal (mirror feedback). This signal is relayed by the *Action recognition* schema and, is basically, a copy of the motor program passed to the *Action recognition* schema itself. The output of this schema is a distributed representation of the object and hand state match (in our implementation the representation is not pre-specified but shaped by the learning process). The idea is to match the object and the hand state as the action progresses during a specific observed reach and grasp. In the current implementation, time is unfolded into a spatial representation of 'the trajectory until now' at the input of the *Object affordance-hand state association* schema, and the Action recognition schema decodes the distributed representation to form the mirror response (again, the decoding is not pre-specified but is the result of the back-propagation learning). In any case, the schema has two operating modes. First is the learning mode where the schema tries to adjust its efferent and afferent weights to ensure the right activity in the *Action recognition* schema. The second mode is the forward mode where it maps the hand state and the object affordance into a distributed representation to be used by the *Action recognition* schema.

The key question for this chapter's modeling will be to account for how learning mechanisms may shape the connections to mirror neuron in such a way that an action in the motor program repertoire of the F5 canonical neurons may become recognized by the mirror neurons when performed by others. In Chapter 5 and Chapter 6 we will present models that can learn a repertoire of grasping actions.

To conclude this section, we note that our modeling is subject to two quite different tests: (i) its overall efficacy in explaining behavior and its development, which can be tested at the level of the schemas (functional units) presented in this article; and (ii) its further efficacy in explaining and predicting neurophysiological data. As we shall see below, certain neurophysiological predictions are possible given the current work, even though the present implementation relies on relatively abstract artificial neural networks.

3.4 Schema implementation

Having indicated the functionality and possible neural basis for each of the schemas that will make up each grand schema, we now turn to the implementation of these three grand schemas. We implement the three grand schemas so that each functions correctly in terms of its input-output relations, and so that the Core Mirror Circuit contains model neurons whose behavior can be tested against neurophysiological data and yield predictions for novel neurophysiological experiments. The Core Mirror Circuit is thus the heart of MNS model that enables us to produce testable predictions (Figure 3.6b), but in order to study it, there must be an appropriate context, necessitating the construction of the kinematically realistic Reach and Grasp Simulator and the Visual Analyzer for Hand State. The latter will first be implemented as an analyzer of views of human hands, and then will have its output replaced by simulated hand state trajectories to reduce computational expense in our detailed analysis of the Core Mirror.

3.4.1 Grand schema 1: reach and grasp

We first discuss the Reach and Grasp Simulator that corresponds to the whole reach and grasp command system shown at the right of the MNS diagram (Figure 3.5). The simulator lets us move from the representation of the shape and position of a (virtual) 3D object and the initial position of the (virtual) arm and hand to a trajectory that successfully results in simulated grasping of the object. In other words the simulator plans a grasp and reach trajectory and executes it in a simulated 3D world (see Chapters 5 and 6 for neural realization of this schema). Trajectory planning (for example Kawato and Gomi 1992; Kawato et al. 1987; Jordan and Rumelhart 1992; Karniel and Inbar 1997; Breteler et al. 2001) and control of prehension (Hoff and Arbib 1993; see Wolpert and Ghahramani 2000 for a review), and their adaptation, have been widely studied. However, our simulator is not adaptive – its sole purpose is to create kinematically realistic actions. A similar reach and grasp system was proposed (Rosenbaum et al. 2001; Rosenbaum et al. 1999) where a movement is planned based on the constraint hierarchy, relying on obstacle avoidance and candidate posture evaluation processes (Meulenbroek et al. 2001). However, the arm and hand model was much simpler than ours as the arm was modeled as a 2D kinematics chain. Our Reach/Grasp Simulator is a non-neural extension of FARS model functionality to include the reach component. It controls a virtual 19 degrees DOF arm/hand (3 at the shoulder, 1 for elbow flexion/extension, 3 for wrist rotation, 2 for each finger joints with additional 2 DOFs for thumb one to allow the thumb to move sideways, and the other for the last joint in the thumb) and provides routines to perform realistic grasps. This kinematics realism is based on the literature of primate reach and grasp experiments (Jeannerod et al. 1995; for human see Hoff and Arbib 1993 and citations therein; for monkey see Roy et al. 2000). During a typical reach to grasp movement, the hand will follow a 'bell-

shaped' velocity profile (a single peaked velocity curve). The kinematics of the aperture between fingers used for grasping also exhibits typical characteristics. The aperture will first reach a maximum value that is larger than the aperture required for grasping the object and then as the hand approaches to the target the hand encloses to match the actual required aperture for the object. It is also important to note that in grasping tasks the temporal pattern of reaching and grasping is similar in monkey and human (Roy et al. 2000). Of course, there are inter-subject and inter-trial variability in both velocity and aperture profiles (Marteniuk and MacKenzie 1990). Therefore in our simulator we captured the qualitative aspects of the typical reach and grasp actions, namely that the velocity profiles have single peaks and that the hand aperture has a maximum value which is larger than the object size (see Figure 3.7, curves $a(t)$ and $v(t)$ for sample aperture and velocity profiles generated by our simulator). A grasp is planned by first setting the operational space constraints (e.g., points of contact of fingers on the object) and then finding the arm-hand configuration to fulfill the constraints. The latter is the inverse kinematics problem. The simulator solves the inverse kinematics problem by simulated gradient descent with noise added to the gradient (see Appendix 11.1.2 for a grasp planning example). Once the hand-arm configuration is determined for a grasp action, then the trajectory is generated by warping time using a cubic spline. The parameters of the spline are fixed and determined empirically to satisfy aperture and velocity profile requirements. Within the simulator, it is possible to adjust the target identity, position and size manually using a GUI or automatically by the simulator as, for example, in training set generation.

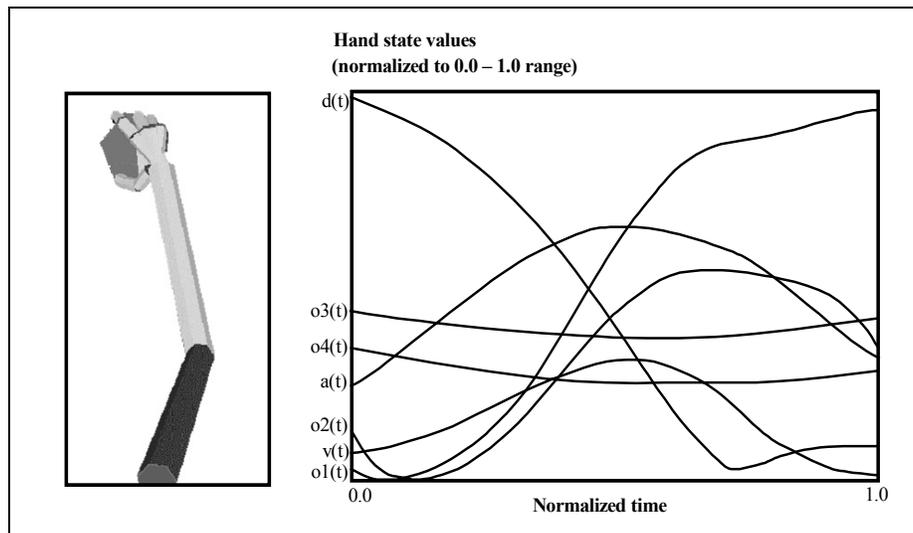


Figure 3.7 (Left) The final state of arm and hand achieved by the reach/grasp simulator in executing a power grasp on the object shown. (Right) The hand state trajectory read off from the simulated arm and hand during the movement whose end-state is shown at left. The hand state components are: $d(t)$, distance to target at time t ; $v(t)$, tangential velocity of the wrist; $a(t)$, Index and thumb finger tip aperture; $o1(t)$, cosine of the angle between the object axis and the (index finger tip – thumb tip) vector; $o2(t)$, cosine of the angle between the object axis and the (index finger knuckle – thumb tip) vector; $o3(t)$, The angle between the thumb and the palm plane; $o4(t)$, The angle between the thumb and the index finger

Figure 3.7 (left) shows the end state of a power grasp, while Figure 3.7 (Right) shows the time series for the hand state associated with this simulated power grasp trajectory. For example, the curve labeled $d(t)$ show the distance from the hand to the object decreasing until the grasp is completed; while

the curve labeled $a(t)$ show how the aperture of the hand first increases to yield a safety margin larger than the size of the object and then decreases until the hand contacts the object.

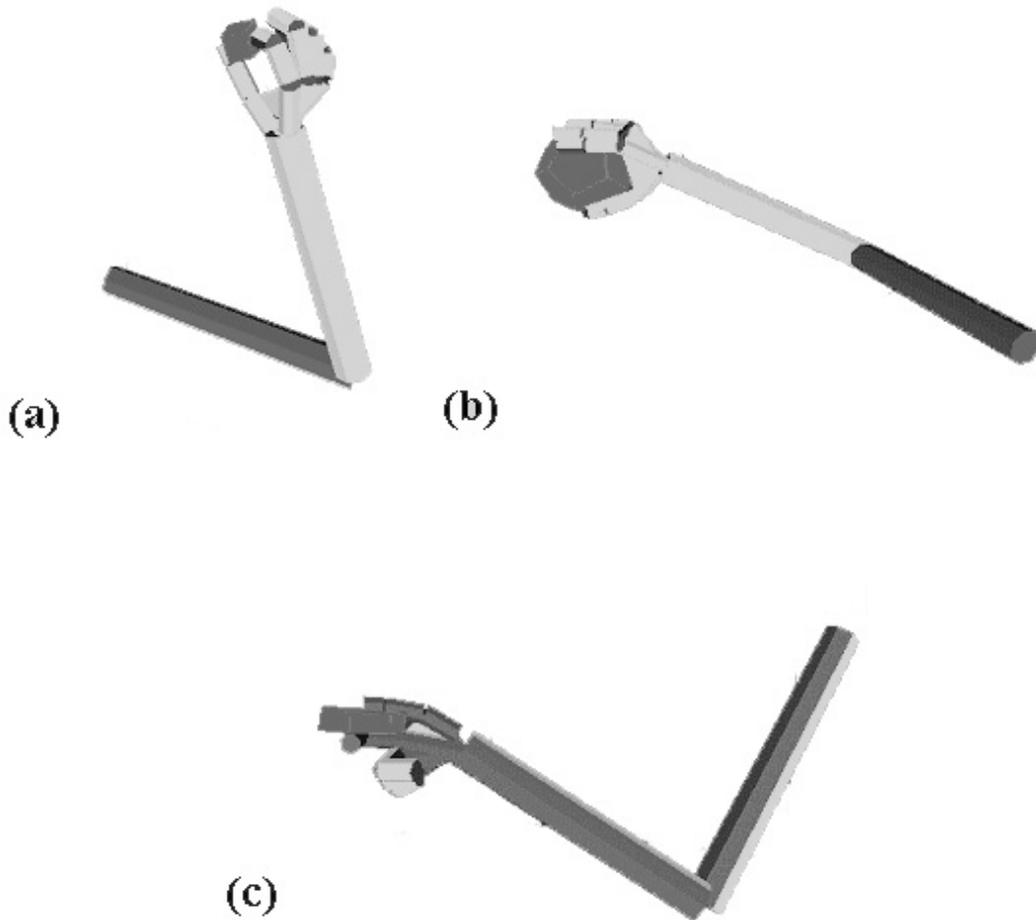


Figure 3.8 Grasps generated by the simulator. (a) A precision grasp. (b) A power grasp. (c) A side grasp

Figure 3.8(a) shows the virtual hand/arm holding a small cube in a *precision grip* in which the index finger (or a larger "virtual finger") opposes the thumb. The power grasp (Figure 3.8(b)) is usually applied to big objects and characterized by the hand's covering the object, with the fingers as one virtual finger opposing the palm as the other. In a side grasp (Figure 3.8(c)), the thumb opposes the side of another finger. To clarify the type of heuristics we use to generate the grasp, Appendix 11.1.2 outlines the grasp planning and execution for a precision pinch.

3.4.2 Grand schema 2: visual analysis of hand state

Visual Analysis of Hand State Schema is a non-neurophysiological implementation of a visual analysis system to validate the extraction of hand parameters from a view of a hand, by recovering the configuration of a model of the hand being seen. The hand model is a three dimensional 14 degrees of freedom (DOF) kinematic model, with a 3-DOF joint for the wrist, two 1-DOF joints (metacarpophalangeal and distalinterphalangeal) for each of the four fingers, and finally a 1-DOF joint for the metacarpophalangeal joint, and a 2-DOF joint for the carpometacarpal joint of the thumb. Note the distinction between 'hand configuration' which gives the joint angles of the hand considered in

isolation, and the 'hand state' which comprises 7 parameters relevant to assessing the motion and preshaping of the hand relative to an object. Thus, the hand configuration provides some, but not all, of the data needed to compute the hand state.

To lighten the load of building a visual system to recognize hand features, we marked the wrist and the articulation points of the hand with colors. We then used this color-coding to help recognize key portions of the hand and used this result to initiate a process of model matching. Thus, the first step of the vision problem was color segmentation, after which the three dimensional hand shape was recovered.

3.4.2.1 Color segmentation and feature extraction

One needs *color segmentation* to locate the colored regions on the image. Gray level segmentation techniques cannot be used in a straightforward way because of the vectorial nature of color images (Lambert and Carron 1999). Split-and-Merge is a well-known image segmentation technique in image processing (Sonka et al. 1993), recursively splitting the image into smaller pieces until some homogeneity criterion is satisfied as a basis for reaggregation into regions. In our case, the criterion is having similar color throughout a region. However, RGB (Red-Green-Blue) space is not well suited for this purpose. HSV (Hue-Saturation-Value) space is better suited since hue in segmentation processes usually corresponds to human perception and ignores shading effects (Russ 1998 Chapters 1 and 6). However, the segmentation system we implemented with HSV space, although better than the RGB version, was not satisfactory for our purposes. Therefore, we designed a system that can learn the best color space.

Figure 3.9(a) shows the training phase of the *color expert* system, which is a (one hidden-layer) feed-forward network with sigmoidal activation function. The learning algorithm is back-propagation with momentum and adaptive learning rate. The given image is put through a smoothing filter to reduce noise in the image before training. Then the network is given around 100 training samples each of which is a vector of ((R, G, B), *perceived color code*) values. The output color code is a vector consisting of all zeros except for one component corresponding to the *perceived* color of the patch. The training builds an internal non-linear color space from which it can unambiguously tell the perceived color. This training is done only at the beginning of a session to learn the colors used on the particular hand. Then the network is fixed as the hand is viewed in a variety of poses.

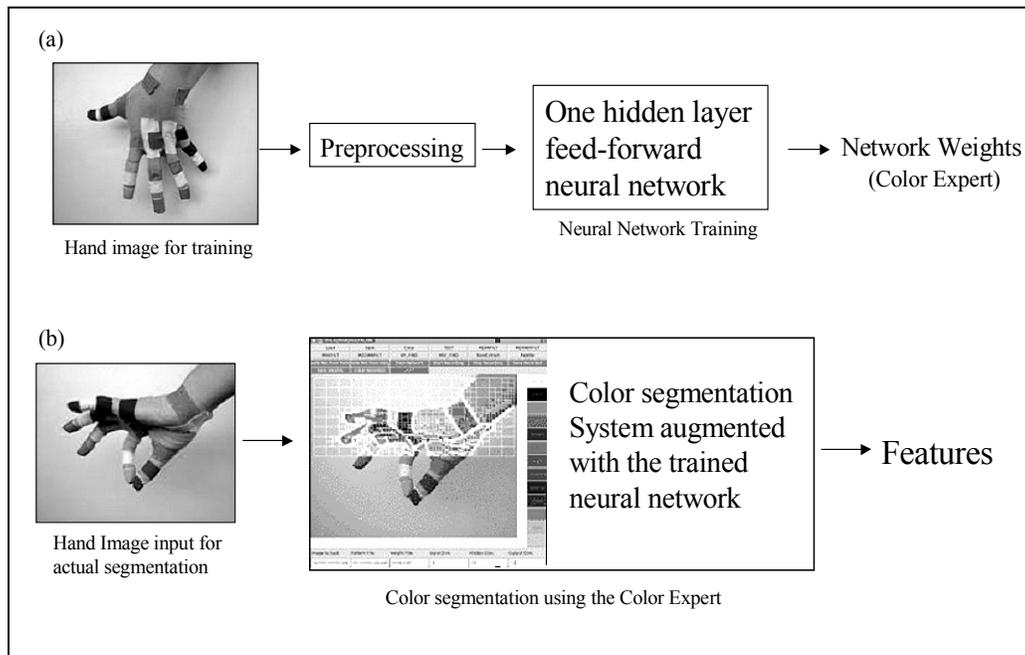


Figure 3.9 (a) Training the color expert, based on colored images of a hand whose joints are covered with distinctively colored patches. The trained network will be used in the subsequent phase for segmenting image. (b) A hand image (not from the training sample) is fed to the augmented segmentation program. The color decision during segmentation is done by consulting to the Color Expert. Note that a smoothing step (not shown) is performed before segmentation

Figure 3.9(b) illustrates the actual segmentation process using the 'color expert' to find each region of a single (perceived) color (see Appendix 11.1.1 for details). The output of the algorithm is then converted into a feature vector with a corresponding confidence vector giving a confidence level for each component in the feature vector. Each finger is marked with two patches of the same color. Sometimes it may not be possible to determine which patch corresponds to the fingertip and which to the knuckle. In those cases, the confidence value is set to 0.5. If a color is not found (e.g., the patch may be obscured), a zero value is given for the confidence. If a unique color is found without any ambiguity then the confidence value is set to 1. The segmented centers of regions (color markers) are taken as the approximate articulation point positions. To convert the absolute color centers into a feature vector we simply subtract the wrist position from all the centers found and put the resulting relative (x,y) coordinate into the feature vector (but the wrist is excluded from the feature vector as the positions are specified with respect to the wrist position).

3.4.2.2 3D hand model matching

Our model matching algorithm uses the feature vector generated by the segmentation system to attain a hand configuration and pose that would result in a feature vector as close as possible to the input feature vector (Figure 3.10). The scheme we use is a simplified version of Lowe's (1991); see Holden (1997) for a review of other hand recognition studies.

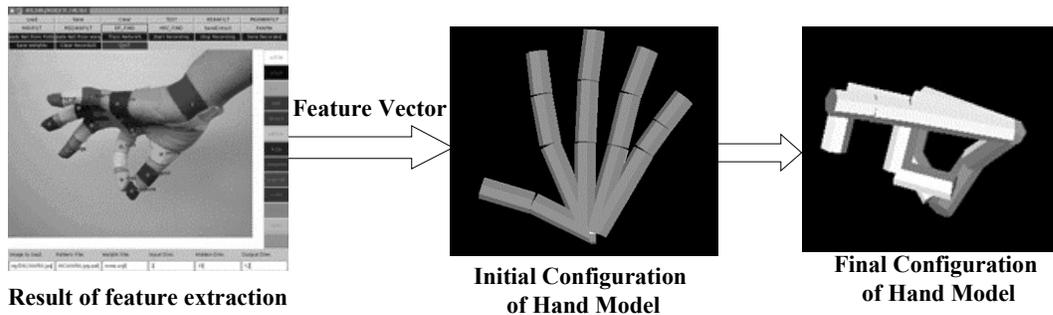


Figure 3.10 Illustration of the model matching system. Left: markers located by feature extraction schema. Middle and Right: initial and final stages of model matching. After matching is performed a number of parameters for the Hand configuration are extracted from the matched 3D model

The matching algorithm is based on minimization of the distance between the input feature and model feature vector, where the distance is a function of the two vectors and the confidence vector generated by segmentation system. Distance minimization is realized by hill climbing in feature space. The method can handle occlusions by starting with ‘don't cares’ for any joints whose markers cannot be clearly distinguished in the current view of the hand

The distance between two feature vectors F and G is computed as follows:

$$D(F, G) = \sqrt{(F_i - G_i)^2 C_i^f C_i^g}$$

where subscripting denotes components and C^f , C^g denotes the confidence vectors associated with F and G . Given this result of the visual processing – our *hand shape recognition* schema – we can clearly read off the following components of the hand state, $F(t)$:

$a(t)$: aperture of the virtual fingers involved in grasping

$o_3(t)$, $o_4(t)$: the two angles defining how close the thumb is to the hand as measured relative to the side of the hand and to the inner surface of the palm (see Figure 3.4). The remaining components can easily be computed once the object affordance and location is known. The computation of the components:

$d(t)$: distance to target at time t , and

$v(t)$: tangential velocity of the wrist

$o_1(t)$: Angle between the object axis and the (*index finger tip – thumb tip*) vector

$o_2(t)$: Angle between the object axis and the (*index finger knuckle – thumb tip*) vector

constitute the tasks of the *hand-object spatial relation analysis* schema and the *hand motion detection* schema. These require visual inspection of the relation between hand and target, and visual detection of wrist motion, respectively. Section 3.5.3 presents justifies the visual analysis of hand state schema by showing MNS performance when the hand state was extracted by the described visual recognition system based on a real video sequence. However, when we turn to modeling the Core Mirror Circuit in the next section, we will not use this implementation of visual analysis of hand state but instead, to simplify computation, we will use synthetic output generated by the reach/grasp simulator to emulate the values that could be extracted with this visual system. Specifically, we use the hand/grasp simulator

to produce both (i) the visual appearance of such a movement for our inspection (Figure 3.7, left), and (ii) the hand state trajectory associated with the movement (Figure 3.7, right). Especially, for training we need to generate and process too many grasp actions, which makes it impractical to use the visual processing system without special hardware as the computational time requirement is too high. Nevertheless, we need to show the similarity of the data from the visual system and the simulator: We have already shown that the grasp simulator generates aperture and velocity profiles that are similar to those in real grasps. Of course, there is still the question of how well the our visual system can extract these features and more importantly how similar are the other components of the hand state that we did not specifically craft to match the real data. Positive evidence will be presented in Section 3.5.3.

3.4.3 Grand Schema 3: core mirror circuit

As diagrammed in Figure 3.6(b), our detailed analysis of the core mirror circuit does not require simulation of the visual analysis of hand state and of reach and grasp so long as we ensure that it receives the appropriate inputs. Thus, we supply the object affordance and grasp command directly to the network at each trial. (Actually, we conduct experiments to compare performance with and without an explicit input which codes object affordance.) For the hand state input, rather than providing visual input to the visual analysis of hand state schema and have it compute the hand state input to the core mirror circuit, we use our reach and grasp simulator to simulate the performance of the *observed* primate – and from this simulation we extract (as in Figure 3.7) both a graphical display of the arm and hand movement that would be seen by the *observing* monkey, as well as the hand state trajectory that would be generated in its brain. We thus use the time-varying hand state trajectory generated in this way to provide the input to the model of the core mirror circuit of the observing monkey without having to simultaneously model its visual analysis of hand State. Thus, we have implemented the core mirror circuit in terms of neural networks using as input the synthetic data on hand state that we gather from our reach and grasp simulator (however see Section 3.5.3 for a simulation with real data extracted by our visual system). Figure 3.13 shows an example of the recognition process together with the type of information supplied by the simulator.

3.4.3.1 Neural network details

In our implementation, we used a feed-forward neural network with one hidden layer. In contrast to the previous sections, we can here identify the parts of the neural network as Figure 3.5 schemas in a one-to-one fashion. The hidden layer of the model neural network corresponds to the *object affordance-hand state association schema*, while the output layer of the network corresponds to the *action recognition schema* (i.e., we identify the output neurons with the F5 mirror neurons). In the following formulation *MR* (mirror response) represents the output of the *action recognition schema*, *MP* (motor program) denotes the target of the network (copy of the output of *motor program (grasp) schema*). *X* denotes the input vector applied to the network, which is the transformed Hand State (and the object affordance). The transformation applied is described in the next subsection. The learning algorithm used is back propagation (Rumelhart et al. 1986) with momentum term. The formulation is adapted from (Hertz et al. 1991).

Activity propagation (forward pass)

$$MR_i = g\left(\sum_j W_{ij} g\left(\sum_k w_{jk} X_k\right)\right)$$

Learning weights from input to hidden layer

$$W_{ij} = W_{ij} + \eta(t)\delta W_{ij} + \mu W_{ij}^{old}, \text{ where}$$

$$\delta W_{ij} = g'\left(\sum_j W_{ij} g\left(\sum_k w_{jk} X_k\right)\right)(MP_i - MR_i)$$

$$W_{old} = W$$

Learning weights from hidden to output layer

$$w_{jk} = w_{jk} + \eta(t)\delta w_{jk} + \mu w_{jk}^{old}, \text{ where}$$

$$\delta w_{jk} = g'\left(\sum_k w_{jk} X_k\right) X_k$$

$$w_{jk}^{old} = w_{jk}$$

The squashing function g we used was $g(x) = 1/(1 + e^{-x})$. η and μ are the learning rate and the momentum coefficient respectively. In our simulations, we adapted η during training such that if the output error was consistently decreasing then we increased η . Otherwise, we decreased η . We kept μ as a constant set to 0.9. W is the $3 \times (6+1)$ matrix of real numbers representing the hidden-to-output weights. w is the $6 \times (210+1)$ ($6 \times (220+1)$ in the explicit affordance coding case) matrix of real numbers representing the input to hidden weights, and X is the $210+1$ ($220+1$ in explicit affordance coding case) component input vector representing the hand state (trajectory) information. (The extra +1 comes from the fact that the formulation we used hides the bias term required for computing the output of a unit in the incoming signals as a fixed input clamped to 1)

3.4.3.2 Temporal to spatial transformation

The input to the network was formed in a way to allow encoding of temporal information without the use of a dynamic neural network, and solved the scaling problem. The input at any time represented the entire input from the start of the action until the present time t . To form the input vector, each of the seven components of the hand state trajectory to time t is fitted by a cubic spline (see Kincaid and Cheney 1991 for a formulation), and the splines are then sampled at 30 uniformly spaced intervals. The hand state input is then a vector with 210 components: 30 samples from the time-scaled spline fitted to the 7 components of the hand-state time series. Note then that no matter what fraction t is of the total time T of the entire trajectory, the input to the network at time t comprises 30 samples of the hand-state uniformly distributed over the interval $[0, t]$. Thus the sampling is less densely distributed across the trajectory-to-date as t increases from 0 to T .

An alternative approach would be to use an SRN (simple recurrent neural network) style architecture to recognize hand state trajectories. However, this raises an extra quantization or segmentation step to convert the continuous hand state trajectories to discrete states. With our approach, we avoid this extra step because the quantization is implicitly handled by the learning process.

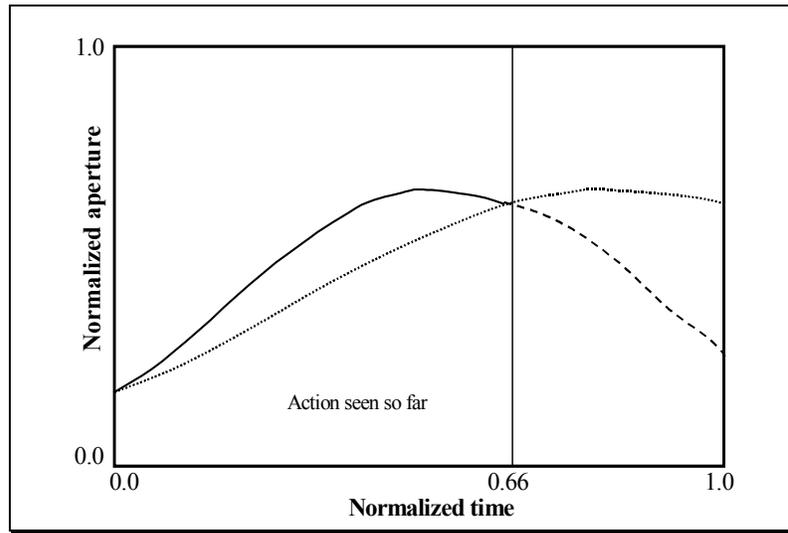


Figure 3.11 The scaling of an incomplete input to form the full spatial representation of the hand state. As an example, only one component of the hand state, the aperture, is shown. When 66 percent of the action is completed, the pre-processing we apply effectively causes the network to receive the stretched hand state (the dotted curve) as input as a re-representation of the hand state information accessible to that time (represented by the solid curve; the dashed curve shows the remaining, unobserved part of the hand state).

Figure 3.11 demonstrates the preprocessing we use to transform time-varying hand state components into spatial code. In the figure, only a single component (the aperture) is shown as an example. The curve drawn by the solid line indicates the available information when 66% of the grasp action is completed. In reality, a digital computer (and thus the simulator) runs in discrete time steps, so we construct the continuous curve by fitting a cubic spline to the collected samples for the value represented (aperture value in this case). Then we resample 30 points from the (solid) curve to form a vector of size 30. In effect, this presents the network with the stretched spline shown by the dotted curve. This method has the desirable property of avoiding the time-scaling problem to establish the equivalence of actions that last longer than shorter ones, as it is the case for a grasp for an object far from the hand compared to a grasp for a closer object. By comparing the dotted curve (what the network sees at $t = 0.66$) with the 'solid + dashed' curve (the overall trajectory of the aperture), we can see how much the network's input is distorted. As the action gets closer to its end, the discrepancy between the curves tends to zero. Thus, our preprocessing gives rise to an approximation to the final representation when a certain portion or more of the input is seen. Figure 3.12 samples the temporal evolution of the spatial input the network receives.

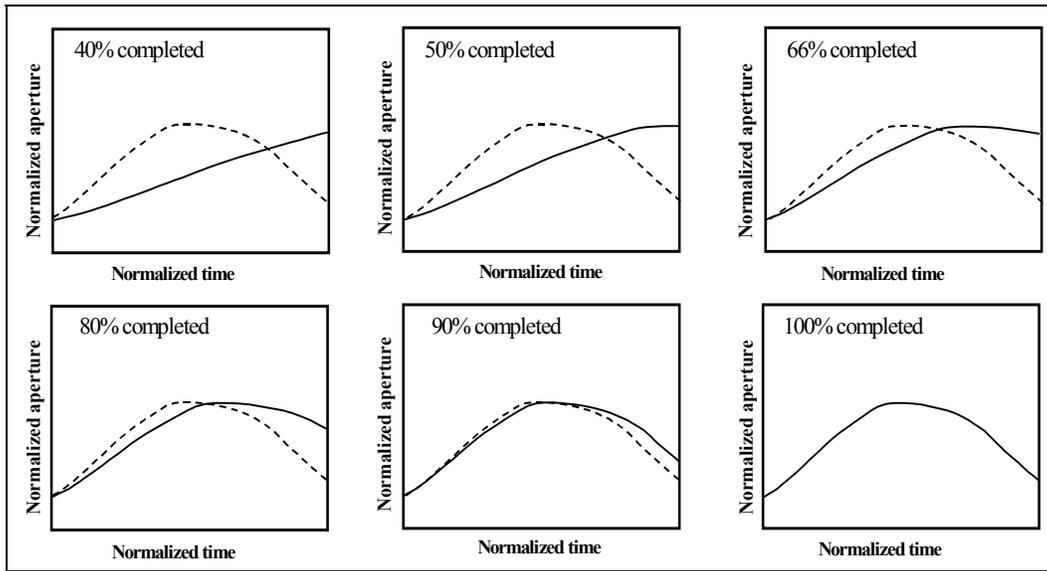


Figure 3.12 The solid curve shows the effective input that the network receives as the action progresses. At each simulation cycle the scaled curves are sampled (30 samples each) to form the spatial input for the network. Towards the end of the action the networks input gets closer to the final hand state

3.4.3.3 Neural network training

The training set was constructed by making the simulator perform various grasps in the following way.

(1) The objects used were a cube of changing size (a generic size cube scaled by a random scale factor between 0.5 and 1.5), a disk (approximated as a thin prism), a ball (approximated as a dodecahedron) again scaled randomly by a number between 0.75 and 1.5. In this particular trial, we did not change the disk size. In the training set formation, a certain object always received a certain grasp (unlike the testing case).

(2) The target locations were chosen from the surface patches of a sphere centered on the shoulder joint. The patch is defined by bounding meridian (longitude) and parallel (latitude) lines. The extent of the meridian and parallel lines was from -45° to 45° . The step chosen was 15° . Thus the simulator made $7 \times 7 = 49$ grasps per object. The unsuccessful grasp attempts were discarded from the training set. For each successful grasp, two negative examples were added to the training set in the following way. The inputs (group of 30) for each parameter are randomly shuffled. In this way, the network was forced to learn the order of activity within a group rather than learning the averages of the inputs (note that the shuffling does not change mean and variance). The second negative pattern was used to stress that the distance to target was important. The target location was perturbed and the grasp was repeated (to the original target position).

Finally, our last modification in the backpropagation training algorithm was to introduce a random input pattern (totally random; no shuffling) on the fly during training and ask the network to produce zero output for those patterns. This way we not only biased the network to be as silent as possible during ambiguous input presentation but also gave the network a higher chance to reach global minima.

It should be emphasized that the network was trained using the complete trajectory of the hand state (analogous to adjusting synapses after the self-grasp is completed). During testing, in contrast, the prefixes of a trajectory were used (analogous to predictive response of mirror neurons while observing a

grasp action). The network thus yielded a time-course of activation for the mirror neurons. As we shall see in the Results section, initial prefixes yields little or no mirror neuron activity, and ambiguous prefixes may yields transient activity of the ‘wrong’ mirror neurons.

We thus need to make two points to highlight the contribution of this study:

(1) It is, of course, trivial to train a network to pair complete trajectories with the final grasp type. What is interesting here is that we can train the system on the basis of final grasp but then observe the whole time course of mirror neuron activity, yielding predictions for neurophysiological experiments by highlighting the importance of the *timing* of mirror neuron activity.

(2) It is commonly understood that the training method used here, namely back-propagation, is not intended to be a model of the cellular learning mechanisms employed in cerebral cortex. This might be a matter of concern were we intending to model the time course of learning, or analyze the effect of specific patterns of neural activity or neuromodulation on the learning process. However, our aim here is quite different: we want to show that the connectivity of mirror neuron circuitry can be established through training, and that the resultant network can exhibit a range of novel, physiologically interesting, behaviors during the process of action recognition. Thus, the actual choice of training procedure is purely a matter of computational convenience, and the fact that the method chosen is non-physiological does not weaken the importance of our predictions concerning the timing of mirror neuron activity.

3.5 Simulation results

In this study, we experimented with two types of network. The first has only the hand state as the network input. We call this version the *non-explicit affordance coding network* since the hand state will often imply the object affordance in our simple grasp world. The second network we experimented with – the *explicit affordance coding network* – has affordance coding as one set of its inputs. The number of hidden layer units in each case was chosen as 6 and there were 3 output units, each one corresponding to a recognized grasp

3.5.1 Non-explicit affordance coding experiments

We first present results with the MNS model implemented without an explicit object affordance input to the core mirror circuit. We then study the effects of supplying an explicit object affordance input.

3.5.1.1 Grasp resolution

In Figure 3.13, we let the (trained) model observe a grasp action. Figure 3.13(a) demonstrates the executed grasp by giving the views from three different angles to show the reader the 3D trajectory traversed. Figure 3.13(b) shows the extracted hand state (left) and the response of the (trained) core mirror network (right). In this example, the network was able to infer the correct grasp without any ambiguity as a single curve corresponding to the observed grasp reaches a peak and the other two units’ output are close to zero during the whole action. The horizontal axis for both figures is such that the onset of the action and the completion of the grasp are scaled to 0 and 1 respectively. The vertical axis in the hand state plot represents a normalized (min=0, max=1) value for the components of the hand state whereas the output plot represents the average firing rate of the neurons (no firing = 0, maximum firing = 1). The plotting scheme that is used in Figure 3.13 will be used in later simulation results as well.

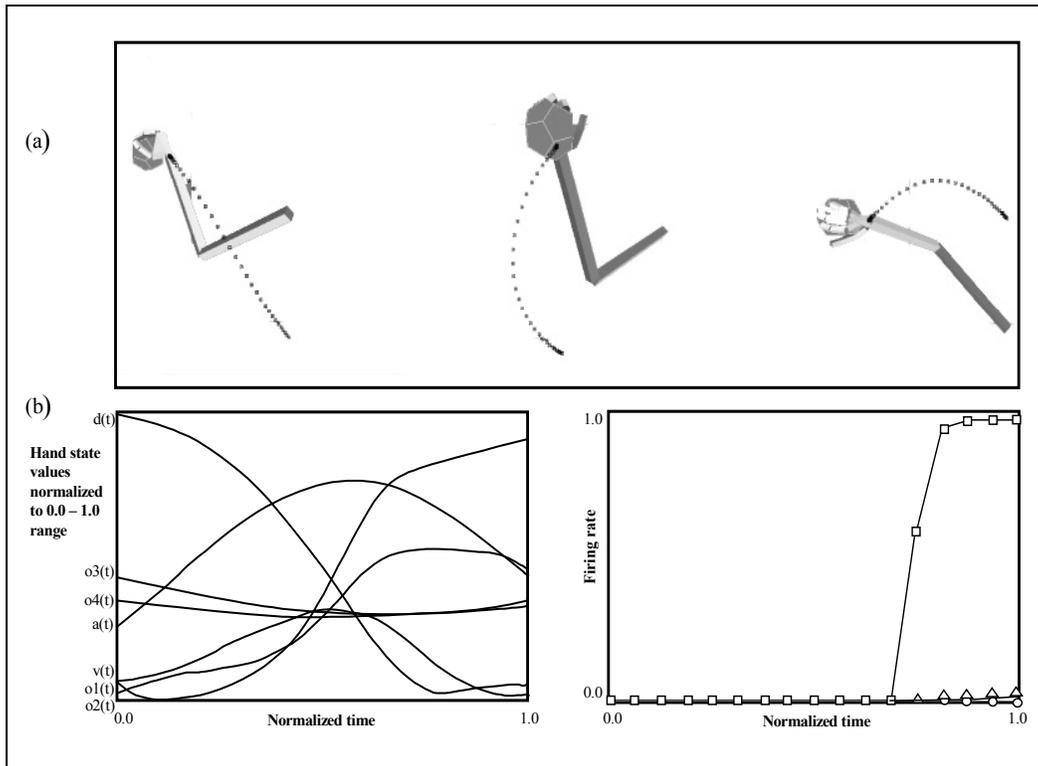


Figure 3.13 (a) A single grasp trajectory viewed from three different angles to clearly show its 3D pattern. The wrist trajectory during the grasp is marked by square traces, with the distance between any two consecutive trace marks traveled in equal time intervals. (b) Left: The input to the network. Each component of the hand state is labelled. (b) Right: How the network classifies the action as a power grasp: squares: power grasp output; triangles: precision grasp; circles: side grasp output. Note that the response for precision and side grasp is almost zero

It is often impossible (even for humans) to classify a grasp at a very early phase of the action. For example, the initial phases of a power grasp and precision grasp can be very similar. Figure 3.14 demonstrates this situation where the model changes its decision during the action and finally reaches the correct result towards the end of the action. To create this result we used the "outer limit" of the precision grasp by having the model perform a precision grasp for a wide object (using the wide opposition axis). Moreover, the network had not been trained using this object for precision grasp. In Figure 3.14(b), the curves for power and precision grips cross towards the end of the action, which shows the change of decision of the network.

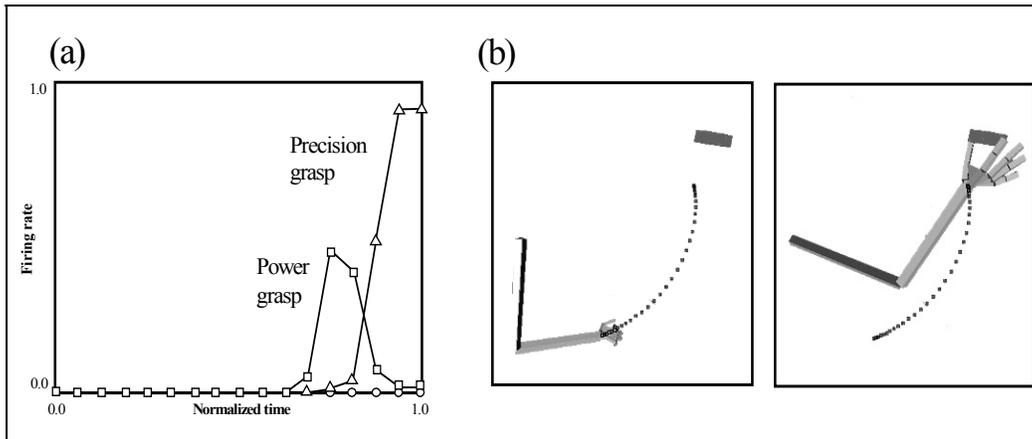


Figure 3.14 Power and precision grasp resolution. The conventions used are as in the previous figure. (a) The curves for power and precision cross towards the end of the action showing the change of decision of the network. (b) The left shows the initial configuration and the right shows the final configuration of the hand

3.5.1.2 Spatial perturbation

We next analyze how the model performs if the observed grasp action does not meet the object. Since we constructed the training set to stress the importance of distance from hand to object, we expected that network response would decrease with increased perturbation of target location.

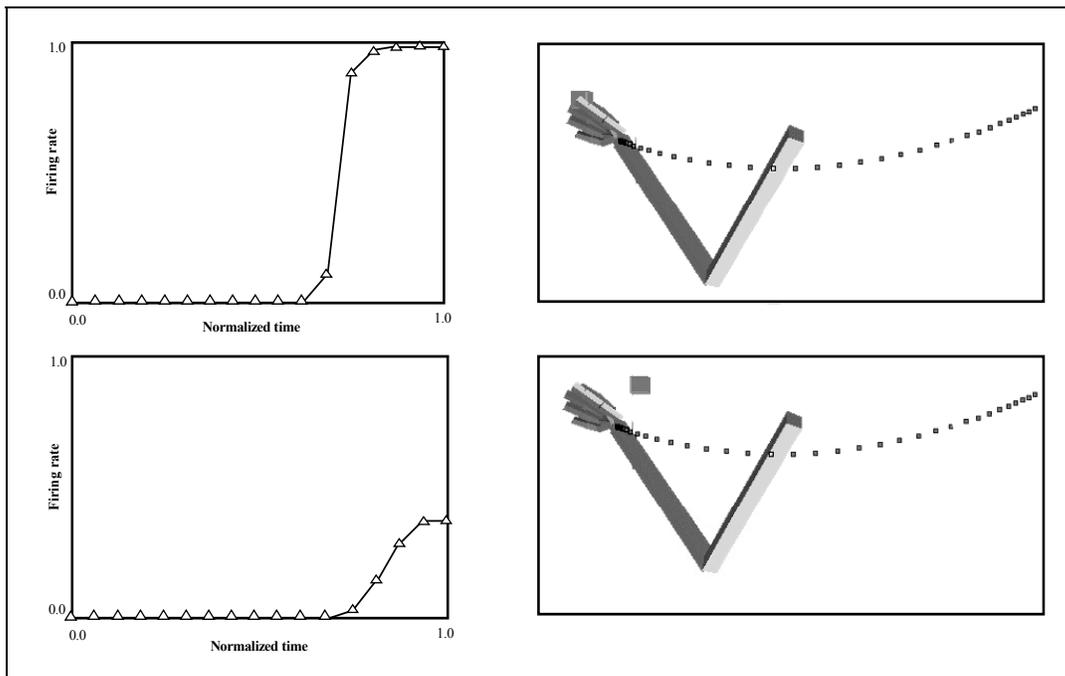


Figure 3.15: (Top) Strong precision grip mirror response for a reaching movement with a precision pinch. (Bottom) Spatial location perturbation experiment. The mirror response is greatly reduced when the grasp is not directed at a target object. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)

Figure 3.15 shows an example of such a case. However, the network's performance was not homogeneous over the workspace: for some parts of the space the network would yield a strong mirror

response even with comparatively large perturbation. This could be due to the small size of the training set. However, interestingly, the network's response had some specificity in terms of the direction of the perturbation. If the object's perturbation direction were similar to the direction of hand motion then the network would be more likely to disregard the perturbation (since the trajectory prefix would then approximate a prefix of a valid trajectory) and signal a good grasp. Note that the network reduces its output rate as the perturbation increases, however the decrease is not linear and after a critical point it sharply drops to zero. The critical perturbation level also depends on the position in space.

3.5.1.3 Altered kinematics

Normally, the simulator produces bell-shaped velocity profiles along the trajectory of the wrist. In our next experiment, we tested action recognition by the network for an aberrant trajectory generated with constant arm joint velocities.

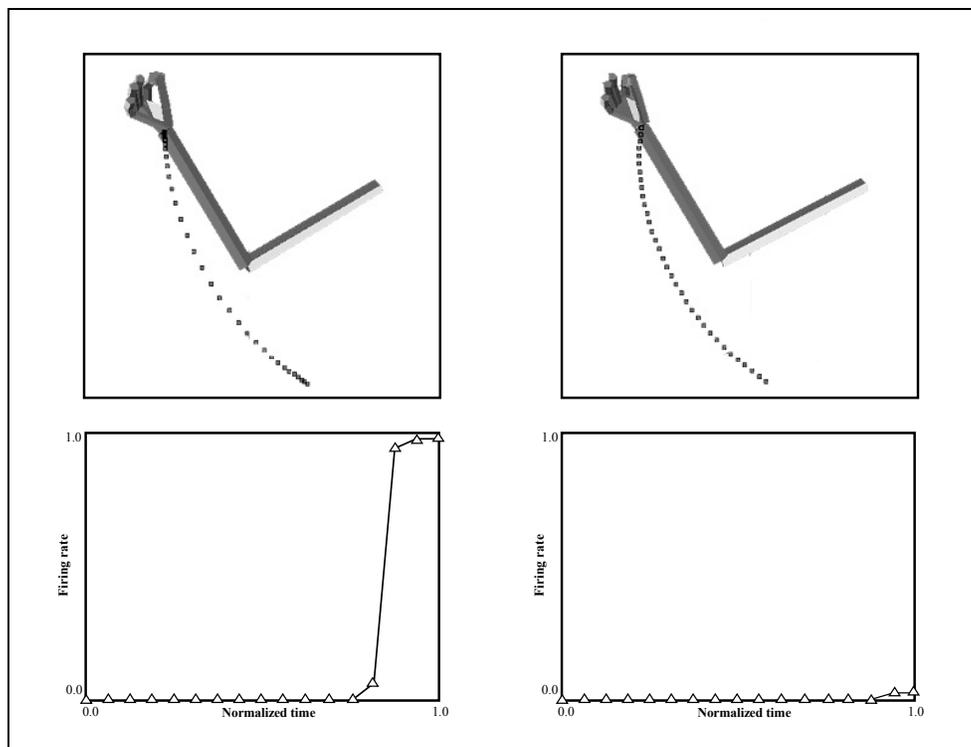


Figure 3.16 Altered kinematics experiment. Left: The simulator executes the grasp with bell-shaped velocity profile. Right: The simulator executes the same grasp with constant velocity. Top row shows the graphical representation of the grasps and the bottom row shows the corresponding output of the network. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)

The change in the kinematics does not change the path generated by the wrist. However, the trajectory (i.e., time course along the path) is changed and the network is capable of detecting this change (Figure 3.16). The notable point is that the network acquired this property without our explicit intervention (i.e. the training set did not include any negative samples for altered velocity profiles). This is because the input to the network at any time comprises 30 evenly spaced samples of the trajectory up to that time. Thus, changes in velocity can change the pattern of change exhibited across those 30 samples. The extent of this property is again dependent on spatial location.

It must be stressed that all the virtual experiments presented in this section used a single trained network. No new training samples were added to the training set for any virtual experiment.

3.5.1.4 Grasp and object axes mismatch

The last virtual experiment we present with non-explicit affordance coding explores the model's behavior when the object opposition axis does not match the hand opposition axis. This example emphasizes that the response of the network is affected by the opposition axis of the object being grasped. Figure 3.17 shows the axis orientation change for the object and the effect of this perturbation on the output of the network. The arm simulator first performed a precision grasp to a thin cylinder. The mirror neuron model's response to this action observation is shown in Figure 3.17, leftmost panel. As can be seen from the plot, the network confidently activated the mirror neuron coding precision grip. The middle panel shows the output of the network when the object is changed to a flat plate but the kinematics of the hand is kept the same. The response of the network declined to almost zero in this case. This is an extreme example – the objects in Figure 3.17 (rightmost panel) have opposition axes 90° apart, enabling the network to detect the mismatch between the hand (action) and the object. With less change in the new axis the network would give a higher response and, if the opposition axis of the objects were coincident, the network would respond to both actions (with different levels of confidence depending on other parameters).

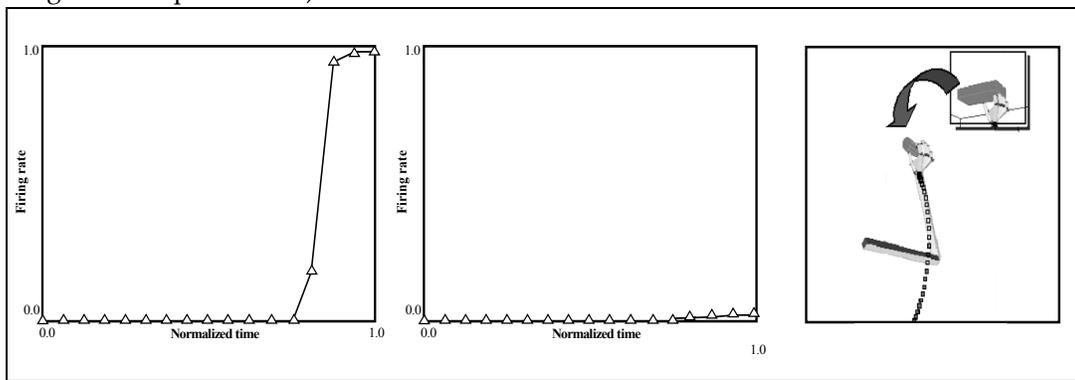


Figure 3.17 Grasp and object axes mismatch experiment. Rightmost: the change of the object from cylinder to a plate (an object axis change of 90 degrees). Leftmost: the output of the network before the change (the network turns on the precision grip mirror neuron). Middle: the output of the network after the object change. (Only the precision grasp related activity is plotted. The other two outputs are negligible.)

3.5.2 Explicit affordance coding experiments

Now we switch our attention to the explicit affordance coding network. Here we want to see the effect of object affordance on the model's behavior. The new model is similar to that given before except that it not only has inputs encoding the current prefix of the hand state trajectory (which includes hand-object relations), but also has a constant input encoding the relevant affordance of the object under current scrutiny. Thus, both the training of the network, and the performance of the trained network will exhibit effects of this additional, affordance, input.

