A Developmental Approach for Pre–Orientation in Humanoid Robots

by

Wouter Aupers

Submitted to the Department of Artificial Intelligence in partial fulfillment of the requirements for the degree of

Doctorandus of Artificial Intelligence

at the

UNIVERSITY OF GRONINGEN

August 2002

Certified by: Prof. Dr. L.R.B. Schomaker
Director Research and Education Department of Artificial Intelligence, University of Groningen
Thesis Supervisor

Accepted by: Dr. P. Hendriks
Researcher/Lecturer at the Department of Artificial Intelligence, University of Groningen

A Developmental Approach for Pre–Orientation in Humanoid Robots

by

Wouter Aupers

Submitted to the Department of Artificial Intelligence on July. 24, 2002, in partial fulfillment of the requirements for the degree of *Doctorandus* of Artificial Intelligence

Abstract

The aim of the research presented in this paper was to show how the developmental perspective for developing humanoid robots could be extended to the pre–grasp orientation of the wrist, based on evidence from development in infants and on the Goodale and Milner hypothesis about dorsal–stream information processing for reaching and pre–grasp orienting in the brain of adults. The model that has been developed features a learning mechanism that associates low–level visual information on the retinal orientation of objects with motor–commands for orienting the wrist, using a posturing task for the hand to learn this association. The model starts with a simple reflex and noisy motor–command generation. Through proprioceptic feedback on the success of the posturing task it gradually builds a more precise mapping of the different motor–commands associated with visually observed orientations of the objects. The model shows, as can be observed in infants, that the orientation behavior of the hand only starts to develop when reaching is precise enough and thus provides proprioceptic feedback on the success of the posturing task for the hand.

Thesis Supervisor: Prof. Dr. L.R.B. Schomaker
Title: Director Research and Education Department of Artificial Intelligence, University of Groningen

Acknowledgments

As part of the requirement for my degree I had to do a project over the course of a half year. I have long had the wish to live abroad for a while, and this fact, combined with the fact that a lot of interesting research in the field of humanoid and biologically inspired robotics is being done outside of the Netherlands, gave me two good reasons to move abroad and live in a foreign country for a while. From December 2001 until may 2002 I have been working on my project for the masters thesis at the LIRA–lab in Genoa on the subject of pre–grasp orientation in a developmental robotic architecture called *Babybot*.

I would like to thank professor Sandini, Lorenzo, Ingrid and all my other colleagues at the LIRA–lab for a wonderful and inspiring time. My special gratitude goes out to professor Sandini, who has given me the opportunity to do the research for my thesis in his lab and also invited me to come along to Lisbon to a scientific meeting. Also Lorenzo deserves a