Due to the simple nature of the objects studied here, the affordance coding used in the present study only encodes the object size. In general, one object will have multiple affordances. The ambiguity then would be solved using extra cues such as the contextual state of the network. We chose a coarse coding of object size with 10 units. Each unit has a preferred value; the firing of a unit is determined by the difference of the preferred value and the value being encoded. The difference is passed through a non-linear decay function by which the input is limited to the 0 to 1 range (the larger the difference, the

smaller the firing rate). Thus, the explicit affordance coding network has 220 inputs (210 hand state inputs, plus 10 units coarse coding the size). The number of hidden layer units was again chosen as 6 and there were again 3 output units, each one corresponding to a recognized grasp.

We have seen that the MNS model without explicit affordance input displayed a biasing effect of object size in the *Grasp Resolution* subsection of Section 5.1; the network was biased toward power grasp while observing a wide precision pinch grasp (the network initially responded with a power grasp activity even though the action was a precision grasp). The model with full affordance replicates the grasp resolution behavior seen in Figure 3.12. However, we can now go further and ask how the temporal behavior of the model with explicit affordance coding reflects the fact that object information is available throughout the action. Intuitively, one would expect that the object affordance would speed up the grasp resolution process (which is actually the case, as will be shown in Figure 3.19).

In the following two subsections we look at the effect of affordance information in two cases: (i) where we study the response to precision pinch trajectories appropriate to a range of object sizes; and (ii) where on each trial we use the same time-varying hand state trajectory but modify the object affordance part of the input. In each case, we are studying the response of a network that has been previously trained on a set of normal hand-state trajectories coupled with the corresponding object affordance (size) encoding.

3.5.2.1 Temporal effects of explicit affordance coding

To observe the temporal effects of having explicit coding of affordances to the model, we choose a range of object sizes, and then for each size drive the (previously trained) network with both affordance (object size) information and the hand-state trajectory appropriate for a precision pinch grasp appropriate to that size of object. For each case we looked at the model's response. Figure 3.18 shows the resultant level of mirror responses for 4 cases (tiny, small, medium, big objects). The filled circles indicate the precision activity while the empty squares indicate the power grasp related activity. When the object to be grasped is small, the model turns on the precision mirror response more quickly and with no ambiguity (Figure 3.18, top two panels). The vertical bar drawn at time 0.6 shows the temporal effect of object size (affordance). The curves representing the precision grasps are shifted towards the end (time = 1), as the object size gets bigger. Our interpretation is that the model gained the property of predicting that a small object is more likely to be grasped with a precision pinch rather than a power pinch. Thus the larger the object, the more of the trajectory had to be seen before a confident estimation could be made that it was indeed leading to a precision pinch. In addition, as we indicated earlier, the explicit affordance coding network displays the grasp resolution behavior during the observation of a precision grip being applied to large objects (Figure 3.18, bottom two panels: the graph labeled big object grasp and to a lesser degree, the graph labeled medium object grasp).

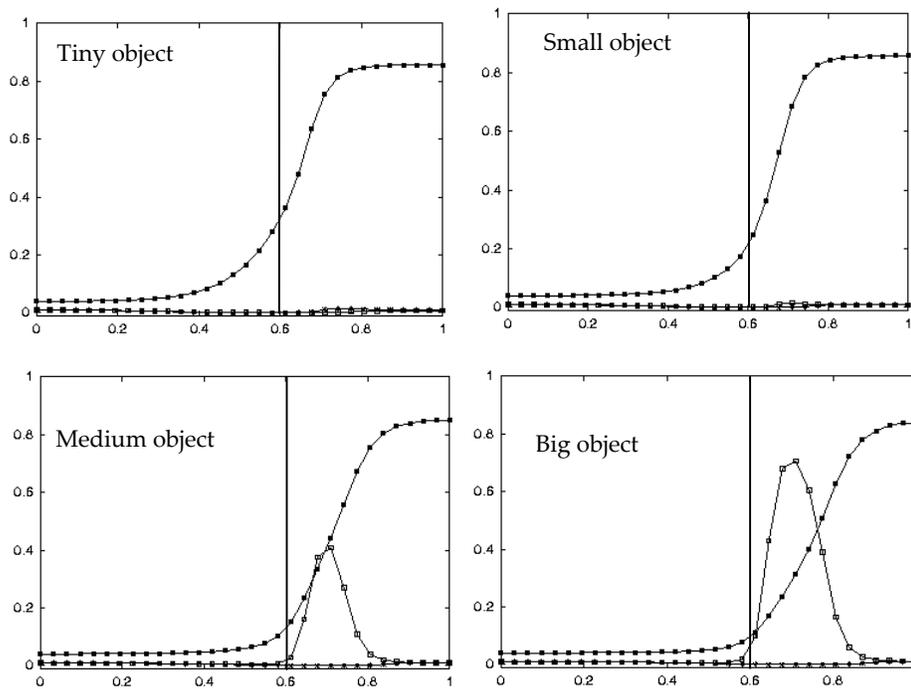


Figure 3.18 The plots show the level of mirror responses of the explicit affordance coding object for an observed precision pinch for four cases (tiny, small, medium, big objects). The filled circles indicate the precision activity while the empty squares indicate the power grasp related activity

We also compared the general response time of the non-explicit affordance coding implementation with the explicit coding implementation. The network with affordance input is faster to respond than the previous one.

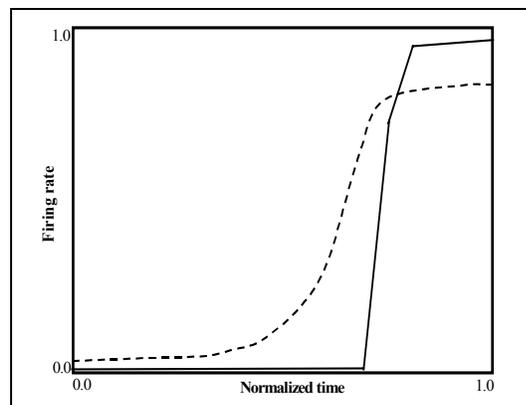


Figure 3.19 The solid curve: the precision grasp output, for the non-explicit affordance case, directed to a tiny object. The dashed curve: the precision grasp output of the model to the explicit affordance case, for the same object

Moreover, it appears that – when affordance and grasp type are well correlated – having access to the object affordance from the beginning of the action not only lets the system make better predictions but also smoothes out the neuron responses. Figure 3.19 summarizes this: it shows the precision

response of both the explicit and non-explicit affordance case for a tiny object (dashed and solid curves respectively).

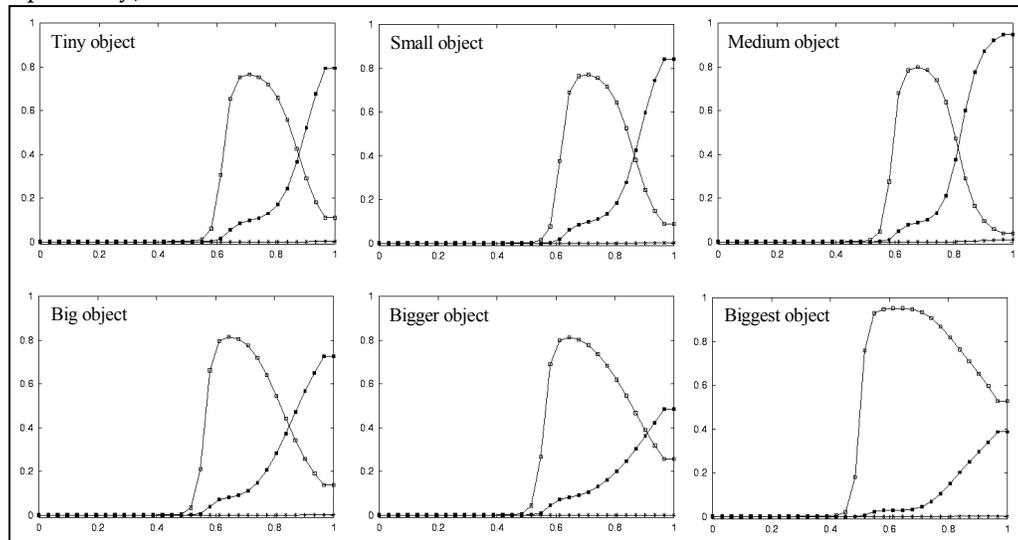


Figure 3.20: Empty squares indicate the precision grasp related cell activity, while the filled squares represent the power grasp related cell activity. The grasps show the effect of changing the object affordance, while keeping a constant hand state trajectory. In each case, the hand-state trajectory provided to the network is appropriate to the medium-sized object, but the affordance input to the network encodes the size shown. In the case of the biggest object affordance, the effect is enough to overwhelm the hand state’s precision bias.

3.5.2.2 Teasing apart the hand state and object affordance components

We now look at the case where the hand state trajectory is incompatible with the affordance of the observed object. In Figure 3.20, the plot labeled *medium object* shows the system output for a precision grasp directed to a medium-sized object whose affordance is supplied to the network. We then repeatedly input the hand state trajectory generated for this particular action but in each trial use an object affordance discordant with the observed trajectory affordance (i.e., using a reduced or increased size of the object). The plots in Figure 3.20 show the change of the output of the model due to the change in the affordance. The results shown in these plots tell us two things. First, the recognition process becomes fuzzier as the object gets bigger because the larger object sizes biases the network towards the power grasp. In the extreme case the object affordance can even overwhelm the hand state and switch the network decision to power grasp (Figure 3.20, graph labeled *biggest object*). Moreover, for large objects, the large discrepancy between the observed hand state trajectory and the size of the objects results in the network converging on a confident assessment for neither grasp.

Secondly, the resolution point (the crossing-point of the precision and power curves) shows an interesting temporal behavior. It may be intuitive to think that as the object gets smaller the network’s precision decision gets quicker and quicker (similar to what we have seen in the previous section). However, although this is the case when the object is changing size from big to small, it is not the case when the object size is getting medium to tiny (i.e., the crossing time has a local minimum between the two extreme object sizes, as opposed to being at the tiny object extreme). Our interpretation is that the network learned an implicit parameter related to the absolute value of the difference of the hand aperture and the object size such that the maximum firing is achieved when the difference is smallest, that is when the hand trajectory matches best with the object. This will explain why the network has quickest resolution for a size between the biggest and the smallest sizes.

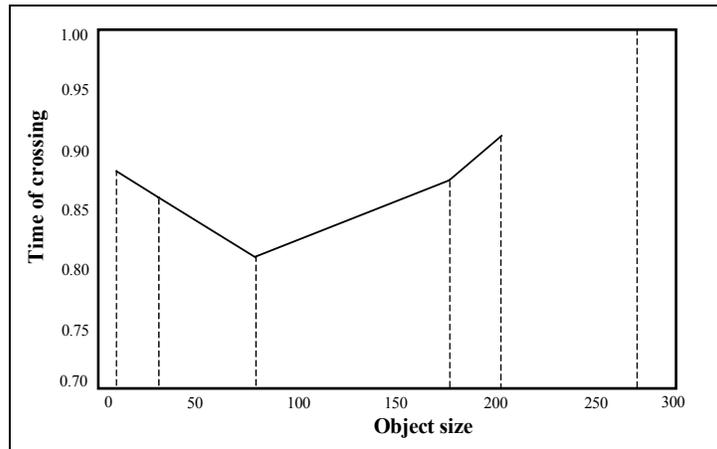


Figure 3.21 The graph is drawn to show the decision switch time versus object size. The minimum is not at the boundary, that is, the network will detect a precision pinch quickest with a medium object size. Note that the graph does not include a point for "Biggest object" since there is no resolution point in this case (see the final panel of Figure 3.19)

Figure 3.21 shows the time of resolution versus object size in graphical form. We emphasize that the model easily executes the grasp recognition task when hand-state trajectory matches object affordance. We do not include all the results of these control trials, as they are similar to the cases mentioned in the previous section.

3.5.3 Justifying the visual analysis of hand state schema

Before closing the results of this chapter, we would like to present a simulation run using a real video input to justify our claim that hand state can be extracted from real video and used to drive the core mirror circuit.



Figure 3.22 The precision grasp action used to test our visual system is depicted by superimposed frames (not all the frames are shown)

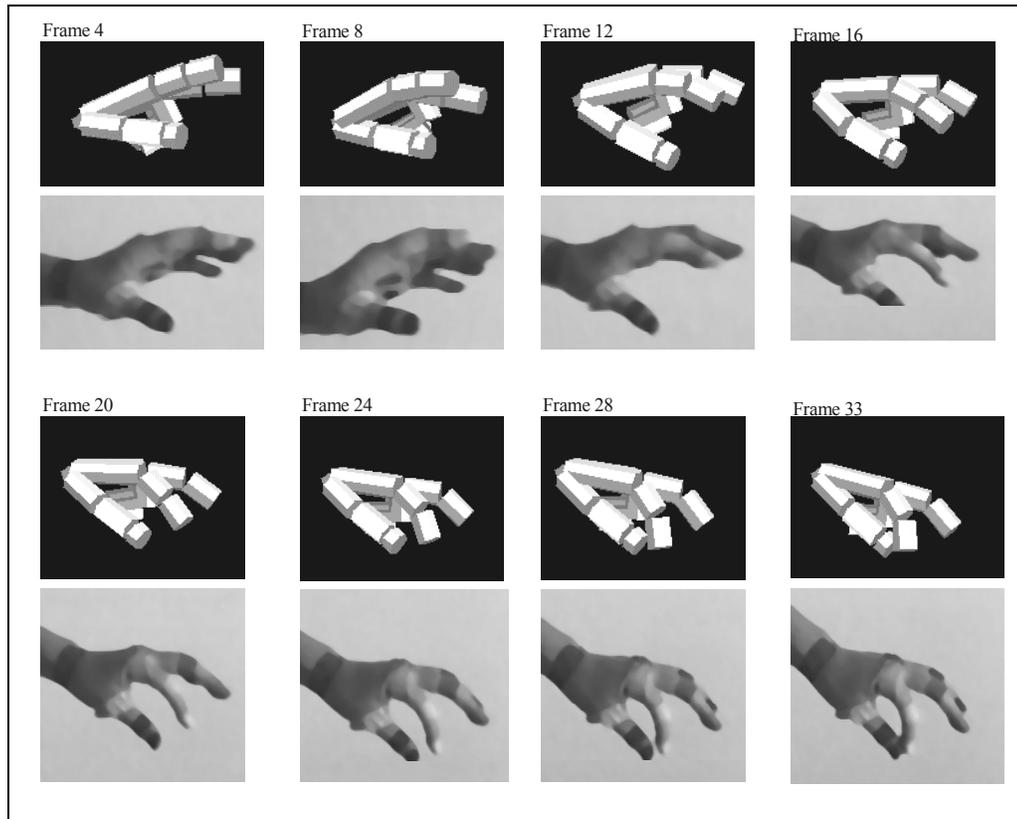


Figure 3.23 The video sequence used to test the visual system is shown together with the 3D hand matching result (over each frame). Again not all the frames are shown

The object affordances are supplied manually as we did not address object recognition in our visual system. However, the rest of the hand state is extracted by the hand recognition system as described in Section 3.4.3. Figure 3.22 depicts the precision grasp action used as input video for the simulation. The result of the 3D hand matching is illustrated in Figure 3.23. The color extraction is performed as described in the Visual Analysis of Hand State section but not shown in the figure. It would be very rewarding to perform all our MNS simulations using this system. However, the quality of the video equipment available and the computational power requirements did not allow us to collect many grasp examples to train the core mirror circuit. Nevertheless, we did test the hand state extracted by our visual system from this real video sequence on the MNS model that has already been trained with the synthetic grasp examples.

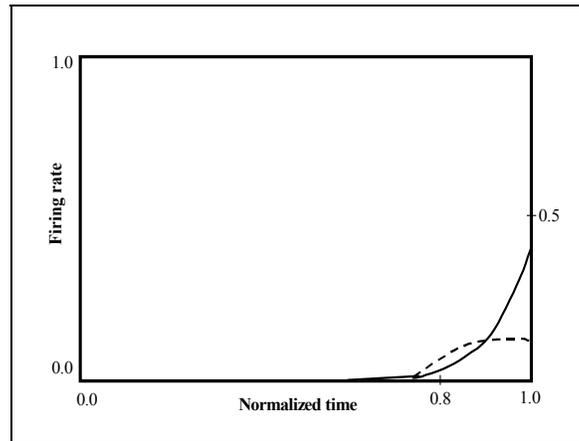


Figure 3.24 The plot shows the output of the MNS model when driven by the visual recognition system while observing the action depicted in Figure 3.22. It must be emphasized that the training was performed using the synthetic data from the grasp simulator while testing is performed using the hand state extracted by the visual system only. Dashed line: Side grasp related activity; Solid line: Precision grasp related activity. Power grasp activity is not visible as it coincides with the time axis

Figure 3.24 shows the recognition result when the actual visual recognition system provided the hand state based on the real video sequence shown in Figure 3.23. Although the output of the network did not reach a high level of confidence for any grasp type, we can clearly see that the network favored the precision grasp over the side and power grasps. It is also interesting to note a similar competition (this time between side and precision grasp outputs) took place as we saw (Figure 3.14) when the grasp action was ambiguous.

3.6 Discussion and predictions

3.6.1 The hand state hypothesis

Because the mirror neurons within monkey premotor area F5 fire not only when the monkey performs a certain class of actions but also when the monkey observes similar actions, it has been argued that these neurons are crucial for understanding of actions by others. Indeed, we agree with the importance of this role and indeed have built upon it elsewhere, as we now briefly discuss. Rizzolatti et al. (1996b) used a PET study to show that both grasping observation and object prehension yield highly significant activation in the rostral part of Broca's area (a significant part of the human language system) as compared to the control condition of object observation. Moreover, Massimo Matelli (in Rizzolatti and Arbib 1998) demonstrated a homology between monkey area F5 and area 45 in the human brain (Broca's area comprises areas 44 and 45). Such observations led Rizzolatti and Arbib (1998) building on Rizzolatti et al. (1996a) to formulate:

The Mirror System Hypothesis: Human Broca's area contains a mirror system for grasping which is homologous to the F5 mirror system of monkey, and this provides the evolutionary basis for language parity - i.e., for an utterance to mean roughly the same for both speaker and hearer. This adds a neural "missing link" to the tradition that roots speech in a prior system for communication based on manual gesture.

Arbib (2001) then refines this hypothesis by showing how evolution might have bridged from an ancestral mirror system to a 'language ready' brain via increasingly sophisticated mechanisms for imitation of manual gestures as the basis for similar skills in vocalization and the emergence of

protospeech. In some sense, then, the present paper can be seen as extending these evolutionary concerns back in time. Our central aim was to give a computational account of the monkey mirror system by asking (i) What data must the rest of the brain supply to the mirror system? and (ii) How could the mirror system learn the right associations between classification of its own movements and the movement of others? In seeking to ground the answer to (i) in earlier work on the control of hand movements (Iberall and Arbib 1990) we were led to extend our evolutionary understanding of the mirror system by offering:

The hand state hypothesis: The basic functionality of the F5 mirror system is to elaborate the appropriate feedback – what we call the *hand state* – for opposition-space based control of manual grasping of an object. Given this functionality, the social role of the F5 mirror system in understanding the actions of others may be seen as an exaptation gained by generalizing from self-hand to other's-hand.

The hand state hypothesis provides a new explanation of the evolution of the ‘social capability’ of mirror neurons, hypothesizing that these neurons first evolved to augment the ‘canonical’ and ‘pure motor’ F5 neurons by providing visual feedback on ‘hand state’, relating the shape of the hand to the shape of the object.

3.6.2 Neurophysiological predictions

We introduced the MNS (Mirror Neuron System) model of F5 and related brain regions as an extension of the FARS model of circuitry for visually-guided grasping of objects that links parietal area AIP with F5 canonical neurons. The MNS model diagrammed in Figure 3.5 includes hypotheses as to how different brain regions may contribute to the functioning of the mirror system. Chapter 6 undertakes the neural implementation of Grasp Learning (area F4, F2 and F5). This chapter focused on the Core Mirror Circuit by aggregating the other functionality into three ‘grand schemas’ – visual analysis of hand state, reach and grasp. Thus we only claim that core mirror circuit is relevant for neurophysiological predictions. We developed the visual analysis of hand state schema to the point of demonstrating algorithms powerful enough to take actual video input of a hand (though we simplified the problem by using colored patches) and produce hand state information. The reach and grasp schema then represented all the functionality for taking the location and affordance of an object and determining the motion of a hand and arm to grasp it (however see Chapter 6 for a detailed neural implementation of this circuit grounded in neurophysiology and infant behavior). As the main aim of this chapter was to analyse the core mirror circuit we showed that if we used the reach and grasp schema to generate an observed arm-hand trajectory (i.e., to represent the reach and grasp generator of the monkey or human being observed), then that simulation could directly supply the corresponding hand-state trajectory, and we thus use these data so that we can analyze the core mirror circuit schema (Figure 3.6(b)) in isolation from the visual analysis of hand state. However note that we have also justified the visual analysis of hand state schema by showing in a simulation that the core mirror circuit can be driven with the proposed vision system without any synthetic data from the reach and grasp schema.

Moreover, the hand state input (regardless of being synthetic or real) was presented to the network in a way to avoid the use of a dynamic neural network. To form the input vector, each of the seven components of the hand state trajectory, up to the present time t , is fitted by a cubic spline. Then this spline is sampled at 30 uniformly spaced intervals; i.e., no matter what fraction t is of the total time T of the entire trajectory, the input to the network at time t comprises 30 samples of the hand-state uniformly distributed over the interval $[0, t]$. The network is trained using the full trajectory of the hand state in a specific grasp; the training set pairs each such hand state history as input with the final grasp type as

output. On the contrary, when testing the model with various grasp observations, the input to the network was the hand state trajectory that was available up to that instant. This exactly parallels the way the biological system (the monkey) receives visual (object and hand) information: When the monkey performs a grasp, the learning can take place after the observation of the complete (self) generated visual stimuli. On the other hand, in the observation case the monkey mirror system predicts the grasp action based on the partial visual stimuli (i.e. before the grasp is completed). The network thus yields a time-course of activation for the mirror neurons, yielding predictions for neurophysiological experiments by highlighting the importance of the *timing* of mirror neuron activity. We saw that initial prefixes will yield little or no mirror neuron activity, and ambiguous prefixes may yield transient activity of the 'wrong' mirror neurons.

Since our aim was to show that the connectivity of mirror neuron circuitry can be established through training, and that the resultant network can exhibit a range of novel, physiologically interesting, behaviors during the process of action recognition, the actual choice of training procedure is purely a matter of computational convenience, and the fact that the method chosen, namely back-propagation, is non-physiological does not weaken the importance of our predictions concerning the timing of mirror neuron activity.

With this we turn to neurophysiological predictions made in our treatment of the Core Mirror Circuit, namely the 'grounding assumptions' concerning the nature of the input patterns received by the circuit and the actual predictions on the timing of mirror neuron activity yielded by our simulations.

Grounding assumptions: The key to the MNS model is the notion of *hand state* as encompassing data required to determine whether the motion and preshape of a moving hand may be extrapolated to culminate in a grasp appropriate to one of the affordances of the observed object. Basically a mirror neuron must fire if the preshaping of the hand conforms to the grasp type with which the neuron is associated; and the extrapolation of *hand state* yields a time at which the hand is grasping the object along an axis for which that affordance is appropriate. What we emphasize here is *not* the specific decomposition of the hand state $F(t)$ into the seven specific components ($d(t)$, $v(t)$, $a(t)$, $o_1(t)$, $o_2(t)$, $o_3(t)$, $o_4(t)$) used in our simulation, but rather that the input neural activity will be a distributed neural code which carries information about the movement of the hand toward the object, the separation of the virtual fingertips and the orientation of different components of the hand relative to the opposition axis in the object. The further claim is that this code will work just as well for measuring how well another monkey's hand is moving to grasp an object as for observing how the monkey's own hand is moving to grasp the object, allowing self-observation by the monkey to train a system that can be used for observing the actions of others and recognizing just what those actions are.

We provided experiments to compare the performance of the Core Mirror Circuit with and without the availability of explicit affordance information (in this case the size of the object) to strengthen our claim that it is indeed adaptive for the system to have this additional input available, as shown in Figure 3.6(b). Note that the "grasp command" input shown in the figure serves here as a training input, and will, of course, play no role in the recognition of actions performed by others.

Also we have given a justification of the visual analysis of hand state schema by showing in a simulation that the core mirror circuit can be driven with the visual system we implemented without requiring the Reach and Grasp simulator provide syntetic data.

Novel Predictions: Experimental work to date tends to emphasize the actions to be correlated with the activity of each individual mirror neuron, while paying little attention to the temporal dynamics of mirror neuron response. By contrast, our simulations make explicit predictions on how a given (hand

state trajectory, affordance) pair will drive the time course of mirror neuron activity – with non-trivial response possibly involving activity of other mirror neurons in addition to those associated with the actual grasp being observed. For example, a grasp with an ambiguous prefix may drive the mirror neurons in such a way that the system will, in certain circumstances, at first give weight to the wrong classification, with only the late stages of the trajectory sufficing for the incorrect mirror neuron to be vanquished.

To obtain this prediction we created a scene where the observed action consisted of grasping a wide object with precision pinch (thumb and index finger opposing each other). Usually this grasp is applied to small objects (imagine grasping a pen along its long axis versus grasping it along its thin center axis). The mirror response we got from our core mirror circuit was interesting. First, the system recognized (while the action was taking place) the action as power grasp (which is characterized by enclosing the hand over large objects; e.g. grasping an apple) but as the action progressed the model unit representing precision pinch started to get active and the power grasp activity started to decline. Eventually the core mirror circuit settled on the precision pinch. This particular prediction is testable and indeed suggests a whole class of experiments. The monkey has to be presented with unusual or ambiguous grasp actions that require a ‘grasp resolution’. For example, the experimenter can grasp a section of banana using precision pinch from its long axis. Then we would expect to see activity from power grasp related mirror cells followed by a decrease of that activity accompanied by increasing activity from precision pinch related mirror cells.

The other simulations we made leads to different testable predictions such as the mirror response in case of a spatial perturbation (showing the monkey a fake grasp where the hand does not really meet the object) and altered kinematics (perform the grasp with different kinematics than usual). The former is in particular a justification of the model, since in the mirror neuron literature it has been reported that the spatial contact of the hand and the object is usually required for the mirror response (Gallese et al. 1996). On the other hand, the altered kinematics result predicts that an alteration of the kinematics will cause a decrease in the mirror response. We have also noted how a discrepancy between hand state trajectory and object affordance may block or delay the system from classifying the observed movement.

In summary, we have conducted a range of simulation experiments – on grasp resolution, spatial perturbation, altered kinematics, temporal effects of explicit affordance coding, and analysis of compatibility of the hand state to object affordance – which demonstrate that the present model is not only of value in providing an implemented high-level view of the logic of the mirror system, but also serves to provide interesting predictions ripe for neurophysiological testing, as well as suggesting new questions to ask when designing experiments on the mirror system.

4 CHAPTER IV: MULTILAYER SUPERVISED HEBBIAN LEARNING AND PROBABILITY CODING

This chapter introduces a learning and data generation model that can be employed in multi-layered circuits. The architecture that we develop in this chapter will be used in the Grasp Learning Models of Chapters 5 and 6. The adaptation of the network weights is performed in a hebbian fashion based on a reinforcement signal. In the general reinforcement learning framework the learning problem is formulated as an agent acting in an environment that returns rewards based on the actions of the agent and state of the environment (Sutton and Barto 1998). By acting, the agent can (and usually does) change the state of the environment. The goal of the agent is to maximize its total reward in the long run, possibly in infinite future (Sutton and Barto 1998). Unlike other optimization based learning methods, reinforcement learning can be implemented by biological circuits. The supervised Hebbian learning that we introduce in this chapter is a special case of the general reinforcement learning. We will use the terms supervised Hebbian and reinforcement learning interchangeably.

4.1 Neural coding

Although there is no agreement on the exact coding neurons employ in the brain we in general adopt a population coding approach where information is represented in the activities of a group of neurons. What we mean by the activity of a neuron is the average firing rate. We do not reject other possibilities such as interval coding where the timing of the spikes carry information but we adopt average firing rate. There is large amount of experimental data showing that a populating coding scheme is employed in the brain (e.g. see Georgopoulos 1986). The term preferred-stimulus is used to indicate the stimulus that causes a neuron to fire maximally. It is observed that neurons would fire for the stimuli which are not the *best*, with a decreasing rate as the stimuli diverges from the preferred one. In most cases, the population code is thought to be encoding a single variable (like the orientation of an edge)(Zemel et al. 1998). However, the population can be involved in encoding more than the *value* of a variable such as the *variance* and *uncertainty* of the variable (Zemel et al. 1998). Indeed there are cases where the probability distributions of the variables are more adequate than the single values as when insufficient information exists to differentiate between two values of the variable (stimuli), or when multiple values underlie the input (former case, place cells in the hippocampus (O'Keefe and Dostrovsky 1971); the latter case directional motion detecting cells in MT (Newsome et al. 1989)).

We adapt the probabilistic interpretation of the population coding (Anderson 1994; Zemel et al. 1998) to represent *multiple values* or *options* given the input stimuli. For example, given a sphere shaped object, what is the most natural representation of a parameter that specifies the direction (e.g. top, down, right, etc.) from which the animal should grasp the object? The most general answer to this question is to have a representation to encode the parameter with its probability distribution. Indeed, in this example, the sphere can be grasped in very many ways. We view each layer in a network as representing the possible choices for the next layer. When the animal needs to act based on the

processed input then it *samples* over the output to generate the movement⁷ possibly by a winner-take-all mechanism.

The ‘usual’ way of using a neural network to perform density estimation is much like a function approximation, where the network is asked to give a single output (the value of the probability density function) given an input (see Bishop 1995). In the architecture we proposed, a neural structure is used to compute a normalized histogram in terms of population activity. Besides biological relevance, by representing the distribution as a population code, we gain the power to use the network’s output as inputs to other networks or layers without an extra decoding step.

4.2 Operation of the proposed network

In the abstract setting, we posit an agent that takes actions and the environment (can include the agent itself) returns rewards or binary signals indicating whether the action was ‘good’ or ‘bad’. Some information of the environment may be available to the agent as input. We denote the output of the agent with $\mathbf{y}(t)$, the input as $\mathbf{x}(t)$ and the reward returned by $r(t)$. We are interested in constructing a network that can be implemented using neural units to achieve: (1) estimation of the probability distribution of the reward given an input, and (2) generating outputs that return rewards approximating the probability distribution of the reward.

We will use $\mathbf{X}(t)$ to denote the input instead of the $\mathbf{x}(t)$ to emphasize that in general the input can be the output of a different layer.

A layer is composed stochastic units (Hertz et al. 1991). The layer generates an output based on the probability distribution represented by $\mathbf{Y}(t)$, the firing potentials of the stochastic units (described below). We introduce layers \mathbf{X}_e , \mathbf{Y}_e as memory traces (a rudimentary form of the ‘eligibility traces’ of Sutton and Barto 1998) for keeping the memory of activities of layers \mathbf{X} and \mathbf{Y} , respectively (we will use ‘memory trace’ and ‘eligibility’ terms interchangeably in this chapter). The connectivity between layers \mathbf{X} and \mathbf{Y} is established through the weight matrix \mathbf{W} . The network operates as the following:

1. Input $X(t)$ is presented
2. $X_e(t) = X(t)$
3. $Y_{net}(t) = g(W(t)X(t))$
where $g(z) = \begin{cases} 0, & z < 0 \\ z, & z \geq 0 \end{cases}$
4. $Y(t) = \frac{Y_{net}(t)}{\sum_{k=1}^N Y_{net}^k(t)}$
5. For $k=1, \dots, N$ $Y_{fire}^k(t) = H(\varphi)$ with probability $Y^k(t)$, where φ non-negative small ($\ll 1$) uniform noise term and H is the step function
6. $k^* = \begin{cases} FTA(Y_{fire}), & \text{with Pr}(1 - \alpha) \\ \text{unif}([1, N]), & \text{with Pr}(\alpha) \end{cases}$ where FTA stands for First-Fire Take All process

⁷ Here we don’t exclude the fact that different biasing mechanisms may reshape the final output

7. $y(t) = \sum_k P^k e^{-\frac{\delta(k,k^*)}{2\sigma^2}}$, where P^k is the preferred action associated with unit k and σ^2 is the variance determining the size of the weighted sum range
8. $Y_e^k(t) = e^{-\frac{\delta(k,k^*)}{2\sigma_e^2}}$, for $k=1,..,N$
9. $W_{ij}(t) = W_{ij}(t-1) + \eta r(t-1) X_e^i(t-1) Y_e^j(t-1)$, where η is the learning rate parameter

The steps 1-2, is standard: the input is propagated via a weight matrix to form the net input. Step 3 applies a threshold operation to eliminate negative net inputs. Step 4 converts the net input into firing potentials by normalization. The function of steps 5 and 6 is not as apparent. Each stochastic unit fires with probability that is equal to its firing potential (computed in Step 4). However, the spike generation timing is variable due to uniform noise (Step 5). Step 6 has two modes of operation. With probability α , a random unit is selected from \mathbf{Y} as a winner (k^*), and with probability $1-\alpha$, a FirstFire-Take-All⁸ process is applied on the stochastic units to find real winner (k^*).

Let us consider the case where no exploratory moves are performed ($\alpha=0$), then what is the distribution of k^* ? The Steps 5 and 6 generate winner k^* s such that $E(Y_{fire}^k | k = k^*) = Y(k)$ holds for all k^9 . In essence we implement an approximate universal data generation (see Leydold and Hormann 2000 for a review) method that is open to neural implementation. Step 7 implements a local population vector computation around the winning unit. The distance metric δ defines how similar the neurons' preferred values are. This metric imposes a topological relation between the units through the learning rule. Step 8 creates eligibility traces to register the activated regions for the next cycles weight update. Step 9 updates the weights in a supervised Hebbian learning fashion (or reinforcement learning update as in Willams' (1992) REINFORCE). The weight update rule creates a normalized histogram of the reward yielding input patterns: **Claim:** Given input \mathbf{X}^m , one of the S distinct input vectors, at time step p , $\mathbf{Y}(k)$ approximates the probability distribution of the reward. **Proof:** We are given the distinct input patterns as $(\mathbf{X}^1, \mathbf{X}^2, .., \mathbf{X}^s)$. We prove the claim with the following assumptions:

1. Memory trace of \mathbf{Y} is concentrated on a single point (let $\mathbf{Z}(t)=\mathbf{Y}_e(t)$ for notional convenience).
2. The rewards returned by the environment are non-negative
3. Number of the training inputs is far larger than the number of distinct input patterns ($p \gg S$)
4. $\alpha=1$

Given a randomly initialized network, When we apply the inputs $\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_{p-1}$ according to the layer update equations (1-9) above we get: (the subscripts denote the time step while the superscripts denote the pattern number.)

⁸ FTA is proposed for spiking neurons analogously to winner take all. Given a set of neurons, and a time interval FTA will return the one which fires first in the given interval. FTA can be implemented using lateral inhibition in biological networks

⁹ Given large number of action generations (p) and large number of output units (S), the expected value of the number of firings of unit k is $p \cdot Y(k)$. At each cycle, the selection of the winner is uniformly distributed (Step 5) over the population of units that are firing. By the law of large numbers the expected number of units that fire can be approximated with Gaussian distribution with some mean (μ) and variance (σ^2). Thus the expected number of a unit will fire and become winner can be approximated as $p \cdot Y(k) / \mu$. Thus the normalized histogram of the winners (k^*) over a large data generation will approximate \mathbf{Y} .

$$Y_{net}(1) = g(W_0 X_1)$$

$$Y_{net}(2) = g((W_0 + \eta r_1 Z_1 X_1^T) X_2)$$

$$Y_{net}(3) = g((W_0 + \eta r_1 Z_1 X_1^T + \eta r_2 Z_2 X_2^T) X_3)$$

...

$$Y_{net}(p) = g((W_0 + \eta r_1 Z_1 X_1^T + \eta r_2 Z_2 X_2^T + \dots + \eta r_{p-1} Z_{p-1} X_{p-1}^T) X_p)$$

Noting that the argument of g is nonnegative and the initial weights' (W_0) contribution can be made arbitrarily small we can write:

$$Y_{net}(p) = W_0 X_p + \eta r_1 Z_1 X_1^T X_p + \eta r_2 Z_2 X_2^T X_p + \dots + \eta r_{p-1} Z_{p-1} X_{p-1}^T X_p$$

If the input patterns that generate the rewards are approximately mutually orthogonal then we can write:

$$Y_{net}(p) \approx \eta \|X_p^s\|^2 \sum_{Z \in V} r^s Z^s \text{ where } s \text{ is from } 1, \dots, S \text{ (} s \text{ indexes the distinct input patterns)}$$

Here, V is the set of Z 's that appear with X^s in the expression for $Y_{net}(p)$. Therefore $Y_{net}(p)$ is the *total reward* for each Z in V . If the reward returns are binary, then $Y_{net}(p)$ becomes (unnormalized) histogram of the *occurrence* of rewards. Thus, step 4 generates the normalized histogram.

In simulations, we observed that *mild* violation of the assumptions used in the proof does deform the final representation in Y . However, the value of α requires some comment. This parameter controls how much the layer is exploring as opposed to exploiting the current probability distribution. The learning rate, η , should be reduced as the layer starts exploiting more. When η and α are chosen constant non-zero values, in the limit, the layer degenerates into a winner-take-all circuit. To see this observe that in the expression of $Y_{net}(p)$ each presentation of input is accompanied by a Z value which is the indicator of the unit (and its neighborhood) that determined the action.

$$Y_{net}(p) = W_0 X_p + \eta r_1 Z_1 X_1^T X_p + \eta r_2 Z_2 X_2^T X_p + \dots + \eta r_{p-1} Z_{p-1} X_{p-1}^T X_p$$

When $\alpha=0$, Z does not depend on the current state of the layer. If $\alpha>0$ then the action generation will pick actions from a neighborhood of the actions that returned reward (with probability α). That is the terms of $Y_{net}(p)$ will be biased. Thus, the histogram computed by the layer will be sharper. In practice, having small α ($\alpha<0.2$), reflects the reward distribution in the layer and provides enough exploitation that the learning is faster than trial and error (i.e. $\alpha=0$ case). The Hebbian update we use in Step 9 although, biologically realistic, does not limit the weight growth. In practice, either a weight normalization or use of an adaptive learning rate that tends to zero is required if the weights need to be bounded.

The analysis when the rewards can be negative is similar. A non-zero $Y_{net}(P)$ can be split into negative and positive reward terms and written in this form:

$$Y_{net}(p) \approx \eta \|X_p^s\|^2 \left(\sum_{Z \in V^+} r^s Z^s - \sum_{Z \in V^-} |r^s| Z^s \right) \text{ where } V^+ \text{ is the set of } Z\text{'s that appear with } X^k \text{ for}$$

which $r^k>0$ and V^- is the set of Z 's that appear with X^k for which $r^k\leq 0$. This simply tells us that given an input X^s at time step p , the firing potential is represented as the difference of two distributions: the positive reward distribution and the negative reward distribution. An intuitive example would be the situation where an action of the agent returns +1 or -1 with equal probability (regardless of state of the environment). Then the learning will generate a Y_{net} distribution where the action is represented as returning 0 reward.

4.3 Testing the proposed architecture

We present two example problems to demonstrate that the architecture we proposed can learn to generate actions that capture the distribution of the reward. In both cases, we take an environment with known reward distribution and we let the agent interact with the environment with the learning rules specified in the previous sections. After the training, we check whether, by interacting with the environment, the network was able to discover the underlying distribution.

In the next chapter we will use this architecture for multivariate action domains, so we chose a two-dimensional action space to make sure that the architecture is adequate for the grasp learning tasks of the next chapters.

4.3.1 Deterministic environment

The test problem is defined as the following. The agent acts by generating a vector (x, y) where each component is from the interval $[-1, 1]$. Then the environment returns a deterministic reward of +1 if $(x^2 + y^2) < 0.25$ and -1 otherwise (Figure 4.1). We used 400 (20x20 grid) neurons to represent the output in the layer. The topology of the layer is reflected through the specification of the distance metric, δ . We defined the distance metric as the usual Euclidean metric over the 2D action space.

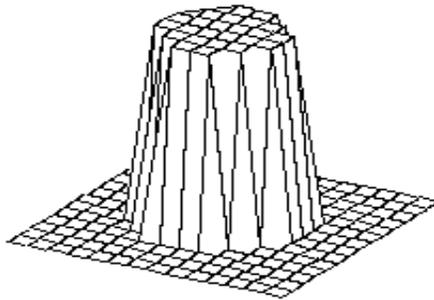


Figure 4.1 The elevated circular region corresponds to the area defined by the equation $(x^2 + y^2) < 0.25$. The environment returns +1 as the reward if the action falls into the circular region, otherwise -1 is returned.

In the learning simulations, the agent generated the actions from its action probability distribution (\mathbf{Y}) with 0.5 probability and performed exploratory trials with 0.5 probability¹⁰ (i.e. $\alpha=0.5$). Figure 4.2 shows the evolution of the network outputs (\mathbf{Y}) subject to the learning rules specified in the previous sections. The firing potential (\mathbf{Y}) is shown as a grid to allow comparison with the environment distribution.

¹⁰ Depending on the task, as the learning progresses α can be adjusted

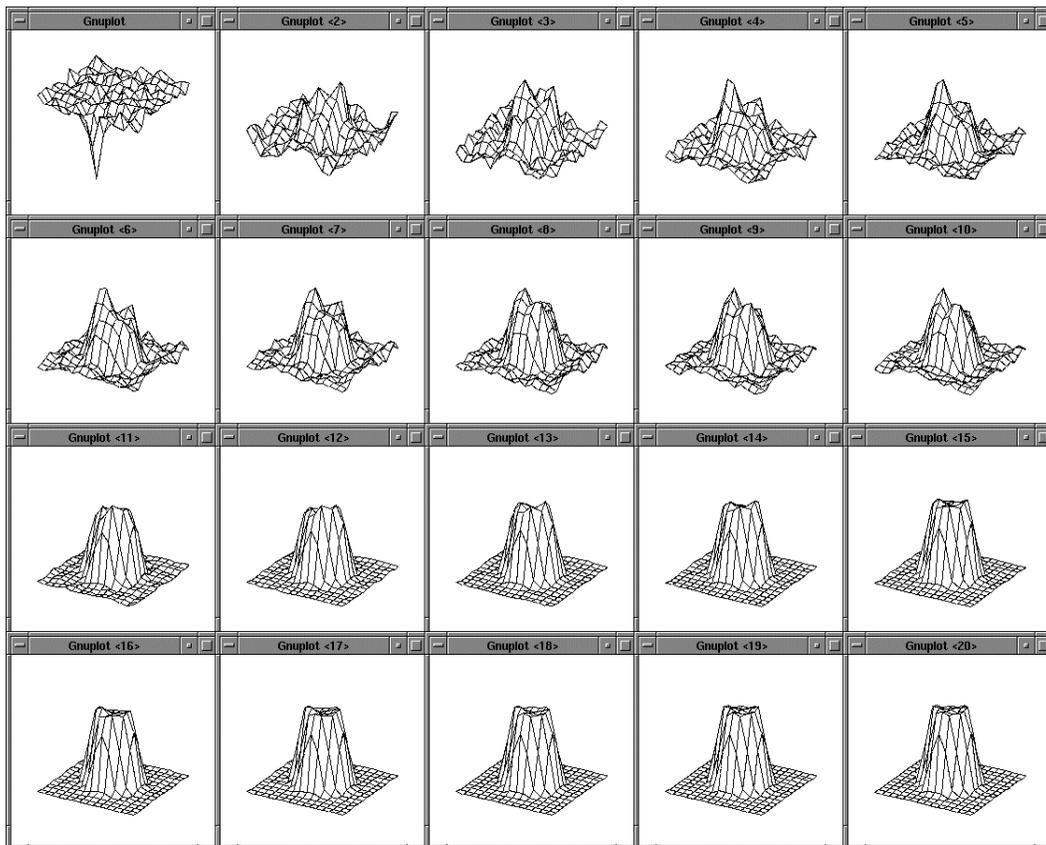


Figure 4.2 The adaptation of the firing potential of the stochastic units are shown as a series of evolving 3D maps. (left to right and top to bottom)

Figure 4.2 only confirms that the firing potentials of the output units capture the reward distribution. We have to test our assertion that the generation of actions by using Steps 5 and 6 of the previous section approximates the distribution of the environment. Figure 4.3, shows that the actions generated with the dynamics we proposed in the previous section leads to a histogram that approximates the environment's reward distribution (compare to Figure 4.1).

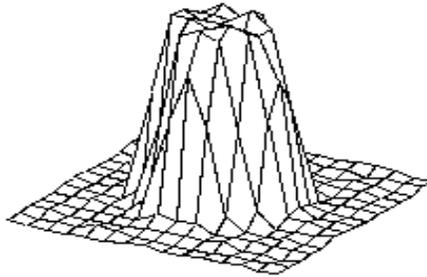


Figure 4.3 The normalized histogram of the actions generated over 60000 samples¹¹. Note that the actions generated captured the environment's reward distribution (see Figure 4.1).

4.3.2 Stochastic environment

The environment we defined above was simple: it was deterministic and the set of actions that would yield reward was convex. Now we confront the architecture we proposed with a more realistic case, where the reward distribution is multimodal and the rewards are returned stochastically. Given an action (x, y) the environment returns a reward of +1 with the probability:

$$p(x, y) = \frac{1}{2\sqrt{2\pi}\sigma^2} \int_{(x,y) \in R} e^{-\frac{(x-0.5)^2+(y-0.5)^2}{2\sigma^2}} + e^{-\frac{(x+0.5)^2+(y+0.5)^2}{2\sigma^2}},$$

where $\sigma^2 = 0.1$ and R is the rectangular grid of size 0.05×0.05 where (x, y) falls in.

The distribution is the summation of two Gaussians centered on $(0.5, 0.5)$ and $(-0.5, -0.5)$, hence double peaked. (Figure 4.4)

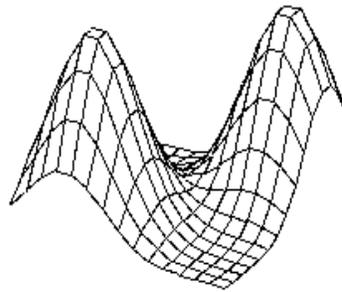


Figure 4.4 The stochastic environment's double peaked reward distribution (see text for the explanation)

¹¹ We arbitrarily generated 60000 actions. The larger the number of samples the better the approximation of underlying distribution.

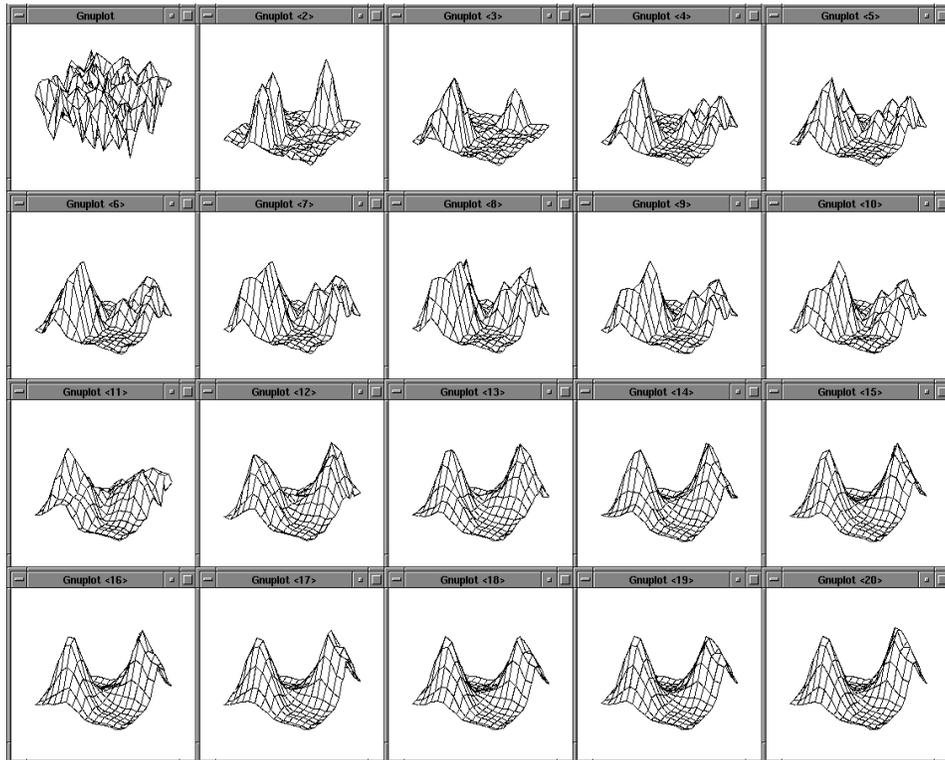


Figure 4.5 Some snapshots showing the phases of learning of the layer in the stochastic environment where the reward distribution has two peaks (see Figure 4.4).

The evolution of the layer's firing potential presented in Figure 4.4 shows differential phases than the unimodal deterministic environment. First, the peaks are sorted out as isolated sharp peaks with different heights. Then the potentials settle on equal height peaks capturing the input distribution (Figure 4.4).

Similar to what we did in the deterministic environment case, we generated actions using the trained layer to test whether the distribution of the generated action matches the input reward distribution. Figure 4.6 shows the normalized histogram of the actions generated. The main structure of the environment was captured (compare to Figure 4.4).



Figure 4.6. The normalized histogram of 60000¹² data points (actions) generated by the trained layer in the stochastic environment depicted in Figure 4.4.

4.3.3 Combining multiple layers

The main aim of introducing the proposed architecture is to combine them for more complex computations. In a hypothetical problem, layers A and B independently compute some parameters. A third layer, layer C uses the output of layer A and B to generate the final output of the combined network. The output of C depends on both A and B. Therefore C has to discovered the action distribution give the inputs from A and C. The combining layer performs this via a synaptic multiplication (we drop the time dependency and use matrix notation for clarity). The subscript e is used to denote the eligibility traces as before. The synaptic input channels of C and the corresponding weights are denoted with C^{synA} , C^{synB} ; and W^{synA} , W^{synB}

Compute the net synaptic inputs:

$$\begin{aligned} C_{net}^{synA} &= g(W^{synA} A_e^T) \\ C_{net}^{synB} &= g(W^{synB} B_e^T) \end{aligned} \quad \text{where } g(z) = \begin{cases} 0, & z < 0 \\ z, & z \geq 0 \end{cases}$$

Compute the normalized synaptic inputs:

$$\begin{aligned} C^{synA} &= \frac{C_{net}^{synA}}{\sum_{k=1}^N C_{net,k}^{synA}} \\ C^{synB} &= \frac{C_{net}^{synB}}{\sum_{k=1}^N C_{net,k}^{synB}} \end{aligned}$$

Combine the normalized synaptic inputs¹³:

¹² We arbitrarily generated 60000 actions. The larger the number of samples the better the approximation of the underlying distribution.

¹³ The * operator performs component-wise multiplication.

$$C = \frac{C^{synA} * C^{synB}}{\sum_{k=1}^N C_k^{synA} * C_k^{synB}}$$

The action generation from layer C, is the same as in the single layer case (Steps 5-8). The synaptic weight update is based on the eligibility of layers. The parameter λ denotes the learning rate for the synaptic connections.

$$W_{new}^{synA} = W_{new}^{synA} + \lambda r C_e A_e^T$$

$$W_{new}^{synB} = W_{new}^{synB} + \lambda r C_e B_e^T$$

Chapters 5 and 6 will present the learning to grasp models. The architecture described here will be used as the mechanism of learning and grasp plan generation. We avoid non-biological computations in the architecture presented here so that we can make realistic prediction with models that use this architecture (Chapter 6). The weight adaptation that is offered in this chapter is based on reinforcement signal and thus biologically plausible. The parameters are represented as local population vectors and hence are open to neurophysiological predictions.

The layer structure we proposed, in essence, computes a histogram of the input patterns that yield rewards. The action that brings reward causes a weight strengthening between the eligibility trace of the unit coding the input stimuli and the eligibility trace of the neuron coding the action. The output trace left is broader than the input (Step 8). This way the weight strengthening not only increases the chance of the neuron that coded the rewarding output to become active next time but also it increases the chance of neighboring output units becoming active as well.

4.4 Summary

We introduced a neural network architecture composed of stochastic units with the following working principle. Each stochastic unit receives a *net input* that is the value after the input vector is passed through a linear transformation (weight multiplication) followed by a threshold operation. Then the units belonging to the same layer are normalized so that their activity sums up to 1. The normalized value of each unit is called the units *firing potential*. Then each unit fires stochastically according to its firing potential. However, the firing *timing* is variable and the unit which fires first (in a given a cycle period) suppress its peers (becomes the winner). The *value* output from the layer is computed as a local population vector summation. The memory traces of which unit fired is kept for the next cycle when the reward signal would be available for updating the weights. If the output generated returned positive reward then the connections that contributed to the firing are enhanced; if the output was a negative reinforcement then the weight that contributed to the output generation are reduced.

This simple reinforcement learning based layer is able to represent the probability distribution of reward conditioned on the input. A very intuitive way to look at the network is to notice that each unit is like a *counter*, counting how many times it is involved in positive-reward situations. The negative weights, the threshold operation and variable reward (i.e. non binary) makes the mathematical analysis nontrivial. Nevertheless, we presented some proofs with certain assumptions and showed the architecture works as we described.

The network is based on biologically realistic computations. Reinforcement learning, population coding, lateral inhibition/shunting (for FTA) and stochastic spike generation (based on the firing potential) are biologically feasible computations. This will enable us to use the architecture in Grasp Learning Model (LGM) of Chapter 6.

In next chapters we will combine layers, each of which is a network of the type we presented here to develop more complex networks. In Chapter 5 we use the multilayer network for infant grasp learning (ILGM). Chapter 6 instantiates the Chapter 5's multilayer grasp learning model as a network of brain regions with respect to monkey neurophysiology and neuroanatomy, which enables us to perform neural level analyzes yielding to neurophysiological predictions.

5 CHAPTER V: INFANT GRASP LEARNING

This chapter presents a computational model of infant grasp learning constrained by infant motor development studies. Key elements of the infancy period, namely elemental motor schemas, the explorative nature of infant motor interaction and inherent motor variability are captured in the model to produce testable predictions and explain how an existing behavior (reaching) yields a more complex behavior (grasping) through exploratory learning.

5.1 Motivation

Many research fields focus on discovering the principles of limb movement control. Neural network researchers address problems of internal representation, learning and execution; Roboticians solve the problems of kinematics and dynamics (Arbib and Hoff 1994) and Computational neuroscientists try to understand principles how the biological systems manage motor control. From a computational point of view, once the fingers to use and/or the wrist position, and the targets of the fingers on the object, are specified kinematics of grasping can be formulated using inverse kinematics methods (Sciavicco and Siciliano 2000). Inverse kinematics methods are based on reducing various kinematics error functions (readers are referred to Arbib and Hoff (1994) for a review of models of reach to grasp).

Although it has been suggested that human grasping involves separate control of digits so as to bring the fingers to their targets as in inverse kinematics methods (Smeets and Brenner 2001; Smeets and Brenner 1999), the general view is that hand transport and finger movements are controlled separately (see Jeannerod et al. 1998). We used the former approach when we studied the Mirror Neuron System Model (MNS) for generating grasp actions to provide visual input stimuli for the MNS model (Chapter 3). However, we will take a very different approach for infant grasp learning.

From the viewpoint of an infant, what does 'error' mean? How is the finger-to-target assignment made? These are the questions that motivate our modeling study of infant grasp learning. This Chapter is built on the notion that infants initially neither have a concept of minimizing 'error' nor know how to match their hands to objects for grasping; however, they do sense the effects of their motor acts which enables them to adjust their behaviors.

5.2 Infant reach and grasp

A child learns its own possibilities of action in the environment (affordances) through exploratory behavior (Jeannerod et al. 1998). By 2-3 months, infants start exploring their bodies as they move in the environment, they babble and touch their own body. They are actively involved in investigating the intermodal redundancies, temporal and spatial relations of their self-perception (Smeets and Brenner 2001). As new skills are acquired, new action domains are opened for the infant. Infants progress from a crude ability of reaching at birth to finer reaching and further grasping ability around four months of age. Infants learn to overcome problems associated with reaching and grasping by interactive searching (Berthier et al. 1996; von Hofsten 1993). The precision grasp appears around 12-18 months of age (Berthier et al. 1999).

To grasp successfully infants have to learn how to control their arms and, further, to match the abilities of their limbs with affordances presented by the environment (Bernstein 1967; Gibson 1969; Gibson 1988; Thelen 2000). Our study focuses on the latter problem; that is discovering how to generate grasps that match the affordances presented by the objects in the environment. Rochat and Morgan (1995) have shown that infants are aware of visual, proprioceptive and haptic consequences of their limb movements.

The onset of reaching and grasping marks a significant achievement in infants' functional interactions with their surroundings. The advent of voluntary grasping of objects is preceded by several weeks in which the infant engages in arm movements and fist swipes in the presence of visible objects (von Hofsten 1984). The term *visually guided reaching* generally refers to the infants having available continuous vision of the hand and target, whereas *visually elicited* reaching refers to the vision of the target, followed by a ballistic hand movement. Clifton et al. (1993) questioned the hypothesis that the earliest accurate reaching behavior is visually guided and appears around 3-5 months (Clifton et al. 1993). They tested seven infants repeatedly between 6 and 25 weeks of age using glowing objects in the light condition and sounding objects for the dark condition. Infants first contacted the object in both dark and light conditions at almost the same ages (mean ages: 12.3 weeks for light, 11.9 weeks for dark condition). Infants first grasped the object in light condition at 16.0 weeks and in the dark at 14.7 weeks. Even though this was not a statistically significant difference, there is other evidence that the self-vision of the hand may retard successful grasp in infants younger than 5 months of age. This could be due to the attentional load that is brought by the existence of two objects in the visual field (Streri 1993). Clifton et al. (1993) argued that since infants could not see their hand or arm in the dark, their early success in contacting the glowing and sounding objects indicated that proprioceptive cues, not sight of the limb, guided their early reaching. This shows that early reaching (and grasping –evidenced in the following text) is performed using an open-loop strategy, because it appears that, the initial localization (glow or the sound of the object) was enough to perform successful movement planning. Reaching in the light developed in parallel with reaching in the dark, suggesting that visual guidance of the hand is not necessary to achieve object contact. It is also noteworthy that infants showed great individual differences. Onset for touch varied between 7 and 16 weeks, while onset for grasp varied between 11 and 19 weeks. The greatest discrepancy (light versus dark conditions) in onset of reach and grasp was 4 weeks. There were three infants out of seven with this discrepancy. Interestingly for all the three infants the behavior occurred earlier in the dark.

An infant, once contacted an object, will occasionally try to grasp it (Clifton et al. 1993). The enclosure reflex disappears around six months of age and it takes four more weeks of infancy to stabilize the grasp (Clifton et al. 1993). It is suggested that the fractionated control of finger movements is not possible since this task requires the cortico-motoneuronal system, which has not fully developed at this age. (Lantz et al. 1996). Before nine months of age, the infant grasp lacks the anticipation of the orientation and size of the object (Rosenbaum 1991). Infants adjust their grasps after touching the objects. In contrast, adults adjust their distance between the thumb and the other fingers according to the size of the object during the hand transport. This holds even though infants younger than nine months old are physically able to vary their grip size, for they can spread their fingers farther apart once they have felt a large object (von Hofsten and Ronnqvist 1988). It appears that in early infancy the fractionated control of fingers is mainly driven by somatosensory feedback.

Butterworth et al. (1997) studied the development of prehension in infants by video recording the grasping behavior of babies from 6 to 20 months of age using objects of different shapes and different sizes. The infants were divided into four groups according to their age. The objects used were wooden

cubes and spheres with sizes of 0.5cm, 1cm, 2cm and 3cm. Butterworth et al. (1997) classified the grips used in the grasps according to two broad categories of power and precision grips. The power grip had subdivisions of ulnar grasp, hand grasp, palm grasp and radial palm grasp while the precision grip had subdivisions of scissor grasp, inferior forefinger grasp, inferior pincer grasp and pincer grasp (Figure 5.1). The main result of this study was that both power and precision grips were observed in infants older than 6 months of age and there is a developmental trend such that in the early second year the occurrence of power grips decreased while the occurrence of precision grips increased, and eventually became the dominant mode of prehension. Furthermore, the size rather than shape (cube vs. sphere) determined the grasp type (Butterworth et al. 1997). The younger infants did not employ a consistent pattern of grasping for small objects whereas the older infants developed preferential grasp patterns for certain sizes (0.5 cm, pincer grasp; 2 cm, forefinger grasp) (Butterworth et al. 1997).

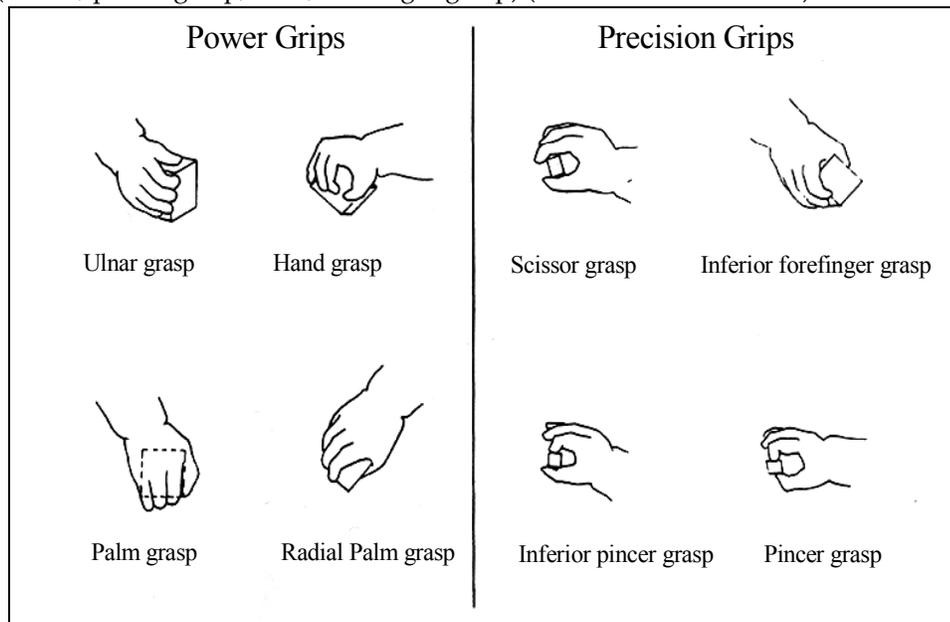


Figure 5.1 Infant grip configurations can be divided in two categories: power and precision grips. Infants tend to switch from power grips to precision grips as they grow (adapted from Butterworth et al. 1997)

Lockman et al. (1984) investigated the development of infants' prehensile adjustments regarding the orientation of objects. They compared the performance of groups of 5 and 9 months old infants (8 infants in each group). The infants were presented with a dowel with two orientations, vertical and horizontal. In the trials, if the infant did not initiate a grasp when the dowel is revealed, the experimenter attracted the attention of the infant to encourage him/her to grasp the dowel. If the infant did not grasp but touched the dowel, the trial was terminated and excluded from the analysis. The hand orientations were analyzed by recording the orientations at four points in time: 1) at the beginning (with the first forward movement of the hand); 2) When the hand passed the midpoint between the first point and the object position; 3) at the time of first contact of the hand with the dowel; 4) and at the end where at least one digit closed around the dowel. The principal finding was that 9 months old infants oriented their hands appropriately earlier during the reach to grasp than did 5 months old infants. Furthermore, the two age groups differed at the last two stages of the grasp (touch and enclosure) whereas the earlier parts of their reaches were similar.

In another related study, 102 infants between 4-8 months old were grouped according to their age and was studied their grasping behavior as a function of object shape and size using seven combinations of objects (Newell et al. 1989). The main finding of this study is that the youngest age group (4 months) required the addition of the haptic system for successful grasp whereas the oldest age group (8 months) mainly relied on information from the visual system to differentiate grip configuration according to the object properties. They found that there was a remarkable similarity between the grasp configurations achieved, irrespective of whether they are visually planned or haptically adjusted.

Butterworth et al. (1997) show that the young infants display a wide variety of grip types that were conventionally not attributed to 6-8 months old infants (Halverson 1931), which is in agreement with our simulations. Lockman et al. (1984) used a single object with different orientations in such a way that the infant's inability to plan according to the target can be detected. An inappropriate hand configuration would cause a big discrepancy between the hand orientation and the object orientation.

Finally, Newell et al. (1989) inform us that the older infants' visually programmed and younger infants' haptically adjusted grasp configurations are very similar. This strongly suggests that the earlier haptic grasping phase serves as the training stimuli for visual grasp planning circuit in infant brain. We use the reviewed studies to constrain and design the Learning to Grasp Model, and evaluate its relevance to infant learning through explicit comparisons.

5.3 Neural maturation versus interactive learning

The corticospinal tract is one of the main neural substrate for independent finger control (Triggs et al. 1998): Firstly, corticospinal projections terminating in the ventral horn innervating hand muscles predict independent finger use for small objects (Bortoff and Strick 1993). Secondly, lesioning the corticospinal tract prevents the development of the independent finger movement in infant animals (Lawrence and Hopkins 1976), and impairs independent finger movements. Human patients recovering from corticospinal lesions initially tend to perform synergistic movement of all the fingers as in a power grip (Denny-Brown 1950; Lassek 1954).

Given these data, it is tempting to theorize that infants' late grasping development is due to neural maturation. This view is embraced by maturational-based theories (Bradley 2000). However correlation studies are not conclusive for it is not known whether other variables account for the observed behavior (Bradley 2000). In contrast, learning-based theories consider the environment and infants' interaction with it as the main factor in shaping infant behavior (Bradley 2000). The model, we propose interacts with its environment (plans and executes grasp actions) and observes the consequences of its actions (grasp feedback) and modify its internal parameters (corresponding to neural connections) such that certain patterns (grasp plans) are selected and refined amongst many other possibilities. Thus, in this sense, our model conforms the learning-based views of motor development.

Infants lack the ability to fully utilize their vision for grasping during the period of early grasping (Newell et al. 1989; Rosenbaum 1991). Nevertheless, infants, as reviewed in the previous section, can grasp objects and even adjusts their grasp actions according to object's visual properties. Thus, we propose that infant grasp learning is mediated by neural circuits specialized for grasp planning, which can function with limited visual analysis of the object, probably position and rough size. The macaque monkey was shown to have a specialized circuit for grasping in which parietal areas extract object affordances (information relevant for grasping) (Taira et al. 1990) and relay those to premotor cortex. Then the premotor cortex with various contextual and intention related bias signals performs grasp

selection/execution task (see Fagg and Arbib 1998 for a modeling study). It is very likely that a similar circuit exists in human (Jeannerod et al. 1995) which is adapted in infancy to provide skillful grasping in adulthood.

We present Infant Learning to Grasp Model (ILGM) in two stages. The first stage is the period when infants are unable to incorporate object affordance into grasp plans while the second phase is when infants start incorporating object information into grasps. Although we favored learning-based theories of motor development, ILGM is compatible with both maturational- and learning-based motor development schools since in either case, with different reasons though, the affordance information cannot be used by the grasp learning circuit in early infancy but it becomes available as development progresses.

In the rest of the chapter, we analyze ILGM via simulation experiments, and present behavioral responses and make comparisons where experimental data is available. When no data is available, we produce useful predictions that can be experimentally tested.

5.4 Infant Learning to Grasp Model (ILGM)

With the term *learning to grasp* we mean to learn how to make motor plans in response to sensory stimuli such that the *open loop* execution of a plan leads to a successful grasp. There is strong behavioral evidence that early grasping is based on open-loop control and does not use visual feedback (Clifton et al. 1993; Clifton et al. 1994; von Hofsten and Ronnqvist 1988; Streri 1993). Chapter 7, studies the visual feedback control of grasping in a simplified setting

We first describe a generic schema level architecture of Learning to Grasp Model. Later sections use the general architecture introduced here to conduct simulation experiments. We propose three computational layers relevant for grasping: Hand Position, Virtual Finger, Wrist Rotation layers. In general, the input to our network is the affordance of the target, which can vary from a single quantity indicating the existence of a graspable object to the full description of the object in terms of its affordances (e.g. size, orientation, etc.). In later sections when running simulations, we engage layers according to experimental requirements. For example to simulate early infant grasp learning, we effectively disable Affordance layer and analyze learning with tactile feedback. The affordances are represented using population-coding scheme and encoded algorithmically (i.e. no visual processing for object recognition and feature extraction is done). We use *Affordance layer* or *Input* terms interchangeably to conform the context. The layers we introduced here encode motor parameters that constitute a minimal set for specifying grasp actions and based on behavioral studies and monkey neurophysiology. We make use of monkey studies by postulating that monkey and human motor development follows similar patterns. In fact the kinematics of reach to grasp movements of the macaque monkey and human is very similar (Roy et al. 2000) and homologous brain structures are involved in motor tasks (Jeannerod et al. 1995).

Adult grasping studies suggests that reach and grasp components of a grasping action is independently programmed (Jeannerod and Decety 1990). This implies that the position of the fingers (hand configuration) on the object determined first, and then based on the hand configuration wrist orientations and arm configuration are determined. However in the infant, reaching component dominates: infants first learn how to reach and then learn to adjust their hands to match according to their ballistic reach (Clifton et al. 1993; Clifton et al. 1994; von Hofsten and Ronnqvist 1988; Streri 1993). Therefore, we suggest that a grasp planning inversion takes place after infants become skillful. The inversion is required because intended manipulations after a grasping action determines how humans grasp the objects. If one were planning to grab a tennis ball with the aim of throwing it far, probably a

power grasp would be preferred. On the other hand, if he/she were going to pass it to a person next to him, a (tripod) precision grasp would be more appropriate.

5.4.1 Layers of infant learning to grasp model

In accordance to infant development, we propose that infants explore the space around the objects (and occasionally touch them) they interact. Thus, infants' early reaches can be considered *object centered*. Early reaching is variable (see Bradley 2000 for a review); the variability must be reduced for successful grasping. We posit the Hand Position layer as specifying the approach direction; that is the object centered (allocentric) position of the hand from where it will approach the target.

Given an approach direction, the orientation of the wrist has to be determined. For example, when approaching a sphere from the bottom side, a large wrist supination is required while approaching from the front a wrist flexion would be required. We posit the Wrist Rotation layer to learn the possible wrist orientations given the allocentric hand position information relayed from Hand Position Layer. Wrist Rotation layer also receives projections from the Affordance layer, because in general, different objects afford different set of approach-direction and wrist-rotation pairs.

The Virtual Fingers layer indicates which finger synergies will be activated given an input. This layer's functionality is fully utilized in adult grasping. When infants start learning to grasp, they first engage all of their fingers in a synergistic way. When infants start to control their digits independently to match object shape, this layer can be engaged in learning the possible virtual finger activations. In the simulations that will be presented in this chapter, the Virtual Fingers Layer is always used to specify a synergistic control of the fingers but is included in the ILGM description for completeness. As will be shown, this does not refrain us from reproducing infant behavior and generating testable predictions.

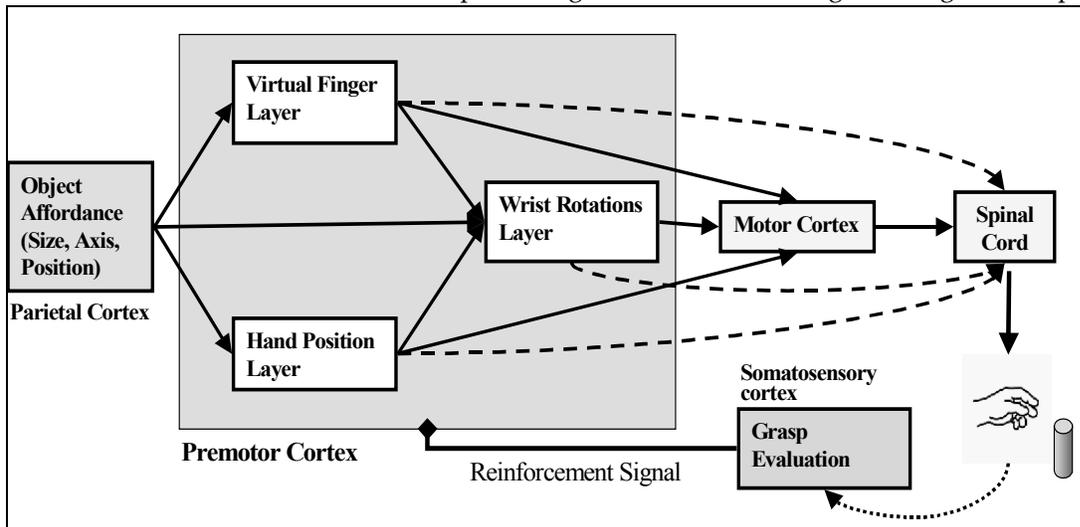


Figure 5.2 The structure of the Infant Learning to Grasp Model. The individual layers are trained based on somatosensory feedback

The layers we proposed are in one-to-one correspondence with Iberall and Arbib's (1990) schemas for grasping. The schemas of *Preshape*, *Approach Vector* and *Orient* correspond to the ILGM layers of *Virtual Fingers*, *Hand Position* and *Wrist Rotation* (Iberall and Arbib 1990).

The layer architecture we introduced in Chapter 3 is instantiated (as a population of neurons) for the parameters, allocentric hand position, virtual fingers and wrist rotations that determine a grasp plan. The key features of the architecture introduced in Chapter 4 are that it is capable of representing multiple choices of actions and it is open to biasing. The model, after learning, will retain a menu of

grasp actions that can be retrieved based on the affordance. The menu then can be biased to satisfy task constraints. Figure 2.1 shows schematically the structure of the ILGM. The task division of brain regions for implementing the schemas left to the next chapter, in which we probe the possible brain localization of grasp learning based on the monkey neurophysiology. For now, we only give an overall view of grasping circuit. Parietal areas extract object affordances and relay them to the premotor cortex, which is the center implicated in grasp programming. In turn, premotor cortex makes a grasp plan and instructs the spinal cord and motor cortex for execution. The result of the plan is integrated in the first somatosensory cortex (SI). Output of somatosensory cortex mediates the adaptation of the grasping circuit.

5.4.2 Functional description of ILGM layers

Hand Position layer determines where the hand will be with respect to the object during grasping. Intuitively, this represents the side from which the hand will grasp the object. The frame of reference for this parameter is allocentric. The coordinate system we choose for this parameter is spherical (see Figure 5.3). Our preliminary simulation showed that this choice is advantageous over a rectangular coordinate system because the spherical representation was less sensitive to errors in radius component while the rectangular coordinate system had the same sensitivity for all the components. The encoding used is as described in Chapter 4: the neural layer represents a probability distribution and the values are read off by local population vector computation. Given an object, the layer represents the feasible positions for the hand. Any item from this position ‘menu’ can be selected by external circuits. However, we simulate the grasping circuit autonomously, by processing the input and generating a grasp program according to the probability distributions represented in ILGM layers (but see Fagg and Arbib 1998 for prefrontal biases for grasp selection).

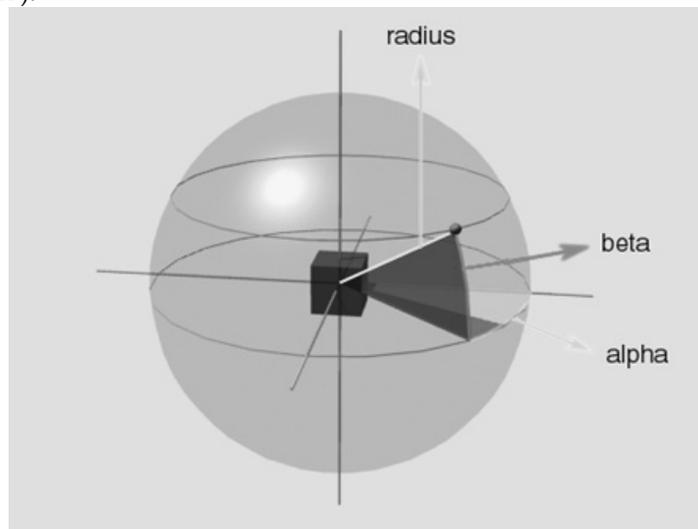


Figure 5.3 Hand Position layer specifies the approach direction of the hand towards the object. The representation is allocentric (centred on the object). Geometrically the space around the object can be uniquely specified with the vector (azimuth, elevation, radius). The Hand Position layer generates the vector by a local population vector computation. The locus of the local neighbourhood is determined by the probability distribution represented in the firing potential of Hand Position layer neurons (see Chapter 4, for details)

Virtual Fingers layer specifies the fingers that will be activated, and with what strength. We use three virtual fingers: thumb, index finger, and the remaining three fingers acting together. The

processing is parallel to the Hand Position Layer's flow. However, this does not mean that this layer has the right to decide the virtual fingers on its own: the virtual fingers that cannot yield grasping are negatively reinforced and hence do not appear in the learned ILGM.

Wrist Rotation layer funnels all the information about the object (affordance) and the Hand Position and Virtual Finger layers' output. Thus, the output of this layer represents the possible wrist orientations given the (1) object related input, (2) output of Hand Position layer (3) output of Virtual Fingers layer in terms of a conditional probability distribution. The parameters generated in this layer determine the movements of wrist extension-flexion (tilt), wrist supination-pronation (bank) and ulnar and radial deviation (heading).

5.5 Joy of grasping

Infants are almost preoccupied with manual manipulation. Infants would play with their own hands; manipulate objects given to their hands, and play with rattles before they can reach to grasp for them as young as 2 months of age (Bayley 1936). Infants, once contacted the object, occasionally would try to grasp it (Clifton et al. 1993). We suggest that the tactile stimuli induced neural signals motivate infants to engage in grasping and holding. However, we do not model the mechanisms of 'joy of grasping' induced by tactile sensation but instead use a physical substitute to emulate the feedback that infant would receive when grasping an object

Through 'joy of grasping' infants explore and learn actions that lead to grasp-like experiences. What we call 'joy of grasping' can be considered as Sporns and Edelman's (1993) *adaptive value* of an action. Sporns and Edelman's (1993) postulate three concurrent steps for sensorimotor learning: (1) The spontaneous movement generation (2) development of the ability to sense the effects of movements, eventually allowing neural selection to be guided by adaptive value (3) actual selection of movements based on the adaptive value. Furthermore, it is argued that selection in the nervous system is mediated mainly via synaptic change (Sporns and Edelman 1993; Sporns et al. 2000) supporting our model's relevance to infant learning.

The model we develop in the chapter will show how sensory feedback shapes infant reaches into grasp actions via explorative learning and produce testable predictions (see Sporns and Edelman 1993 for a simple reach learning architecture based on similar principles).

5.5.1 Mechanical grasp stability

A successful grasp requires that the object stays stable in the hand (must not drop or move) (MacKenzie and Iberall 1994), which is physically defined as the following (Fearing 1986)

- (1) The net force (F_{net}) acting on the object must be zero

$$F_{net} = \sum_1^M N_i (F_i \bullet N_i) \text{ where } F_i \text{ denotes the force applied by the fingers at contact } i \text{ and } N_i$$

denotes the normal of the surface that is involved in the contact. M denotes number of contact points of the hand on the object.

- (2) The net torque (T_{net}) acting on the object must be zero

$$T_{net} = \sum_1^M (P_i - P_c) \times N_i (F_i \bullet N_i) \text{ where } P_i \text{ is the contact position and } P_c \text{ is the center of the mass of}$$

the object.

- (3) For any force acting on the object, the angle between the directions of the force with the surface normal must be less than a certain angle ϕ . This angle is defined through the finger and contact surface dependent coefficient μ with the relation:

$$\phi = \tan^{-1} \mu$$

The constant μ satisfies the property that if F_n and F_t are the normal and tangential components of an applied force then there will be no slip if

$$\mu F_n > |F_t|$$

(4) The magnitude of the grasping force should be adaptable to prevent any displacement due to an external force.

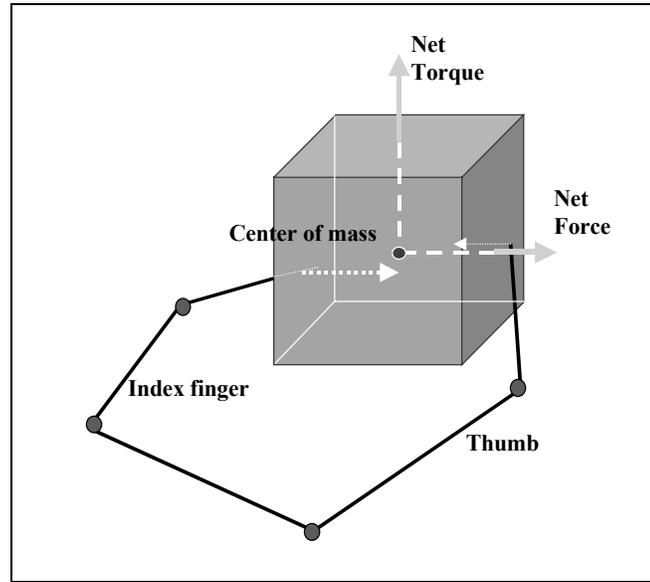


Figure 5.4: The grasp stability we used in the simulations is illustrated for a hypothetical precision pinch grip (note that this is a simplified, the actual hand used in the simulations has five fingers)

5.5.2 Implementing the grasp stability

The parameters generated by Virtual Fingers layer determine how much each virtual finger is activated. By converting these values to forces exerted on the object during contact, ILGM, in theory, can discover the force values that will stabilize the object. However, in our preliminary simulation studies, we have seen that learning to grasp required excessive grasp trials to discover the force balance¹⁴. As mentioned before we do not intend to model the details of the tactile feedback system our aim is to compute a value that captures the joy (behavioral reward) of the infant during grasping or a neural signal (neural level reward) that indicates a stable grasp. Therefore, we concentrated on the question of whether, given an object, the simulated hand's contact configuration could afford a stable grasp. We formulated the problem as a constrained minimization problem with the cost function E .

$$E(F_i) = \alpha |F_{net}| + (1 - \alpha) |T_{net}| + \beta \sum_{i=1}^M \cos^{-1} \left(\frac{F_i \cdot N_i}{|F_i| |N_i|} \right)$$

$$|F_i| > f_{\min}, \forall i = 1..M$$

The values α and β determine relative contributions of the individual costs terms to the total cost function, E . The first two terms capture the grasp stability conditions (1) and (2); the last term captures

¹⁴ In fact, human hand is endowed with compensatory spinal feedback loops to stabilize grasping in addition to the feed-forward mechanisms to specify finger forces (see Rothwell 1994; Johansson and Westling 1987a, 1987b).

the grasp stability condition (3)¹⁵. The value f_{min} is an arbitrary positive constant to avoid the degenerate solution ($F_i=0$).

5.6 Learning approach direction with palm orienting behavior

To test the hypothesis that goal directed reaching could be the basis of infants' early grasp learning we mimicked the infants' elemental hand behaviors: we implemented palm orienting behavior and enclosure reflex in the 3D hand/arm which we developed for Reach and Grasp schema of Chapter 3.

During the neonatal period when an infant reaches towards an object, his/her hand is usually open (von Hofsten 1982). Although this behavior disappears before the second month of age; it resumes around the third month of postnatal age (see Bradley 2000). It has been suggested that enclosure reflex constitutes the first stage of grasping (Twitchell 1970; see Streri 1993 for a review). In our model we wanted to test whether reflex based grasps can be factored into grasp plans so that instead of a randomly directed reach (before learning stage), the infant can use the reinforced directions that yield appropriate contact with the object so that enclosure of the hand yields a stable grasp.

The simulation is set up as the following. The model generates a grasp plan. In this restricted case, the grasp is determined by the (allocentric) hand position parameter which is a triplet describing a point with respect to the target. The value is read off from the population activity of the neurons in the Hand Position layer (see Chapter 4 for details of layer operations). Initially, the value is random since we initialize the weights to small random values. The Hand Position layer captures the infant's variable reaches towards the visual targets. The task of learning is then to narrow down the variability to account for only appropriate approach directions (for a given object at a certain location).

After a grasp plan (approach direction) is generated, the hand starts executing the reach action specified by the Hand Position parameter. During the reach, the wrist is rotated so that the palm normal always points to the target object. From literature, we know that infants (except a one month period) open their hand while they are reaching for objects (von Hofsten 1982) and orient their palms towards the them (see Streri 1993 pages 46-47).

As the model generates grasp plans and executes them, it receives rewards from those plans that yield grasp or close-to-grasp actions. In the implementation, the model receives negative reward for plans not yielding stable grasps. However, when the grasp error, E is small (e.g. the grasp is close to stable) the model receives a positive reward that is proportional to e^{-E} . Thus, an almost-grasp program is encouraged which enables the model to have a higher chance of producing similar plans. This makes our approach goal directed and different from pure trial and error learning.

5.6.1 Simulation results

Figure 5.5, right panel, illustrates a discovered grasp (i.e. approach direction) by the model. The model, by interacting with the object, learned that it is possible to grasp the object from the bottom side. The grasp plan our model generates is defined by a triplet (α, β, r) (α stands for azimuth; β for elevation and r for radius) as shown in Figure 5.3. If we sum over the radius we can represent the grasp plan as a pair and plot it as shown in Figure 5.5, left panel. We can think of this plot as the normalized histogram of (α, β) pairs generated by the model over many trials. The plot tells us that the model prefers grasp directions from $\alpha > 0^\circ$ (below the object) and $\beta < 90^\circ$ (right and/or front of the object). The position of the object allowed the arm to grasp the object from bottom and front/right. For example, a trial with an approach to the object from the top does not yield a stable grasp (the length of the arm limits the

position of the hand. Those positions that could not be reached do not yield a stable grasp when the hand contacts the object).

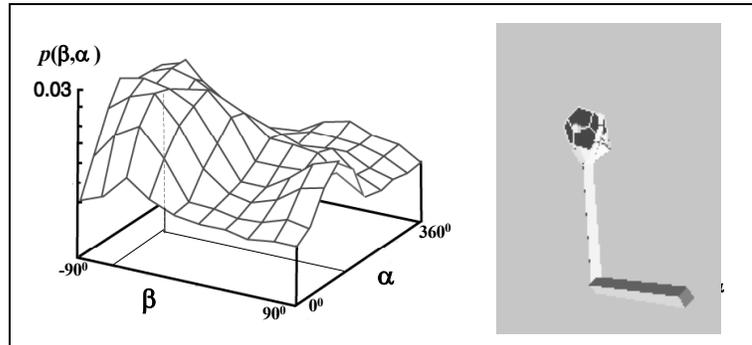


Figure 5.5 The trained model's Hand Position layer is shown as a 3D plot. One dimension is summed to reduce the 4D map to a 3D map. Intuitively the map says: 'when the object is above the shoulder and in front grasp it from the bottom'

We also tested the model when the object is located at a different place in the workspace.

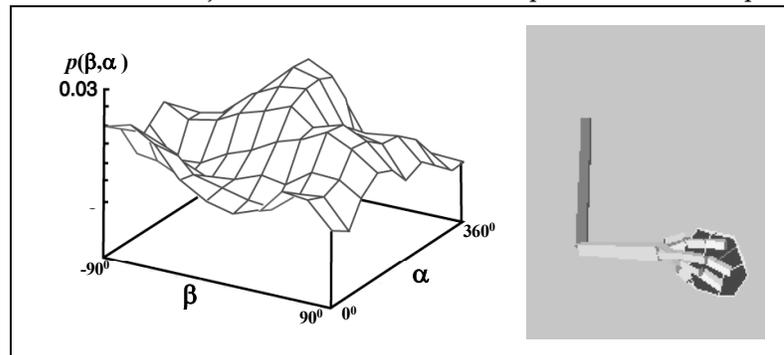


Figure 5.6: The output of the trained model's target position layer is shown as a 3D plot. One dimension is summed to reduce the 4D map to a 3D map. The object is on the left side of the (right handed) arm. Intuitively, the map says 'when the object is on the left side grasp it from the right side of the object'

This time the object is placed rather low in the workspace on the left side of the (right-handed) arm. Here we observe that the model discovered that a grasp directed to the right hand side of the object is likely to result in a stable grasp. Figure 5.6 shows the results of this simulation. Figure 5.7 shows the distribution of the Hand Position layer during training. Initially the probability distribution of approach directions was set randomly. The distribution gets a regular pattern via learning and the reward yielding regions gain higher levels of activity.

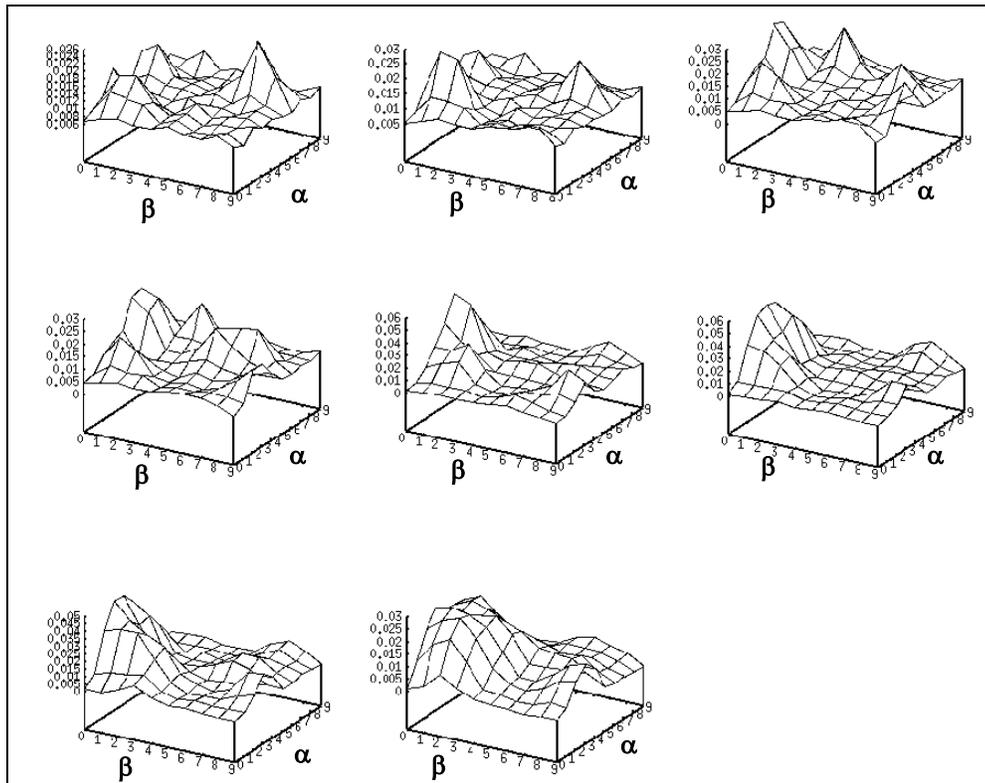


Figure 5.7 The learning evolution of the distribution of the Hand Position layer is shown as a 3D plot. Note that the 1000 neurons shown represent the probability distribution of approach directions. Initially, the layer is not trained and responds in a random fashion to the given input. As the learning progresses, the neurons gain specificity for this object location.

5.6.2 Conclusions and predictions

The simulation results showed that goal directed reaching with palm orienting behavior is enough for generating power grasps. The model predicts that infants will narrow their reach variability when presented with objects that they can interact with. Furthermore, the variability reduction will be more pronounced for the space where they could grasp objects.

5.7 Is infant palm orienting learned or innate? Learning the wrist orientation

In this section, we deprive the simulator from auto palm-orienting behavior to test whether the infant palm-orienting behavior can be mediated via learning, rather than being innate. Although there are accounts that infants usually orient their hands during reaching so to increase the likeliness of a contact between palm and the object (von Hofsten 1982), to our knowledge, there is no account of innateness of this behavior. The period we are modeling corresponds to early grasping in development when the object affordance is not available to the grasping circuit, either because of maturational shortcomings or because the complexity of learning holds back the motor development (in accordance to maturational- and learning-based theories of motor development, respectively).

In terms of learning task, ILGM has to discover the distribution of wrist movements (supination-pronation, extension-flexion, ulnar/radial deviation) for multiple approach directions. Presented with the object, Hand Position layer produces the distribution of possible approach directions. The selection is relayed to Wrist Rotation layer. Then, Wrist Rotation layer computes the distribution of feasible wrist

orientations conditioned on the generated hand position. Virtual Finger layer generates synergistic parameters effectively learning a single parameter to dictate the enclosure speed of the hand as a whole.

5.7.1 Simulation results

The model learned to generate parameters to perform power and precision grasps, and many variations of the two. The most abundant grip generated was the power grip and its variations. Precision type grasps were less frequently generated.

In earlier studies, it was thought that infants were unable to demonstrate precision grips during grasping (Halverson 1931). The newer studies showed that, when infants are tested in proper conditions, perform precision grips occasionally (Butterworth et al. 1997; Newell et al. 1989). ILGM simulation results are in accordance with this finding. The fact that power grasp is inherently easier manifests itself in both infants and ILGM simulations: As long as the object is brought in contact with the palm, an enclosure (as in palmar reflex) is likely to produce power grasp. Figure 5.8, shows the grasping of a cube with a power grasp plan generated by learned ILGM.

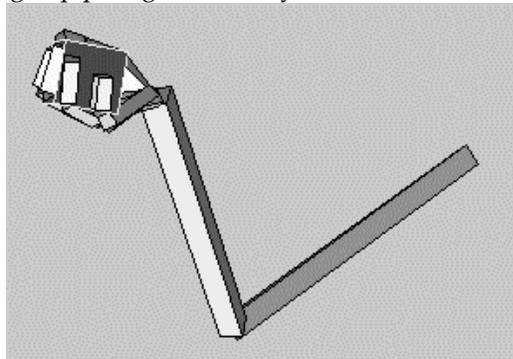


Figure 5.8 ILGM planned and performed a power grasp after learning. Note the supination (and to a lesser extent extension) of the wrist required to grasp the object from the bottom side

The learned precision grasp varieties were mainly involved the engagement of fingers other than thumb and index finger. Figure 5.9 shows two examples for this. Usually the object is secured between three or four fingers, thumb opposing the center of remaining fingers. This emergent grasping behavior is in accordance with the theory of virtual fingers and opposition spaces (Iberall and Arbib 1990) and human tripod grasping (Baud-Bovy and Soechting 2001).

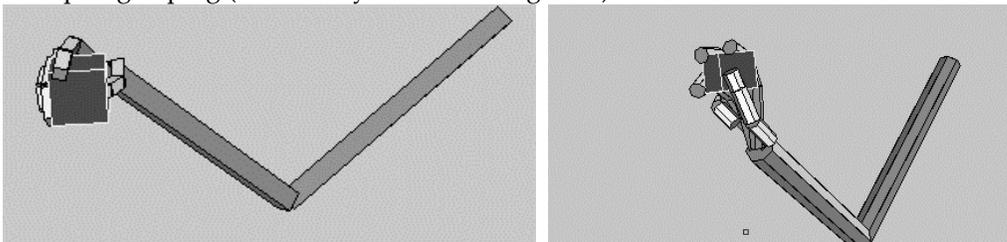


Figure 5.9 Two learned precision grips (left: three fingered; right four fingered) are shown. Note that the wrist configuration for each case. ILGM learned to combine the wrist location with the correct wrist rotations to secure the object

Figure 5.10 show examples of two finger precision grasps which were less frequently generated than the three or four fingered precision grasps. These results show that even without object affordance, precision type grips can emerge from a circuit adapted using supervised Hebbian learning, which can be employed by infants.

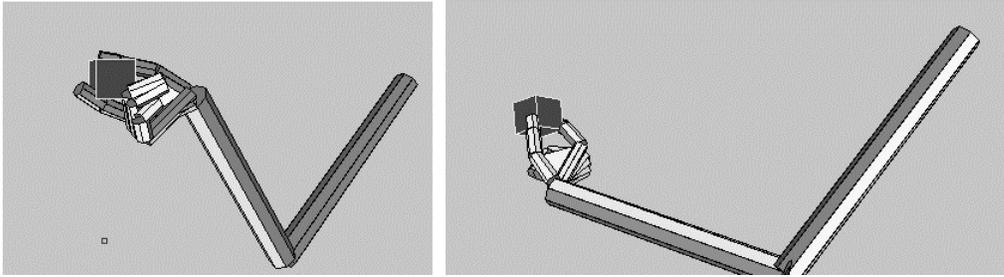


Figure 5.10 ILGM was able to generate two fingered precision grips. However these were less than the three or four finger grips

5.7.2 Conclusions and predictions

The result shows that even without object affordance, performing reaches directed towards the object from various allocentric positions and associating the grasp-yielding wrist orientations with allocentric positions, and a variety of grips, including the precision grip, can be learned predicting the results of the study of Butterworth et al. (1997). Thus, infant can learn to select the ‘right’ grasp from the grasp ‘menu’ based on internal motives or environmental constraints. There are two observations that have to be made: Firstly, the secondary learning mentioned here is easier to master. Infants only have to associate the object properties (and the context) to the correct configuration that they have already achieved. Secondly, no extrinsic to intrinsic transformation is required because a posture that yields a successful grasp has already been discovered. In computational terms, supervised learning can take place.

5.8 Task constraints shape infant grasping

The normative developmental phases of infant prehension starts with palmar reflex, followed by power grips; and finally, ends with dexterous finger-thumb opposition precision grips (Halverson 1931; see Newell et al. 1989 for a review). However, this early view has been challenged by advanced recording techniques and careful experimental conditions (Newell et al. 1989; Butterworth et al. 1997). If a variable (e.g. infant grip configuration) is dynamically context specific, the experimental approach may be too artificial to reveal its effect on motor control (Bradley 2000). For example, the postural requirements could have been a factor in masking infant grasp abilities in the earlier infant motor development studies (Newell et al. 1989; Bradley 2002). Task constraints may be viewed as including the goal of the task or the rules that constrain the response dynamics (Newell 1986). Some examples of task constraints are the object properties such as size and shape (Newell et al. 1989). To verify the hypothesis that infant grasping is shaped by task constraints we designed a simulation experiment with a physical constraint as the following.

In the earlier simulations of ILGM we presented the object without constraints, hanging in space, thus the model could grasp the object from all anatomically possible directions. However, this is not usually the case for the infant. For example, infants often interact with objects that lay on flat surfaces such as the floor or table. A small object on a horizontal table is grasped best by opposing index finger and thumb with an approach from top.

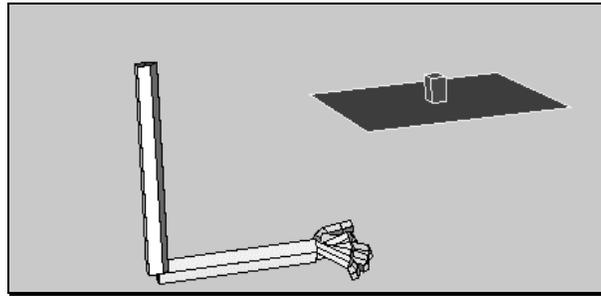


Figure 5.11 The *cube on the table* simulation set up. ILGM interacts with the object with the physical constraint that it has to avoid collision with the table

We simulated the situation by presenting ILGM with a small cube placed on a horizontal plane (Figure 5.11). The plane constituted a physical obstacle for many grasp attempts. Thus, when the simulated hand/arm collided with the table, a negative reward was returned to ILGM. We let ILGM interact with the *cube on the table* and analyzed the acquired grasping behavior.

5.8.1 Simulation results

ILGM with the *cube on the table* task condition was unable to acquire power grasps (whole hand prehension). The grasp attempts of the model that would result in a whole hand prehension resulted in negative rewards as the fingers always collided with the table surface. Thus, the grasp plan parameters yielding power grasps were not represented in the grasp repertoire after learning.

The approach directions for avoiding the collision were learned perfectly. LGM always attempted to grasp the cube from top. Figure 5.12 shows typical precision grips executed by ILGM after learning. The grasp ‘menu’ acquired was composed of grasps with wrist positions above the object. The contact points on the cube showed variability (see Figure 5.12).

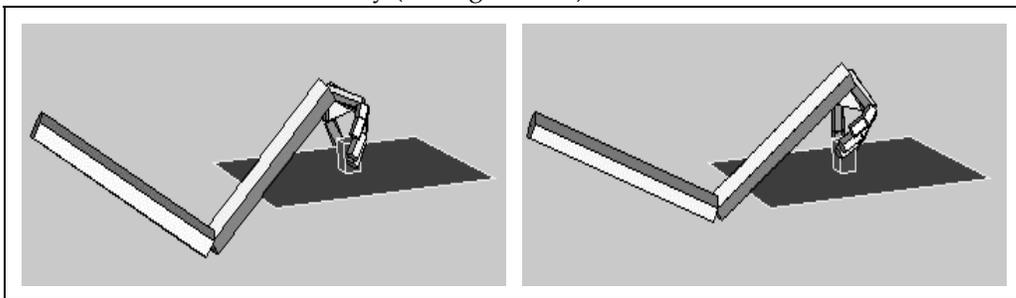


Figure 5.12 ILGM learned a ‘menu’ of precision grips with the common property that the wrist was placed well above the object. The orientation of the hand and the contact points on the object showed some variability. Two example precision grips are shown in the figure

Many of the precision grips that were learned involved supportive fingers other than the thumb and the index finger. However, two finger precision grips were also acquired (see Figure 5.13).

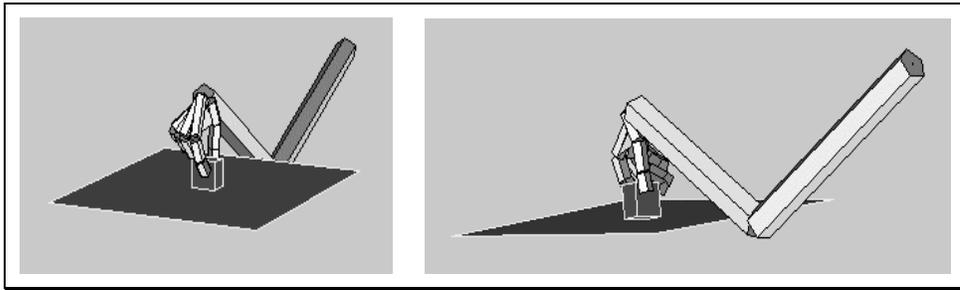


Figure 5.13. ILGM acquired thumb opposing index finger precision grips

Most of the precision grips learned correspond to inferior forefinger grasp (Figure 5.12) and inferior pincer grasp (Figure 5.13), according to classification of Butterworth et al. (1997) (see Figure 5.1). One of the interesting observations is that ILGM assimilated the object affordances into the grasp 'menu' it learned. By comparing Figure 5.12 and Figure 5.13, we see that the opposition axes used for grasping were 90 degrees apart. In both grasps of Figure 5.12, the thumb was placed on the left surface of the cube whereas in Figure 5.13, it was placed on the surface that is parallel and closer to the presumed infant.

5.8.2 Conclusions and predictions

Clearly, it is not possible to grasp a small object with the whole-hand grip without further manipulating (e.g. dropping or raking) the object or without deforming the hand to a pincher grasp (readers are encouraged to grasp a pellet from a table without using precision grip). Infants, during early grasping phase (4-6 months of age), certainly contact the surface before they can grasp the objects as they use tactile senses for grasping (Newell et al. 1989). However, we predict that infants' ability to grasp small objects from a hard, flat surface using whole-hand precision would be a rare, if not impossible, occasion. On the other hand it has been shown that infants are able to use various precision grips (Newell et al. 1989; Butterworth et al. 1997; Corbetta et al. 2000) which is in full accordance with our findings. It appears that older infants are more interested in small objects compared to younger infants. ILGM explains why, as infants grow older, prefer to approach objects from top (Fagard 2000) with the answer that from-top approach is the most natural way to grasp small objects in a constrained condition like the one we presented (i.e. small object on the hard flat surface).

The simulation results of *cube on the table task* (combined with earlier simulation results) have nontrivial consequences. To be precise, ILGM simulation results:

- Predict that infant even without object affordance input, that is during the age of early grasping, can perform precision grasp
- Show that task constrains shape motor development and support the view that development of precision grips is mediated by task constraints.
- Show that object affordances could be represented in infants' grasp repertoire in spite the fact that they are unable to access/extract object affordance information. This result is very important because it is a proof that tactile learning can train/modulate the visuomotor learning in infancy. A corollary prediction is that infants learn to extract affordance of the objects they can grasp.
- Predict (in relation to the preceding item) that object (visual) affordances would be heavily represented in the motor and sensory areas of human cerebral cortex for objects that we manipulate often. For example, presented with a drill, carpenter's neural circuits would extract more elaborate affordances than a fisherman's would.

5.9 Affordance input matters

Until now, we simulated infant grasp learning assuming that vision of the object is not used to adjust grasp plans according to object features. Therefore, we can associate the earlier ILGM infancy period of two months to six months of age. In this section, we introduce affordance input to ILGM and associate the model with infancy period of nine months and after.

This section is reserved for simulating the study of Lockman et al. (1984) and compare ILGM results with their infant data. Lockman et al. (1984) used infants of 5 and 9 months of age and compared their performance in orientating their hands to a dowel presented in horizontal and vertical position (see 5.2 section for more details). To replicate their experimental condition, we let the affordance of the object (orientation) to be relayed to ILGM preferentially for 5 and 9 months of age. We presented a cylinder to ILGM analogous to a dowel. With the (realistic –see section 5.2) assumption that information about the axis orientation (i.e. affordance of the object) is not available to the grasp planning circuit during early infancy (5 months of age) and that it becomes available later in development (9 months of age), we effectively disabled the orientation encoding¹⁶ when simulating the younger infants' grasp learning whereas when simulating older infants' learning we enabled the orientation coding in the Affordance layer. We refer to the former case as the *poor-vision* and the latter as the *full-vision* condition.

We predicted that the infants without affordance input would not be able to factor the object orientation into their motor plans. To make the experiment a little bit more interesting, we also included a third, diagonal orientation condition in addition to the existing horizontal and vertical orientation conditions.

Lockman et al. (1984) used orientation difference between the hand and the dowel as a measure of how much infants adapted their hand orientations to the target. The experimenters scored the difference between the hand orientation of infants and the dowel. Lockman et al. (1984) used 0 for full match and 4 for maximal mismatch (i.e. the difference between the orientations of the hand and the dowel was closer to 90 than 67.5 degrees). Each of the sixteen infants performed eight grasps totaling 128 grasps. The grasp definition used by Lockman et al. (1984) however, was relaxed: a finger wrapping the dowel would be counted as a grasp. In our simulation, when ILGM was learning we used our grasp stability measure; when data collecting for the analysis, to be compatible with Lockman et al. (1984), we included cases where the grasp stability was not achieved.

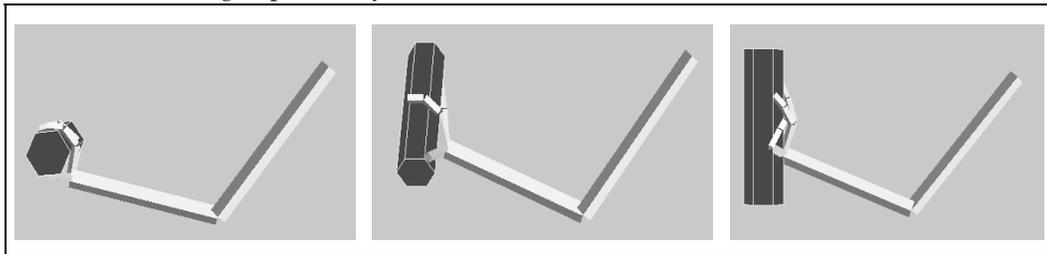


Figure 5.14 The three cylinder orientations and grasp attempts by the poor vision condition.

Figure 5.14, centre grasp does not satisfy our grasp stability criterion but conforms to the definition of Lockman et al. (1984). The right two grasps do not satisfy the grasp stability criterion

¹⁶ by clamping it to a random value

5.9.1 Simulation results

Figure 5.14 shows the orientations and the cylinder we used in the simulation. The grasp actions shown were performed by ILGM in the poor-vision condition. With our grasp stability measure, we observed that in 10 trials the horizontal cylinder could be grasped six times, The vertical cylinder could be grasped four times and the diagonal cylinder could not be grasped at all. Using the convention of Lockman et al. (1984), we saw that the horizontal cylinder could be grasped eight times, the vertical cylinder five times and the diagonal cylinder seven times. The numbers indicate that without affordance input (orientation) the grasp learning was not satisfactory.

5.9.2 Comparison of ILGM with Lockman et al. (1984)

We first compare the averaged data from Lockman et al. (1984) and data from our simulator over multiple runs (128). We compare the performance of the infants and the simulator for the vertical oriented¹⁷ cylinder (see Figure 5.15, right panel). Figure 5.15, left panel shows average orientation match score versus reach progression for infants and for the simulation. The diamonds over dashed line indicates the 5 months old infants' performance. The 9 months old infants adjusted their hands better than the younger ones as indicated with diamonds over solid line.

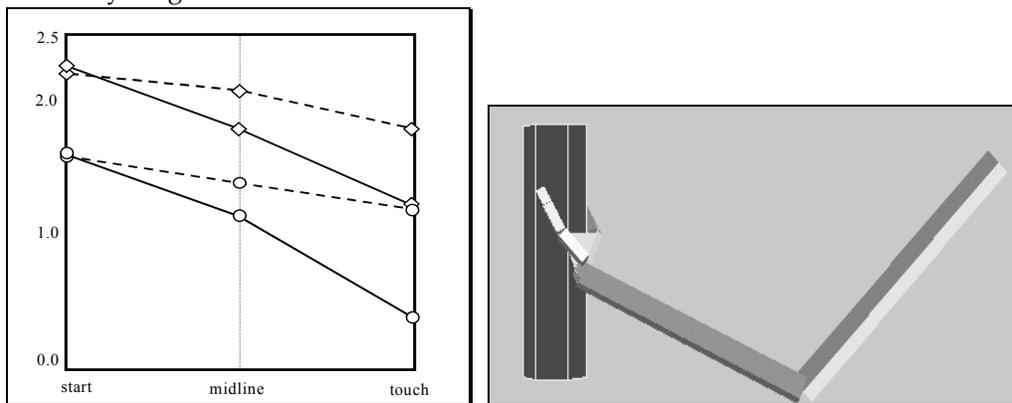


Figure 5.15 The orientation match of the hand and the cylinder is illustrated. Dashed line with diamonds: 5 months old infants; Solid line with diamonds: 9 months old infants; Dashed line with circles: ILGM with no affordance; Solid line with circles: ILGM with affordance (infant data from Lockman et al. (1984)). Right panel illustrates the object orientation used for the simulation and for the infants in this comparison

The simulated ILGM data is shown with the same line style, but data points are marked with circles. Although the absolute scores differ between simulation and infant case, the performance improved in the full-vision case similar to the performance improvement for 9 months old infants. Moreover, the performance increment of ILGM and Lockman et al.'s (1984) infants were comparable.

5.9.3 ILGM kinematics analysis (five months of age)

Now we try to infer what the model learned. When we analyzed individual errors made by the simulator, we observed six typical error curves (Figure 5.16). We grouped the error curves into rows to differentiate the mode of operation learned. The columns, from left to right, correspond to horizontal cylinder, diagonal cylinder and vertical cylinder. Note that top-left graph is flat and shows almost zero error whereas the other two graphs in the same row show flat curves with higher error.

¹⁷ We excluded the horizontal case, because for the horizontal cylinder ILGM learned an underarm grasp that is not observed in infants. The issue will be discussed later in the text.

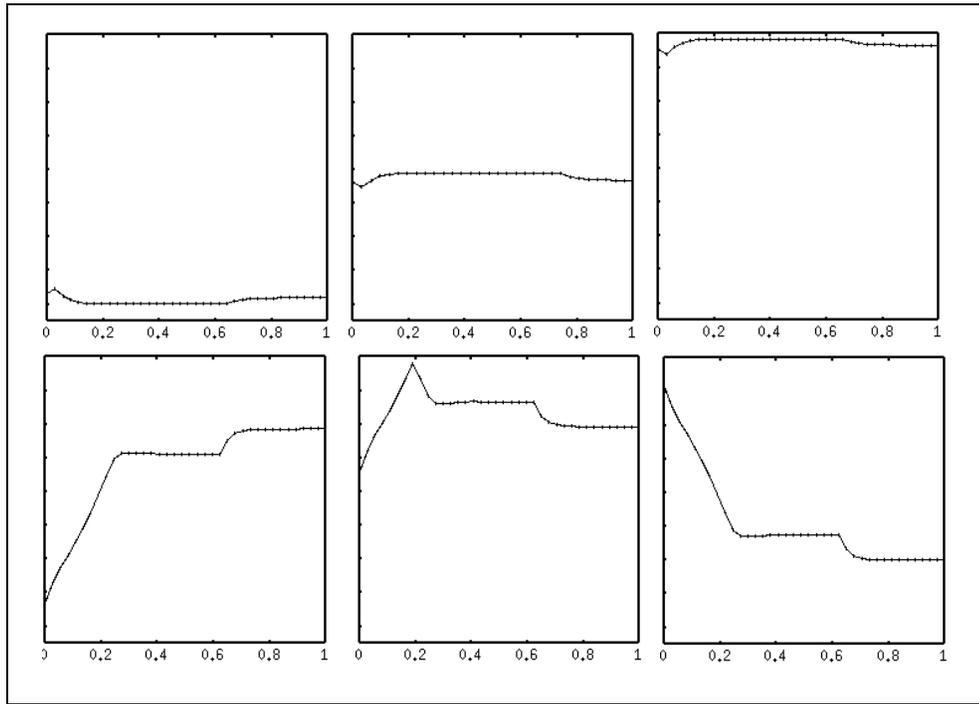


Figure 5.16 The hand orientation and cylinder orientation difference curves for individual trials. The columns from left to right correspond to horizontal, diagonal and vertical orientations. The upper row *flat* class of error curves, lower row *non-flat* class for error curves (see text for explanation)

Lockman et al. (1984) found that infants start their reaches with horizontal orientation so that their initial errors for the horizontal dowel are low. Based on this finding the initial configuration of the hand was set to a horizontal posture in our simulations. In accordance with Lockman et al. (1984) observations, ILGM made more corrections for the vertical cylinder case (bottom-right panel in Figure 5.16) and less for the horizontal case. From this we can infer that in the upper row trials of Figure 5.16, ILGM model used a grasp plan appropriate for horizontal orientation. As can be seen from the bottom-right panel, the model can occasionally perform a vertical cylinder adaptation too (remember that the architecture of ILGM allows representation of multiple grasp plans). However, the model cannot differentiate the two grasping strategies. Thus, the model learns a strategy to increase its chance to make successful grasps. Since the initial hand posture is close to horizontal, it is intuitive that the dominant mode of grasp planning becomes the one best suited for horizontal cylinder, as it requires less correction (Lockman et al. 1984). The model replicates this observation as the horizontal cylinder could be grasped most frequently.

5.9.4 ILGM kinematics analysis (nine months of age)

ILGM in the full-vision condition learns to perform grasps similar to 9 months of age infants. We present the plots of difference of hand orientation and cylinder orientations in Figure 5.17

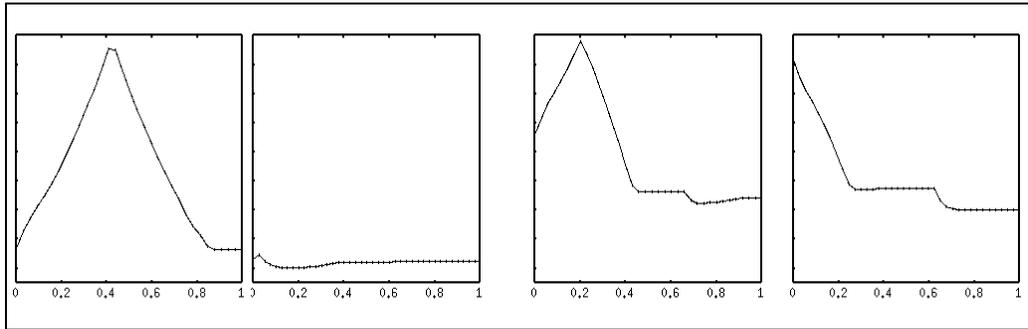


Figure 5.17 The hand orientation and cylinder orientation difference curves while ILGM was executing four types of grasp in the full-vision condition. Left two figures are two typical error curves for the horizontal cylinder. Note that the two horizontal case error patterns reflect the two possible grasps: from the bottom and from the top. The third and fourth are typical error curves for the diagonal and vertical cylinders respectively

The left two panels show the difference in orientations for the horizontal cylinder case. The flat curve corresponds to the *easy* grasp that could be observed in poor-vision condition also. The other high curvature one corresponds to a bottom grasp of the cylinder as shown in the left panel of Figure 5.18. To our knowledge, there is not report of bottom whole hand grasping during infancy. (Lockman et al. (1984) also did not observe this kind of grasp). This may be due to the general experimental set up; or the inability of infants to work against gravity. In our simulation, we did not address the dynamics of the arm. Therefore, both grasp actions are equal in their reward value. One could incorporate a penalty term for energy use in the ILGM reward computation (which evaluates how well the a grasp action has been). This would tilt the balance of grasp choice towards the *easy* grasp. It would be interesting to find out the developmental course of this particular grasp and the underlying reasons why it does not appear early in development. One explanation could be that gravity helps infants to contact with the object earlier as infants engages in arm movements and fisted swipes in the presence of visible objects (von Hofsten 1984). Thus, the infant has more experience approaching objects from the top and front. When approaching from the bottom the gravity works against the infant since the infant has to counteract gravity to grasp the object. One notable property of this bottom grasp is that it did not appear when we simulated ILGM for replicating 5 months of age infant behaviour; because ILGM had to find a strategy that will work for three orientations and hence had no exploration potential to discover alternative grasps. Fagard (2000) found that hand orientation at object contact changed with age. Horizontal orientation (the *easy* grasp) decreased and vertical ones increased from 5 months old infants to 12 months old infants¹⁸.

¹⁸ Orthogonal to this finding, the precise grasping from above became more predominant in infants older than 9 months of age.

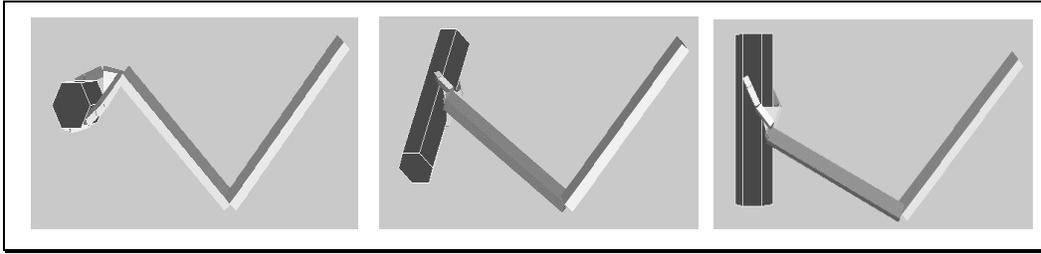


Figure 5.18 The grasps performed after ILGM learned the association between the wrist rotations and the object affordance (orientation)

5.10 Summary and conclusion

This chapter presented the Infant Learning to Grasp Model (ILGM) and presented simulations yielding nontrivial predictions. We first showed that even having a very limited set of behaviors and visual input, explorative learning could yield grasping behavior and reduce the variability of grasping actions. Then we studied the palm orienting behavior and test whether it could be acquired through interactive learning. The results of this study showed that object affordance was not a prerequisite for learning the set of correct wrist orientations. ILGM learned a ‘grasp menu’ including precision grips, which could be retrieved based on the approach direction.

In our simulations, the side grasp was not discovered. Although infants do not exhibit side grasps (opposition of thumb to side of index finger), we would like to point out the possible reason. The neural architecture we introduced in Chapter 4 is capable of learning any reward yielding grasp plan. However, the physics modeling of the arm and the objects is not realistic. We did not model contact forces. Therefore, thumb opposing the side of the finger cannot be counteracted by the side of the finger. Except for the palm, the forces can only be exerted in the directions determined by joint angles of each finger.

It should be emphasized that, the point of our modeling was not to give a realistic rigid dynamics model of the arm/hand but rather propose a grasping architecture that can learn to program grasps as long as a correct reward signal is given by the environment. The learning we demonstrated, and the variety of grasps we could generate met our aim.

During development, infants have to deal with constraints and find ways to act within the limitations of the environment and the context. We investigated the possibility that task constraints may play a role in shaping infant’s grasping behavior. We simulated a situation where a small cube was placed on a table. The model was asked to interact with this simple environment. The grasping configurations learned by the model reflected the task constraints. The model could not acquire any whole hand prehension grasps but acquired grasps that reach the cube from the top avoiding a collision with the table.

Finally we analyzed what affordance may add to ILGM by simulating the experimental set up of Lockman et al. (1984). With this simulation we showed not only the improvement in grasp execution (measured as the orientation match between infant’s hand and the target dowel) but also the similarity of the improvement pattern was comparable to Lockman et al.’s (1984) results, which indicates that ILGM captured the behaviour of 5 and 9 months age infants via differential affordance access.

Combining the summarized simulations we explicitly state that ILGM:

- Predicts that infants, even without object affordance input (i.e. during the age of early grasping), can perform precision grasp

- Shows that task constrains shape motor development and supports the view that development of precision grip is mediated by task constraints.
- Shows that object affordances could be represented in infants' grasp repertoire in spite the fact that infants' are unable to access/extract object affordance information.
- Predict that the distribution of object (visual) affordance in motor and sensory areas in human cerebral cortex would reflect the frequency of their manipulation.

The last two results are very important because it shows that tactile learning can train and modulate the visuomotor learning during infancy. By projecting this statement back to infants, we predict that infants learn to extract visual affordances of the objects they interact with.

5.11 Discussion

The simulations with ILGM explored both the early grasping period when infants are unable to factor object affordance into grasp plans and the period when they start using visual input for grasp planning. We captured the infancy period by limiting the visual input available to the model.

From a maturational-based theory point of view, ILGM should correlate more with a phylogenetically older grasping circuit with limited access to visual areas. One candidate for such a circuit is the primary motor cortex that receives object location information from the superior colliculus. As the cortical control becomes dominant while the infant is growing, motor cortex leaves the grasp-planning task to higher cortical areas, probably to premotor cortex, in adulthood.

From viewpoint of learning-based theory motor development, infant and adult grasping circuits are identical but infants have to sort out the flux of information they receive and organize them into useable schemas by interactive learning. In contrast to maturational-based theories, learning-based theories suggest that development of motor abilities is the consequence of learning by trial and error to control motor schemas that are genetically determined in their rudimentary forms (Bradley 2000). Thus, grasp planning and execution mechanism can manifest itself only after the infant has interacted enough with the environment. This is the period required for infant to master his/her sensorimotor skills, which starts from birth and continues until the first year of life. The movements of newborns are usually treated as unintentional, purposeless or reflexive muscle activities (van der Meer et al. 1995), probably with a bias from maturational-based theories. However, there exists strong evidence that infants engage in learning and exploring actively as early as 10 days of age. van der Meer et al. (1995) recorded spontaneous arm-waving movements of newborns while they were allowed to see only the arm they were facing, only the opposite arm on a video monitor, or neither arm. The newborns' hands were pushed downward in the direction of the toes. When the newborns could see their arms, either directly or via the video screen, they opposed the forces and moved normally, effectively preferring to have feedback on their movements. The findings indicate that newborns control their arm movements as long as they can receive visual feedback even when they have to oppose external forces. This shows that babies at a very early age start exploring and collecting data indicating that their visuomotor learning hardware is functioning soon after birth.

6 CHAPTER VI: NEUROPHYSIOLOGICAL VIEW OF LEARNING TO GRASP

In Chapter 5, we presented the schema level model of grasp learning (ILGM) without spelling out the brain regions that contribute to learning. In this chapter, we constrain ILGM with neurophysiological and neuroanatomical data to pin down the brain regions involved in grasp learning. Our analysis leads to two alternative hypothesis of primate grasp circuit. We present evidence for both hypothesis and analyze the one that is best explained by experimental data through simulation. In particular, we propose the Affordance-based Grasp Learning model (LGM) which meets the neurophysiological and neuroanatomical constraints and functional requirements derived from ILGM. Thus, ILGM is functionally equivalent to LGM. We will refer the reader to Chapter 5, when the material to be presented is already introduced there unless it is necessary to have an overlap to keep the Chapter self-contained.

Our analysis of grasp learning in terms of neurophysiology complements the Mirror Neuron System model we presented in Chapter 3. The simulation results enable us to explain the mechanism of grasp learning in terms of brain circuits and show how adaptation shapes the visuomotor transformation that enables primates to select and execute suitable grasps based on the object affordances leading to units with properties similar to F5 canonical neurons. Through simulation experiments, we make explicit predictions, which can be tested experimentally and used for refining and validating (or invalidating) the model we propose.

6.1 Grasp learning circuit and mirror neurons are complementary networks

One of the basic assumptions of the Mirror Neuron System (MNS) model (Chapter 3) was that self-observation of grasping was the training stimuli adapting parietal and premotor circuits of MNS model for action recognition. Using schema methodology (Arbib et al. 1998), the grasp learning and execution were encapsulated under the Reach and Grasp Schema, which included a set of algorithmic routines implementing reach and grasp execution based on techniques from robotics. Learning to Grasp Model (LGM) of this chapter serves also, as the biologically realistic schema, substituting the engineered Reach and Grasp Schema of MNS. Figure 6.1 highlights the relevant regions for grasp learning using the schema level view of MNS that we have presented in Chapter 3.

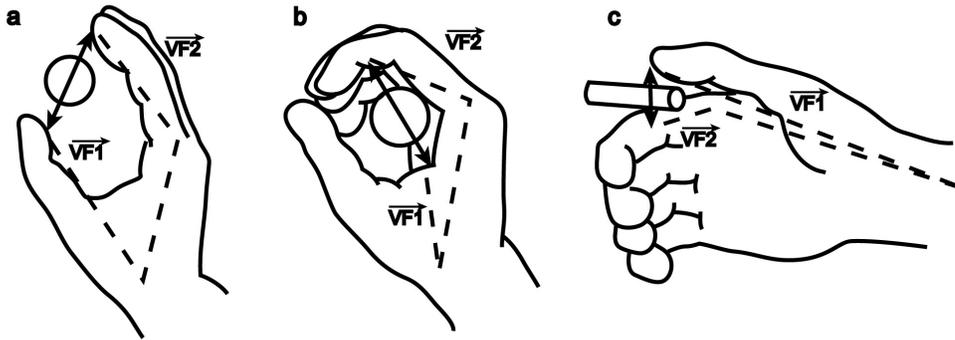


Figure 6.2 Left: precision grasp (pad opposition); Middle: Power grasp (palm opposition); Right: Side grasp (side opposition). Each of the 3 grasp types here is defined by specifying two 'virtual fingers', VF1 and VF2, which are groups of fingers or a part of the hand such as the palm which are brought to bear on either side of an object to grasp it. The specification of the virtual fingers includes specification of the region on each virtual finger to be brought in contact with the object. A successful grasp involves the alignment of two "opposition axes": the *opposition axis in the hand* joining the virtual finger regions to be opposed to each other, and the *opposition axis in the object* joining the regions where the virtual fingers contact the object (adapted from Iberall and Arbib 1990)

From a robotics viewpoint, we can model the grasping process by extending the techniques available for trajectory planning by formulating the grasp planning as an inverse kinematics problem. Indeed, we used this approach when we studied the Mirror Neuron System Model to generate grasp actions for providing visual input stimuli for the MNS model (Chapter 3). However, in this chapter, we take a learning approach and propose a model that learns to generate successful grasp plans via exploration and selection. By doing so we aim at satisfying these goals:

1. Shed light on the possible organization of the primate premotor circuit involved in grasp planning by offering biologically realistic learning rules and structures.
2. Complement our earlier work on the Mirror Neuron System, by substituting the engineered grasp-planning module of Chapter 3 with biologically realistic and self-organized grasp-planning circuit
3. Form a solid basis for hypotheses about Mirror Neuron System development and visuomotor learning in parietal and premotor circuits, which can be experimentally tested and further investigated with modeling studies following the structure we propose.

6.3 Neural correlates of infant reach and grasp

The corticospinal tract is one of the main neural substrate for independent finger control (Triggs et al. 1998). Firstly, corticospinal projections terminating in the ventral horn innervating hand muscles predict independent finger use for small objects (Bortoff and Strick 1993). Secondly, lesioning the corticospinal tract prevents the development of the independent finger movement in infant animals (Lawrence and Hopkins 1976), and impairs independent finger movements. Human patients recovering from corticospinal lesions initially tend to perform synergistic movement of all the fingers as in a power grip (Denny-Brown 1950; Lassek 1954). More evidence is presented by Olivier et al. (1997). Hinde and Rowell (1964) observed that dexterous grooming was not observed in infant macaques until 6 months of age. Lawrence and Hopkins (1976) reported that in rhesus monkey, the earliest signs of skillful hand use appear around 2-3 months and stabilizes as a mature pattern at 7-8 months. Galea and Darian-Smith (1995) reported that performance on a reach-and-grasp test (with infant macaques) reached adult levels

around 6 months. More importantly, this correlated with the emergence of an adult-like distribution of cortical motor areas contributing to the corticospinal tract. Armand et al. (1997) showed that at birth CM projections from primary motor cortex are very weak or not present. It is not very clear how much information these weak projections may carry at an early age. However Flament et al. (1992) reported that the earliest EMG responses to transcranial magnetic stimulation as in the adult could not be obtained before 2-3 months of age.

The division of function between premotor and motor cortex projecting to corticospinal tract is not well studied. The primary motor cortex has circuitry to facilitate grasping once a proper contact with the object is established (Rothwell 1994). Certain primary motor cortex neurons that control the finger muscles have cutaneous receptive fields on the skin that likely encounter obstacles and may be stimulated when a movement is caused by the neurons' activity (Rothwell 1994). Similarly, the joint receptors, tendon organs and muscle spindles have afferent organization such that if the microstimulation of a certain primary motor cortex patch produces a movement in one direction then the passive movement of the *same joint* in the *opposite* direction is likely to excite that area of the cortex (Rothwell 1994). This means that, at least for spindles, if a muscle is passively stretched, the afferents of the spindles will activate primary motor cortex neurons, which would in turn produce contraction of the same muscle (Rothwell 1994). Thus, the intrinsic wiring of the primary motor cortex can enable grasping and holding of the object once it is touched. Fogassi et al. (2001) showed with reversible inactivation studies that the (precise) grasping behavior is compromised when a muscimol injection was made to a certain part of area F5. Area F5 is divided into two main sectors based on cytoarchitectonics (Rizzolatti et al. 1988): the F5 sector lying on the cortical convexity (F5 convexity) and the part buried in the arcuate sulcus (F5 bank). Both sectors have neurons that respond to visual stimuli. The visual neurons of F5 bank respond to the presentation of three-dimensional objects, usually, in a congruent way with their motor responses in terms of grip type and size (Rizzolatti et al. 1988; Murata et al. 1997b). The visual neurons in F5 convexity fire when the monkey observes an individual performing certain actions involving object interaction (Gallese et al. 1996; Rizzolatti et al. 1996a). The former neurons are named 'canonical neurons', whereas the latter ones are named 'mirror neurons'. Although the muscimol injection to the region of F5 mirror neurons did not impair grasping behavior (only a slowing down was observed), an injection made to F5 bank impaired (precise) grasping (Rizzolatti et al. 1996a). The hand was not adjusted according to the object size and shape. Nevertheless, the monkeys could perform the grasp after they touched the object. Interestingly, the grasp for large objects appeared almost unaffected (Murata et al. 1997b). In one of the monkeys tested, the use of the contralateral hand was largely impaired with F5 bank injection. This monkey often refused to make reaching-to-grasp movements toward small objects and, when it made them, grasping was clumsy and the hand shape was inappropriate for the object size and shape. A large sphere and a large cylinder, however, were grasped almost normally (Fogassi et al. 2001). The other monkey had similar deficits but to a lesser extent. In both monkeys, the ipsilateral hand performance was also compromised (when the muscimol injection was strong). It is important to note that, finger dexterity was not abolished when F5 bank was inactivated because after contact with the object the grasps could be completed. This is very similar to the description of grasp performance during early grasping phase of human infancy (von Hofsten and Ronnqvist 1988). Fogassi et al. (2001) also injected muscimol in the hand region of motor cortex (F1). The result was a strong impairment in the capacity of the grasping of the hand contralateral to the injection site. Unlike F5 injection, the ipsilateral hand was not impaired (Fogassi et al. 2001). Both monkeys with F1 injection, could reach the tray that holds the object, but with a stereotypical flat hand configuration after which they retrieved (i.e. used their hands like a rake) rather than grasped the objects. The hand

section of area F1 receives rich projections from area F5 (Matelli 1986). We can argue that F5 modulates F1 hand related neurons so as to engage them in a precision grasp in the following way. We know that neither area F1 nor area F5 alone is enough to perform a precision grip (Fogassi et al. 2001) and area F5 neurons become preferentially active when the animal performs grasping actions with varying degrees of grip specificity (Rizzolatti et al. 1988; Murata et al. 1997b; Rizzolatti et al. 2000). Thus, we suggest that area F1 needs to be modulated by area F5 in order to carry on the dexterous grasp actions. This proposal is also supported by human imaging studies. Ehrsson et al. (2000) showed, using fMRI, that when the human subjects performed power grasp, the primary motor cortex of the contralateral hemisphere showed increased activity whereas when the subjects performed precision grasp both hemispheres were activated. Importantly, the dominant activity was observed in the (ipsilateral) ventral premotor cortex, which is the homologue of monkey F5 (Gallese et al. 1996; Rizzolatti and Arbib 1998). Ehrsson et al. (2000) suggested that the control of fingertip actions with a precision grasp is mediated with neural circuits that are different than the circuits involved in power grasp, which are phylogenetically older. Further support for this proposal comes from the studies of Muir and Lemon (1983). Some primary motor cortex cells that project to motoneurons that mobilize hand muscles are active during precision grasp execution but not for a power grasp execution, albeit the same muscles can be mobilized in both grasps. This indicates that there exists a skilled grasping circuit, which uses some part of primary cortex for its exclusive function. Then, how much visual information is available to phylogenetically older (primary motor cortex) grasping circuit and how well can it learn sensory motor associations (as opposed to the fixed motor plans)? Shen and Alexander (1997a) showed that the neurons in motor cortex participate in sensory and/or associative (context-dependent) processing of spatial information relevant to visually guided reaching movements. They have found many neurons in motor cortex that showed behavior-correlated discharge that depended on the visuospatial target of the monkey's instructed reach, irrespective of the limb trajectory used (as expected, there were also a substantial proportion of neurons with limb-dependent activity as well). This means that the primary motor cortex neurons indeed can encode motor plans based on visual cues that are not tied to certain muscle groups.

6.4 Primate grasp development hypotheses

We now present two alternative hypotheses of primate grasp development that will guide us in locating the specific brain areas involved in grasp learning in the following sections.

6.4.1 Hypothesis I: two coexistent grasping circuits

The first hypothesis is that the early power-like grasp is mainly controlled by area F1 and F4 (for the reach component), and involvement of area F5 is not substantial. The review in the previous section suggests the existence of a phylogenetically older lower level grasp circuit that does not require premotor regions and can work with limited visual analysis of the object, probably position and rough size. The power grasp appears early in development but it takes longer to mature a precision type of grip (Lockman et al. 1984; Fagard 2000). In addition, infants lack the ability to fully utilize vision for grasping before the precision grip becomes part of their grasp repertoire (Rosenbaum 1991). Ehrsson et al. (2000) suggested that the control of fingertip actions with a precision grasp is mediated with different neural circuits than the phylogenetically older circuits for power grasp.

Hypothesis I explains the emergence of skillful finger use as the following. The visual control of dexterous finger use (e.g. precision grip) is learned by area F5 by associating the performance of the lower level grasp circuit with the visual analysis of the object performed by parietal cortex (AIP). Area F5 modulates grasp selection and the activity in the lower level grasp circuit. The modulation and

selection are based on the high level goals of the individual (prefrontal influences based on context), affordances extracted by parietal cortex, the lower level grasp plans and the actual grasping performance (success or failure).

6.4.2 Hypothesis II: single grasping circuit

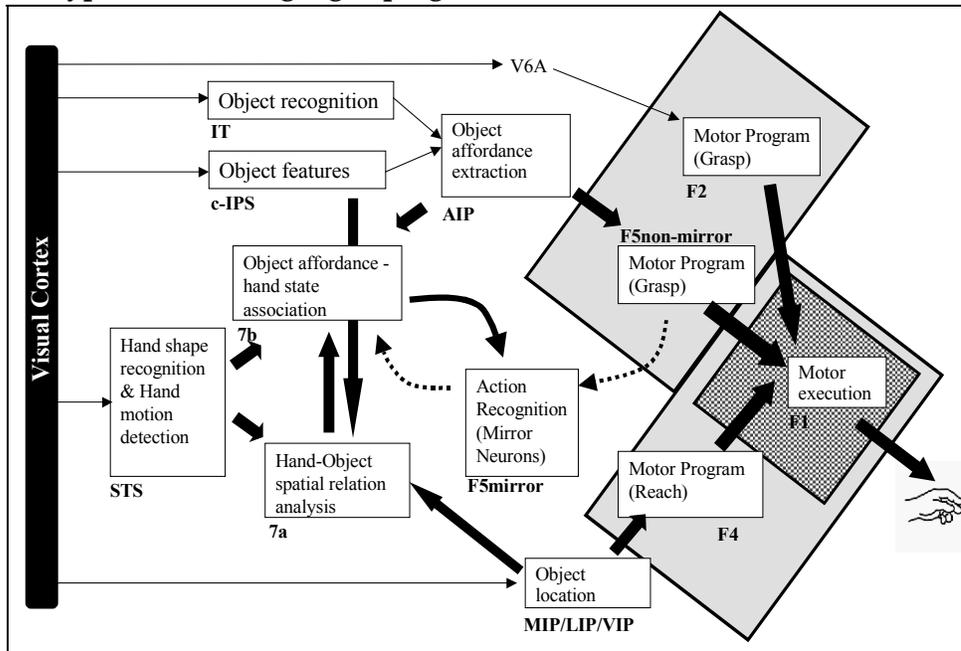


Figure 6.3 The two possible organization of learning to grasp circuit are shown. According to Hypothesis I, two grasping circuits exist; the phylogenetically older one located in area F1 (hatched background) and the newer one in the premotor cortex (solid background). According to Hypothesis II, F1 is involved in only executing the premotor cortex instructed movements. LGM is based on the latter hypothesis. The details of LGM are shown in Figure 6.4. Note that we introduced area F2 for complementing the MNS structures. The visual input to area F2 originates from MIP (not shown) and V6a

The alternative hypothesis is that learning to grasp circuit is distributed over a large area in the premotor cortex including area F5, with access to direct object affordance input from AIP; and area F1 is responsible for the execution of the plan instructed by area F5 (see Figure 6.3). According to this hypothesis, the inability of infants to perform adult-like grasps can be explained by the underdevelopment of AIP for affordance extraction and the inability of infants to control their limbs, or the computational complexity of AIP-F5 learning. The former corresponds to the maturational-based theory of motor development; while the latter corresponds to learning-based theories (see Chapter 5).

We will introduce Affordance-based Learning to Grasp Model (LGM) based on Hypothesis II, which means that LGM will be the model of grasp related visuomotor circuit of monkey premotor cortex. LGM, being a simulated neural realization of Hypothesis II, will yield testable neurophysiological predictions.

6.5 Affordance-based learning to grasp model (LGM)

With the term *learning to grasp* we mean to learn how to make motor plans in response to sensory stimuli such that the *open loop* execution of a plan leads to a successful grasp. There is strong behavioral evidence that early grasping is based on open-loop control and does not use visual feedback (Clifton et al. 1993; Clifton et al. 1994; von Hofsten and Ronqvist 1988; Streri 1993). Further, adult practiced

movements move from a visual feedback control strategy to an open loop control strategy evidenced by postural invariance studies in grasping movements. Desmurget et al. (1998) studied a prehension task requiring subjects to grasp a cylindrical object presented at different locations with changing orientations. The effect of initial arm posture was investigated. The results showed that individual subjects had stereotypical grasping patterns resulting in fixed postures of the arm, which varied systematically as a function of initial posture and object location and orientation (Desmurget et al. 1998). Furthermore Grea et al. (2000) showed that the final posture to be reached is planned in advance and used as a control variable by the central nervous system. This was true even when the object jumped to a new location during the transport phase of the reach. The new position of the object could determine the final arm posture with the same precision as a stationary target (Grea et al. 2000).

However, we do not claim that visual feedback is not used for reaching and grasping. On the contrary, we suggest that when the task demands cannot be satisfied with existing motor schemas, the visual feedback control becomes necessary. In Chapter 7, though in a simplified setting, we study the visual feedback control of grasping. Current chapter addresses the issue of acquiring a grasp repertoire that can generate suitable grasping based on the object affordances.

The link between open-loop grasp execution and the visual feedback based grasp execution of Chapter 7 is established through AIP-F5 learning in LGM. As will be shown LGM learning yields units that show object selective responses similar to F5 canonical neurons. The favored hypothesis of Chapter 7 is that F5 canonical neurons, based on the object properties, gate F5 visual servo circuits. Thus, Chapter 7 without loss of generality, will concentrate on a visual servo circuit specialized for precision grip learning and execution, which is presumably selected by a population of F5 canonical neurons.

6.5.1 Localizing learning to grasp model in primate cortex

We propose three computational layers relevant for grasping: Hand Position, Virtual Finger, Wrist Rotation layers, as we functionally justified in Chapter 5. The affordance of the target can vary from a single variable indicating the existence of a graspable object to the full description of an object in terms of its affordances, and is encoded in the Affordance layer. The affordances are represented using the population-coding scheme and encoded algorithmically (i.e. no visual processing for object recognition and feature extraction is performed). We use *Affordance layer* or the *input* terms interchangeably to conform the context

The parameters (hand position, virtual fingers and wrist rotations) encoded in LGM layers are abundantly represented in the premotor and motor cortices. Thus, it is not always possible to constrain the localization of the layers with a high level of confidence. Nevertheless, we can minimize the number of alternatives based on relevant literature as follows.

The *wrist rotation* parameters are represented in the primary motor cortex (area F1) in terms of direction (i.e. independent of the muscle groups activated) (Takei et al. 1999). The ventral premotor cortex (area F5) neurons are involved in extrinsic coding of hand direction (Takei et al. 2001). Area F2 has control over wrist movements and is organized similarly to area F1 in terms of somatotopic organization (Fogassi et al. 1999); thus Wrist Rotation layer can be associated with F2 as well as with F1. Premotor cortex (area F4) can potentially encode hand location with respect to the object (Fogassi et al. 1992; Fogassi et al. 1996), as it is the target of ventral intraparietal area (VIP) (Geyer et al. 2000), which is involved in egocentric target representation (Duhamel et al. 1997). The VIP-F4 circuit, therefore, can play an important role in reach and grasp planning in monkeys (Rizzolatti et al. 1998). Thus, it is tempting to associate the Hand Position layer with area F4 and/or area F5. Recalling the object selective motor properties of F5 neurons (see Chapter 2), we can posit F5 in performing the task of Virtual Finger layer by instructing which configuration to use for a given object. In addition, Cisek and Kalaska (2002)

presented evidence that the primate premotor cortex can simultaneously represent discrete directional signals related to multiple alternative reaching actions. In this study, a monkey was asked to reach for possible targets, which were cued by a nonspatial 'go' stimulus. During the first period, while the monkey was waiting for the go signal, two directional signals coexisted in the activity of neurons in dorsal premotor cortex encoding the reach directions toward the two potential targets (Cisek and Kalaska 2002). When the 'go' signal was given, the activity encoding the non-cued direction disappeared and the remaining signal predicted the monkey's reach choice (Cisek and Kalaska 2002). In a similar study by Hoshi and Tanji (2000), an additional cue was introduced to instruct the monkey to use its left or right hand when pointing one of the two targets. The results indicated that alternative motor plans were represented in the dorsal premotor cortex before the action is uniquely determined. Note that one of the main motivations behind the neural architecture we developed in Chapter 4 was to be able to encode multiple action plans.

In the light of the above review and in accordance with Hypothesis II, we propose that area F5 works closely with areas F4 and F2 to create a feasible grasp plan based on the object affordances relayed via AIP (see Figure 6.4). The grasp plan is then, relayed to the primary motor cortex (F1) and the spinal cord for execution. The tactile feedback of the grasp is assimilated in the first somatosensory cortex (SI), which mediates learning in parietal and premotor connections.

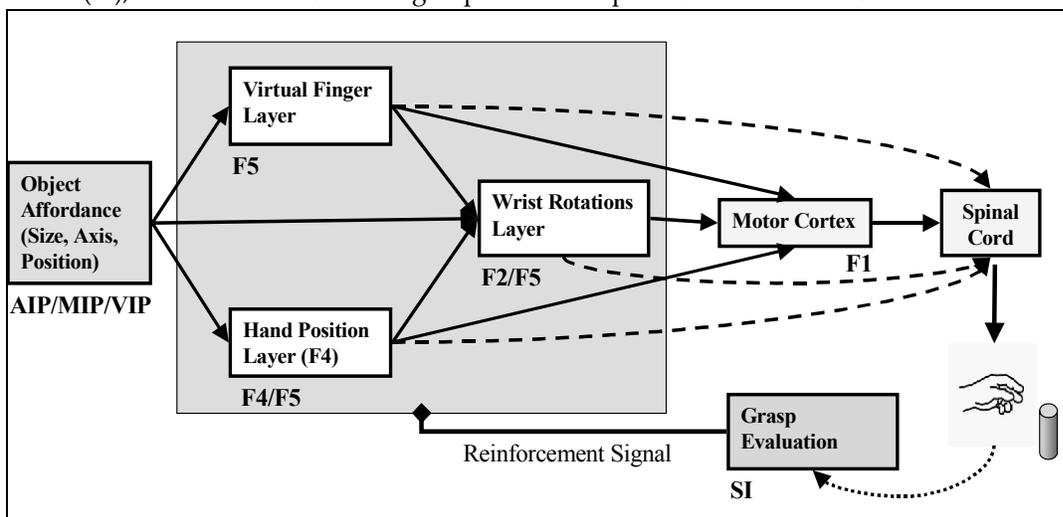


Figure 6.4 The Learning to Grasp Model. F5 is implicated in all grasp related parameters. Dashed connections indicate the direct corticospinal projections of premotor areas. Area F5 works with area F2 and F4 to transform visual affordances signalled by parietal areas into a grasp plan. The grasp plan is then, relayed to primary motor cortex (F1) and spinal cord for execution. The tactile feedback of the action is integrated in the first somatosensory cortex (SI), which mediates the adaptation of the parietal-premotor and inter-premotor connections

The involvement of area F5 in multiple grasp parameter coding finds support by the finding that most F5 neurons were selective for grasping movements, but there were also reach related as well as wrist rotation selective neurons (Rizzolatti et al. 1988). The multiplicity of motor representations (Wu et al. 2000; Rizzolatti et al. 1988; Gentilucci et al. 1988) makes it impossible to rule out other alternatives such as the view that F5 being exclusively involved in mapping object affordances to finger configurations. This, in fact, points out the importance of our modeling. With our model, we aim to

motivate neurophysiologists to challenge the model predictions with experiments, which will help uncover the detailed functional roles of the premotor areas in grasp related visuomotor transformations.

6.5.2 What does cerebral cortex know about a grasp?

When a human infant or a monkey touches, many receptors in the hand transmit many signals related to the contact such as the skin indentation and slip (Rothwell 1994; Johansson and Westling 1987b; Johansson and Westling 1987a; Salimi et al. 1999b). The literature on the mechanisms of mechanoreceptors and the transmission of their signals to the cortical areas is vast. Our intention is not to model these mechanisms, but rather convince the reader (and ourselves) that a reinforcing signal indicating a successful grasping is available to the learning-to-grasp circuit located in the premotor cortex. Thus, we present here a brief but relevant data from the literature on the role of the somatosensory cortex in representing the sensations of the hand, including grasping. The interested reader is referred to other literature (e.g. Akoev et al. 1988; Willis and Coggeshall 1991) for detailed information on mechanoreceptors and their functional organization.

The primary somatosensory cortex (SI) has somatic representation of the fingers and shows differential activity during grasping phases (Gardner et al. 1999; Ro et al. 2000). In one study, anterior SI was found to be dominated almost exclusively by neurons with cutaneous receptive fields (88%) and posterior SI neurons were found to receive tactile inputs (51%) and deep inputs from muscle and joints (41%) (Debowy et al. 2001). Furthermore, Debowy et al. (2001) showed that during prehension somatosensory cortex units signaled the *formation of hand and object as a functional unit* in combination with other hand actions. Debowy et al. (2001) classified SI neurons, among others, as approach, contact, contact-grasp, grasp-lift, manipulation and grasp inhibited neurons. In terms of tuning, they have found neurons with grasp tuning, approach tuning, hold tuning and contact tuning.

Martin et al. (2000) examined the effects of blocking neural activity in cat sensory motor cortex (muscimol infusion) during early postnatal development on prehension. Grasping occurred on only 14.8% of trials with the limb contralateral to the infusion. In addition, the grasping was replaced by raking without distal movements. This data suggests that the normal development of skilled motor behavior requires activity in the sensory motor cortex during early postnatal life.

Salimi et al. (1999) examined the receptive field properties of somatosensory cortex neurons in monkeys during a precision grip task. The majority of the receptive fields found was cutaneous and covered less than one digit. Two types of neurons were described: dynamic and static. The dynamic neurons, showed a brief increase in activity beginning near grip onset, which quickly reduced even the pressure to the receptive field continued (Salimi et al. 1999a). Some of the dynamic neurons responded to both skin indentation and release (Salimi et al. 1999a). The static neurons had higher activity during the stationary holding phase of the task (Salimi et al. 1999a).

Based on the brief review above, we postulate that grasp success is signaled as a population activity in the somatosensory cortex (SI)¹⁹. In order to simulate and test the proposed grasp-learning hypothesis we need to emulate the grasp success signal to drive the learning in LGM. We base the emulation on the physical definition of grasp stability as we did in Chapter 5. Of course, we do not claim that such computation is performed in the primate brain. We are using the tenets of schema methodology (Arbib et al. 1998) to substitute a biological *reinforcement or success* schema ('joy of grasping' of Chapter 5) with the engineered version for the sake of analysis.

¹⁹ The form of the signal is not the main issue here, what we are after is a neural representation of *holding* (stable grasping). As a first approximation, we posit somatosensory cortex for this role without assigning an emotional meaning to the signal. Prefrontal cortex may be involved in such motivation-based signal.

6.5.3 Simulation level description of LGM layers

We have already introduced LGM layers in Chapter 5 as schemas. Now, we summarize the functional description of these layers to prepare the reader for neural level analyses.

Hand Position layer determines where the hand will be with respect to the object during grasping. Intuitively, this represents the side from which the hand will grasp the object. The frame of reference for this parameter is allocentric. The coordinate system we choose for this parameter is spherical. The encoding used is as described in Chapter 4. Thus the neural layer represents a probability distribution of hand position values and the specific output is generated by a local population vector computation.

Virtual Fingers layer specifies which fingers will be activated, and with what strength. We use three virtual fingers: thumb, index finger, and the remaining three fingers acting together. The processing is parallel to the Hand Position Layer's flow. However, this does not mean that this layer has the right to decide the virtual fingers on its own. The virtual fingers that cannot yield grasping are negatively reinforced; and hence do not appear in learned LGM.

Wrist Rotation layer combines all the information about the object (affordance) and the Hand Position and Virtual Finger layers' output. Thus, the output of this layer represents the possible wrist orientations given the (1) object related input (affordances), (2) output of Hand Position layer (3) output of Virtual Fingers layer in terms of a conditional probability distribution. The parameters generated in this layer determine the movements of wrist extension-flexion (tilt), wrist supination-pronation (bank); and ulnar and radial deviation (heading).

A typical scenario for grasp execution would be as the following:

- A small box is presented in the workspace of the arm to the left of midline of the body.
- Hand Position layer computes the distribution of feasible approach directions and a selection is made according to the distribution, for example as 'from top'
- The Virtual Finger layer works similarly. Let us assume that a selection is made such that index and thumb fingers are activated such that their trajectories coincide (a precision pinch).
- Then, the Wrist Rotation layer combines the Virtual Finger and Hand Position layers' parameters with the Affordance layer to compute a probability distribution for the applicable (e.g. to the precision pinch approaching to the object from top) wrist orientations. The grasp plan is complete when the final selection is made from the Wrist Rotation layer.

6.5.4 Why LGM is relevant: good model versus bad model

Microstimulation studies ensure that there are neurons in the primary motor cortex and premotor cortex that control finger digits, wrist movements, and reaching. This chapter provides a learning mechanism to adapt the connectivity of premotor regions so that they act cooperatively to yield feasible grasp plans. One of the main motivations of using a minimal set of grasp parameters was to test whether learning by interacting with the environment, can shape neurons to have properties that we did not manually encode. The emergent neuron properties such as the object selectivity are very important because they justify that the structure and the learning proposed are adequate to capture the learning for the grasp related visuomotor transformation in the primate since we bootstrapped the grasp learning from a minimal set of elemental/postnatal abilities and behaviors summarized as the following.

1. Motor abilities: infants are able to move their wrist, and fingers
2. Visuomotor abilities: infants reach for visual targets
3. Behaviors: infants explore the space with their hands through variable movements

4. Reflexes: infants are born with reflex behaviors helping them to shape their (visuo)motor abilities (e.g. enclosure reflex)

This is why we did not assume, for example, a layer of *grasping neurons* specialized for different hand apertures. If we did so, those neurons and the objects with compatible sizes would be trivially associated via trial and error learning, which would not add anything to our knowledge. However, with our approach we show that starting from a basic set of abilities/behaviors in accordance to infant development, complex neural properties emerge which yield predictions that can be experimentally tested. If the model is validated by comparing experimental findings to predictions of the model, we can suggest that primate brain follows a similar strategy to develop visuomotor abilities by interacting with the environment. Thus, a validated model will (1) bring new insights to primate visuomotor transformation and (2) enable us to make sound predictions with new simulation experiments. As will be presented in later sections, LGM leads to the emergence of neurons with nontrivial properties such as object preference, similar to those of F5 canonical neurons in monkey premotor cortex.

6.6 Wrist orientation-learning revisited: neural level analysis

Chapter 5 showed that Infant Learning to Grasp Model (ILGM) was able to learn how to adjust wrist orientation and approach-direction via explorative learning. The model learned to generate parameters to perform power and precision grasps, and many variations of the two. The most abundant grip generated was the power grip and its variations. Precision type grasps were less frequently generated. In terms of learning task, the ILGM had to discover the distribution of wrist movements (supination-pronation, extension-flexion, ulnar/radial deviation) for multiple approach directions. Presented with an object, Hand Position layer produced the distribution of possible approach directions. The selection was relayed to Wrist Rotation layer. Then, Wrist Rotation layer computed the distribution of feasible wrist orientations conditioned on the generated hand position. Virtual Fingers layer dictated the enclosure speed of the hand.

Noting that ILGM being functionally equivalent to LGM, analyzed the behavioral aspects of grasp learning, we now present a neural level analysis of the wrist-orientation learning of Chapter 5.

6.6.1 Neural level analysis

In this section, we analyze the activities of LGM layers. We will use two graphical representations: one for the probability distribution represented by LGM units and one for the memory (or eligibility) traces. The former map represents how likely the units would fire given the input and the context while the latter represents the 'generated parameter' at the particular instant. A neurophysiological analogy for the probability distribution would be a multi-electrode recording experiment. The normalized firing histogram of the neuron population over many trials (with fixed input and the same experimental conditions) would be very much comparable to what we call the probability distribution. Thus the probability distribution graphs represent (1) population level activity of the neurons in a layer and (2) individual neurons' preferred stimulus. The eligibility trace maps show the activity of neurons that are involved in parameter generation in a single trial. Figure 6.5 shows the activity of the Hand Position and Wrist Rotation layers (left two panels) as well as the memory traces (center two panels) for the grasp shown in the right panel. Note that the preferred values of Hand Position and Wrist Rotation layers are three-dimensional parameters. To visualize them we average the activities over an axis. For Hand Position we average over *radius* axis. For Wrist Rotation layer we average over the *ulnar/radial deviation* axis. These two axes are chosen because they are found to be least influential in generating successful grasps.

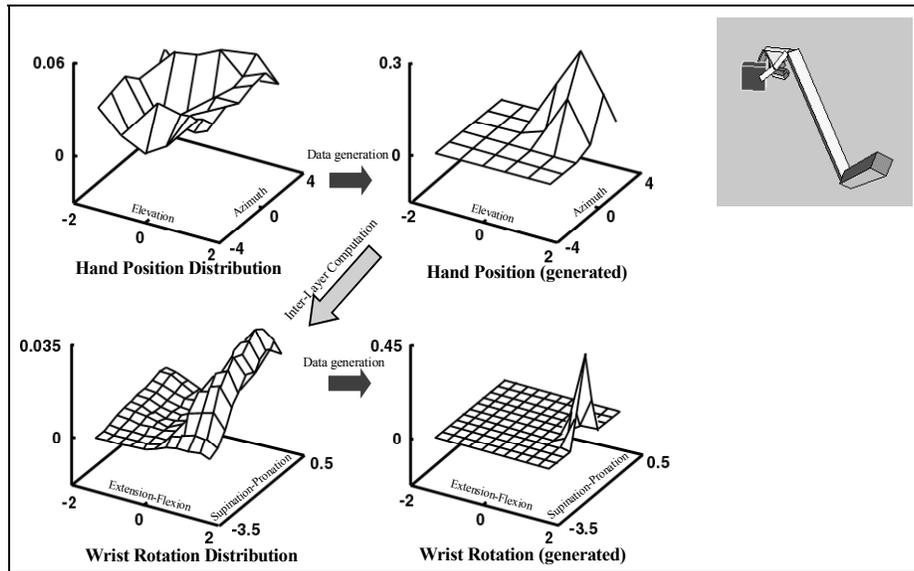


Figure 6.5 The top-left shows the Hand Position layer output summed over the radius (approach direction is encoded in spherical coordinates) as a 3D plot. The top-center panel shows the sample generated from the Hand Position distribution. Bottom-left shows the Wrist Rotation layer output summed over the heading axis as a 3D plot. The bottom-center panel shows the parameters picked from the Wrist Rotation layer distribution. Note that Wrist Rotation layer distribution depends on (i.e. represents a conditional distribution) the sample picked from the Hand Position layer. The rightmost panel shows the executed grasp

The allocentric position encoded by top-center map roughly corresponds to a location higher than the object, which is also at the behind and right side of the object. This location determines the approach direction of the hand (hand reaches for the object from that location). Given this approach the possible wrist rotations that yield stable grasps are computed by Wrist Rotation layer. The bottom-left panel shows the firing potential of units in this layer. Having a population level activity, we can interpret the strategy ILGM learned. The Wrist Rotation probability distribution roughly says, as long as the wrist is flexed between 30 to 80 degrees (tilted downward) then the grasp will succeed with the hand approach direction encoded in the memory trace. Of course, there is no way to tell in advance, which of the combinations will yield a precision pinch. Our general notion is among the performed grasps, there will be precision grips, which will be picked and tuned by the premotor cortex.

6.6.2 LGM represents a ‘menu’ of grasps in terms of neural activity

Now we look at another grasp plan made by the same LGM (no new training). As we mentioned earlier Wrist Rotation layer learns feasible wrist rotations based on the eligibility trace of the Hand Position layer. This time Hand Position layer generated a different approach (Figure 6.6, top-center panel), which changed the landscape of the Wrist Rotation layer’s distribution (bottom-left graph). The wrist rotation parameter generation is shown in the bottom-center panel and the resulting grasp is shown in the right panel.

One of the important observation is that LGM was able learn a ‘menu’ of grasps which is open to biasing. The bottom-left panel of Figure 6.6 explicitly shows that two distinct sets of wrist rotations could be generated given the approach direction encoded in the memory trace of the hand position parameter (Figure 6.6, top-center graph).

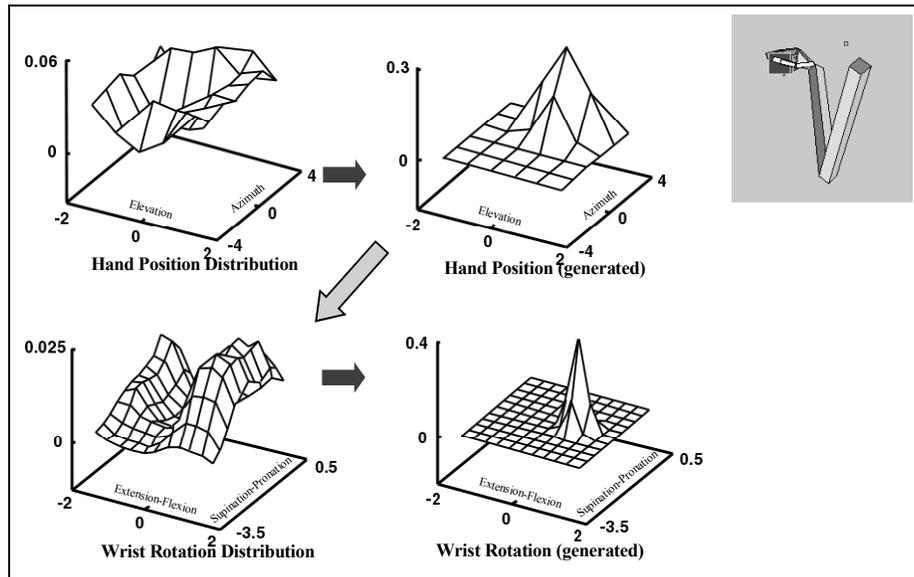


Figure 6.6 Using the same LGM used for Figure 6.5, another grasp plan is generated (left four panels). The resulting grasp is shown on the right. By comparing the grasp plan shown on the left four panels with of Figure 6.5's grasp plan we see how the selection of a different approach direction (see the centre-top panels of both figures) changed the Wrist Orientation distribution

Now we demonstrate the two options from the 'grasp menu'. Figure 6.7, shows clearly the variability of the grasp that can be encoded in LGM. In the upper panels the Wrist Rotation layer specified maximum wrist extension with pronation (top-center panel). The resulting grasp is shown on the upper-right panel. In the second trial, the generated wrist rotation instructed maximum supination and small extension. The resulting grasp was very different although the hand approached the object from the same location (the Hand Position layer is not shown).

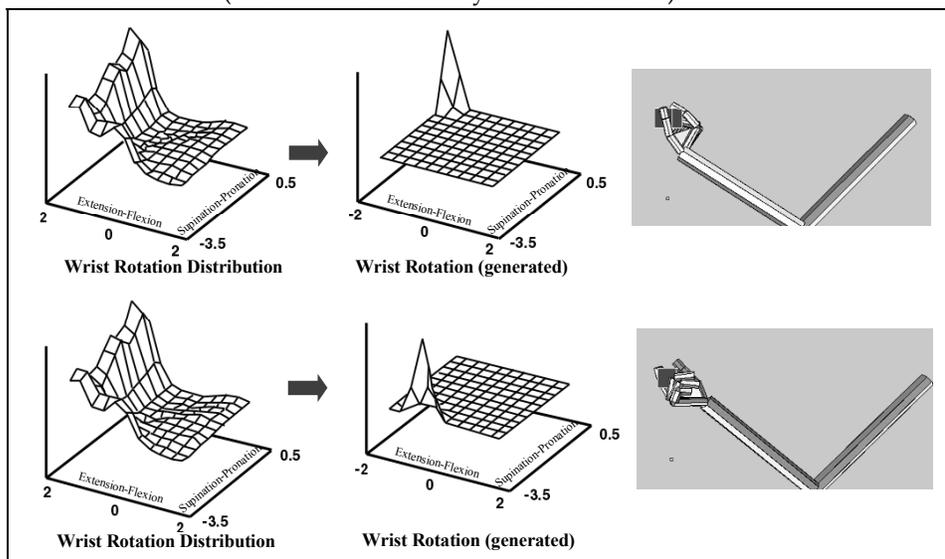


Figure 6.7: Two very different grasp generation from the same LGM. Upper panel: Grasping with maximum wrist extension with some pronation. Lower panel: Grasping with maximum wrist supination and small wrist

extension. Note that the Wrist Layer probability map is the same since the approach direction was chosen the same (the small dots in the right most panels).

6.6.3 Predictions and discussion

LGM simulation in particular predicts that the premotor cortex neurons that are involved in encoding limb movement parameters (e.g. wrist rotations) must be modulated by other movement parameters (e.g. the direction of approach).

By generalizing grasp actions to general movements, we predict that complex movements cannot be determined by a single layer. For example, the neurons controlling digits and wrist rotations are segregated in area F2 (Fogassi et al. 1999); area F4 controls proximal while area F5 distal movements (Rizzolatti et al. 1988; Gentilucci et al. 1988). Somatotopic organization of multiple motor areas is an evident strategy of the primate motor cortex (see reviews: Wu et al. 2000; Geyer et al. 2000). Thus, the general prediction we make is that the activities in these segregated regions *must be modulated* by their related peers (e.g. hand-arm). For an experimenter the neurons of this nature would appear to have ‘gain fields’ or appear to be modulated by other behavioral contexts.

Let us take a hypothetical but common experimental set up (for example see Georgopoulos et al. 1982). The experimenter wants to find the neurons encoding wrist extension-flexion in monkey. The monkey places its arm on a table and performs wrist extension and flexion movements. The experimenter locates the motor region that (e.g. F2) correlates very well with the monkey’s movement. Can he say that this region encodes wrist extension/flexion? With the insights gained from our simulation studies, we answer ‘no’. The activity could be well depend on other motor parameters, such as posture. Indeed, this kind of modulation has been shown to exist in the primary motor cortex (Sergio and Kalaska 1997) with reaching movements using different arm postures. The finding has also been extended to the dorsal premotor cortex confirming that similar modulation exists in higher motor areas (Scott et al. 1997).

In fact, Arbib and Hoff (1994) noted the important distinction between neural activity that *controls* movement versus neural activity that *correlates* with movement. We propose that to discover motor circuits in the cortex the correlation studies (correlating behavior to neural firing) is not enough. Simultaneous recording from anatomically connected regions must be required to understand the computational elements underlying the modulation or gain field phenomenon

6.7 Object axis selectivity: neural level analysis

Lockman et al. (1984) used subjects of 5- and 9-months of age and compared their performance in orientating their hands to a dowel presented in horizontal and vertical position (see Figure 6.8). To replicate their experimental condition, we let the affordance of the object (orientation) to be relayed to Learning to Grasp Model in accordance with our Hypothesis II (see Figure 6.3). We presented a cylinder to LGM analogous to a dowel. To implement the premises of our hypothesis that the information about the axis orientation (i.e. affordance of the object) is not available to the grasp planning circuit during early infancy (5 months of age) and that it becomes available later in development (9 months of age), we effectively disabled the orientation encoding for the simulation of younger infants’ grasp learning whereas to simulate older infants’ case we enabled the orientation coding in their affordance input (see Chapter 5). The former case was referred as the *poor-vision* and the latter as the *full-vision* condition.

With this simulations we had two goals: (1) to compare the simulation results with of Lockman et al.’s (1984) and (2) to analyze the internal representation of the learned model. Chapter 5 compared the model results with experimental data fulfilling (1). Now we present the internal representation analysis

of LGM for both affordance and no affordance case representing 9 months and 5 months of age infants respectively.

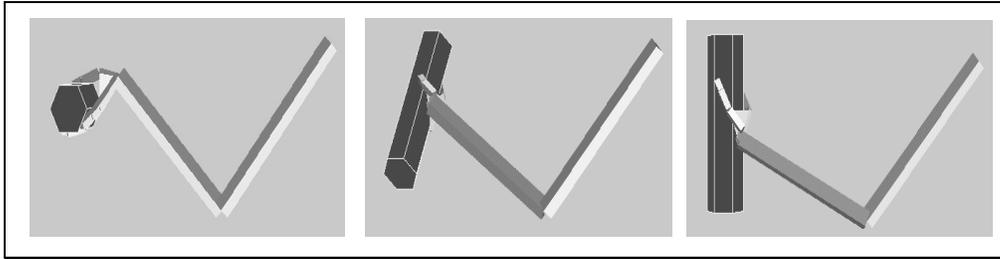


Figure 6.8 The grasps performed after LGM learned the association of hand rotations with the object orientation input (full vision condition). Note that the left panel shows a bottom side grasp. All of the shown grasp configurations satisfied grasp stability criterion

6.7.1 Neural level analysis

By analyzing the kinematics error curves of the orienting behavior, we concluded that there were two modes of grasp planning (Chapter 5); but how correct is this? We answer this question by examining the population level representations emerged by LGM learning.

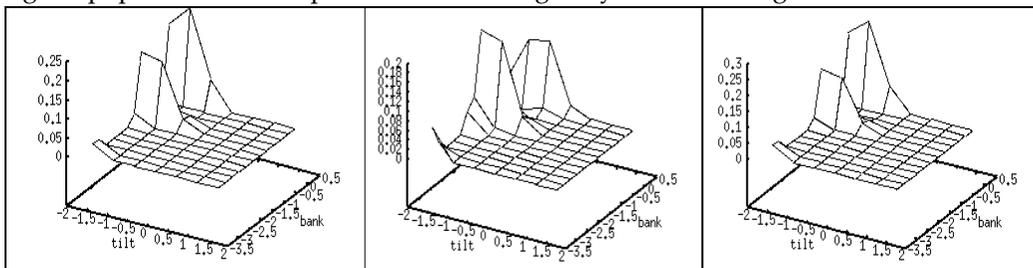


Figure 6.9 In the poor-vision case, the hand rotation neurons in LGM show the same response for horizontal (left panel), diagonal (centre panel) and vertical (right panel) object presentations because of the lack of axis orientation input

We are interested in deciphering the grasp planning strategy learned by LGM in poor-vision condition. Figure 6.9 shows the Hand Rotation neurons of LGM: (from left to right) for horizontal, diagonal and vertical cylinder presentation. The Hand Rotation distribution is almost identical since the axis orientation information is not accessible to LGM in the poor-vision case. The Hand Rotation Layer distribution confirms our inference from the kinematics: there are two peaks of neuron activity, which corresponds to the vertical and horizontal orientations of the hand during grasping.

However, note that the two possibilities could not be specialized for the different orientations and hence represented in all three cases. The peak at $(-\pi/2, 0)$ indicates a wrist rotation that makes the angle between the backside of the hand and the arm 90 degrees (full wrist extension), which is required for grasping a centered horizontal object from the front side. On the other hand, the other peak at $(-\pi/2, \pi/2)$ indicates an additional supination of the hand, which is required for grasping vertical cylinder from the front side. These values corresponds to the grasp configurations we kinematically observed in Chapter 5.

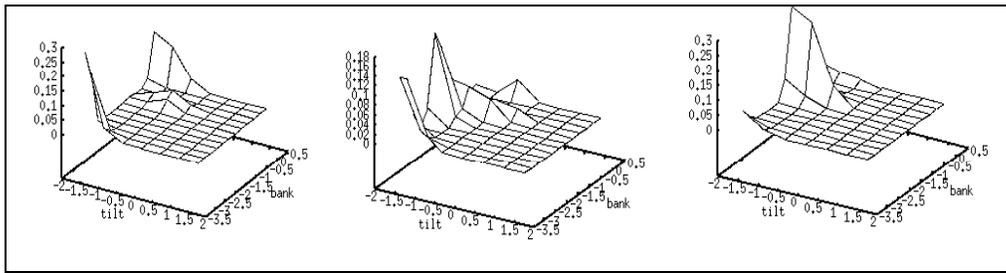


Figure 6.10 When LGM has access to axis orientation information the Hand Rotation neurons represent different plans in response to horizontal (left panel), diagonal (centre panel) and vertical (right panel) object presentations

In the full-vision case, we supply the orientation of the cylinder as an input to the LGM. Figure 6.10 shows the Hand Rotation neuron's distribution for different orientations, for a direct comparison with poor-vision case Figure 6.9. Now the neurons show *preferential activity* for different cylinder orientations. Also we can see that in the horizontal orientation case the LGM have two alternative plans as evidenced by the two peaks at the edges of the Hand Rotation neuron distribution shown in Figure 6.10, left panel. Note that these two alternatives correspond to grasping the horizontal cylinder from top and bottom (left panel, Figure 6.10).

Similar multiple plan representation is also observed in the diagonal case (Figure 6.10, center panel). The vertical cylinder case has a single grasp plan representation (Figure 6.10, right panel). We can relate the Hand Rotation neuron activities of poor-vision and full-vision cases. We can see that the poor-vision distributions are superimpositions of the three maps shown in Figure 6.10 with varying degrees of inhibition on areas where a common activity was not observed. For example there is activity around $(-\pi/2, -\pi/2)$ in all three neuron responses (Figure 6.10), and hence the peak appeared in the responses shown in Figure 6.9. However, the peak of activity $(-\pi/2, -\pi)$, which is shared by horizontal and diagonal orientations is not pronounced in Figure 6.9, because these neurons are not active when the object is vertically oriented (Figure 6.10, right panel)

6.7.2 Conclusions and neurophysiological predictions

Based on the wrist rotation and hand position maps formed via learning we can make predictions concerning the neural response properties of the premotor cortex. Up to now, we have used the term *Affordance* or *Input* to LGM without specifying from where this information can be relayed to premotor regions. Now we relate the emergent properties of the simulated neurons to affordances relayed.

The anterior part of the lateral bank of the intraparietal sulcus (area AIP) is implicated in extracting visual properties of objects relevant for grasping (Sakata et al. 1997a; Sakata et al. 1998; Sakata et al. 1995; Murata et al. 1996). There is strong anatomical reciprocal connection between area F5 and area AIP (Matelli 1994, Sakata, 1997). Furthermore some AIP visually activated neurons, show tuning according to the orientation of the longitudinal axis or the plane (surface) of flat objects (Sakata et al. 1999). Therefore, the axis information can be extracted by area AIP and channeled to area F5. The lateral bank of the intraparietal sulcus (c-IPS area) is involved in three-dimensional analysis of objects (Sakata et al. 1997a; Sakata et al. 1999). Some of these binocular visual neurons are selective for the orientation of the axis of the objects (AOS neurons). The c-IPS neurons may be the basis of AIP neuron properties (Sakata et al. 1997a; Sakata et al. 1999).

Analogous to the AIP-like object representation found in area F5 (Murata et al. 1997a), our model predicts that a premotor area that is connected with AIP or area F5 must have visually triggered neurons that are selective for *object orientations*. Furthermore, we suggest that the area must have control

over *wrist rotations*. Area F5 and F2 of the dorsal premotor cortex, satisfies the conditions we have listed. Fogassi et al. (1999) found that F2 neurons responded to visual stimuli and showed by microstimulation that the wrist joints were controlled within area F2. To our knowledge there is no study probing the object orientation selectivity of neurons in area F2 and F5. We suggest that the canonical neurons may have tuning fields according to the object orientation similar to the object type neurons with orientation tuning in area AIP (Sakata et al. 1999).

To summarize, we predict that a sub-population of F5 canonical neurons or F2 neurons that are involved in wrist control must have selectivity or tuning according to objects' grasp relevant properties. More specifically, we claim that a closer examination of F5 canonical neurons and/or area F2 neurons would reveal *orientation selectivity* either in isolation or as a tuning mechanism over object selectivity.

6.8 Object size selectivity

The goal of this simulation is to investigate the emergent neural properties when Learning to Grasp Model is presented with objects of different sizes. The object size is accessible to LGM via the Affordance layer and algorithmically coded (no object recognition and feature extraction is performed).

6.8.1 Simulation results

We used cubes of different sizes as objects. After the model interacted with the object, we identified the grasps learned. For large object case, the model often generated power grasps, and some variations. For small objects model generated precision grasps. The number of different precision grasp types was less compared to the number of different type of power grasping observed for big objects. In other words, LGM formed a large 'menu' of power grasps but gained a small 'menu' of precision grasps²⁰. After the behavioral observation, we analyzed the unit level activity of Hand Position layer. Wrist Rotation layer's output was modulated by Hand Position layer's activity (see the analyses in section 6.6).

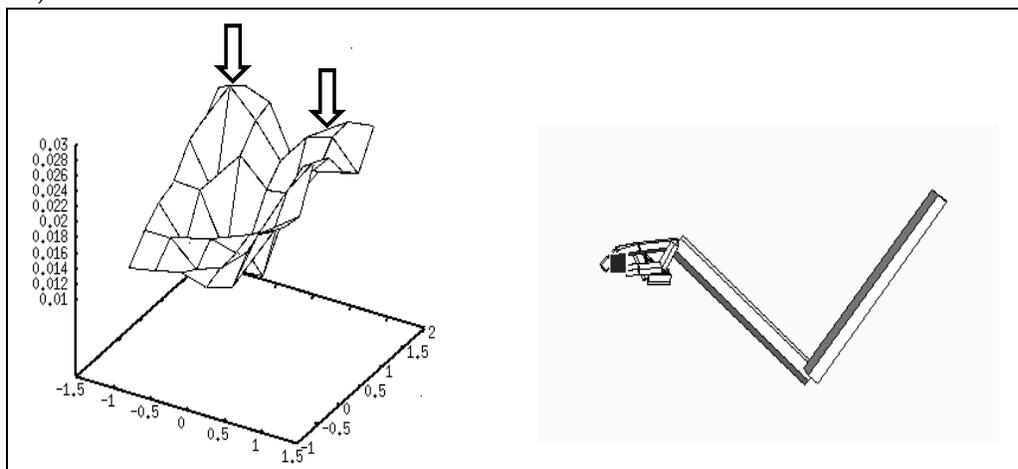


Figure 6.11 The small object presentation produced two peaks of activity in the Hand Position layer corresponding to the probability distribution of approach directions. The right panel shows the executed grasp when the data generation was localized in the area pointed by the leftmost arrow.

Figure 6.11 shows a learned precision pinch and the activity of Hand Position layer. The small object activated two peaks as indicated by the arrows in the figure. If a neurophysiologist recorded the activity of the units in Figure 6.11 (each point on the 3D mesh represents a neuron with its average

²⁰ Grasps with different wrist orientations are counted as different.

activity as the z-axis), he would notice over many trials that those units are the most active ones. To claim the specificity he still would need to compare the very same neuron with different object sizes. We do the same and show that the activity is object-size specific (compare with Figure 6.12 and Figure 6.13).

Figure 6.12 shows a power-like (radial palm grasp see Figure 5.1) grasp with the corresponding Hand Position layer activity. This time the hypothetical experimenter would notice the unit indicated with the arrow fired maximally for the presentation of this object. Note that the activity locus is different from the small object presentation (see Figure 6.11).

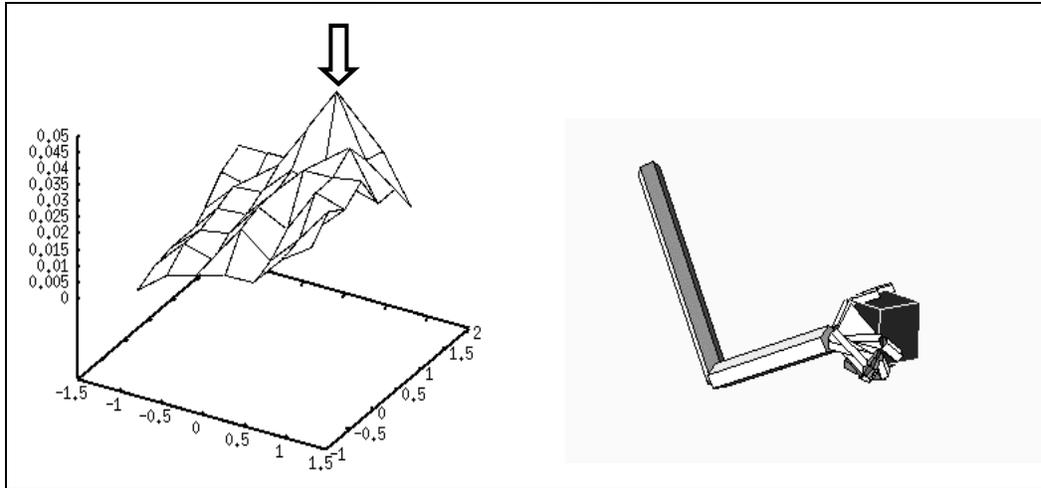


Figure 6.12: A large cube was grasped by securing the object between the thumb and the other fingers (right panel). The Hand Position layer activity is shown on the left panel. The neuron with largest activity is marked with an arrow

As the last simulation we present the unit level activity when the model presented with the largest object. The Hand Position layer showed a clear peak (Figure 6.13). In order to emphasize that these simulated neurons do not show unspecific activity, but rather are selective for object sizes, we compare the locus of the neurons that the arrows point to in each figure (Figure 6.13; Figure 6.12; and Figure 6.11)

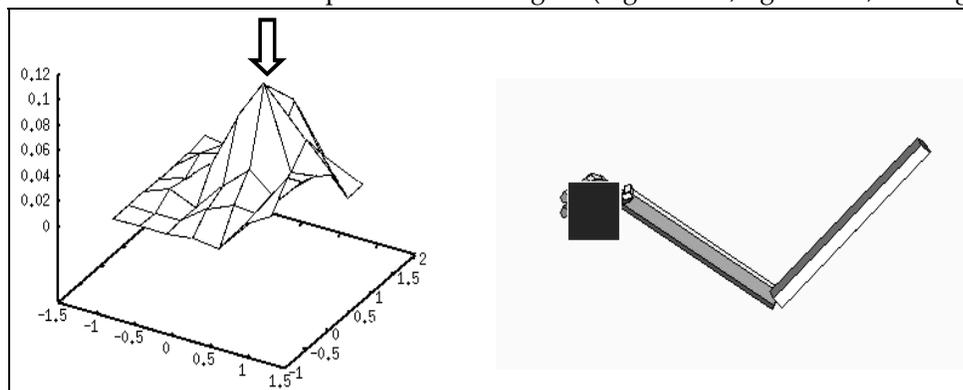


Figure 6.13 The largest object presentation and grasping. The Hand Position reflects a single reach direction as indicated with an arrow

Figure 6.14 shows the three activities superimposed after the axis are aligned. The maximum activity regions are also linked to the corresponding object presentation. Notice that the peak activity loci are not the same indicating that the neurons gained *selectivity to object size*.

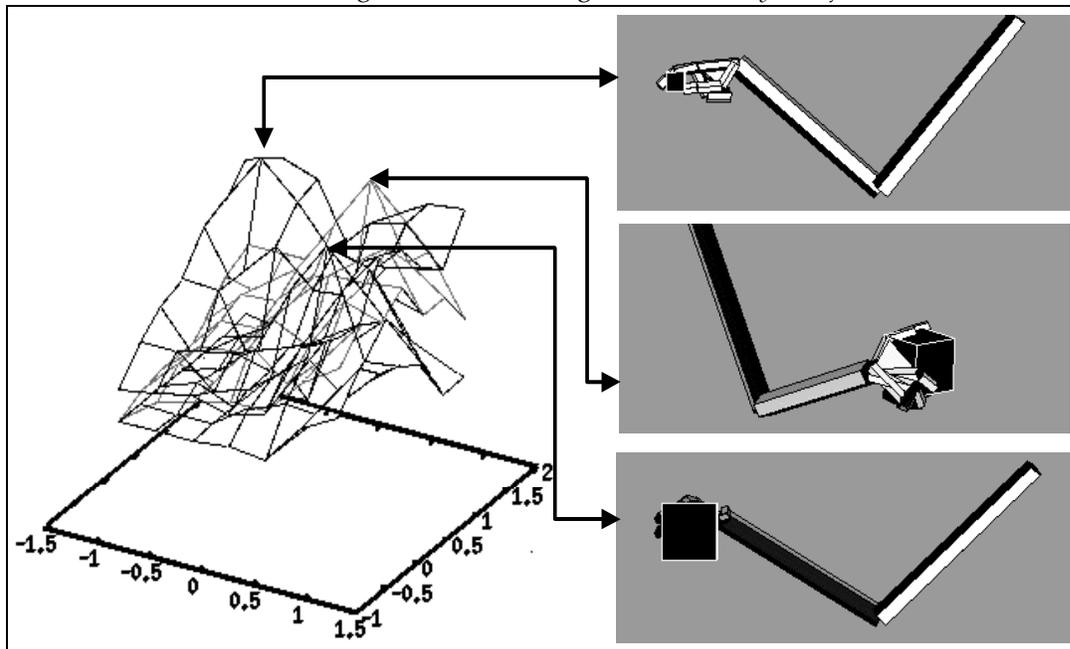


Figure 6.14: The Hand Position layer activity is superimposed to demonstrate that the maximum activity loci are separated for each object indicating *selectivity for object size*.

6.8.2 Conclusions and neurophysiological predictions

Even though we did not engage Virtual Finger layer in its full capacity (it controlled the fingers synergistically) the Hand Position layer activity of LGM showed object specificity. Therefore, we suggest new experiments to test whether F5 canonical object specificity is due to some other motor parameter encoding. Although microstimulation of F5 neurons produces complex movements they do not produce a grasping behavior, so the activity during grasping do not necessarily specify complete grasp plan but rather a subset of motor parameters required for grasping. This is analogous to the LGM layers differential contribution to the grasp plan. Therefore, the challenge to the experimentalists is to find out whether F5 canonical activity alone specifies a grasp plan or not. Our prediction is negative because, for example, in LGM simulations, the wrist rotations generated did not affect the activity of Hand Position neurons (analogous to F5 canonical neurons' object related activity) but the grasp plan was only complete with the Wrist Rotation layer's contribution. Thus, F5 canonical neurons can be just a part of the grasp planning circuit.

To decipher the functional structure of the premotor grasping circuit, this hypothesis must be either ruled out or validated. If it turns out that F5 neurons are specifying only a subset of grasp parameters, new experiments must be designed to uncover the neuron properties of the areas connected area F5. Since our model predicts that if some neurons specify a grasp plan partially (by being responsible for one parameter of the plan), the learning must shape the connections between the connected regions to enable the cooperative computation that can yield compatible grasp parameters. This is a nontrivial observation. if a hypothetical experimenter, investigating the encoding of grasp parameters, finds out that

1. Area A is responsible for wrist orientation
2. Area B is responsible for controlling the reach direction
3. If the activities of A and B are true control variables of wrist rotation and reach direction

Then, we claim that

1. Regions A and B must be anatomically connected (the connection may be a multi staged one, such as A connected to C, C connected to B)
2. The synaptic strength of the connections between A and B encode the grasp selection strategy of the grasping circuit
3. If the parameters are, in the physical sense, coupled then the dependency of parameters must be captured in the connection between A and B

The claim (3) is based on the observation that the only way to produce coherent grasp plans is to capture the physical coupling inherent in the movement. For example given a sphere, there are infinitely many approach directions, which can be used for grasping (assume A encodes approach direction). Similarly, there are infinitely many wrist orientations (assume B encodes wrist orientation). The crucial observation is that, not all pairs of approach direction and wrist orientation are compatible. For grasping, given A, B has to be determined or given B, A has to be determined because the parameters are coupled. With this setting, we further claim that:

(1) If it is found that the connection between A and B is unidirectional (say from A to B) then (from the view point of the experimenter) the activity of B neurons must be tuned by the population activity A

(2) Recording simultaneously from A and B can reveal the organization of visuomotor grasping circuit as the following. If activity of A predicts the activity of B, better than the other way round (i.e. activity of B predicts the activity of A) the underlying principle is that area A selects some of the grasp parameters and then B, based on the selection of A, generates the remaining grasp parameters

In terms specific predictions, LGM predicts that F5 canonical neurons, to specify a grasp plan, must have 'gain fields' based on other brain areas (which, behaviourally would appear to be based on other movement parameters). A good example for such systematic relation (unfortunately, to our knowledge only behaviourally) is the effect of arm posture on reach-related activity of motor cortex neurons (Caminiti et al. 1990; Caminiti et al. 1991; Sergio and Kalaska 1997). In these studies, it has been found that the reach encoding neurons have gain fields based on arm posture or position of the arm in the workspace. These findings, combined with multiple action representations in the motor areas (Cisek and Kalaska 2002) strongly indicate that the underlying circuit for reach generation have similar properties as LGM.

The structure (the dependence relations of layers) we offered in LGM may not be the only possibility. To uniquely determine the structure, neurophysiological studies must go beyond correlation studies. LGM postulates the dependency of wrist rotations on the virtual fingers and approach direction. Therefore, according to LGM, virtual fingers and approach direction predicts the wrist orientation. Then LGM predicts that the wrist orientation coding neurons will have 'gain fields' based on the virtual fingers selected and the approach direction determined. Whether LGM structure is right can be tested by experiments similar to reach and posture coupling experiments (Caminiti et al. 1990; Caminiti et al. 1991; Sergio and Kalaska 1997; Scott et al. 1997).

An alternative structure for LGM would be to posit that the wrist rotations with virtual fingers are determined first then based on those the approach direction is determined. This can be a feasible alternative because the wrist orientations are important in determining the manipulative freedom of the hand once the object is secured. For example when grasping the knob of a door with the intention of opening, we intentionally choose a 'hard' wrist orientation to have a large manipulative freedom (i.e.

turning the knob to open the door). However, if we assume that the intended future action plans are also relayed to wrist rotation layer (together with affordance of the object), LGM structure can account for such anticipatory grasp planning as well.

LGM presents a rich set of hypotheses that can be verified or falsified with clean-cut experiments. We now, suggest that simultaneous recording experiments for premotor regions of F2, F5 bank (canonical neurons), F4 and wrist rotation encoding neurons of motor cortex (F1) must be performed to understand the underlying principles of grasp related visuomotor computation. A feasible experimental setup would pick two pairs from (F2, F4, F5, F1) at a time and record from the two sites while the monkey is performing grasp actions directed to objects located at different positions with changing orientations. The simultaneous pair-wise recording of areas F2, F4, F5 and F1 would require 6 sessions (F2- F4, F2-F5, F2 F1, F4-F5, F4-F1, F5-F1). We claim that, with rigorous analyses of these recordings with reference to hand kinematics and object affordances (i.e. the location and orientation of the object, and the monkey's wrist rotations and approach direction) the structure of the cooperative computation underlying grasping can be revealed.

6.9 Generalization: learning to plan based on object location

Up to now, we presented simulations where the objects were located at a fixed position in the space. In general, wrist rotations and approach directions depend on the position of the object as well as intrinsic object properties. Here we concentrate on the object location, but the analysis is valid for object properties (as long as they are represented as population coded activities and relayed to LGM) such as the orientation of the object. In fact we had already used multiple orientation learning when simulating Lockman et al.'s (1984) experiments (sections 5.9 and 6.7). However, there, we did not perform extensive analysis to show the generalization of the learning to novel orientations.

A given wrist rotation and approach direction that yield a stable grasp, in general, does not necessarily yield a stable grasp when the object is moved to an other position as approach direction and wrist orientation parameters are coupled for a successful grasp. For example, we don't grasp objects that are located on the left of midline as the same way as we do when they are on the right side. In addition, reach component of different grasp actions yield different absolute wrist orientations because the absolute wrist orientations depend on the configuration of the arm. Thus, the grasp plan has to take into account the location of the object. In this section we presents results showing that LGM can learn to generate stable grasp plans for different locations and generalize well for the locations that has not been experienced before. The generalization property is important from a neural network perspective. A generalizing network does not need to *memorize* all possible grasp plans for each location and hence does not fail for novel situations.

6.9.1 Simulation results

We trained LGM model by randomly²¹ placing a sphere in the workspace and letting the model interact with the object. The egocentric location of object is encoded algorithmically in the Affordance layer using population coding. Note that there is evidence that objects are encoded in egocentric reference frames in the parietal cortex (Siegel 1998; Colby and Goldberg 1999) which projects to premotor cortex (Geyer et al. 2000).

After 10000 grasp attempts, the model acquired the ability to make grasp plans for objects located in the workspace. Figure 6.15 illustrates the learning achieved as superimposed images of completed grasps. Note that the object locations were not used while the model was learning to grasp. The lower

²¹ The positions were generated randomly on the surface of a invisible sphere centered on the shoulder

three grasp actions are anatomically hard to achieve but they were included for demonstrating the range of actions that LGM learned to grasp at. As in earlier simulations, the illustration in Figure 6.15 only depicts a possible grasp for each location out of many alternatives. We reran the trained model on exactly the same object locations as in Figure 6.15. A set of different grasping configurations is selected to show that LGM could both generalize and represent multiple grasp plans for each target location (see Figure 6.16).

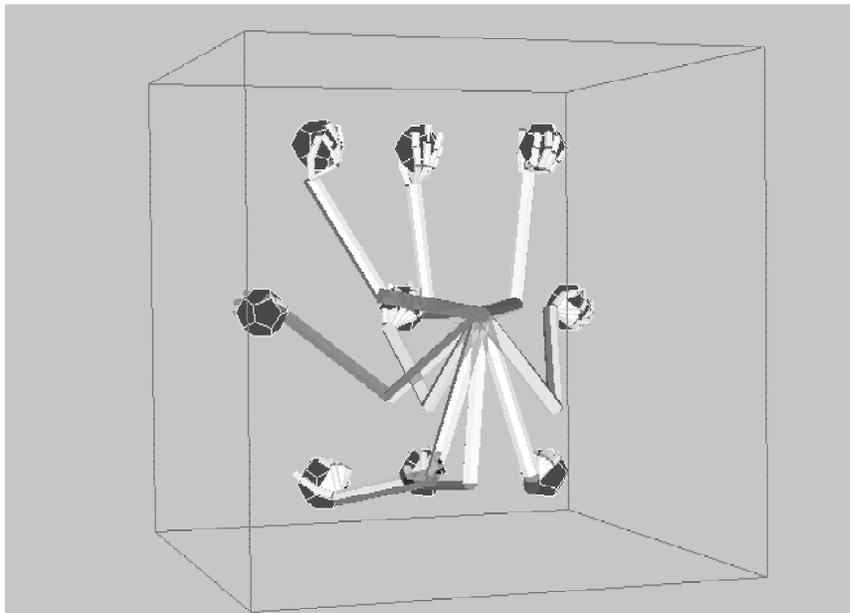


Figure 6.15 The trained Learning to Grasp Model executed grasps to objects located at nine different locations in the workspace. The grasp locations were not used in the training. All of the grasps shown were stable

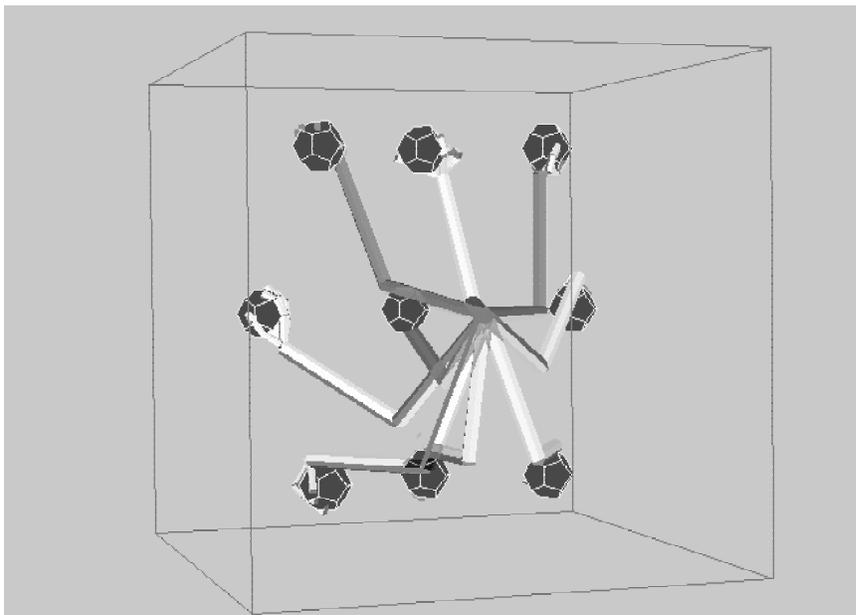


Figure 6.16 The same model used in generating Figure 6.15 was used to generate a different set of grasps. Again all the grasps were stable.

6.9.2 Summary and Conclusion

The crucial feature of LGM we implemented is that it can produce a variety of grasp plans based on input. Furthermore, the layer structure is well suited for further refining or biasing (e.g. contextual or motivational biases). For example, one can use different biasing when grasping an apple for eating or for placing it in a shopping basket. The reinforcement framework also gives the flexibility to include soft constraints in grasp evaluation such that the grasps that are not favorable (due to discomfort, excess energy consumption, etc.) are not represented or represented with low probability. Thus the reinforcement signal ('joy of grasping' of Chapter 5 and 'neural grasp stability' representation of this chapter) can incorporate the *anatomical and environmental constraints* which are important in shaping grasp development (Newell et al. 1989) or the *adaptive value* of Sporns et al.(1998)

We summarize the computational ingredients of grasp learning we proposed with LGM by tracing the neural level computations through an example.

We will use Figure 6.17 to illustrate how the same input condition can give rise to different grasp plans. First, the object is presented (bottom center). The object location is encoded in an egocentric reference frame using population coding (we use a spherical coordinate system) (the center plot). The object location is transformed into approach direction distribution (the top-center) by Hand Position layer. After approach direction generation, the center stream branches into two. On the left, the approach direction is generated as (approximately) from bottom, indicated by the top-left plot. On the right, the approach direction is even lower but it is from the backside of the object as can be read-off from the top-right plot.

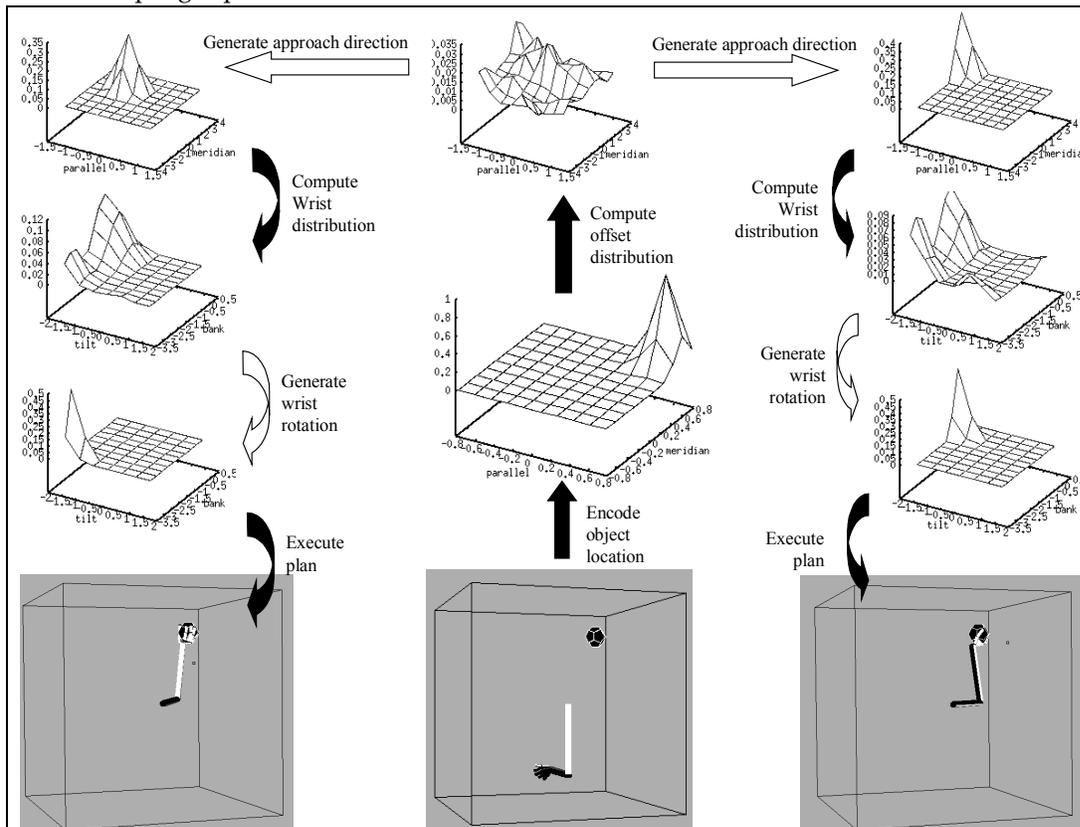


Figure 6.17 The internal mechanisms of representing and generating multiple grasp plans are shown. Solid arrows (except object encoding) denote learned connections while empty arrows indicate data generation. The

flow of operation starts with the presentation of the object (the bottom centre) and follows the arrows. At the top-centre, the data generation can yield multiple approach directions. The two possible approach directions are shown creating two streams (left column and right column), each of which yields different grasp execution (bottom pictures of left and right column).

Based on this two approach directions two different wrist rotation distributions are computed by the Wrist Rotation layer as indicated by the downward solid arrows on each side. This is followed by wrist rotation generation, which results in different wrist joint rotations in the left and right streams as indicated with the downward empty arrows. The grasp actions that are instructed by the left and right grasp plans would yield the completed grasps shown on the bottom-left and bottom-right columns respectively.

7 CHAPTER VII: BIOLOGICALLY REALISTIC F5 VISUAL SERVO CIRCUITS FOR GRASPING AND EMERGENCE OF MIRROR NEURONS

It has been argued that mirror neurons forms the basis of understanding other's actions (Gallese et al. 1996; Rizzolatti et al. 1996a; Rizzolatti et al. 2001a; Umiltà et al. 2001). The goal of this chapter is to bring an alternative view that this cognitive function may be secondary to a role for the mirror neurons in providing visual control signals for manipulation. Figure 7.1 shows the focus of this chapter. MNS model is redrawn and the regions of interest are marked with the gray background rectangle. In Chapter 3, we introduced the idea that the mirror neurons might be involved in visual feedback for manipulation, but in implementation, we emphasized how the self-action observations yield mirror neurons. In Chapter 6, we have studied how grasp alternatives can be formed and selected based on the object affordance and showed the emergent object feature selective properties of the simulated premotor neurons. In this chapter, for the sake of tractability, we use planar arm/hand model and focus on precision grasps. However, it should be emphasized that the simplification we make does not reduce the value of the message we wish to communicate to experimentalists. The point of this chapter is not to show the object selective properties of the F5 neurons, but to study the temporal aspects of grasping. Specifically, we show that units that encode visually defined grasp errors can yield activities similar to mirror neurons and suggest experiments to validate and challenge the visual feedback control hypothesis of mirror neurons.

7.1 Motivation

First, we present a biologically realistic *feedback* circuit composed of leaky integrators that can visually servo the hand to achieve grasping. Then we augment the circuit with a *feed-forward* controller that is composed of pattern matching neural units and present two alternative hypotheses that associate the visual control signals with mirror neurons. We demonstrate that the visual feedback and feed-forward grasping system can work with lower level motor control circuits by implementing a position and velocity (PD) controller which receives its desired trajectory from the visual grasp control model. Finally, we introduce a method to compare the controller unit activities with real mirror neuron recording data and suggest explicit experiments to validate or invalidate our proposal that mirror neurons are involved in visual control of grasping movements.

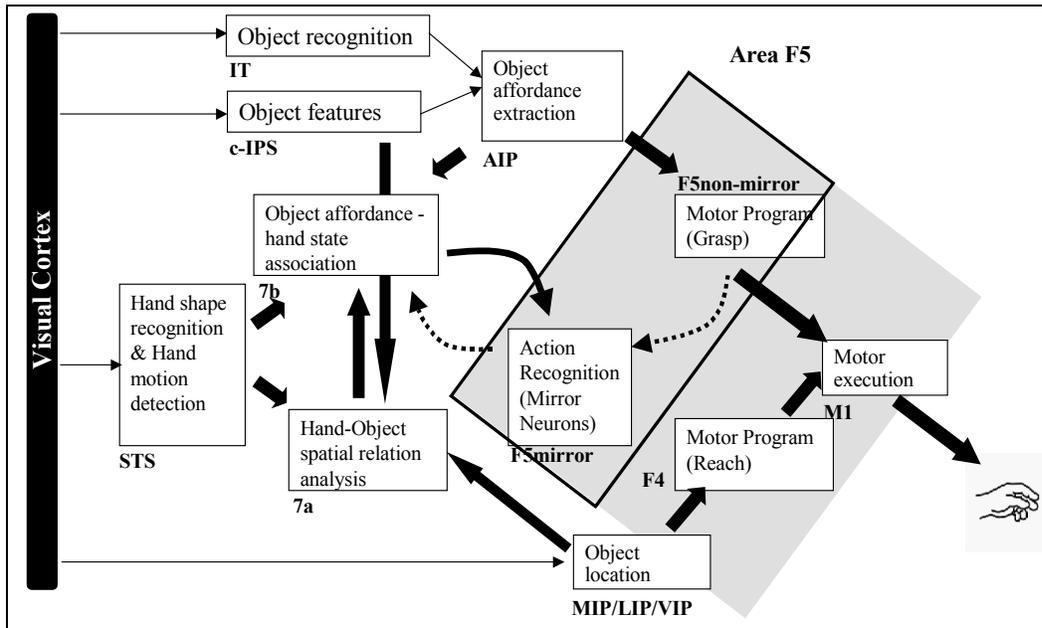


Figure 7.1 The MNS model repartitioned to show the focus of this chapter. The grey background marks area of interest

7.2 The link between the mirror neuron system and grasp learning

To our knowledge, up to now, all research on the mirror system has focused on the operation of a mirror system with the implicit assumption that it is developed for higher-level tasks (such as understanding and imitation), without exploring the possible biological precursors. MNS (Chapter 3) showed how observation of self-action may serve as the learning stimulus for shaping the mirror neuron system but did not address the issue of why the brain might contain such learning hardware. The standard answer is that it is there to help the animal recognize the actions of others by means of some similarity to its own actions. However, we wish to explore the hypothesis that the mirror system can best be understood through exaptation of a system for visual feedback control for manual actions.

For a reaching task, the simplest visual feedback is some form of signal of the distance between object and hand. This may suffice for grabbing bananas, but for peeling a banana, feedback on the shape of the hand relative to the banana, as well as force feedback become crucial. The parameters that are needed for such visual feedback have the ingredients of the hand state we used in MNS. We do not claim that the hand state we defined exists in the brain but we do claim that such hand-object relations must be represented in the primate brain. There are studies where neurons sensitive for observation of hands approaching to the points of attention or fixation have been found (Siegel and Read 1997). Also there are studies suggesting that the monkey parietal area have an allocentric representation of object locations, which simply indicates that the distance between objects can be encoded in parietal cortex (Murata 2000).

With Learning to Grasp Model (LGM) (Chapter 5), we showed that it is possible to *learn to generate* grasping movements based on available affordances. LGM is capable of generating precision grips, but it does not have the machinery to visually servo the fingers to their targets on the object. In other words LGM learns to generate grasp plans for open-loop control. We suggest that a visual feedback system develops in area F5 augmenting the LGM grasp machinery. The fundamental idea is that LGM

bootstraps F5 visual servo system by presenting examples of suitable grasps that trains the feedback system. In this chapter, we present two possible organization of the manual visual servo circuit and perform simulations using one.

Here we do not use the control term strictly. It should be understood that the command sent from area F5 are higher level than the control command sent to actuators of a robot. In fact, in robotics the proposed area F5 output would be considered as a *trajectory plan* for grasping which specifies the kinematics aspects of the grasp but leaves the dynamics to the lower levels. In the next sections, we address the dynamics of the reach and grasp by giving a 2D model of grasping. It is known that F5 projects directly via the corticospinal tract to motoneurons that control finger muscles (Dum and Strick 1991), these connections are not enough to perform a grasping action as the lesion of the primary motor cortex (F1) completely disrupts the grasp execution (Fogassi et al. 2001). Therefore, our assumption that the premotor controller outputs are higher-level signals (in the sense that they require a subordinate layer to interpret them) is in fact supported by neurophysiology.

7.2.1 Two Visual Control Hypotheses

Currently, there is no hand kinematics data synchronized with F5 firing. Therefore, it is not possible to reliably assert the roles of different neurons in area F5. Nevertheless, we propose two likely control structures for area F5. The first one (Figure 7.2) postulates that F5 canonical neurons mediate the controller development by priming or gating mechanism such that controllers are differentially associated with different affordances relayed by AIP. F5 neurons that are recruited by F5 canonical neurons form multiple feedback and feed-forward controller pairs (motor schemas or modules) and partition the task space with inter-module competition and F5 canonical guidance.

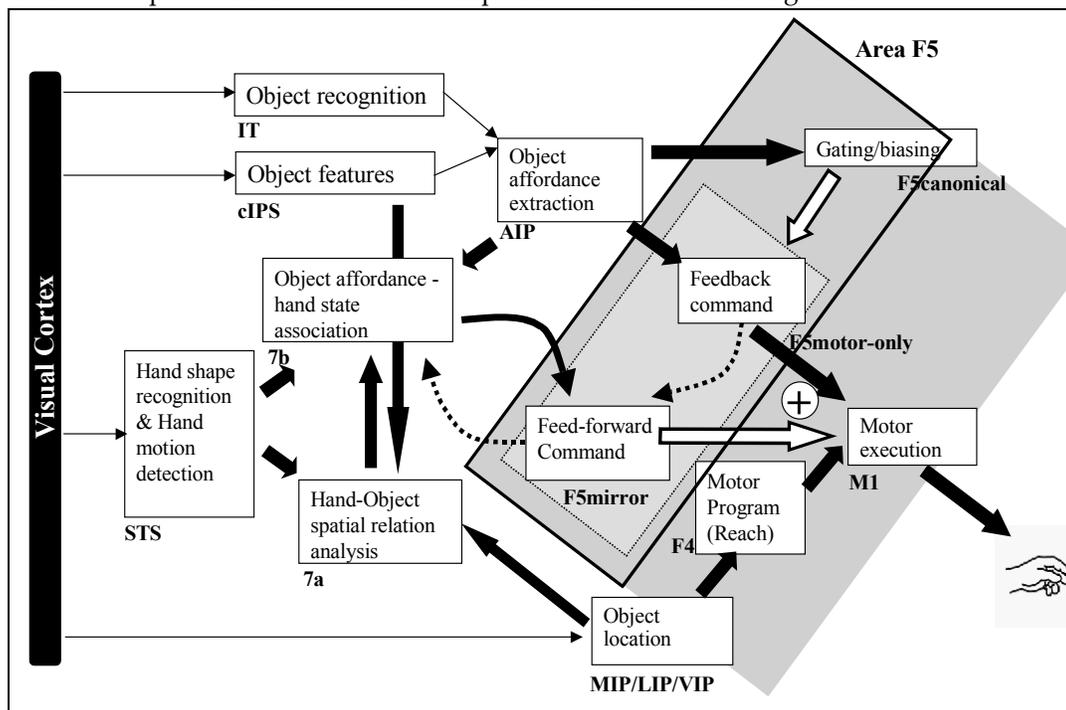


Figure 7.2 One alternative visual control structure for manipulation is shown within the MNS framework. The mirror neurons generate feed-forward commands

The complexity of each module depends on the number of modules that share the task. If the system to be controlled is complex, it is beneficial to have modularity, as a better overall control

performance can be achieved with controllers specialized for each separate task than with a single super-complex controller for all the tasks. Wolpert and Kawato (1998; Haruno et al. 2001) proposed the *multiple paired forward and inverse models* for this kind of control scheme. Figure 7.2 expands the MNS model in accordance to this view. Note that the F5 neurons are now, split into three: canonical, mirror and motor-only. The motor-only neurons send transient feedback commands to correct ongoing grasping movements while the F5 mirror neurons provide the motor command to achieve the final hand configuration or intermediate configurations as the subgoals of a grasping task. In the second alternative control structure (Figure 7.3), feed-forward command is generated by canonical neurons, and mirror neurons are implicated in feedback control of grasping.

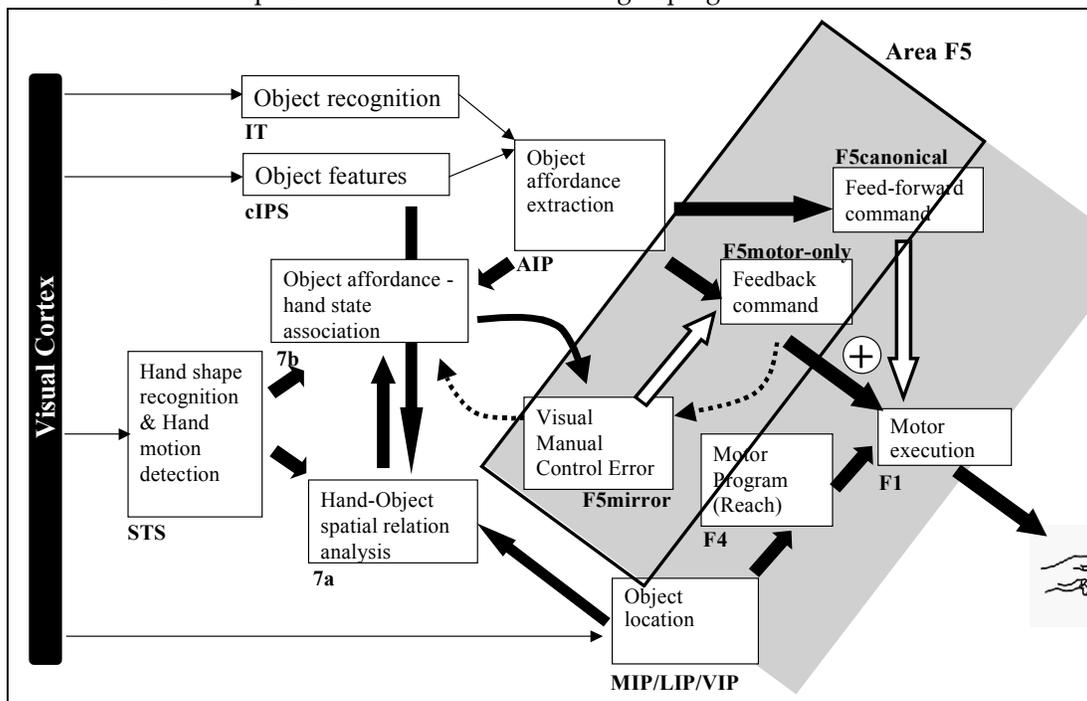


Figure 7.3 Another alternative visual control structure for manipulation is shown within the MNS framework (compare with Figure 7.2). The mirror neurons generate feedback commands

In both alternatives, the visual control task is to

1. Compute a *visually defined error* based on the current state and the desired state
2. Generate a command to reduce the error

From a developmental point of view, the questions one might ask is what the error and the desired state for the newborn are. Can we assume that from birth the circuits that compute visual error (e.g. distance to target) exists in the primate brain? We argue that the answer is negative and postulate that an LGM-like circuit trains the manual visual feedback control circuit. In intuitive terms, the animal learns how a successful grasp looks like by observing the performance of LGM-like circuit. Assuming the animal learns to extract desired state in visual terms by observing LGM,²² we can study how neural circuits in area F5 may use this signal to function as a manual visual servo circuit. For example, for the precision pinch schema, the desired state could be the contact of index and thumb fingers on the object

²² A desired state, given an object, could be simply a look up table of visual descriptions of a hand grasping the object

surface. The manually defined Hand State (Chapter 3) that monitors the relation of the hand to the object could serve as the desired state. However, in the general setting, we assume that the desired state is learnt from LGM based grasps.

7.2.2 Mirror neurons in feed-forward control (alternative I)

Figure 7.4, illustrates the basic manual visual feedback control circuit we propose for the structure that posit mirror neurons as feed-forward command units. The canonical neurons act as gating units; they select the appropriate feedback and feed-forward pairs based on the objects to be grasped. Based on the F5 neurophysiology, we can suggest that there exists a single controller module (schema) for different type of grasps (e.g. precision, power etc.).

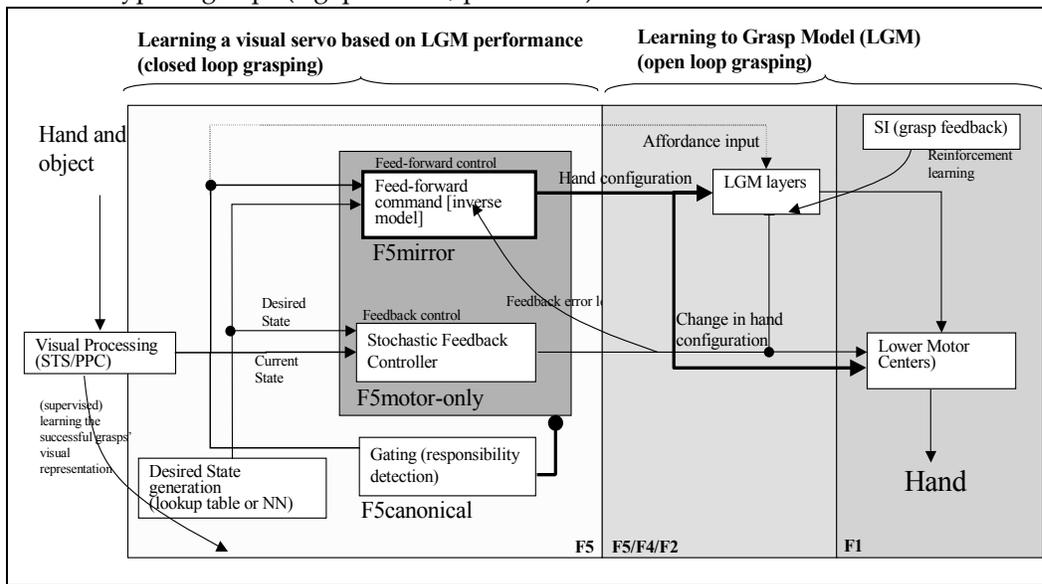


Figure 7.4 The feedback and feed-forward control view of the F5 grasping circuit, alternative I: F5mirror neurons learn to generate feed-forward command. The desired state is assumed to be available and is converted to a correction motor command by F5motor-only units using stochastic gradient descent. F5canonical neurons gate the feed-forward and feedback pairs.

The feedback controller receives the desired state and the current state. The aim of the servo circuit is to generate signals (for simplicity we will assume the signals determine change in hand configuration), such that the new state gets closer to the desired state. Note that we don't define the state. For MNS model, we defined a common state (hand state of Chapter 3) which could be used for all the grasps that are accounted, which may not be a parsimonious choice for the brain. Here each module may learn its own state. In intuitive terms, it means that each module, for its input can use the aspects of the action they are in charge of controlling. The key point is that the state depends on the *vision* of the *hand* and the *object*. The simplest of such a state is the distance between the hand and the object, which is a good choice for a pointing task servo circuit, but a bad choice for a grasping servo circuit since it is clearly not enough to control a hand to perform a grasp action based on a single scalar parameter. There is psychophysical evidence that such visual servo systems do exist in man. Ghahramani et al. (1996) in a psychophysical experiment, limited the visual feedback of finger position at one or two locations in the workspace, where a discrepancy was introduced between the actual and visually perceived finger position. The remapping induced changes in pointing task were largest near the locus of remapping and

decreased away from it. This pattern of pointing disturbance suggests that visual feedback is used by the reach related visuomotor circuit of human during pointing.

7.2.3 Mirror Neurons in feedback control (alternative II)

Figure 7.5, illustrates the basic manual visual feedback control circuit we propose for the structure that posits mirror neuron involvement in feedback control. The outputs of mirror neurons are converted to motor commands by F5motor-only neurons. The feed-forward command is generated by canonical neurons. Note that the output of mirror neurons cannot be interpreted as a feedback command in a trivial manner. Most mirror neurons increase activity during the approach phase where the *grasp error* is decreasing. It can be argued that such neurons do not affect the overt movement because their targets in area F1 are inhibited. Another explanation could be that the mirror neuron activity works via inhibition, however this is unlikely because the microstimulation of F5 triggers movement. We propose that mirror neurons, at any given instant, keep an error map based on their preferred action (e.g. precision grasp) and visual stimuli.

Thus, the mirror neurons provide the *feedback error* (for their preferred actions), on which a feedback command can be generated. It is up to some other neuron population (e.g. F5 motor-only neurons) to produce a corrective command based on the population of mirror outputs relevant for the object to be manipulated.

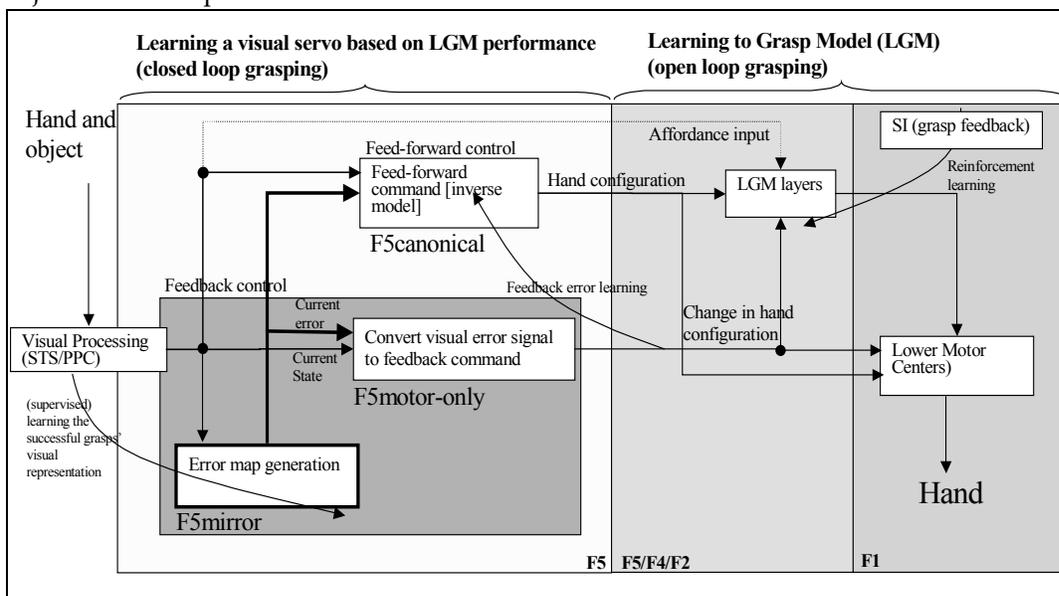


Figure 7.5. The feedback and feed-forward control view of the F5 grasp circuit, alternative II: F5mirror neurons learn to compute the error. The error is then converted to a correction motor command by F5motor-only units. F5canonical generates the feed-forward command signal.

7.2.4 The target of implementation

We based our choice on the following observation. If F5 mirror neurons were lesioned in 'alternative I' (the visual servo circuit that employs mirror neurons as feed-forward elements -Figure 7.4), the model would fall back to feedback only mode and could perform the action but with the lack of synergy of having the target configuration (the feed-forward signal). Analogously, it has been shown that inactivation of F5 mirror neurons in the monkey does not abolish the grasp but merely slows it. On the other hand, when F5 canonical neurons are inactivated the grasps that require precision cannot be executed. This agrees with 'alternative I' that F5 canonical neurons select the modules to be used in the

grasping tasks where precision is required (so that a feedback control is necessary). Therefore, we favor F5 control structure that assigns the mirror neurons the role of feed-forward motor command generation²³(Figure 7.4).

7.2.5 The visual servo task

In the next section, we propose a biologically realistic visual feedback circuit that can work as an autonomously visual servo circuit for *reaching and grasping*. It must be noted that, the computational background for this task is well established in robotics, although we do not use it here because of biological implausibility. The required change in controller output (change in hand configuration) in order to get closer to the desired state can be achieved using resolved motion rate control techniques from robotics (Klein et al. 1995; Whitney 1969). The computation requires the determination of an input dependent matrix to convert the (desired state-current state) vector direction to the correct gradient to be followed by the control output. Once the module learns this relation, the hand configuration changes for the precise positioning of fingers becomes possible. The visual servo circuit that will be presented follows the gradient stochastically. However, stochastic or deterministic, the gradient following requires *constant monitoring of the hand and object* to calculate the instant configuration changes.

It can be beneficial to determine the target configuration with *one-shot* computation (feed-forward control). Adopting the 'alternative I' (discussed in section 7.2.4), we involve mirror neurons in performing the feed-forward control function; that is given an object and hand in action the mirror neurons report the hand configuration required to achieve a grasping configuration. Once the mirror system has learned an action the animal can act faster, more accurate and in a more robust-to-perturbation fashion using both feed-forward and feedback control outputs.

7.2.6 The feed-forward model learning

The feed-forward controllers are trained by observing the successful grasps performed by the feedback controllers. The Reach and Grasp schema of Chapter 3 used inverse kinematics techniques for implementing the feedback controller. There the desired states were defined differently for each grasp type. The desired state information was defined algorithmically. With the LGM, we are one step ahead. While LGM learns and performs exploratory grasp plans, the feedback controller(s) can learn the desired states (successful grasp plans, even if rare) and the relation of input (visual) states to the output (motor) commands. Then, the learning in the feed-forward controllers can be accomplished in a biologically plausible way using feedback error learning (Gomi and Kawato 1993) since the error is readily available as the output of the corresponding feedback controller. One alternative to feedback learning is using 'distal learning' (Jordan and Rumelhart 1992) approach where first the forward model is trained and then the inverse model is trained using the error that is propagated backward through the forward model. However, distal learning is less biological because of the error back propagation. This relates our discussion to the *multiple paired forward and inverse model* architecture introduced by Wolpert and Kawato (1998; Haruna et al. 2001).

7.3 Implementation: F5 manual visual control circuit for 2D arm

This section presents the visual control circuit implementation outlined in Figure 7.4. The next section, focuses on the feedback controller module of Figure 7.4 and present a biologically plausible model in the sense that it can be implemented by leaky integrator units. We sidestep the complexity of

²³ However the activation of mirror neurons, as spelled in the next sections, is based on an explicit error map that coincides with the alternative structure's (Figure 7.5), with the exception that the error metric is manually chosen, not learned as it would be in the alternative structure.

the 3D arm that we used in earlier chapters by switching to a 2D arm model, but we require the system to learn the reach component as well as the grasping component. In Chapter 6, we have seen that learning the affordance-grasp associations produced selective populations for different grasps. Here without loss of generality, we limit ourselves to precision pinch grasps and study the control aspects based on kinematics properties of the reach and grasp movements.

Specifically we

1. Present a visual servo circuit using leaky integrators that can reach and grasp (section 7.3.1)
2. Propose a memory based feed-forward module and expand (1) into a feedback and feed-forward control system (section 7.3.4)
3. Present simulation results demonstrating that (1) is effective
4. Present simulation results with a dynamics model of the arm showing that (1) and (2) improves controller performance

The results (3) and (4) are meant to convince the reader that the proposed visual controller based on leaky integrators (feedback) and memory-based neural units (feed-forward) is adequate as a visual servo circuit. The readers interested in the main results concerning the mapping between the mirror neurons and the controller can skip to section 7.4 where we compare feed-forward units activity with real mirror neuron recordings.

7.3.1 A leaky integrator model for F5 manual visual feedback circuit

For a feedback controller system, a desired behavior or an error signal showing the deviation from the desired behavior is required. For the grasping the desired behavior, is grasping the object. In this model, we assume that the visual error signal is available for defining the how close the hand is to its final grasping configuration. A neural circuit can learn to generate the error signals based on the observed successful grasp examples. However, we do not model the learning of error signal generation and concentrate on how the error signal can be used to perform feedback based grasps.

From a computational point of view, the visual error signal has to be converted into motor error to deterministically control the behavior of the system. In this model, we take a stochastic gradient approach where the system determines its control output based on the commands it has sent in the previous time steps. One of the goals of the model is to give a simple but autonomous neural circuit that is open to neural implementation. Therefore, we avoid computations requiring programming constructs such as if-then, but instead use differential equations to describe the system. Figure 7.6 shows the schema level view of the feedback controller. The visual processing encapsulates the process of extracting a visual error based on the vision of the hand and the object. Lower Motor Centers encapsulates the functionality involved in transforming the motor signal sent by the feedback controller into neural signals sent to muscles. In later sections, we augment and detail the circuit to include the feed-forward module.

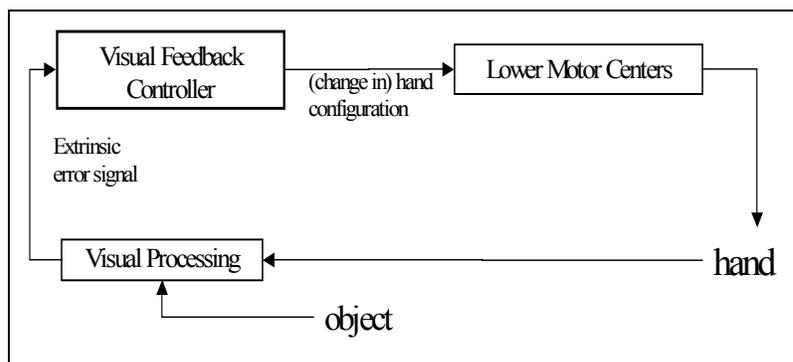


Figure 7.6 The schema level view of the feedback controller. The visual processing encapsulates the process of extracting an error based on the vision of the hand and the object. Lower Motor Centers encapsulates the functionality involved in transforming the motor signal into actual commands sent to muscles

The feedback controller sees outside world through its input (the extrinsic error signal) and affects the environment through its output (the change in behavior) signal. For the task of grasping, we define the output as the change of joint angles in the arm-hand.

The error signal is taken as the distance of the fingers used in grasping to the target location on the object. In general, an object affords multiple grasps. However, the multiple affordances are not a concern here since Chapter 5 showed how a menu of affordances can be formed and selected (see also Fagg and Arbib 1998 for prefrontal influences on grasp selection) based on the object affordances. The hand/arm model is a planar one and comprised of three links, a thumb with 1DOF an index finger with 2DOF. We only consider precision type grasp with varying aperture sizes²⁴, which is defined by two points on the object as targets for the fingertips (see Figure 7.8). The error signal²⁵ is then, a four-dimensional (two planar coordinates) and the output is a five-dimensional (five joint angle changes) vector. The computational elements we use are biologically realistic (e.g. integration, summation and shunting). We use leaky integrators as neural units. The representation we use is rate coding. The output of a leaky integrator unit defines single parameter in contrast to the population coding we used before. However, we can use single units without loss of generality because in theory we can expand a unit's activity to a population activity and apply the single unit equations to a population with proper weight coefficients. Figure 7.7 presents the detailed circuitry of the feedback controller module. The figure represents the differential equations and the operations in a schematic form. The output of the feedback controller is the change in joint angles ($\Delta\mathbf{q}$), and the error signal that drives the circuit is the sum of the distances of the fingers to their targets ($\mathbf{e}(\mathbf{t})$). The working principle is based on stochastic gradient descent. The network tries a random move; if the move was efficient in reducing the error then it is more likely to make a similar move. If it was a bad move, the movement is backed up. Although the stochastic gradient method is algorithmically very simple, it requires some care for implementing it as a set of differential equations. The merit in implementing an algorithm in terms of leaky integrators is to show that the computation can be implemented in the brain. Once this is shown, the algorithm can be encapsulated and used as a computational block (schema) in designing other brain circuits much like

²⁴ The simplifications we make in this chapter does not decrease the value of the message we would like to communicate to experimentalists. The point of this chapter is not to show the object selective properties of the F5 neurons, but to study the temporal aspects of grasping.

²⁵ The feedback circuit uses the Euclidean norm of the error signal while the vector form is used by the feed-forward module

the winner-take-all circuit (see Arbib 1989, chapter 4.4 for a leaky integrator implementation and mathematical analysis).

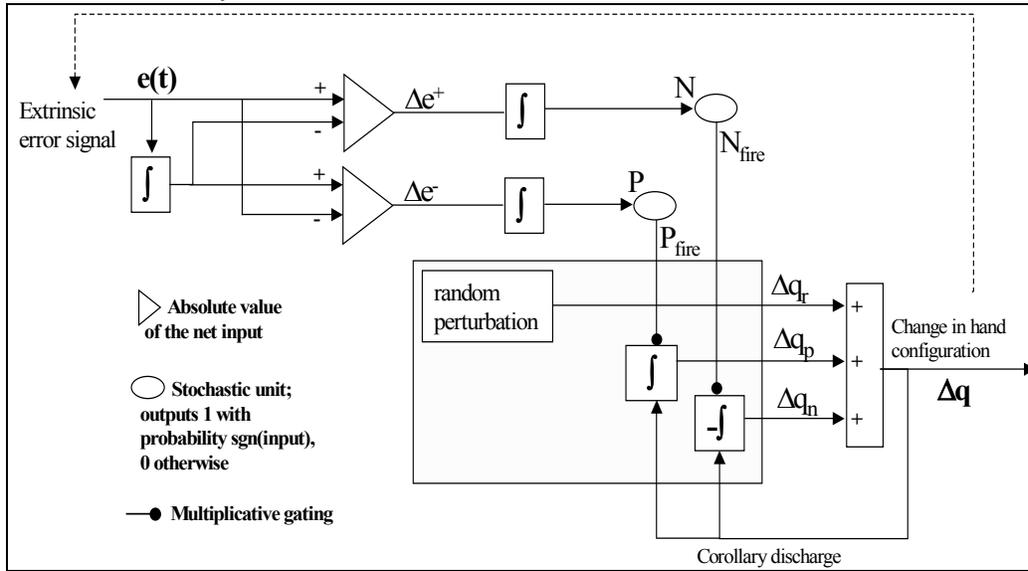


Figure 7.7 The leaky integrator implementation of the feedback circuit that solves the inverse kinematics problem for precision grasping. See text for the explanation

The key stochastic elements of the circuit are the firing of the units P and N (*random perturbation* can be considered as a background noise). The units P and N are leaky integrators where P follows the value of $\Delta e^-(t)$, whereas N follows $\Delta e^+(t)$. The variables $\Delta e^-(t)$ and $\Delta e^+(t)$ keep track of the change in error. If the recent commands increased the error $\Delta e^-(t)$ charges up, whereas if the recent commands decreased the error $\Delta e^-(t)$ charges up.

$$\tau \frac{dP(t)}{dt} = -P(t) + \Delta e^-(t)$$

$$\tau \frac{dN(t)}{dt} = -N(t) + \Delta e^+(t)$$

Both parameters $\Delta e^-(t)$ and $\Delta e^+(t)$ are non-negative and have zero resting potential level. The change in error is monitored with auxiliary parameter $e^{\text{delayed}}(t)$ which follows the error with a lag determined by the time constant κ .

$$\kappa \frac{de(t)^{\text{delayed}}}{dt} = -e(t)^{\text{delayed}} + e(t)$$

$$\Delta e^+ = \frac{1}{2} |e(t) - e(t)^{\text{delayed}}| + \frac{1}{2} (e(t) - e(t)^{\text{delayed}})$$

$$\Delta e^- = \frac{1}{2} |e(t) + e(t)^{\text{delayed}}| - \frac{1}{2} (e(t) - e(t)^{\text{delayed}})$$

At any instant, $P_{\text{fire}}(t)$ and $N_{\text{fire}}(t)$ are either 0 or 1 corresponding to the firing of the units P and N respectively. The probability of each unit's firing is determined as the following.

$$\Pr(N_{fire}(t) = 1) = \frac{1}{1 + e^{-\beta N(t)}}$$

$$\Pr(P_{fire}(t) = 1) = \frac{1}{1 + e^{-\beta P(t)}}$$

β is the steepness parameter determining how quick the circuit is likely to respond to a change in error. This parameter is not critical; however, we empirically found that setting it to a high value (~ 100) results in fast convergence. One noticeable fact is that, unlike many stochastic units (Hertz et al. 1991), the probability is always greater than 0.5. The output of \mathbf{P} and \mathbf{N} units gate the three channels $\Delta\mathbf{q}_r$, $\Delta\mathbf{q}_p$, $\Delta\mathbf{q}_n$ that sum up to give the net $\Delta\mathbf{q}$. These are defined as:

$$\Delta q_r(t) = \text{random} - \text{perturbation}$$

$$\Delta q_p(t) = P_{fire}(t)\Delta q_{old}(t)$$

$$\Delta q_n(t) = N_{fire}(t)\Delta q_{old}(t)$$

The $\Delta\mathbf{q}_{old}$ keeps a history of the net output of the controller ($\Delta\mathbf{q}$) with a time constant of ς .

$$\varsigma \frac{d\Delta q_{old}(t)}{dt} = -\Delta q_{old}(t) + \Delta q(t)$$

The net output is simply given by the summation of three sources. Note however that $\Delta\mathbf{q}_n$ is non-negative and therefore its effect is reflected with a negation.

$$\Delta q(t) = \Delta q_r(t) + \Delta q_p(t) - \Delta q_n(t)$$

The net output intuitively interpreted as the following. If both \mathbf{P} and \mathbf{N} fired at the same time, claiming that the error both increased and decreased, the net output is mainly determined randomly (by $\Delta\mathbf{q}_r$). If \mathbf{N} fired but \mathbf{P} did not fire, meaning that it is likely that the error was increasing due to recent commands, then we try to undo the last commands ($-\Delta\mathbf{q}_n$). On the other hand, if the error was decreasing (\mathbf{P} fired but \mathbf{N} did not fire) then we try to repeat what we did before. The final case is when neither \mathbf{N} nor \mathbf{P} fires, which results in a random decision ($\Delta\mathbf{q}_r$).

7.3.2 Simulation: visual feedback control with leaky integrators

We implemented the proposed circuit using Matlab. For this section, the arm was implemented as a kinematics chain without dynamics. The reason was we wanted to avoid the coupling of the dynamics of the arm with the feedback circuit's internal dynamics. However, in later sections we introduce dynamics as well. For solving the differential equations that define the circuit, we used Euler integration with time step 0.01. The command is sent to (and in the kinematics case, applied immediately to the arm's configuration) at each 0.1 time units.

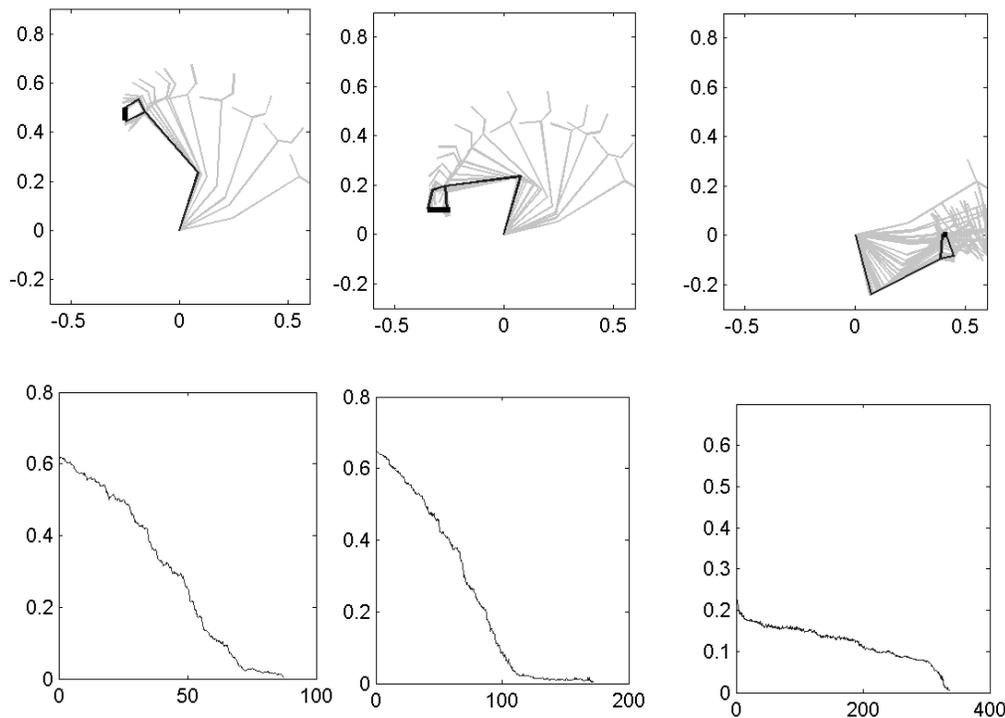


Figure 7.8 Three grasping tasks executed by the feedback circuit proposed shown on the upper half of the figure. The change of arm/hand configuration during the execution is illustrated by snapshots of the arm/hand. Each hand figure is accompanied (lower half) by the error plot. The grasp execution is stopped (success) when the sum of finger distances to their target was less than 2mms.

The Figure 7.8 show some example grasps performed together with the error plot. Each tick in the time axis of the error corresponds to 0.1 time units. For example, the leftmost grasp required less than 100 commands while the rightmost *harder* grasp required more than 300 commands. The rightmost grasp is harder because on the way to the target there are local minima to be overcome.

7.3.3 Simulation: feedback and lower motor centers

From a control theory point of view, it is important to ask the question how the *lower motor centers* will work with the feedback circuit we proposed. To answer this question we first need to implement a dynamics model of the arm and then design a motor controller, which tracks a given trajectory. Simplest such a controller is a PD controller with suitable gains. Note that we implicitly assume that PD controller is encapsulated in the lower level grasp schema that we have postulated to receive the output of the visual servo output signals. Figure 7.9 shows the control system we implemented. For the arm, we assumed that hand's contribution is negligible and set the upper limb length as 0.25 meters and weight as 4kgs, and the distal limb length as 0.35 meters and weights as 3kgs. Assuming cylindrical links and homogeneous mass with no gravity the 2D arm dynamics can be given by (Sciavicco and Siciliano 2000):

$$\begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} = M \begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix} + F \begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \end{bmatrix}$$

Where \mathbf{M} is the inertia matrix and \mathbf{F} is the matrix of Coriolis and centrifugal forces. Note that we don't include a gravity term since we assumed no gravity. q_1 and q_2 are the joint angles of the shoulder and

the elbow respectively. τ_1 and τ_2 denotes the corresponding torques. Single and double dot notation refers to usual first and second time derivatives. The inertia matrix and the matrix of Coriolis and centrifugal forces are given by:

$$M = \begin{bmatrix} \frac{m_1 l_1^2 + m_2 l_2^2}{12} + \frac{m_1 l_1^2 + m_2 (4l_1^2 + l_2^2 + 4l_1 l_2 \cos(q_2))}{4} & \frac{m_2 l_2^2}{12} + \frac{m_2 (l_2^2 + 2l_1 l_2 \cos(q_2))}{4} \\ \frac{m_2 l_2^2}{12} + \frac{m_2 (l_2^2 + 2l_1 l_2 \cos(q_2))}{4} & \frac{m_2 l_2^2}{12} + \frac{m_2 l_2^2}{4} \end{bmatrix}$$

$$F = \frac{m_2 l_1 l_2 \sin(q_2)}{2} \begin{bmatrix} -\dot{q}_2 & -(\dot{q}_1 + \dot{q}_2) \\ \dot{q}_1 & 0 \end{bmatrix}$$

The PD controller loop has to be faster than the visual feedback controller loop. Otherwise, the PD controller will not be able to track the desired joint angles specified by the visual feedback controller. In our simulations, we used 250 Hz. for the PD controller loop and 10 Hz for the visual feedback servo.

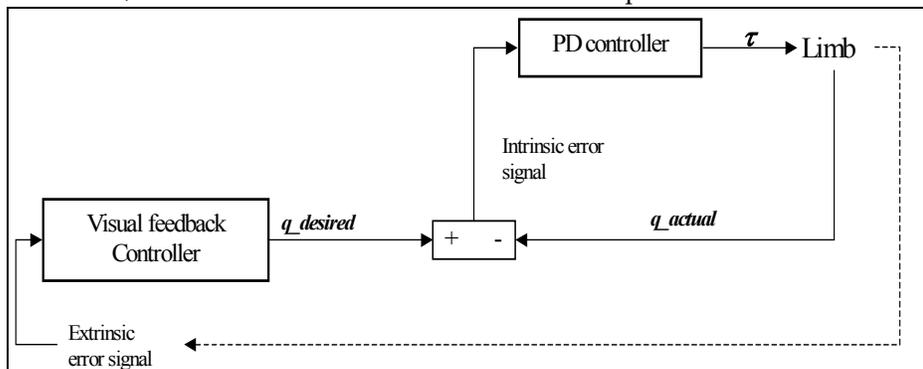


Figure 7.9. The Visual feedback circuit generating desired trajectories for the 'lower level motor centre' (implemented as a PD controller)

The cycle rate of the PD controller becomes more important when tracking high-speed trajectories. The PD gains can be selected to reduce tracking error, but the cycle rate of the controller brings a limitation on that also. In this simulation, the position gain was chosen as 30 and the velocity gain was chosen as 10. In Figure 7.9, this speed effect can be observed. The bottom half of the Figure 7.10 shows the tracking error and the followed path when the action took 2 seconds whereas the upper half shows similar graphs when the action took 0.5 seconds. When the action was slow, the combination of visual feedback circuit and the lower motor centers can work well (Figure 7.10, lower half). However, for higher speeds the circuit is not very effective for precise guiding of the hand (Figure 7.10, upper half).

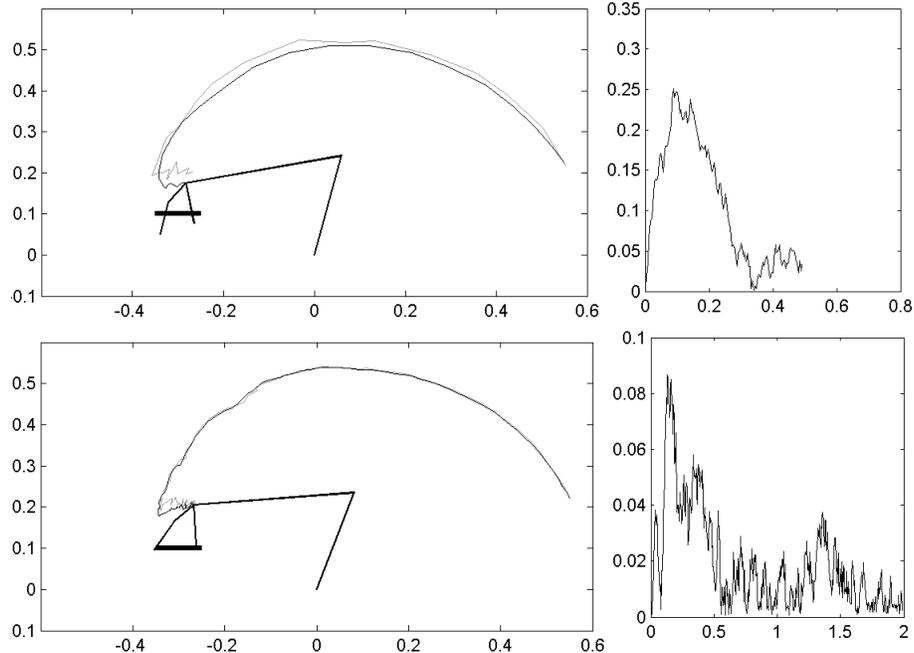


Figure 7.10 The slow (2 seconds) (lower half) and fast (0.5 seconds) (upper half) performance of the ‘visual feedback servo’ + ‘PD controller’ system is shown. The right hand side graphs show the tracking error (of the wrist) versus time. In the slow case, the object can be grasped but in the fast case, it is missed

It is desirable to generate desired trajectories in ahead without requiring a correction afterwards. In our simple visual feedback controller the system has to make a move; if it is not good it has to undo it and perform another move and so on. The desired trajectory produced in this way is not very smooth and would suffer from the feedback delays in a real robotics or biological system. The system we had could work pretty well (Figure 7.10) because we assumed the feedback signal has no delays. Having only the feedback signal as the corrective mechanism has severe consequence on the controller performance when there is delay in feedback loop. To overcome this we can use a feed-forward controller and issue the right command without requiring the feedback signal. The perturbations that cannot be accounted for (external perturbations for example) have to be still taken care of by the feedback controller. In the next section we augment the visual feedback servo circuit with a feed-forward module addressing one of the important elements of the visual control of grasping circuit we proposed (Figure 7.4) and show how the feed-forward units can behave as mirror neurons.

7.3.4 F5 Feed-forward visual control and mirror neurons

In this section, we augment the feedback controller of the previous section with a feed-forward module. The feed-forward module is formed by neural units that are selective to visual grasping errors. During any grasping action a number units activate based on their match on the object-hand relation. The one that fires maximally specifies the current feed-forward command. We define the feed-forward command in the visual control framework and introduce the mirror neurons as candidate for feed-forward command formation.

We present our design of the circuit and present simulations performed using the planar arm/hand we used in the previous section. Importantly we take a bold step and try to establish a link between F5 mirror neurons by comparing the activity of real F5 neurons with the simulated feed-forward units.

7.3.4.1 Mirror neurons, inverting actions ?

The natural way to construct a feed-forward controller is to *invert* the controlled plant. If the plant is characterized by $x=f(u)$, where x is the behavior of the system and u is the command used to manipulate the plant, then if we can compute²⁶ $f^{-1}(x)$ we will have a perfect feed-forward controller, because the analytic relation $u=f^{-1}(x)$ unambiguously tells us what command u is required to achieve a desired behavior. Now let us motivate why F5 mirror neurons can be involved in such an inverse computation. We know that F5 mirror neurons that are involved in encoding motor plans are activated by the observation of similar actions (Gallese et al. 1996; Rizzolatti et al. 1996a). An inverse computation in a visuomotor task is intuitively defined by: ‘given a visually defined desired behavior what is the required motor representation?’ Thus, if the mirror neuron activity represents some part of a motor plan then we can view them as elements of a feed-forward control system. Assuming a visual feedback controller exist as we described in the previous section, then even without F5 mirror neurons we can expect to have visual control of grasping with limited degradation (especially in slow movements the delays in feedback becomes less severe). Indeed, in the study of (Fogassi et al. 2001) even though the grasping movements were not abolished when F5 mirror neurons were inactivated, there was a *slowing* in the grasping movements but the hand reaching and preshaping were intact. Although we argue in favor of F5 mirror region’s being an inverse model we need to emphasize that a single mirror neuron itself cannot be an inverse model for a visual control task since the single neuron activity of a neuron is much too variable to encode a precise action. The activity of mirror neurons can be strictly selective for the type of the grip used (e.g. precision or power grasp) or broad (Gallese et al. 1996; Rizzolatti et al. 1996a). If the mirror neurons encode motor plans, the plan must be represented in a distributed fashion.

Now we turn back to our visual controller and specify the possible input output parameters for the feed-forward controller, which we advocated F5 mirror neurons for. In the simple grasping world we defined, a feed-forward command would specify the (change of) arm configuration (the joint angles) given the visual information about the object and the hand. The nature of the visual information deserves some comment. The intrinsic grasp related object properties (affordance) are relayed to area F5 via AIP, which is reciprocally connected with F5 (Matelli 1986). The object location information can be relayed to area F5 via other parietal regions such as VIP (Duhamel et al. 1998; Colby et al. 1993a) or premotor regions such as F4 (Fogassi et al. 1992; Fogassi et al. 1996). For grasping it is likely that area F4 and F5 work together with F4 being involved in the reach component and F5 in grasp component. Also there is evidence that F2 may be involved in controlling wrist rotations (Raos et al. 1998). However, in our simple grasp world we assume without loss of generality, that our inverse model specifies the full arm configuration without explicitly making a task division over mentioned possible premotor regions.

7.3.4.2 Inverse kinematics and requirement for local representations

From a computational point of view the task of feed-forward command is simply stated as finding the inverse mapping of $f(u)$, which is the forward function describing how the plant behaves with given command u . Our case $f(.)$ function is referred as the *forward kinematics* mapping and f^{-1} is called the inverse kinematics mapping. Mathematically speaking $f^{-1}(.)$ may not exist because $f(.)$ need not be one-to-one, that is $f(u_1)=f(u_2)$ does not necessarily imply that $u_1=u_2$. In this case, the manipulator (e.g. arm/hand) is called *redundant*. Even our simple arm/hand is redundant. To see this, simply note that the wrist can move without breaking the finger contacts on the object during a two-finger pinch. In the literature there are many techniques to solve the inverse kinematics problem (Flash and Sejnowski 2001;

²⁶ The inverse may not exist, which is one of the difficulties in learning a inverse model.

Klein et al. 1995; Whitney 1969). In a redundant system, the inverse kinematics problem can be solved by introducing an extra optimization criterion such as minimum energy and shortest distance. The commonly used *resolved motion rate control* (Whitney 1969) methods rely on the matrix called the Jacobian of the forward kinematics which is composed of the partial derivatives $\left[\partial \vec{f} / \partial q_i\right]$ of the forward mapping function. The Jacobian transpose, pseudo-inverse of the Jacobian (with null space optimisation), and the extended Jacobian (Klein et al. 1995) methods are widely used in robotics and computer graphics for mapping the extrinsic coordinates into joint angles (Sciavicco and Siciliano 2000). However, we are interested in solving the task using neuron like units rather than pure engineering techniques. There is also a plentiful of robotics neural network literature on learning inverse kinematics functions (Tang and Wang 2001; Oyama et al. 2001a; D'Souza et al. 2001; Driscoll 2000). The key point in solving the inverse kinematics for a redundant manipulator is to employ a modular neural network (Oyama et al. 2001b) or to use locally specialized learning methods such as Locally Weighted Regression (LWR) (Atkeson and Schaal 1995; Schaal et al. 2002). The reason is that usually the inverse kinematics function is not convex, thus the average of two solution points may not be a valid solution point. A homogeneous non-local neural network (e.g. a back-propagation feed-forward neural network) averages over two solution points when queried with a non-trained point (unless a negative data point in the region of query is manually added to the training set to *disable* the 'wrong' averaging). The simplest locally weighted algorithm is the memory-based LWR (Atkeson 1992; Atkeson 1989; Atkeson and Schaal 1995).

7.3.4.3 Memory based visual feed-forward control

The feed-forward module for the visual servo circuit uses the memory-based learning with nearest neighbour (Atkeson and Schaal 1995). The reach model of Rosenbaum et al. (1995) showed that summing over postures is not a feasible strategy (Rosenbaum et al. 1999). Thus, when querying a point we do not apply a weighted sum as in LWR but simply return the nearest match with a winner-take-all circuit. This ensures that we never have a wrong answer to a query. In their planar kinematics grasp simulator Rosenbaum et al. (1999) used the same approach with the additional explicit constraint satisfaction criterion.

The learning takes place in this fashion: when the hand is performing a grasp action servo-ed by the feedback controller new F5 feed-forward units are allocated when the hand configuration is not similar to any configuration encountered before. The instantiation of the configuration is based on joint angles (and visual gating signals). As a new configuration is stored its corresponding visual representation is stored along with it.

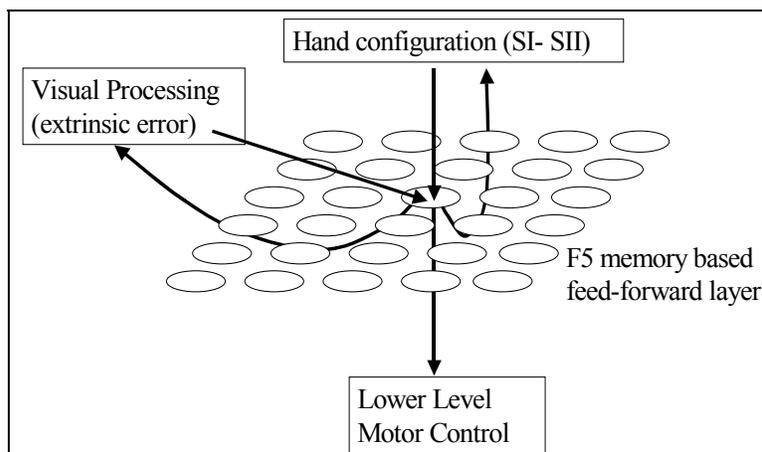


Figure 7.11 The F5 mirror neurons viewed as the memory based feed-forward controller. The arrows below the sheet of neurons indicate outputs while the arrows coming above the sheet indicate inputs

A F5 memory unit is activated to the extent that the *hand configuration* and the *visual error* fall into its receptive field. The activity signals what change must be done given the extrinsic error and the current hand configuration establishing an inverse kinematics circuit. In this model, the lower level applies a winner-take-all process on the activity of F5 feed-forward neurons. The result is then used for movement execution (Figure 7.11).

To our knowledge, there is no experimental data²⁷ as to how a grasp plan is represented in premotor cortex. Therefore, we are not claiming that the brain uses exactly these mechanisms; but rather we are proposing a *distributed* feed-forward module that can supply a feed-forward signal complementing the feedback signal. However, when we show later in this chapter that it is possible to observe unit activities that are similar to mirror neuron activities, we suggest that the mirror neurons may be involved in encoding visually defined grasp errors. It should be understood that our modeling in this chapter is not intended to be a proof that the premotor cortex works as we modeled, but rather an important message to experimentalists that quantitative experiments are necessary to pin down the mechanism of grasp planning and execution.

Now let us look into the details of our representation for F5 mirror neurons. There are two kinds of inputs to F5 mirror units: (1) the somatosensory cortex input that signals the configuration of the hand/arm. We capture this input as a vector of joint angles of the arm/hand: $s=[q1,q2,q3,q4,q5]^T$ (shoulder angle, elbow angle, thumb angle, index finger angle, index finger second metacarpal angle), and (2) inputs from the visual centers described as the following. The visual input to feed-forward units is a five dimensional vector $x=[dx_1,dy_1,dx_2,dy_2,\gamma A]$. The first four components are the parameters are the *errors* signaled by the feedback loop, namely the distances between fingertips and their contact targets on the object. The last component represents the affordance of the object. In this model, we use only the size of the object as the affordance. The parameter γ controls the relative importance of the error and the object affordance input. Note that γA parameter plays the role of a soft gating network. As we advertised earlier, F5 canonical neurons may implement an explicit gating network based on visual object properties (Murata et al. 1997a) for biasing F5 mirror units. In that case, only the feedback error signals would suffice, which is in full accordance with the learning structure we offered in Figure 7.4. However, for the sake of implementation we embed the gating mechanism into the error signal. It must

²⁷ Except that the qualitative description of selectivity of canonical and mirror neurons.

be emphasized that x is purely defined within an extrinsic space. Thus, it can be applied to self-action as well as the actions of others. However when learning (storing of hand configurations) the recruitment of new units to store a configuration is based only on intrinsic parameters (s vector). The circuit continually monitors the arm it controls (i.e. online learning); if the maximum activation based on s input of the units is below a threshold a new unit is recruited. The hand configuration activity s_i of unit i , is computed simply as the length of the vector between s and s_i ($|s-s_i|$) where s is the current hand configuration and s_i is the stored configuration i . Figure 7.12 shows the acquired configurations during the simulations of this chapter for six objects. In a general system for satisfactory performance many more units are required (Rosenbaum et al. 1995; Rosenbaum et al. 1999; Rosenbaum et al. 2001).

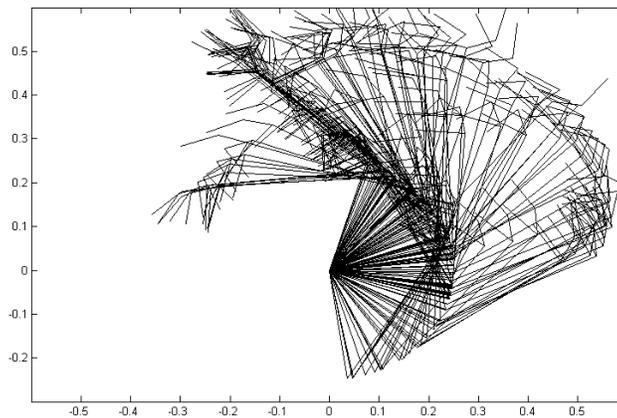


Figure 7.12 The arm configurations that were acquired during 6 object grasping actions are shown. Each of the superimposed configurations is represented by a unit in the feed-forward layer

We mentioned (1) configuration dependent activation and (2) vision based activation. Now we need to explain the mechanism of inverse mapping, that is, how the given visual information is mapped to a motor plan. At any instant, during a grasp action, the vision based activity and the hand configuration based activity lights up some candidate configurations. The task is to pick one of the units as the best candidate and return its stored hand configuration as the next target (as the joint angles). In a general framework, a sophisticated way of combining activated units would be employed as in LWR methods. We speculate that the brain circuits must have a reinforcement type of learning circuit to learn to pick a chain of hand configurations that will satisfy extrinsic constraints (obstacle avoidance) and intrinsic constraints (such as minimum energy). However for simplicity we offer a heuristic for selecting a configuration that will enable the system to generate reasonable trajectories so that we can test the system's control performance and look at the properties of the units that are acquired. Our heuristic is to pick the unit that maximizes the activity function: $g_i = e^{-\left(\alpha\|(dx_1, dy_1, dx_2, dy_2)\| + \lambda\|q - q_i\|\right)}$ where α and λ are constants controlling the relative contributions of the intrinsic and extrinsic signals. Here there is a subtle issue: the minimum of g may not be a right plan because the error vector $[dx_1, dy_1, dx_2, dy_2]$ is relative to the given object. To circumvent this problem there are two possible mechanisms that can be applied: (1) having a gating mechanism which biases the correct subspace of the error space based on the object affordance (this is the task we offered for F5 canonical neurons), (2) having a *forward model* of the arm/hand kinematics that predicts the error vector if the unit were selected and use this predicted error $[dx_1, dy_1, dx_2, dy_2]'$ in the computation of $g_i = e^{-\left(\alpha\|(dx_1, dy_1, dx_2, dy_2)'\| + \lambda\|q - q_i\|\right)}$. It is suggested that both gating and forward model strategy is employed in the human brain from cerebellar motor prediction to

social prediction (Miall and Wolpert 1996; Wolpert et al. 1998; Wolpert et al. 2001). For our simulations the winning unit is selected using the latter approach (we use a non-neural forward model).

7.3.5 Simulation: trajectory planning and controller performance

Although it is a standard result that use of feed-forward signals improves performance we need to first show that the distributed feed-forward system we proposed does not degrade but improves the performance as expected.

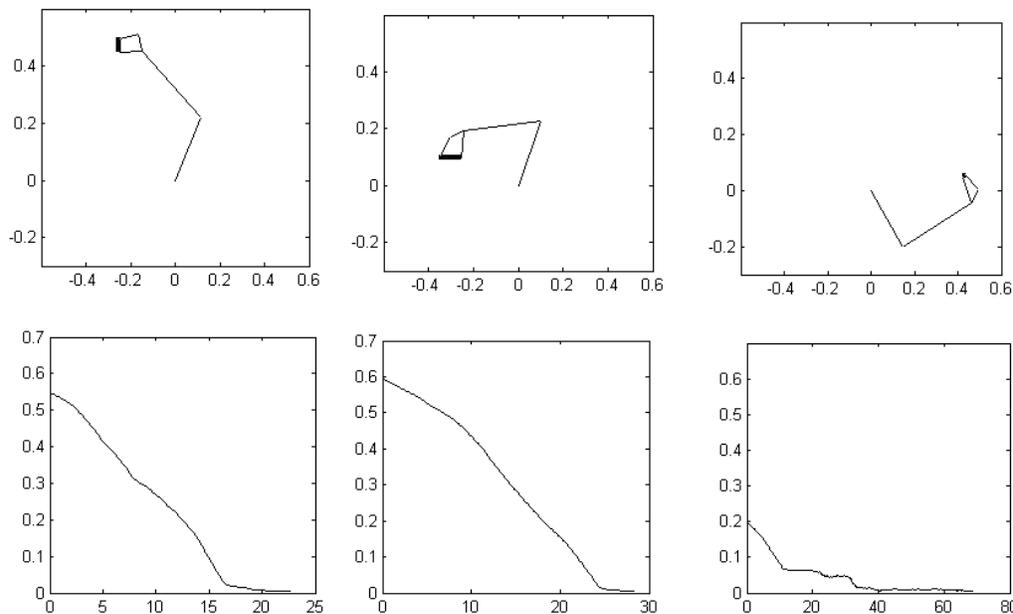


Figure 7.13 The trajectory generation with feedback and feed-forward control is illustrated for comparison with Figure 7.8 (feedback-only system). In the lower panel the error graphs are plotted as error versus iteration. The error is the sum of squared distances of the fingertips to their targets. The rightmost object was not grasped in the training (a novel object/location). Thus the system could not make use of the feed-forward signal, approximately after iteration 25 and switched to feedback only mode, resulting in slower positioning of the fingers on the target locations

For comparison purposes, we tested the speed of the desired trajectory generation. Figure 7.13 shows the three grasps that are performed using the feedback and a trained feed-forward controller. The discrepancy of the errors at time=0 for Figure 7.13 and Figure 7.8 is due to the initial finger configurations of the hand (but the arm started from exactly the same configuration). In fact, the initial error for the middle grasp is larger in Figure 7.13.

The feed-forward control introduction to the system reduced the time required to output a desired trajectory command by about four times. This means that when we actually connect the system to the lower level motor servo we can achieve 4 times frequency that of a feedback-only system could achieve. This effectively gives a higher range of the PD gains for the lower level motor servo.

The system learned online while it was performing grasps using the feedback and feed-forward model learned so far. The leftmost two grasp objects in Figure 7.13 were included in the training; thus, the error curve goes down very fast since the feed-forward command can take care of most of the task (as it acquired an inverse map during training). However, the rightmost object was presented the first time to the model. The feed-forward module could make use of the earlier experiences until step 25, only partially between 25, and 40. After that point, the feedback servo worked alone. Thus, the

generation of the trajectory was slower. We did not use the lower level motor servo to be able to make a fair comparison with Figure 7.7 of feedback-only system. Figure 7.14 shows the complete systems', i.e. the feed-forward, feedback and lower level feedback motor servo performance all working together.

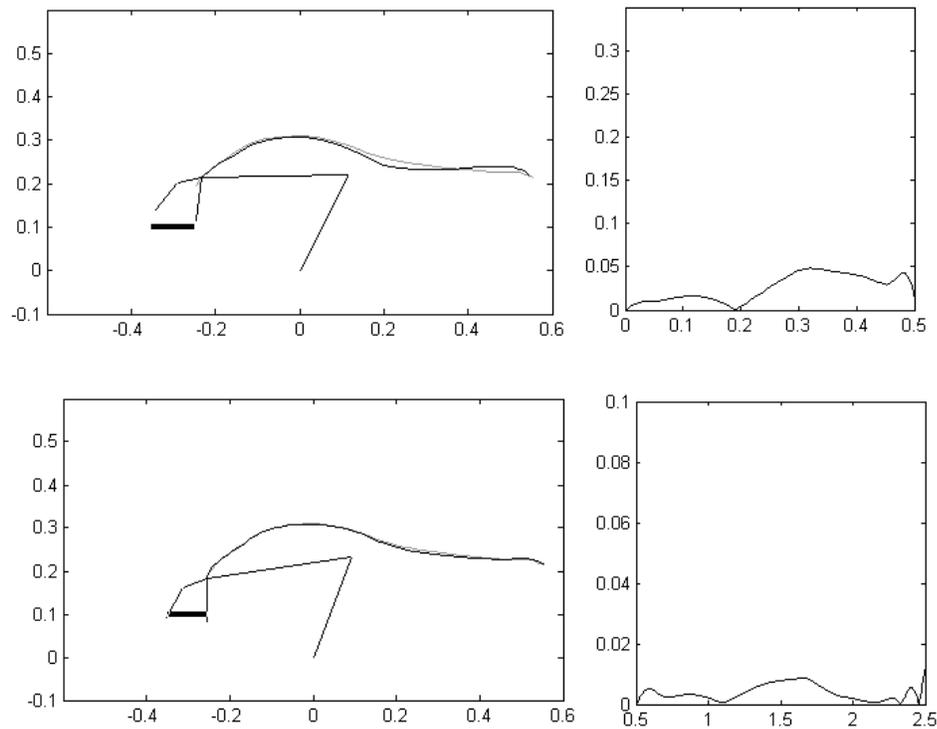


Figure 7.14 The feedback, feed-forward and lower level motor servo and the dynamic arm was simulated all together. Upper half: The grasp lasted 0.5 seconds. Lower half: the grasp lasted 2 seconds. The fast movement error reduced with a factor of 6 while the slow movement reduced with a factor of 10 in terms (compare with Figure 7.10)

The plots in Figure 7.14 use the same scales as Figure 7.10, for direct comparison. A high degree of improvement is observed as expected. The slower motion trajectory can be followed with 1/10 times and the fast motion with 1/6 times the error made with only feedback controller. The reason for the poor performance of the feedback only control case (Figure 7.10) is that the generated desired trajectories are not very smooth, which causes the PD controller to overshoot and follow the trajectory harder. In contrast when using the feed-forward controller, the generated trajectories are smooth. Note that both PD controllers of the lower level motor servo used the same PD gains (30 for position and 10 for derivative gain). The values were manually optimized for the feedback-only system and it was not possible to increase the gains further to achieve better tracking. In Figure 7.14 even with the feed-forward controller it looks like the system did not achieve a satisfactory grasp. However, note that we required the simulator to complete the grasp in the allocated time; the system completes the grasp after the allocated time with a negligible increase in the error (not shown).

We mentioned earlier that having memory-based feed-forward controller may allow explicit planning using the controller representation. Here we demonstrate this fact by a simulation. An obstacle in the workspace for example, can be encoded as simple inhibition over the feed-forward units that

encode hand configurations that would result in a collision with the obstacle. Figure 7.16 shows two grasp examples in the top row where the effect of the obstacle is zero (the obstacle is drawn for reference). In the lower level the obstacle inhibits the cost function associated with the units that would bring the hand in the obstacle. Considering the limited range of motions available to move the arm from its initial position to the target, we can suggest that the representation we proposed allows generation of non-trivial trajectories that can be automatically acquired (in contrast to set by heuristics) once the system is equipped with a learning mechanism that can do the planning (e.g. reinforcement learning).

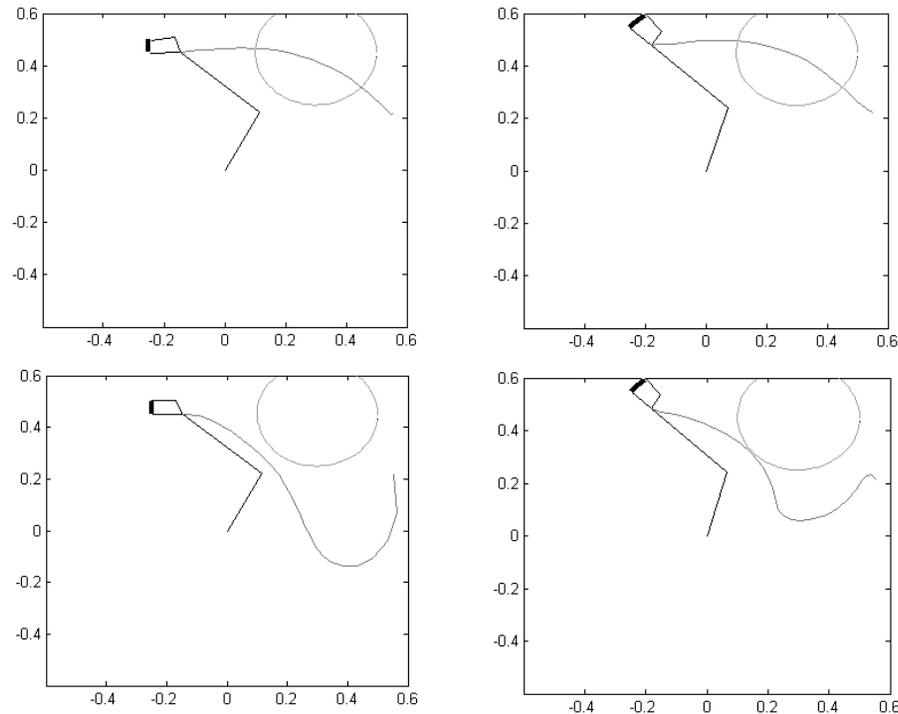


Figure 7.15 The top row demonstrates two trajectory-planning examples for grasping without obstacle. The bottom row demonstrates how new trajectories are formed by introduction of an obstacle as a local inhibition on the feed-forward controller units

7.4 Feed-forward unit activity and mirror neurons

In this section, we focus on the unit level activities of the feed-forward controller and introduce a method to compare the electrophysiological recordings from monkey mirror neurons with our feed-forward units activities.

First we look at the population level activity of the feed-forward (F5 mirror neurons) that we introduced in section 7.3.4.3 during various grasp actions. Figure 7.16 shows the unit activities of four grasps as area level plots. Each plot consists of 157 neurons acquired during the learning phase (the rows). The columns represent the time (note that in general the grasp actions take different time steps for different grasps). The left edge of each graph is aligned with the start of the grasp actions. The right edge is aligned with the completion of the grasps. We can see that different units are activated at different times depending on the grasp action. It must be emphasized that this map is based on only *visual* information and mainly dominated by the error patterns that occur during a grasp action. It is likely to have similar errors during different grasps unless the map is modulated by object affordances. We used the object size as a soft modulator by including it in the error computation. How much object

related information is actually encoded in such a map in the primate is a topic of future research for experimentalists. Our predictions and experiment suggestions for object specificity is presented in Chapter 6. Our goal here is to show that a *visual error-based activation* of units can lead to similar firing patterns of mirror neurons.

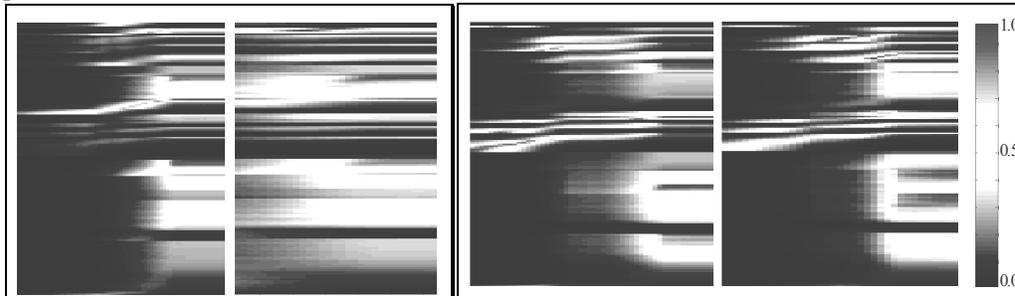


Figure 7.16 The feed-forward unit activations for four grasp observations shown as unit versus time. Each graph consists of 157 neurons acquired during the learning phase (the rows). The columns represent the time

The activity of a feed-forward unit corresponds to *average firing rate* of a real neuron. We can map the average firing rates during an action (observation activity only) to actual firings (neuron spikes) using a Poisson distribution model of a neuron. The Poisson distribution is the extension of the binomial distribution to the continuous case. Under the Poisson encoding model the probability of a neuron generating r spikes during $(t, t+\delta t)$ for encoding a parameter x is given by (Snippe 1996; Sanger 1996; Zemel et al. 1998):

$$\Pr(r | x) = e^{-\langle f(x) \rangle} \frac{\langle f(x) \rangle^r}{r!}$$

where $\langle \cdot \rangle$ denotes the average firing rate and $f()$ is the response function (or tuning function) of the neuron which actually determines the average firing rate activity one computes using a non-spiking neural circuit model²⁸. Zemel et al. (1998) suggest that a typical tuning function is proportional to a Gaussian:

$$f(x) = \beta e^{-\frac{(x-x_{pref})^2}{2\sigma^2}}$$

where x_{pref} denotes the input that best activates the neuron; x is the current input to the neuron and β is a proportionality constant and σ^2 is the variance determining the receptive field size. We used this form of activity function for F5 feed-forward controller units. We set $\beta=1$ and empirically determined a single σ^2 (0.1) value for all the neurons in our simulations.

The feed-forward units activation (the average firing rate) is determined by the activation function we have defined in section 7.3.4.3. However, note that intrinsic contribution is set to zero during an action observation (i.e. the activity is purely visual during observation). This can be accomplished in the primate brain by shunting of intrinsic input when the animal is not engaged in any motor act. Noting that in section 7.3.4.3 we defined $x=[dx_1, dy_1, dx_2, dy_2, \gamma A]^T$ we can write the probability of a feed-forward units firing based on the units preferred stimulus x_{pref} and the current stimulus x as the following.

²⁸ δt is absorbed in the expectation operator $\langle \cdot \rangle$. Formally $\langle f(t) \rangle = \int_t^{t+\delta t} g(t) dt$ where $g(t)$ is the instantaneous firing rate.

$$\Pr(\text{firing} | \bar{x}) = e^{-e^{-(\bar{x}-\bar{x}_{pref})^T(\bar{x}-\bar{x}_{pref})/(2\sigma^2)}} e^{\frac{-(\bar{x}-\bar{x}_{pref})^T(\bar{x}-\bar{x}_{pref})}{2\sigma^2}}$$

Now we can convert our average firing rates plotted in Figure 7.16 into spikes and make (qualitative) comparisons with real mirror neuron data. The mirror neuron data available to us does not allow comparisons based on kinematics of the observed actions. Therefore, instead we try to spot feed-forward units that have firing profiles similar to those available to us.

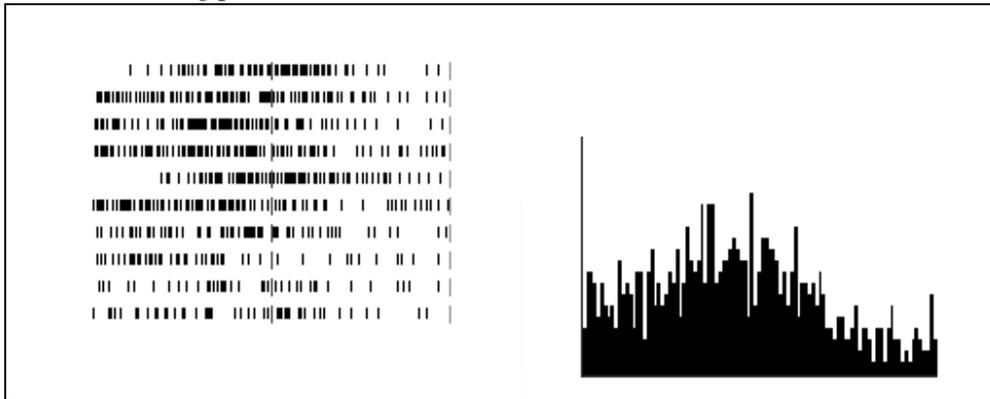


Figure 7.17 A mirror neuron recorded during a grasp observation. On the left the raster plots; on the right the histogram. The recording data shown spans 2 seconds. In addition, the hand start to move approximately at time = 1 second indicated by the vertical bar at the centre of the raster panel (Rizzolatti and Gallese 2001)

Figure 7.17 shows a real recording from Rizzolatti lab (Rizzolatti and Gallese 2001; Umilta et al. 2001), which is displayed using the tool we have developed (Oztop 2000). The recording shown spans two seconds and the end of the rasters are aligned with the experimenter's touch of the object. We did not include the holding phase of the recordings for the comparison because different processes may be involved in holding, which we did not model. Figure 7.18 (lower panel) shows one of the feed-forward module unit's activity that is produced using the Poisson model described above. The generated spikes are shown as raster plots and histogram for a direct comparison. The raster plots corresponds to actual trials in the real experimentation (10 trials), however in the simulation we collected the average firing rate information via a single trial and run the Poisson spike generation multiple times (25 runs) The histogram bin width is selected as 20ms for both experimental and the simulation case. The mirror neuron in Figure 7.17 (same neuron is also shown in Figure 7.18, upper panel) shows an interesting behavior: first, the activity of neuron rises but before contact with the object, around half way of the grasp the activity reduces again. This type of behavior is not very well understood. If this mirror neuron were involved in understanding the meaning of actions as suggested in the literature (Rizzolatti et al. 2001a; Rizzolatti et al. 2000; Umilta et al. 2001), then one is tempted to ask the question, why the firing is inhibited prior to contact. To our knowledge, there is no satisfactory answer to the question and no published explanation of such mirror activity. Here we demonstrate that such an activity can be generated as a by-product of the visual servo circuit (the feed-forward module) that we implemented.

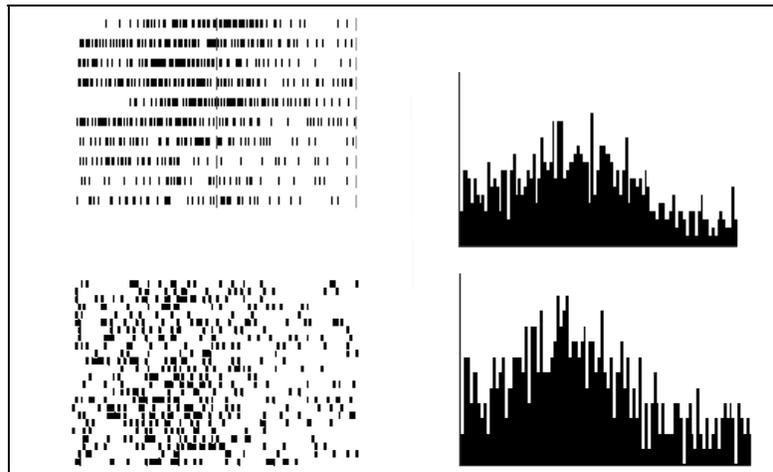


Figure 7.18 Top row: Real mirror neuron recording during a precision grasp. Bottom row: One of the feed-forward controller unit's responses to vision of a grasping action. In the left panels, each raster row corresponds to a trial (Poisson spike generation for the model). The right panels show the histograms. The rasters aligned according to the contact of the hand with the object

We now switch our attention to a more common mirror neuron that increases its activity as the action observed is progressing towards completion, which is in support for the mirror neuron involvement in understanding the meanings of actions (Rizzolatti et al. 2001b). We can find many units in the feed-forward module, which mimics the activity of this type of mirror neurons with different rise time profiles.

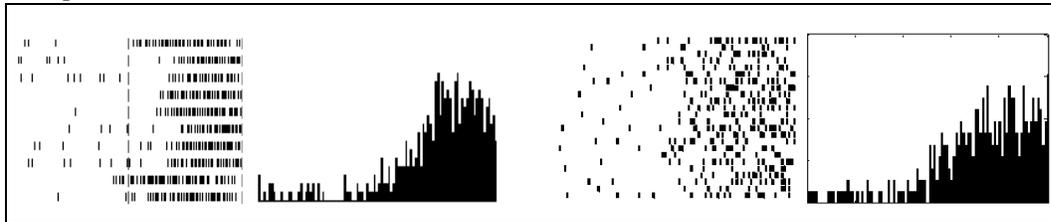


Figure 7.19. The similarity of a real neuron and model unit is demonstrated. Left two panels real mirror neuron rasters and histogram. Right two panels are the model generated rasters and histogram. A slow increasing activity is observed in both cases

In Figure 7.19, we picked one, which matches a real mirror firing profile (also shown in the same figure). Note that no parameter fitting was done; the parameters of the model and the Poisson spike generation process were the same as in the previous example.

Next, we present an interesting mirror neuron profile, which we could not replicate without changing the receptive field size parameter (σ^2). This neuron was silent as the background activity and stayed silent when the experimenter started the grasping action (Figure 7.20, left two panels). Before contact, it started firing vigorously and stopped before holding phase. As the very first neuron we have presented we suggest that this neuron must be involved in some processes other than understanding because of the similar reasoning. When we tried to find this kind of activity in our feed-forward control module at first we could not find such sharp profile units in any grasp action observation. However, when we reduced the receptive fields of the units we could get similar activity (Figure 7.20, right two panels) although the timing of such activity could not be replicated (the length of the silent time after the burst). However, this could be because we require the grasp to be completed for the simulator case.

The real mirror activity shown is only until the contact of the hand to the object. Inspecting Figure 7.21 indeed reveals that the final portions of the grasping action do not change the population activity. Nevertheless, it is certainly too daring to claim that the mirror neurons use the same error space we used for the simple grasp world we implemented. However, we would like to suggest that mirror neurons can be involved some visual error computation for visual guidance of the hand during grasping.

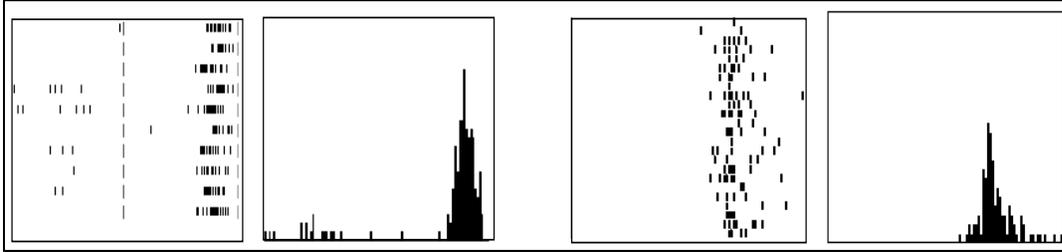


Figure 7.20 Left: a sharp mirror neuron activity, which could only be replicated with our simulator by reducing the receptive field. Right: Similar response profile obtained from one of the feed-forward module units

When we look at the population activity of the feed-forward units with the smaller receptive field we see that the main characteristics are preserved, but the tuning of the units become stricter (Figure 7.21). In the general case, each unit may have their own receptive field sizes, which can adapt with experience. The unit shown in Figure 7.20 is marked with an ellipse in Figure 7.21. The unit was activated strongly for a short period of time and declined its activity rapidly. A group of other neurons became active after the decline.

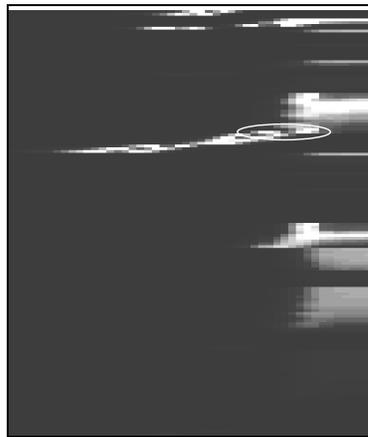


Figure 7.21 The population activity of feed-forward units with smaller receptive fields. The feed-forward unit we used to match the real mirror firing profile in Figure 7.20 is marked with an ellipse

8 CHAPTER VIII: CONCLUSIONS

The scope for the dissertation was set in Chapter 3 where (1) the visual feedback control of manual manipulation hypothesis of F5 neurons and (2) the view that infant grasp learning precedes the mirror neuron development were introduced, and (3) the assertion that self-grasp observation trains the mirror neurons during infancy was made. Our study took a schema based approach to address (1), (2) and (3).

First (3) was modeled in Chapter 3, focusing on the Core Mirror Circuit. The context for the circuit was provided by the functional schemas we implemented without intending to address neural localization. This enabled us to concentrate on the Core Mirror Circuit thoroughly and make testable predictions.

The view (2) motivated us in developing the Learning to Grasp Models. Chapter 5 developed a behavioral level model of infant grasp learning, while Chapter 6 localized the circuit in monkey brain. Both chapters produced nontrivial predictions that can be experimentally tested.

Finally, Chapter 7 attempted to realize (1) through simplified but nevertheless adequate (according to the F5 canonical priming hypothesis) simulation studies. It was shown that units encoding kinematics based errors, could yield mirror neuron like responses. This is in full accordance with the suggestion in Chapter 3 that, the degree of conformity of the hand movement and preshape to the object affordance must be represented as a distributed neural code. This code in its simplest form is a kinematics error function which we have elaborated in Chapter 7.

We concluded our modeling tour of action recognition and visuomotor transformation by handing neurophysiologists an extensive set of testable predictions and hypotheses, and equipping the modelers who are interested in understanding the cortical mechanism underlying the action recognition and visuomotor learning, with an excellent starting guide.

8.1 Mirror neurons

Chapter 3 studied how self-executed grasp actions can adapt the parietal and premotor connections to shape F5 neurons into mirror neurons. We used schema approach for the Visual Analysis and Reach and Grasp components of the Mirror Neuron System (MNS). Therefore, the neurophysiological predictions pertaining to MNS model could be gathered only from the Core Mirror Circuit, whose output corresponded to the mirror neurons (see Figure 3.5).

The key to the MNS model was the notion of *hand state* as encompassing data required to determine whether the motion and preshape of a moving hand may be extrapolated to culminate in a grasp appropriate to one of the affordances of the observed object. A mirror neuron had to fire if the preshaping of the hand conformed to the grasp type with which the neuron was associated. We emphasized that the specific decomposition of the hand state into the specific components that we had used in our simulation was not the crucial issue. In fact, we suggested that a distributed neural code that carries information about the movement of the hand toward the object, the separation of the virtual fingertips and the orientation of different components of the hand relative to the opposition axis in the object must be represented in monkey parietal or premotor cortices. The crucial property of the code we predict is that it would work just as well for measuring how well another monkey's hand is moving to

grasp an object as for observing how the monkey's own hand is moving to grasp the object. This allows the self-observation of the monkey to train a system that can be used for observing the actions of others and recognizing just what those actions are.

Experimental work to date tends to emphasize the actions to be correlated with the activity of each individual mirror neuron, while paying little attention to the temporal dynamics of mirror neuron response. Our simulations make explicit predictions on how a given (hand state trajectory, affordance) pair will drive the time course of mirror neuron activity. For example, we have shown that a grasp with an ambiguous prefix (e.g. using precision grip for a large object) drives the mirror neurons in such a way that the system at first gives weight to the wrong classification, with only the late stages of the trajectory sufficing for the incorrect mirror neuron to be vanquished (*grasp resolution*).

The *grasp resolution* prediction is testable and indeed, suggests a whole class of experiments. The monkey has to be presented with unusual or ambiguous grasp actions. For example, the experimenter can grasp a section of banana using precision pinch from its long axis. Then we would expect to see activity from power grasp related mirror neurons followed by a decrease of that activity accompanied by an increasing activity from precision pinch related mirror neurons. Our other simulations lead to different testable predictions such as the mirror response in case of a spatial perturbation (showing the monkey a fake grasp where the hand does not really meet the object) and altered kinematics (perform the grasp with different kinematics than usual). We have also noted how a discrepancy between hand state trajectory and object affordance may block or delay the system from classifying the observed movement.

In summary, we have conducted a range of simulation experiments – on grasp resolution, spatial perturbation, altered kinematics, temporal effects of explicit affordance coding, and analysis of compatibility of the hand state to object affordance – which demonstrate that MNS model is not only valuable in providing an implemented high-level view of the logic of the mirror system, but also provides interesting predictions that are ripe for neurophysiological testing, as well as suggesting new questions when designing experiments on the mirror system.

Chapter 7 investigated the *manual visual feedback* hypothesis asserted in Chapter 3 through a simplified hand/arm system. We concentrated on a single grasp type. However, this does not weaken the impact of the analyses we conducted since Chapter 6 demonstrated, through explicit simulations, that grasp learning gave rise to emergent neural properties comparable to F5 canonical neurons, which are thought to be involved in gating mechanisms, selecting the feedback circuit best suited for the object to be grasped. To be more specific, after LGM had mastered grasp learning, when we investigated the neural responses of LGM layers we found units that are selectively active for specific object sizes and orientations. Thus, we concluded that these were the *canonical neurons* of LGM. In parallel, Chapter 7 hypothesized that F5 canonical neurons implement a gating or priming mechanism based on the object affordance. For example, given a small object the LGM canonical neuron with small object preference would become active (as part of its grasp planning task, as shown in Chapter 6) priming the visual feedback circuit (a subpopulation of F5 purely motor and mirror neurons) that is best suited for the control for small object grasping (e.g. the visual servo circuit that is specialized for precision grip).

Although not modeled, a learning mechanism is proposed for the gating mechanism as well. The canonical neurons not only prime the F5 visual feedback circuit for grasp execution but also they prime the feedback circuits for learning ensuring that the neurons in a particular circuit will *learn to control specific grasps* that are directed to objects for which the canonical neuron has selectivity.

8.2 Grasp learning: infant development

ILGM simulations of Chapter 5 showed that even without object affordance, a variety of grips, including the precision grip, can be learnt by actively interacting with the environment predicting the results of the study of Butterworth et al. (1997).

We showed that task constraints are influential in shaping the grasp repertoire of infants via the *cube on the table task*, which together with other simulation results enabled us to make useful predictions. To be precise, Chapter 5

1. Predicted that infants in the period of early grasping, can perform precision grasps
2. Showed that task constraints shape motor development supporting the view that development of precision grips is mediated by task constraints.
3. Showed that object affordances could be represented in infants' grasp repertoire in spite the fact that the grasp planning circuit was unable to access/extract object affordance information. This result is important because it implies that tactile learning can train the visuomotor learning in infancy.
4. Predicted (as a corollary to 3) that object (visual) affordances would be heavily represented in the motor and sensory areas of cerebral cortex for objects that are manipulated often.

8.3 Grasp learning: monkey neurophysiology

The simulation results of emergent object selectivity in LGM and the analysis of what it takes to combine movement parameters into a coherent whole for a successful grasp enabled us to produce nontrivial testable predictions.

Specifically, we predicted that area F2 must be having visually triggered neurons that are selective for *object orientations*. Furthermore, we suggested that the area must have control over *wrist rotations*. We further predicted that F5 canonical neurons would have tuning fields according to the object orientation, similar to AIP visual dominant neurons (Sakata et al. 1999). Therefore, we asserted that a closer examination of area F2 and F5 canonical neurons would reveal *orientation selectivity* either in isolation or in the form of gain fields over existing object selectivity.

We generalized the above prediction and proposed a systematic way of interpreting neural activities to amend the shortcomings of analyses based on behaviour and neural firing correlation paradigm.

If some population of neurons were to specify a grasp plan partially (by being responsible for one parameter of the plan), a learning mechanism had to exist for shaping the connections between regions that specify the complementary grasp parameters. This implies that in order to specify a grasp plan F5 canonical neurons must have 'gain fields' based on other brain regions' activity (the 'gain fields' would, behaviorally appear to be based on movement parameters). Let us review the hypothetical experimenter we introduced in Chapter 6 to present our proposed systematic approach.

If the experimenter finds out that area A is responsible for some aspect (wrist orientation) of a complex movement (reach-to-grasp) and that area B is responsible for controlling some other aspect (reach direction) of the movement and if the activities of A and B are true control variables determining the movement then we claim that

- (1) The region A and B must be anatomically connected (the connection may be a multi staged one, such as A connected to C, C connected to B)
- (2) The synaptic strength of the connection between A and B encodes the movement generation (grasping) strategy of the grasping circuit

(3) If it is found that the connection between A and B is unidirectional (say from A to B) then (from the viewpoint of the experimenters) the activity of B neurons must be *tuned* by the population activity A.

(4) The verification of (3) can be used to decipher the organization of visuomotor transformation circuit by recording simultaneously from A and B. If activity of A predicts the activity of B better than the other way round (i.e. activity of B predicts the activity of A) then the underlying principle is that area A selects movement parameters, and B, based on the selection of A, generates other movement parameter(s) completing (or further specifying) the movement plan.

We hereby explicitly suggest a methodology to uncover the roles of premotor areas in reach-to-grasp movements. For a true understanding of premotor function, we strongly suggest that the premotor regions of F2, F5 bank (canonical neurons), F4 and wrist rotation encoding neurons of motor cortex (F1) must be investigated with *simultaneous recording* experiments. A feasible experimental design would be to set up sessions for pair-wise recording of areas F2, F4, F5 and F1, in which the monkey would be asked to reach and grasp for objects at various locations with changing orientations. We claim that the rigorous analyses of neural firing, kinematics of the hand and the object condition will reveal the mechanisms of grasp learning and planning by showing the structural and functional dependency of premotor regions in specifying grasp parameters.

8.4 Grasp learning: neural architecture

The crucial feature of LGM inherited from the architecture we developed in Chapter 4 was that it could represent and produce a variety of grasp plans conditioned on input, effectively learning a repertoire of grasp plans. The reinforcement framework we adapted in Chapter 4 gives LGM the flexibility to include soft constraints in grasp evaluation such that the grasps that are not favorable (due to discomfort, excess energy consumption, etc.) are eliminated or represented with low probability. Thus the reinforcement signal of Chapter 4, (*joy of grasping* of Chapter 5 and *neural grasp stability* of Chapter 6) can incorporate the *anatomical and environmental constraints* which are important in shaping grasp development (Newell et al. 1989). The reinforcement signal can be viewed as the *adaptive value* of action that shapes the movement repertoire through selection (Sporns and Edelman 1993).

The simulation results of Chapter 5 and 6 showed that the architecture we introduced was adequate for grasp learning by producing testable prediction for both infant development and for monkey neurophysiology. We showed the neural network competency of LGM by demonstrating its generalization ability over novel input conditions while retaining the property of representing multiple actions for a given input.

9 CHAPTER IX: FUTURE WORK

9.1 Simultaneous learning in MNS model and LGM

While infants learn how to shape their reaching movements into grasping actions, they also observe the visual stimuli they generate. The main hypothesis of the MNS model was that the visual stimuli from the successful grasps were used to train the mirror neurons. To this date, there is no study pinpointing any stage in development when the mirror neurons start to function fully. Our suggestion is that the MNS learning and grasp-learning overlap. The amount of the overlap may have important consequences on the organization of the two circuits. A very early learning in MNS may result in mirror neurons that respond to premature grasping actions, as the training stimuli will be based on the premature grasping attempts. A very late learning may result in delayed visuomotor development in the infant. The effect of the overlap amount must be studied via simulation studies to produce predictions to be tested by the experimenters. These studies will stimulate research on the development of mirror neurons, which is a lacking component of current mirror neuron research.

9.2 Going beyond grasping: learning the ‘hand state’

The MNS model was based on grasping movements and excluded some actions such as tearing and twisting for which mirror activity has been observed. The MNS model must be extended with the help of new neurophysiological studies to account for other hand manipulation actions. In essence, the *hand state* we introduced must be extended to account for other actions. Although the first step would be manually crafting the representation for the hand state to cover the excluded hand manipulation actions, the major research direction would be to model the mechanisms of how the brain learns to extract such a representation from the visual world. We have already outlined one alternative by claiming that the hand state is used as the visual feedback for hand manipulation actions. Thus, a computational model showing how a visual feedback mechanism may be bootstrapped by biologically realistic neural circuits, would provide a general theory of mirror neuron development as well as explaining the existence of mirror neuron system. Chapter 7 presented alternative feedback mechanisms for grasping focusing on a single grasp type, namely the precision grasp. The models there viewed with LGM (Chapters 5 and 6) should serve as a starting point for the interested researcher. As shown in Chapters 5 and 6, the infant/model learns how to make open-loop reach and grasp movements via goal directed trial and error learning. Then the circuits involved in visual feedback control are built based on the observation of the self-executed ballistic grasps.

9.3 Tactile feedback for grasping

Infants’ early reaching and grasping movements are ballistic and usually require correction after contact with the object. Infants use the tactile feedback to conform their hands to the shape of the object reducing the number of trials they need to perform for adapting their ballistic reach and grasp planning. Learning to Grasp Models that we have presented in Chapters 5 and 6 did not include the tactile correction phase and required large number of grasping actions for learning. The modeling of tactile feedback based grasping is not trivial and requires (1) a detailed physical modeling of the hand and the

objects in terms of friction and contact forces and (2) a realistic model of mechanoreceptor signals at the spinal and cortical levels. When combined with learning to grasp models of Chapters 5 and 6, the tactile feedback modeling research will provide a theory of how the tactile exploration can train a visuomotor mapping circuit, and generate predictions on the grasp development at the cortical and spinal level.

9.4 Learning to extract the right affordance

The neurophysiological findings assure that the visual affordances of objects are represented in the monkey brain. We extended the monkey studies to humans by assuming that object affordances were available to the grasp programming centers and mirror neuron system without addressing the issue of how and why the affordances were extracted from the vision of objects. The *how* part is one of the major problems of computer vision, namely feature extraction and object recognition. The *why* part is the question of selecting the *relevant* features of an object for the motor task. The natural extension of the learning studies presented in this thesis is to model learning of *what feature is relevant* given a set of features. The research on how to learn to associate relevant features of objects with the movement parameters will present a more thorough view of the neural mechanisms of the visuomotor transformation that takes place for hand manipulation actions.

9.5 More realistic models of the limb and the brain regions

The arm/hand model we used was a kinematics chain, which was articulated by specifying the angles of each joint. A more realistic arm/hand would include muscles that are controlled with neural signals from the spinal cord and the cortex. Such an arm/hand system would allow research on the effects of corticospinal pathway immaturity during infancy.

Using the schema methodology, we focused on various brain regions and implemented them with various levels of biological realism. A major contribution would be to replace the schemas that have no neural implementation with realistic neural circuits. For example, the visual analysis of the hand was implemented using computer vision techniques with no neural implementation. In addition, the affordances of the presented objects were encoded without actually performing feature extraction. Thus, the two visual processes, the visual analysis of the hand and the object, stand as future research challenges.

While modeling the core mirror circuit of MNS model, we noted that the core mirror circuit could be adapted using any supervised learning technique. The range of the neurophysiological predictions would be greatly extended by implementing the supervised learning in terms of realistic neural circuits. In particular, the temporal to spatial conversion required for the current implementation could be avoided.

Depending on future experiments that address the timings of mirror neuron activities, a revision of the MNS model might be required. The current modeling presented in this dissertation confronted experimentalists with a range of nontrivial predictions that should stimulate experimenters in conducting new experiments for testing the timing and population level activity of the neurons involved in mirror neuron system and grasp planning, which in turn, must be factored into the revisions of the models presented in this dissertation.

9.6 Sensitivity analyses for the simulation parameters

In many of the presented simulation experiments, the parameter selection was based on empirical tests. The dissertation studied a broad range of functionality with paying less attention to the effect of

parameter selection on learning. It was empirically observed that the parameters were not very sensitive in the sense that the models would not collapse with small variations of the parameters. However, the result of learning may show some differences according to the variations of the parameters. For example when studying the learning to grasp, we did not analyze the dependency of the acquired 'grasp menu' properties to the simulation parameters. This study requires establishing a method to quantify the properties of the 'grasp menu'. Then using an extensive set of simulations, the dependency can be formulated. For example, one property of a 'grasp menu' can be chosen as the ratio between the number of power grasps and the precision grasps acquired.

10 REFERENCES

- Akoev GN, Alekseev NP, Krylov BV (1988) *Mechanoreceptors, their functional organization*. Springer-Verlag, Berlin ; New York
- Andersen RA, Asanuma C, Essick G, Siegel RM (1990) Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology* 296: 65-113
- Andersen RA, Snyder LH, Bradley DC, Xing J (1997) Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience* 20: 303-330
- Anderson CH (1994) Basic elements of biological computational systems. *Int J Mod Phys C* 5: 313-315
- Arbib MA (1981) Perceptual structures and distributed motor control. In: Brooks VB (ed) *Handbook of physiology, section 2: The nervous system, vol ii, motor control, part 1*. American Physiological Society, pp 1449-1480
- Arbib MA (1989) *The metaphorical brain 2 : Neural networks and beyond*. Wiley, New York, N.Y.
- Arbib MA (2001) The mirror system, imitation, and evolution of language. In: Nehaniv C, Dautenhahn K (eds) *Imitation in animals and artifacts*. The MIT Press
- Arbib MA, Erdi P, Szentâagotai J (1998) *Neural organization : Structure, function, and dynamics*. MIT Press, Cambridge, Mass.
- Arbib MA, Hoff B (1994) Trends in neural modeling for reach to grasp. In: Bennett KMB, Castiello U (eds) *Insights into the reach to grasp movement*. North-Holland, Amsterdam ; New York, pp xxiii, 393
- Armand J, Olivier E, Edgley SA, Lemon RN (1997) Postnatal development of corticospinal projections from motor cortex to the cervical enlargement in the macaque monkey. *J Neurosci* 17: 251-266
- Atkeson CG (1989) Using local models to control movement. In: Touretzky D (ed) *Advances in neural information processing systems 1*. Morgan Kaufmann, San Mateo, CA, pp 157-183
- Atkeson CG (1992) Memory-based approaches to approximating continuous functions. In: Casdagli M, Eubank S (eds) *Nonlinear modeling and forecasting*. Addison Wesley, Redwood City, CA, pp 503-521
- Atkeson CG, Schaal S (1995) Memory-based neural networks for robot learning. *Neurocomputing* 9: 243-269
- Barnes CL, Pandya DN (1992) Efferent cortical connections of multimodal cortex of the superior temporal sulcus in the rhesus monkey. *Journal of Comparative Neurology* 318: 222-244

- Battaglia-Mayer A, Ferraina S, Mitsuda T, Marconi B, Genovesio A, Onorati P, Lacquaniti F, Caminiti R (2000) Early coding of reaching in the parietooccipital cortex. *Journal of Neurophysiology* 83: 2374-2391
- Baud-Bovy G, Soechting JF (2001) Two virtual fingers in the control of the tripod grasp. *J Neurophysiol* 86: 604-615
- Bayley N (1936) *The California infant scale of motor development, birth to three years*. University of California Press, Berkeley, Calif.,
- Bernstein NA (1967) *The coordination and regulation of movements*. Pergamon Press, Oxford
- Berthier NE, Clifton RK, Gullapalli V, McCall DD, Robin DJ (1996) Visual information and object size in the control of reaching. *J Motor Behav* 28: 187-197
- Berthier NE, Clifton RK, McCall DD, Robin DJ (1999) Proximodistal structure of early reaching in human infants. *Exp Brain Res* 127: 259-269
- Bishop CM (1995) *Neural networks for pattern recognition*. Clarendon Press ;Oxford University Press, Oxford New York
- Bonda E, Petrides M, Ostry D, Evans A (1996) Specific involvement of human parietal systems and the amygdala in the perception of biological motion. *J Neurosci* 16: 3737-3744
- Bortoff GA, Strick PL (1993) Corticospinal terminations in 2 new-world primates - further evidence that corticomotoneuronal connections provide part of the neural substrate for manual dexterity. *J Neurosci* 13: 5105-5118
- Bota M (2001) *Neural homologies: Principles, databases and modeling*. In: *Neurobiology*. University of Southern California, Los Angeles
- Boussaoud D, Ungerleider LG, Desimone R (1990) Pathways for motion analysis - cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque. *J Comp Neurol* 296: 462-495
- Bradley NS (2000) Motor control: Developmental aspects of motor control in skill acquisition. In: Campbell SK, Vander Linden DW, Palisano RJ (eds) *Physical therapy for children*. Saunders, Philadelphia, pp xvi, 1006
- Bradley NS (2002) Postural control is crucial in determining infant abilities. In, Los Angeles
- Bremmer F, Graf W, Ben Hamed S, Duhamel JR (1999) Eye position encoding in the macaque ventral intraparietal area (vip). *NeuroReport* 10: 873-878
- Breteler MDK, Gielen SCAM, Meulenbroek RGJ (2001) End-point constraints in aiming movements: Effects of approach angle and speed. *Biol Cybern* 85: 65-75
- Buccino G, Binkofski F, Fink GR, Fadiga L, Fogassi L, Gallese V, Seitz RJ, Zilles K, Rizzolatti G, Freund HJ (2001) Action observation activates premotor and parietal areas in a somatotopic manner: An fmri study. *Eur J Neurosci* 13: 400-404
- Butterworth G, Verweij E, Hopkins B (1997) The development of prehension in infants: Halverson revisited. *Brit J Dev Psychol* 15: 223-236

- Caminiti R, Genovesio A, Marconi B, Mayer AB, Onorati P, Ferraina S, Mitsuda T, Giannetti S, Squatrito S, Maioli MG, Molinari M (1999) Early coding of reaching: Frontal and parietal association connections of parieto-occipital cortex. *European Journal of Neuroscience* 11: 3339-3345
- Caminiti R, Johnson PB, Burnod Y, Galli C, Ferraina S (1990) Shift of preferred directions of premotor cortical cells with arm movements performed across the workspace. *Exp Brain Res* 83: 228-232
- Caminiti R, Johnson PB, Galli C, Ferraina S, Burnod Y (1991) Making arm movements within different parts of space: The premotor and motor cortical representation of a coordinate system for reaching to visual targets. *Journal of Neuroscience* 11: 1182-1197
- Cavada C, Goldman-Rakic PS (1989) Posterior parietal cortex in rhesus-monkey .2. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal-lobe. *J Comp Neurol* 287: 422-445
- Cavada C, Goldman-Rakic PS (1989) Posterior parietal cortex in rhesus monkey: I. Parcellation of areas based on distinctive limbic and sensory corticocortical connections. *J Comp Neurol* 287: 393-421
- Cisek P, Kalaska JF (2002) Simultaneous encoding of multiple potential reach directions in dorsal premotor cortex. *Journal of Neurophysiology* 87: 1149-1154
- Clifton RK, Muir DW, Ashmead DH, Clarkson MG (1993) Is visually guided reaching in early infancy a myth. *Child Dev* 64: 1099-1110
- Clifton RK, Rochat P, Robin DJ, Berthier NE (1994) Multimodal perception in the control of infant reaching. *J Exp Psychol Human* 20: 876-886
- Colby CL, Duhamel JR (1991) Heterogeneity of extrastriate visual areas and multiple parietal areas in the macaque monkey. *Neuropsychologia* 29: 517-537
- Colby CL, Duhamel JR (1996) Spatial representations for action in parietal cortex. *Cognitive Brain Research* 5: 105-115
- Colby CL, Duhamel JR, Goldberg ME (1993a) Ventral intraparietal area of the macaque - anatomic location and visual response properties. *J Neurophysiol* 69: 902-914
- Colby CL, Duhamel JR, Goldberg ME (1993b) Ventral intraparietal area of the macaque: Anatomic location and visual response properties. *J Neurophysiol* 69: 902-914
- Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. *Annual Review of Neuroscience* 22: 319-349
- Constantinidis C, Steinmetz MA (1996) Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology* 76: 1352-1355
- Constantinidis C, Steinmetz MA (2001) Neuronal responses in area 7a to multiple-stimulus displays: I. Neurons encode the location of the salient stimulus. *Cerebral Cortex* 11: 581-591
- Corbetta D, Thelen E, Johnson K (2000) Motor constraints on the development of perception-action matching in infant reaching. *Infant Behav Dev* 23: 351-374

- Crammond DJ, Kalaska JF (1996) Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. *Exp Brain Res* 108: 45-61
- Debowy DJ, Ghosh S, Ro JY, Gardner EP (2001) Comparison of neuronal firing rates in somatosensory and posterior parietal cortex during prehension. *Experimental Brain Research* 137: 269-291
- Decety J, Grezes J, Costes N, Perani D, Jeannerod M, Procyk E, Grassi F, Fazio F (1997) Brain activity during observation of actions. Influence of action content and subject's strategy. *Brain* 120: 1763-1777
- Denny-Brown D (1950) Disintegration of motor function resulting from cerebral lesions. *Journal of Nervous and Mental Disease* 112: 1-45
- Desmurget M, Grea H, Prablanc C (1998) Final posture of the upper limb depends on the initial position of the hand during prehension movements. *Exp Brain Res* 119: 511-516
- di Pellegrino G, Wise SP (1991) A neurophysiological comparison of three distinct regions of the primate frontal lobe. *Brain* 114: 951-978
- Dipellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G (1992) Understanding motor events - a neurophysiological study. *Exp Brain Res* 91: 176-180
- Dong WK, Chudler EH, Sugiyama K, Roberts VJ, Hayashi T (1994) Somatosensory, multisensory, and task-related neurons in cortical area 7b (pf) of unanesthetized monkeys. *Journal of Neurophysiology* 72: 542-564
- Driscoll JA (2000) Comparison of neural network architectures for the modeling of robot inverse kinematics. In: *IEEE Southeastcon*, pp 44-51
- D'Souza A, Vijayakumar S, Schaal S (2001) Learning inverse kinematics. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol 1, pp 298 -303
- Duhamel JR, Bremmer F, BenHamed S, Graf W (1997) Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* 389: 845-848
- Duhamel JR, Colby CL, Goldberg ME (1998) Ventral intraparietal area of the macaque: Congruent visual and somatic response properties. *J Neurophysiol* 79: 126-136
- Dum RP, Strick PL (1991) The origin of corticospinal projections from the premotor areas in the frontal lobe. *J Neurosci* 11: 667-689
- Ehrsson HH, Fagergren A, Jonsson T, Westling G, Johansson RS, Forssberg H (2000) Cortical activity in precision- versus power-grip tasks: An fmri study. *J Neurophysiol* 83: 528-536
- Eskandar EN, Assad JA (1999) Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nature Neuroscience* 2: 88-93
- Fadiga L, Fogassi L, Gallese V, Rizzolatti G (2000) Visuomotor neurons: Ambiguity of the discharge or 'motor' perception? *Int J Psychophysiol* 35: 165-177
- Fadiga L, Fogassi L, Pavesi G, Rizzolatti G (1995) Motor facilitation during action observation - a magnetic stimulation study. *J Neurophysiol* 73: 2608-2611

- Fagard J (2000) Linked proximal and distal changes in the reaching behavior of 5- to 12- month-old human infants grasping objects of different sizes. *Infant Behav Dev* 23: 317-329
- Fagg AH, Arbib MA (1998) Modeling parietal-premotor interactions in primate control of grasping. *Neural Networks* 11: 1277-1303
- Fearing RS (1986) Simplified grasping and manipulation with dexterous robot hands. *IEEE Journal of Robotics and Automation* 2: 188-195
- Ferraina S, Battaglia-Mayer A, Genovesio A, Marconi B, Onorati P, Caminiti R (2001) Early coding of visuomanual coordination during reaching in parietal area pec. *Journal of Neurophysiology* 85: 462-467
- Ferraina S, Garasto MR, Battaglia-Mayer A, Ferraresi P, Johnson PB, Lacquaniti F, Caminiti R (1997a) Visual control of hand-reaching movement: Activity in parietal area 7m. *European Journal of Neuroscience* 9: 1090-1095
- Ferraina S, Johnson PB, Garasto MR, Battaglia-Mayer A, Ercolani L, Bianchi L, Lacquaniti F, Caminiti R (1997b) Combination of hand and gaze signals during reaching: Activity in parietal area 7 m of the monkey. *Journal of Neurophysiology* 77: 1034-1038
- Flament D, Hall EJ, Lemon RN (1992) The development of cortico-motoneuronal projections investigated using magnetic brain-stimulation in the infant macaque. *J Physiol-London* 447: 755-768
- Flash T, Sejnowski TJ (2001) Computational approaches to motor control. *Curr Opin Neurobiol* 11: 655-662
- Fogassi L (1999) Mirror-like neurons in area 7b. In, personal communication. Parma, Italy
- Fogassi L, Gallese V, Buccino G, Craighero L, Fadiga L, Rizzolatti G (2001) Cortical mechanism for the visual guidance of hand grasping movements in the monkey - a reversible inactivation study. *Brain* 124: 571-586
- Fogassi L, Gallese V, Dipellegrino G, Fadiga L, Gentilucci M, Luppino G, Matelli M, Pedotti A, Rizzolatti G (1992) Space coding by premotor cortex. *Exp Brain Res* 89: 686-690
- Fogassi L, Gallese V, Fadiga L, Luppino G, Matelli M, Rizzolatti G (1996) Coding of peripersonal space in inferior premotor cortex (area f4). *J Neurophysiol* 76: 141-157
- Fogassi L, Gallese V, Fadiga L, Rizzolatti G (1998) Neurons responding to the sight of goal-directed hand/arm actions in the parietal area pf (7b) of the macaque monkey. In: 28th Annual Meeting of Society for Neuroscience, Los Angeles
- Fogassi L, Raos V, Franchi G, Gallese V, Luppino G, Matelli M (1999) Visual responses in the dorsal premotor area f2 of the macaque monkey. *Exp Brain Res* 128: 194-199
- Galea MP, Darian-Smith I (1995) Postnatal maturation of the direct corticospinal projections in the macaque monkey. *Cereb Cortex* 5: 518-540
- Gallese V (2002) Majority of hand related f5 neurons are purely motor related. In, Los Angeles
- Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. *Brain* 119: 593-609

- Gallese V, Murata A, Kaseda M, Niki N, Sakata H (1994) Deficit of hand preshaping after muscimol injection in monkey parietal cortex. *Neuroreport* 5: 1525-1529
- Gardner EP, Ro JY, Debowy D, Ghosh S (1999) Facilitation of neuronal activity in somatosensory and posterior parietal cortex during prehension. *Exp Brain Res* 127: 329-354
- Gentilucci M, Fogassi L, Luppino G, Matelli M, Camarda R, Rizzolatti G (1988) Functional organization of inferior area 6 in the macaque monkey. I. Somatotopy and the control of proximal movements. *Exp Brain Res* 71: 475-490
- Georgopoulos AP (1986) Neuronal population coding of movement direction. *Science* 233: 1416-1419
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience* 2: 1527-1537
- Geyer S, Matelli M, Luppino G, Zilles K (2000) Functional neuroanatomy of the primate isocortical motor system. *Anat Embryol* 202: 443-474
- Ghahramani Z, Wolpert DM, Jordan MI (1996) Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci* 16: 7085-7096
- Ghosh S, Brinkman C, Porter R (1987) A quantitative study of the distribution of neurons projecting to the precentral motor cortex in the monkey (m. Fascicularis). *Journal of Comparative Neurology* 259: 424-444
- Gibson EJ (1969) *Principles of perceptual learning and development*. Prentice-Hall, Englewood Cliffs, NJ
- Gibson EJ (1988) Exploratory behavior in the development of perceiving, acting and acquiring of knowledge. *Annual Review of Psychology* 39: 1-41
- Gibson JJ (1966) *The senses considered as perceptual systems*. Houghton Mifflin, Boston,
- Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70: 216-220
- Gomi H, Kawato M (1993) Neural-network control for a closed-loop system using feedback-error-learning. *Neural Networks* 6: 933-946
- Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual salience in monkey parietal cortex. *Nature* 391: 481-484
- Grafton ST, Arbib MA, Fadiga L, Rizzolatti G (1996) Localization of grasp representations in humans by positron emission tomography .2. Observation compared with imagination. *Exp Brain Res* 112: 103-111
- Grafton ST, Fadiga L, Arbib MA, Rizzolatti G (1997) Promotor cortex activation during observation and naming of familiar tools. *Neuroimage* 6: 231-236
- Grea H, Desmurget M, Prablanc C (2000) Postural invariance in three-dimensional reaching and grasping movements. *Exp Brain Res* 134: 155-162

- Halverson HM (1931) An experimental study of prehension in infants by means of systematic cinema records. *Genetic Psychology Monographs* 10: 107-285
- Hari R, Forss N, Avikainen S, Kirveskari E, Salenius S, Rizzolatti G (1998) Activation of human primary motor cortex during action observation: A neuromagnetic study. *Proceedings of the National Academy of Sciences of the United States of America* 95: 15061-15065
- Haruno M, Wolpert DM, Kawato M (2001) Mosaic model for sensorimotor learning and control. *Neural Comput* 13: 2201-2220
- Hertz J, Palmer RG, Krogh A (1991) Introduction to the theory of neural computation. Addison-Wesley Pub. Co., Redwood City, Calif.
- Hinde RA, Rowell TE (1964) Behavior socially living rhesus monkeys in their first six months. *Proc Zool Soc Lond* 143: 609-649
- Hoff B, Arbib MA (1993) Models of trajectory formation and temporal interaction of reach and grasp. *J Motor Behav* 25: 175-192
- Holden EJ (1997) Visual recognition of hand motion. In: *Computer Science*. University of Western Australia
- Hoshi E, Tanji J (2000) Integration of target and body-part information in the premotor cortex when planning action. *Nature* 408: 466-470
- Iacoboni M, Woods RP, Brass M, Bekkering H, Mazziotta JC, Rizzolatti G (1999) Cortical mechanisms of human imitation. *Science* 286: 2526-2528
- Iberall T, Arbib MA (1990) Schemas for the control of hand movements: An essay on cortical localization. In: M.A. G (ed) *Vision and action: The control of grasping*. Ablex, Norwood, NJ
- Jeannerod M, Arbib MA, Rizzolatti G, Sakata H (1995) Grasping objects - the cortical mechanisms of visuomotor transformation. *Trends Neurosci* 18: 314-320
- Jeannerod M, Decety J (1990) The accuracy of visuomotor transformation: An investigation into the mechanisms of visual recognition of objects. In: Goodale MA (ed) *Vision and action : The control of grasping*. Ablex Pub. Corp., Norwood, N.J., pp viii, 367
- Jeannerod M, Paulignan Y, Weiss P (1998) Grasping an object: One movement, several components. *Novartis Foundation Symposium* 218: 5-16; discussion 16-20
- Johansson RS, Westling G (1987a) Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. *Experimental Brain Research* 66: 141-154
- Johansson RS, Westling G (1987b) Significance of cutaneous input for precise hand movements. *Electroencephalography & Clinical Neurophysiology - Supplement* 39: 53-57
- Johansson RS, Westling G, Backstrom A, Flanagan JR (2001) Eye-hand coordination in object manipulation. *Journal of Neuroscience* 21: 6917-6932
- Jordan MI, Rumelhart DE (1992) Forward models - supervised learning with a distal teacher. *Cognitive Sci* 16: 307-354

- Kakei S, Hoffman DS, Strick PL (1999) Muscle and movement representations in the primary motor cortex. *Science* 285: 2136-2139
- Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. [see comments.]. *Nat Neurosci* 4: 1020-1025
- Kandel ER, Schwartz JH, Jessell TM (2000) *Principles of neural science*. McGraw-Hill Health Professions Division, New York
- Karniel A, Inbar GF (1997) A model for learning human reaching movements. *Biol Cybern* 77: 173-183
- Kawato M, Furukawa K, Suzuki R (1987) A hierarchical neural neural network model for control and learning of voluntary movement. *Biological Cybernetics* 57: 169-185
- Kawato M, Gomi H (1992) A computational model of 4 regions of the cerebellum based on feedback-error learning. *Biological Cybernetics* 68: 95-103
- Kincaid D, Cheney EW (1991) *Numerical analysis : Mathematics of scientific computing*. Brooks/Cole, Pacific Grove, Calif.
- Klein CA, Chujenq C, Ahmed S (1995) A new formulation of the extended jacobian method and its use in mapping algorithmic singularities for kinematically redundant manipulators. *Ieee T Robotic Autom* 11: 50-55
- Krams M, Rushworth MF, Deiber MP, Frackowiak RS, Passingham RE (1998) The preparation, execution and suppression of copied movements in the human brain. *Exp Brain Res* 120: 386-398
- Lambert P, Carron T (1999) Symbolic fusion of luminance-hue-chroma features for region segmentation. *Pattern Recogn* 32: 1857-1872
- Lantz C, Melen K, Forssberg H (1996) Early infant grasping involves radial fingers. *Dev Med Child Neurol* 38: 668-674
- Lassek A (1954) *The pyramidal tract: Its status in medicine*. Charles C. Thomas, Springfield, Illinois
- Lawrence DG, Hopkins DA (1976) The development of motor control in the rhesus monkey: Evidence concerning the role of corticomotoneuronal connections. *Brain* 99: 235-254
- Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *Journal of Comparative Neurology* 428: 112-137
- Leydold J, Hormann W (2000) Universal algorithms as an alternative for generating non-uniform continuous random variates. In: *Proc International Conference on Monte Carlo Simulation*
- Lockman J, Ashmead DH, Bushnell EW (1984) The development of anticipatory hand orientation during infancy. *Journal of Experimental Child Psychology* 37: 176-186
- Lowe DG (1991) Fitting parameterized 3-dimensional models to images. *Ieee T Pattern Anal* 13: 441-450
- Luppino G, Matelli M, Camarda RM, Gallese V, Rizzolatti G (1991) Multiple representations of body movements in mesial area-6 and the adjacent cingulate cortex - an intracortical microstimulation study in the macaque monkey. *J Comp Neurol* 311: 463-482

- Luppino G, Murata A, Govoni P, Matelli M (1999) Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas aip and vip) and the ventral premotor cortex (areas f5 and f4). *Exp Brain Res* 128: 181-187
- MacKay WA (1992) Properties of reach-related neuronal activity in cortical area 7a. *Journal of Neurophysiology* 67: 1335-1345
- MacKenzie CL, Iberall T (1994) *The grasping hand*. North-Holland, Amsterdam ; New York
- Maioli MG, Squatrito S, Samolsky-Dekel BG, Sanseverino ER (1998) Corticocortical connections between frontal periarculate regions and visual areas of the superior temporal sulcus and the adjoining inferior parietal lobule in the macaque monkey. *Brain Res* 789: 118-125
- Marconi B, Genovesio A, Battaglia-Mayer A, Ferraina S, Squatrito S, Molinari M, Lacquaniti F, Caminiti R (2001) Eye-hand coordination during reaching. I. Anatomical relationships between parietal and frontal cortex. *Cerebral Cortex* 11: 513-527
- Marteniuk RG, MacKenzie CL (1990) Invariance and variability in human prehension: Implications for theory development. In: Goodale MA (ed) *Vision and action: The control of grasping* 163-180, Norwood, NJ: Ablex
- Matelli M (1984) Interconnections within the postarcuate cortex (area 6) of the macaque monkey. *Brain Res* 310: 388-392
- Matelli M (1986) Afferent and efferent projections of the inferior area 6 in the macaque monkey. *J Comp Neurol* 251: 281-298
- Matelli M, Luppino G, Fogassi L, Rizzolatti G (1989) Thalamic input to inferior area-6 and area-4 in the macaque monkey. *J Comp Neurol* 280: 468-488
- Matelli M, Luppino G, Murata A, Sakata H (1994) Independent anatomical circuits for reaching and grasping linking the inferior parietal sulcus and inferior area 6 in macaque monkey. In: *Society for Neuroscience, vol 20*, p 404.404
- Matelli M, Luppino G, Rizzolatti G (1991) Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. *J Comp Neurol* 311: 445-462
- Maunsell JH (1995) The brain's visual world: Representation of visual targets in cerebral cortex. *Science* 270: 764-769
- Miall RC, Wolpert DM (1996) Forward models for physiological motor control. *Neural Networks* 9: 1265-1279
- Motter BC, Mountcastle VB (1981) The functional properties of the light-sensitive neurons of the posterior parietal cortex studied in waking monkeys: Foveal sparing and opponent vector organization. *Journal of Neuroscience* 1: 3-26
- Motter BC, Steinmetz MA, Duffy CJ, Mountcastle VB (1987) Functional properties of parietal visual neurons: Mechanisms of directionality along a single axis. *Journal of Neuroscience* 7: 154-176
- Muir RB, Lemon RN (1983) Corticospinal neurons with a special role in precision grip. *Brain Research* 261: 312-316
- Murata A (2000) Parietal neurons discriminating allocentric fixation locations. In, Kyoto, Japan

- Murata A, Fadiga L, Fogassi L, Gallese V, Raos V, Rizzolatti G (1997a) Object representation in the ventral premotor cortex (area f5) of the monkey. *J Neurophysiol* 78: 2226-2230
- Murata A, Fadiga L, Fogassi L, Gallese V, Rizzolatti G (1997b) Visuomotor properties of grasping neurons of inferior area 6. *Pflug Arch Eur J Phy* 434: 73-73
- Murata A, Gallese V, Kaseda M, Sakata H (1996) Parietal neurons related to memory-guided hand manipulation. *J Neurophysiol* 75: 2180-2186
- Murata A, Gallese V, Luppino G, Kaseda M, Sakata H (2000) Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area aip. *J Neurophysiol* 83: 2580-2601
- Neal JW, Pearson RC, Powell TP (1990) The ipsilateral cortico-cortical connections of area 7b, pf, in the parietal and temporal lobes of the monkey. *Brain Research* 524: 119-132
- Newell KM (1986) Motor development in children: Aspects of coordination and control. In: Wade MG, Whiting HTA (eds) *Motor development in children : Aspects of coordination and control*. Nijhoff, Boston, pp 341-360
- Newell KM, Scully DM, McDonald PV, Baillargeon R (1989) Task constraints and infant grip configurations. *Developmental Psychobiology* 22: 817-831
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341: 52-54
- Nishitani N, Hari R (2000) Temporal dynamics of cortical representation for action. *Proceedings of the National Academy of Sciences of the United States of America* 97: 913-918
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34: 171-175
- Olivier E, Edgley SA, Armand J, Lemon RN (1997) An electrophysiological study of the postnatal development of the corticospinal system in the macaque monkey. *J Neurosci* 17: 267-276
- Oyama E, Agah A, Chong NY, Maeda T (2001a) Inverse kinematics learning by modular architecture neural networks with performance prediction networks. In: *IEEE International Conference on Robotics and Automation*, vol 1
- Oyama E, Agah A, MacDorman KF, Maeda T, Tachi S (2001b) A modular neural network architecture for inverse kinematics model learning. *Neurocomputing* 38: 797-805
- Oztop E (2000) *Neurobench*. In, Los Angeles
- Pandya DN, Seltzer B (1982) Intrinsic connections and architectonics of posterior parietal cortex in the rhesus monkey. *Journal of Comparative Neurology* 204: 196-210
- Perret DI, Harries MH, Benson PJ, Chitty AJ, Mistlin AJ (1990a) Retrieval of structure from rigid and biological motion: An analysis of the visual responses of neurons in the macaque temporal cortex. In: Blake A, Troscianko T (eds) *Ai and the eye*. John Wiley and Sons Ltd., pp 181-199
- Perret DI, Mistlin AJ, Harries MH, Chitty AJ (1990b) Understanding the visual appearance and consequence of hand actions. In: Goodale MA (ed) *Vision and action: The control of grasping* 163-180, Norwood, NJ: Ablex, pp 163-180

- Raos V, Franchi G, Fogassi L, Gallese V, Luppino G, Matelli M (1998) Functional organization of area f2 in the monkey. *Eur J Neurosci* 10: 87-87
- Ratcliff G (1991) Brain and space: Some deductions from clinical evidence. In: Paillard J (ed) *Brain and space*. Oxford University Press, Oxford England; New York
- Rizzolatti G, Arbib MA (1998) Language within our grasp. *Trends Neurosci* 21: 188-194
- Rizzolatti G, Arbib MA (1999) From grasping to speech: Imitation might provide a missing link - reply. *Trends Neurosci* 22: 152-152
- Rizzolatti G, Camarda R, Fogassi L, Gentilucci M, Luppino G, Matelli M (1988) Functional organization of inferior area 6 in the macaque monkey. II. Area f5 and the control of distal movements. *Exp Brain Res* 71: 491-507
- Rizzolatti G, Fadiga L (1998) Grasping objects and grasping action meanings: The dual role of monkey rostraventral premotor cortex (area f5). In: *Sensory guidance of movement, novartis foundation symposium* 218. Wiley, Chichester, pp 81-103
- Rizzolatti G, Fadiga L, Gallese V, Fogassi L (1996a) Premotor cortex and the recognition of motor actions. *Cognitive Brain Res* 3: 131-141
- Rizzolatti G, Fadiga L, Matelli M, Bettinardi V, Paulesu E, Perani D, Fazio F (1996b) Localization of grasp representations in humans by pet .1. Observation versus execution. *Exp Brain Res* 111: 246-252
- Rizzolatti G, Fogassi L, Gallese V (2000) Mirror neurons: Intentionality detectors? *Int J Psychol* 35: 205-205
- Rizzolatti G, Fogassi L, Gallese V (2001a) Neurophysiological mechanisms underlying the understanding and imitation of action. *Nat Rev Neurosci* 2: 661-670
- Rizzolatti G, Gallese V (2001) Mirror neuron electrophysiological recording data (raw). In. *Personal Communication*
- Rizzolatti G, Gallese V, Fogassi L, Keysers C (2001b) Experimental setup of mirror neuron recordings. In. *Personal Communication*
- Rizzolatti G, Luppino G, Matelli M (1998) The organization of the cortical motor system: New concepts. *Electroen Clin Neuro* 106: 283-296
- Ro JY, Debowy D, Ghosh S, Gardner EP (2000) Depression of neuronal firing rates in somatosensory and posterior parietal cortex during object acquisition in a prehension task. *Experimental Brain Research* 135: 1-11
- Robinson CJ, Burton H (1980a) Organization of somatosensory receptive fields in cortical areas 7b, retroinsula, postauditory and granular insula of m. Fascicularis. *Journal of Comparative Neurology* 192: 69-92
- Robinson CJ, Burton H (1980b) Somatic submodality distribution within the second somatosensory (sII), 7b, retroinsular, postauditory, and granular insular cortical areas of m. Fascicularis. *Journal of Comparative Neurology* 192: 93-108

- Rochat P, Morgan R (1995) Spatial determinants in the perception of self-produced leg movements by 3-5 month-old infants. *Developmental Psychology* 31: 626-636
- Rosenbaum DA (1991) *Human motor control*. Academic Press, San Diego
- Rosenbaum DA, Loukopoulos LD, Meulenbroek RG, Vaughan J, Engelbrecht SE (1995) Planning reaches by evaluating stored postures. *Psychological Review* 102: 28-67
- Rosenbaum DA, Meulenbroek RGJ, Vaughan J, Jansen C (1999) Coordination of reaching and grasping by capitalizing on obstacle avoidance and other constraints. *Exp Brain Res* 128: 92-100
- Rosenbaum DA, Meulenbroek RJ, Vaughan J, Jansen C (2001) Posture-based motion planning: Applications to grasping. *Psychol Rev* 108: 709-734
- Rothwell JC (1994) *Control of human voluntary movement*. Chapman & Hall, London ; New York
- Roy AC, Paulignan Y, Farne A, Jouffrais C, Boussaoud D (2000) Hand kinematics during reaching and grasping in the macaque monkey. *Behav Brain Res* 117: 75-82
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning internal representations by error propagation. In: Rumelhart DE, McClelland JL, group aP (eds) *Parallel distributed processing, vol 1: Foundations*, pp 151-193
- Rushworth MF, Nixon PD, Passingham RE (1997) Parietal cortex and movement. I. Movement selection and reaching. *Experimental Brain Research* 117: 292-310
- Russ JC (1998) *The image processing handbook*. CRC Press, Boca Raton, FL
- Sakata H, Taira M, Kusunoki M, Murata A, Tanaka Y (1997a) The tins lecture - the parietal association cortex in depth perception and visual control of hand action. *Trends Neurosci* 20: 350-357
- Sakata H, Taira M, Kusunoki M, Murata A, Tanaka Y, Tsutsui K (1998) Neural coding of 3d features of objects for hand action in the parietal cortex of the monkey. *Philos T Roy Soc B* 353: 1363-1373
- Sakata H, Taira M, Kusunoki M, Murata A, Tsutsui K, Tanaka Y, Shein WN, Miyashita Y (1999) Neural representation of three-dimensional features of manipulation objects with stereopsis. *Exp Brain Res* 128: 160-169
- Sakata H, Taira M, Murata A, Gallese V, Tanaka Y, Shikata E, Kusunoki M (1997b) Parietal visual neurons coding three-dimensional characteristics of objects and their relation to hand action. In: Their P, Karnath HO (eds) *Parietal lobe contributions to orientation in 3d space*. Springer-Verlag, Heidelberg
- Sakata H, Taira M, Murata A, Mine S (1995) Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey. *Cereb Cortex* 5: 429-438
- Salimi I, Brochier T, Smith AM (1999a) Neuronal activity in somatosensory cortex of monkeys using a precision grip. I. Receptive fields and discharge patterns. *Journal of Neurophysiology* 81: 825-834
- Salimi I, Brochier T, Smith AM (1999b) Neuronal activity in somatosensory cortex of monkeys using a precision grip. Iii. Responses to altered friction perturbations. *Journal of Neurophysiology* 81: 845-857

- Sanger TD (1996) Probability density estimation for the interpretation of neural population codes. *J Neurophysiol* 76: 2790-2793
- Schaal S, Atkeson CG, Vijayakumar S (2002) Scalable techniques from nonparametric statistics for real time robot learning. *Applied Intelligence* (in press)
- Sciavicco L, Siciliano B (2000) *Modelling and control of robot manipulators*. Springer, London ; New York
- Scott SH, Sergio LE, Kalaska JF (1997) Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. *J Neurophysiol* 78: 2413-2426
- Sergio LE, Kalaska JF (1997) Systematic changes in directional tuning of motor cortex cell activity with hand location in the workspace during generation of static isometric forces in constant spatial directions. *J Neurophysiol* 78: 1170-1174
- Shadlen MN, Newsome WT (1996) Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America* 93: 628-633
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J Neurophysiol* 86: 1916-1936
- Shen LM, Alexander GE (1997a) Neural correlates of a spatial sensory-to-motor transformation in primary motor cortex. *J Neurophysiol* 77: 1171-1194
- Shen LM, Alexander GE (1997b) Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *J Neurophysiol* 77: 1195-1212
- Siegel RM (1998) Representation of visual space in area 7a neurons using the center of mass equation. *Journal of Computational Neuroscience* 5: 365-381
- Siegel RM, Read HL (1997) Analysis of optic flow in the monkey parietal area 7a. *Cerebral Cortex* 7: 327-346
- Smeets JB, Brenner E (1999) A new view on grasping. [see comments.]. *Motor Control* 3: 237-271
- Smeets JB, Brenner E (2001) Independent movements of the digits in grasping. *Experimental Brain Research* 139: 92-100
- Snippe HP (1996) Parameter extraction from population codes: A critical assessment. *Neural Comput* 8: 511-529
- Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. [see comments.]. *Nature* 386: 167-170
- Snyder LH, Batista AP, Andersen RA (2000) Intention-related activity in the posterior parietal cortex: A review. *Vision Research* 40: 1433-1441
- Sonka M, Hlavac V, Boyle R (1993) *Image processing, analysis, and machine vision*. Chapman & Hall Computing, London ; New York
- Sporns O, Edelman GM (1993) Solving bernstein's problem: A proposal for the development of coordinated movement by selection. *Child Development* 64: 960-981

- Sporns O, Edelman GM, Meijer OG (1998) Bernstein's dynamic view of the brain: The current problems of modern neurophysiology (1945). *Motor Control* 2: 283-305
- Sporns O, Tononi G, Edelman GM (2000) Connectivity and complexity: The relationship between neuroanatomy and brain dynamics. *Neural Networks* 13: 909-922
- Stein JF (1991) Space and the parietal association areas. In: Paillard J (ed) *Brain and space*. Oxford University Press, Oxford England; New York
- Streri A (1993) *Seeing, reaching, touching : The relations between vision and touch in infancy*. Harvester Wheatsheaf, London ; New York
- Sutton RS, Barto AG (1998) *Reinforcement learning : An introduction*. MIT Press, Cambridge, Mass.
- Taira M, Mine S, Georgopoulos AP, Murata A, Sakata H (1990) Parietal cortex neurons of the monkey related to the visual guidance of hand movement. *Exp Brain Res* 83: 29-36
- Tang WS, Wang J (2001) A recurrent neural network for minimum infinity-norm kinematic control of redundant manipulators with an improved problem formulation and reduced architecture complexity. *Ieee T Syst Man Cy B* 31: 98-105
- Thelen E (2000) Motor development as foundation and future of developmental psychology. *Int J Behav Dev* 24: 385-397
- Triggs WJ, Yathiraj S, Young MS, Rossi F (1998) Effects of task and task persistence on magnetic motor-evoked potentials. *Journal of Contemporary Neurology* 1998 (electronic journal: <http://scholar.lib.vt.edu/ejournals/JCN/ncn-mirror/articles/003/Triggs.pdf>)
- Twitchell TE (1970) Reflex mechanisms and the development of prehension. In: Connolly KJ (ed) *Mechanisms of motor skill development*. Academic Press, London, New York
- Umiltà MA, Kohler E, Gallese V, Fogassi L, Fadiga L, Keysers C, Rizzolatti G (2001) I know what you are doing: A neurophysiological study. *Neuron* 31: 155-165
- van der Meer AL, van der Weel FR, Lee DN (1995) The functional significance of arm movements in neonates. *Science* 267: 693-695
- von Hofsten C (1982) Eye-hand coordination in the newborn. *Developmental Psychology* 18: 450-461
- von Hofsten C (1984) Developmental changes in the organization of prereaching movements. *Developmental Psychology* 20: 378-388
- von Hofsten C (1993) The structuring of neonatal arm movements. *Child Dev* 64: 1046-1057
- von Hofsten C, Ronnqvist L (1988) Preparation for grasping an object: A developmental study. *Journal of Experimental Psychology: Human Perception & Performance* 14: 610-621
- Whitney DE (1969) Resolved motion rate control of manipulators and human prostheses. *IEEE Transactions on Man-Machine Systems* 10: 47-53
- Williams RJ (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8: 229-256
- Willis WD, Coggeshall RE (1991) *Sensory mechanisms of the spinal cord*. Plenum Press, New York

- Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: Corticocortical connectivity and combinatorial computations. *Annual Review of Neuroscience* 20: 25-42
- Wolpert DM, Ghahramani Z (2000) Computational principles of movement neuroscience. *Nat Neurosci* 3 Suppl: 1212-1217
- Wolpert DM, Ghahramani Z, Flanagan JR (2001) Perspectives and problems in motor learning. *Trends Cogn Sci* 5: 487-494
- Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317-1329
- Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. *Trends Cogn Sci* 2: 338-347
- Wu CW, Bichot NP, Kaas JH (2000) Converging evidence from microstimulation, architecture, and connections for multiple motor areas in the frontal and cingulate cortex of prosimian primates. *J Comp Neurol* 423: 140-177
- Zemel RS, Dayan P, Pouget A (1998) Probabilistic interpretation of population codes. *Neural Comput* 10: 403-430

11 APPENDIX

11.1 Mirror neuron system model (MNS)

The system was implemented using Java programming language on a Linux operating system. The grasp simulator can be accessed using the URL <http://www-clmc.usc.edu/erhan/models/MNS>. The simulation applet at this URL also includes a simplified version of the MNS model. The applet enables the users to test the action recognition ability of the model.

11.1.1 Color segmentation

The segmentation system as a whole works as follows:

1. Start with N rectangles (called nodes), set thresholds for red, green and blue variances as rV , gV , bV
2. For each node calculate the red, green, blue variance as rv , gv , bv
3. If any of the variance is higher than the threshold ($rv > rV$ or $gv > gV$ or $bv > bV$) then split the node into four equal pieces and apply step 2 and step 3 recursively
4. Feed in the mean red, green and blue values in that region to the Color Expert to determine the color of the node.
5. Make a list of nodes that are of the same color (add node to the list reserved for that color).
6. Repeat 2-5 until no split occurs.
7. Cluster (in terms of Euclidean distance on the image) the nodes and discard the outliers from the list (use the center of the node as the position of the node). The discarding is performed either when a region is very far from the current mean (weighted center) or it is not "connected" to the current center position. The connectedness is defined as follows. The regions A and B are connected if the points lying on the line segment joining the centers of A and B are the same color as A and B. Once again, the Color Expert is used to determine the percentage of the correct (colors of A and B) colors lying on the line segment. If this percentage is over a certain threshold (e.g. 70%) then the regions A and B are taken as "connected". (This strategy would not work for a "sausage-shaped" region, but does work for patches created by the coloring we used in the glove.)
8. For each pruned list (corresponding to a color) find the weighted (by the area of the node) mean of the clusters (in terms of image coordinate).
9. Return the cluster mean coordinates as the segmented regions center.

So we do not exactly perform the merge part of the split-merge algorithm. The return values from this procedure are the (x,y) coordinates of the center of color patches found. Another issue is how to choose the thresholds. The variance values are not very critical. A too small value increases computation time but does not affect the number of colors extracted correctly (though the returned coordinates may be shifted slightly). To see why intuitively, one can notice that the center of a rectangle and the centroid of the centers of the quarter rectangles (say after a split operation) would be the same. This means that if a region is split unnecessarily (because the threshold variances were set to very small

values) it is likely to be averaged out with our algorithm since it is likely that the four split rectangles will have the same color and will be connected (with our definition of connectedness)

11.1.2 Reach and grasp schema precision grasp planning and execution

- Determine the opposition axis to grasp the object.
- Compute the two (outer) points A and B at which the opposition axis intersects the object surface. They serve as the contact points for the virtual fingers that will be involved in the grasp.
- Assign the real fingers to virtual fingers. The particular heuristic we used in the experiments was the following. If the object is on the right [left] with respect to the arm then the thumb is assigned to the point A if A is on the left of [at a lower level than] B otherwise the thumb is assigned to B. The index finger is assigned to the remaining point.
- Determine an approximate target position C, for the wrist. Mark the target for the wrist on the line segment connecting the current position of the wrist and the target for the thumb a fixed length (determined by the thumb length) away from the thumb target.
- Solve the inverse kinematics for only the wrist reach (ignore the hand).
- Solve the inverse kinematics for grasping. Using the sum of distance squares of the finger tips to the target contact points do a random hill climbing search to minimize the error. Note that the search starts with placing the wrist at point C. However, the wrist position is not included in the error term.
- The search stops when the simulator finds a configuration with error close to zero (success) or after a fixed number of steps (failure to reach). In the success case the final configuration is returned as the solution for the inverse kinematics for the grasp. Otherwise failure-to-reach is returned.

Execute the reach and grasp. At this point the simulator knows the desired target configuration in terms of joint angles. So what remains to be done is to perform the grasp in a realistic way (in terms of kinematics). The simplest way to perform the reach is to linearly change the joint angles from the initial configuration to the target configuration. But this does not produce a bell shaped velocity profile (nor exactly a constant speed profile either because of the non-linearity in going from joint angles to end effector position). The perfect way to plan an end-effector trajectory requires the computation of the Jacobian. However we are not interested in perfect trajectories as long as the target is reached with a bell-shaped velocity profile. To get the effect it is usually sufficient to modify the idea of linearly changing the joint angles little bit. We simply modulate the change of time by replacing time with a third order polynomial that will match our constraints for time (starts at 0 climbs up to 1 monotonically). Note that we are still working in the joint space and our method may suffer from the non-linearity in transforming the joint angles to end effector coordinates. However, our empirical studies showed that a satisfactory result, for our purposes, could be achieved in this way.

11.2 Learning to grasp models (ILGM and LGM)

The simulation environment is developed using Java programming language. The simulation environment consists of a 3D kinematics arm model (see Chapter 3) and routines implementing the LGM circuit. Note that from an implementation point of view the learning to grasp circuits of Chapters 5 and 6 are essentially the same. In this section, we present a brief overview of the simulation environment. The further details are available at the URL <http://www-clmc.usc.edu/erhan/models/LGM>.

The main behavior of the simulation system is determined by the resource file 'KolParameters.res'. Many simulation parameters can be set within this file. However, some of the parameters are hard coded in the Java class files comprising the simulation system. The 3D positions and vectors are defined using a spherical reference frame. The PAR, MER and RAD tags are used to indicate elevation, azimuth and radius components respectively.

Object position range parameters define the range of the object position in the workspace. In the following segment of text, the first two correspond to the range of the elevation; the following two corresponds to the range of azimuth and finally the last two corresponds to the range of radius.

```
minPAR      -45
maxPAR      45
minMER      -45
maxMER      45
minRAD      700
maxRAD      1300
```

Object location coding length parameters define how many units to allocate for representing each position component in the affordance layer. In the below setting, the object location will be represented with 100 units.

```
obj_locPAR_code_len  10
obj_locMER_code_len  10
obj_locRAD_code_len  1
```

Object location coding variance defines the variance for population encoding of the object location.

```
obj_encode_var      1
```

Object axis orientation range parameters define the minimum and maximum allowed tilt of the object around the z-axis (in the frontal plane). The number of units allocated for encoding the tilts (in three coordinate axes) is hardwired in the file Motor.java

```
minTILT      0
maxTILT      90
```

Base learning rate parameter is used as the common multiplier for all the learning rates in the grasp learning circuit.

```
eta          0.5
```

LGM layer length parameters define the number of units to allocate for the layers generating the motor parameters. BANK, PITCH and HEADING tags indicate the supination-pronation, wrist extension-flexion and radial/ulnar deviation movements respectively. (The Virtual Fingers layer in these simulations is not engaged fully as it is represented as a layer of ten units encoding the synergistic enclosure speed of the hand.)

```
hand_rotBANK_code_len  9
hand_rotPITCH_code_len 9
hand_rotHEADING_code_len 1
```

The tags locMER, locPAR and locRAD indicate the approach direction vector components.

```
hand_locMER_code_len  7
hand_locPAR_code_len  7
hand_locRAD_code_len  1
```

Learning session parameters define the behavior of the simulator during learning. For a learning session, the simulator makes MAXBABBLE number of reach/grasp attempts. For each approach-direction, the simulator makes MAXROTATE grasping attempts. After MAXREACH reaches are done, the next input condition is selected (e.g. the object is placed in a different position). After each weightSave reach/grasp trials the learned connections are written to disk. Note that MAXBABBLE only limits the maximum number of attempts the simulator will make. A particular simulation may be stopped at any instant. The saved connection weights then can be used for testing the performance at a later time. Reach2Target parameter indicates which part of the hand should be used in reaching for the object. The possible values are {INDEX, MIDDLE, THUMB} x {0,1,2} where 2 indicates the tip.

```
MAXREACH      5
MAXROTATE     7
MAXBABBLE     10000
weightSave    4500
Reach2Target  INDEX1
```

Grasp stability parameters define the acceptable grasps in terms of physical stability. costThreshold specifies the allowable inaccuracy in grasping. Ideally, the cost of grasping (a measure of the instability, see Chapter 5.5) should be small indicating that the grasp is successful. Empirically, a value between 0.5 and 0.8 gives a good result for the implemented cost function. If the distance of the touched object to the palm is less than palmThreshold and the movement of the object due to finger contact is towards the palm then the palm is used as a virtual finger to counteract the force exerted by the fingers. The negReinforcement parameter specifies the level of punishment returned when a grasp attempt fails. Empirically values greater than -0.1 and less than 0 result in good learning. Generally, a large negative reinforcement overwhelms the positively reinforced plans before they have chance to get represented in the layers.

```
costThreshold  0.8
palmThreshold  150
negReinforcement -0.05
```

Exploration and exploitation parameters specify how often to use the learned distribution to generate grasp plans. A value of 1 means always use random parameter selection, while a value of 0 means always generate parameters from the current distribution of the layer. The tag 'rot' indicates the Wrist Rotation layer while 'off' indicates the Hand Position Layer.

```
rotRandomness  1      # 1=full random 0=from the pdf
offRandomness  1      # 1=full random 0=from the pdf
```

Now we present the simulation parameters used for learning to grasp experiments presented in the thesis. The default values of the parameters given in the descriptions above will not be repeated.

Simulation parameters for sections 6.6.1 and 6.6.2, and Figure 5.8

```
hand_rotBANK_code_len  11
hand_rotPITCH_code_len 11
hand_rotHEADING_code_len 6
hand_locMER_code_len   6
hand_locPAR_code_len   6
hand_locRAD_code_len   1
MAXREACH                N/A
MAXROTATE               55
MAXBABBLE               50000 (stopped at 16000)
```

weightSave	16000
Reach2Target	MIDDLE0
costThreshold	0.75
palmThreshold	125
negReinforcement	-0.1
rotRandomness	0.95
offRandomness	0.95

Simulation parameters for section 5.7 (except Figure 5.8)

hand_rotBANK_code_len	11
hand_rotPITCH_code_len	11
hand_rotHEADING_code_len	6
hand_locMER_code_len	6
hand_locPAR_code_len	6
hand_locRAD_code_len	1
MAXREACH	N/A
MAXROTATE	55
MAXBABBLE	50000 (stopped at 10000)
weightSave	5000
Reach2Target	INDEX0
costThreshold	0.75
palmThreshold	125
negReinforcement	-0.1
rotRandomness	0.95
offRandomness	0.95

Simulation parameters for sections 5.9 and 6.7

hand_rotBANK_code_len	12
hand_rotPITCH_code_len	7
hand_rotHEADING_code_len	1
hand_locMER_code_len	1
hand_locPAR_code_len	1
hand_locRAD_code_len	1
MAXREACH	1
MAXROTATE	25
MAXBABBLE	200000 (stopped at 20000)
weightSave	2500
Reach2Target	MIDDLE0
costThreshold	0.85
palmThreshold	150
negReinforcement	-0.1
rotRandomness	1
offRandomness	1

Simulation parameters for sections 5.8

hand_rotBANK_code_len	9
hand_rotPITCH_code_len	9
hand_rotHEADING_code_len	1
hand_locMER_code_len	7
hand_locPAR_code_len	7
hand_locRAD_code_len	1
MAXREACH	1
MAXROTATE	7
MAXBABBLE	200000 (stopped at 45000)
weightSave	4500
Reach2Target	INDEX2
costThreshold	0.80

palmThreshold	150
negReinforcement	-0.05
rotRandomness	1
offRandomness	1

Simulation parameters for section 6.9

maxRAD	1200
hand_rotBANK_code_len	10
hand_rotPITCH_code_len	7
hand_rotHEADING_code_len	5
hand_locMER_code_len	10
hand_locPAR_code_len	10
hand_locRAD_code_len	1
MAXREACH	10
MAXROTATE	30
MAXBABBLE	200000 (stopped at 50000)
weightSave	10000
Reach2Target	MIDDLE0
costThreshold	0.75
palmThreshold	128
negReinforcement	-0.1
rotRandomness	1
offRandomness	1

Simulation parameters for section 6.8

hand_rotBANK_code_len	9
hand_rotPITCH_code_len	9
hand_rotHEADING_code_len	1
hand_locMER_code_len	7
hand_locPAR_code_len	7
hand_locRAD_code_len	1
MAXREACH	10
MAXROTATE	10
MAXBABBLE	200000 (stopped at 50000)
weightSave	4500
Reach2Target	INDEX1
costThreshold	0.8
palmThreshold	150
negReinforcement	-0.05
rotRandomness	0.95
offRandomness	0.95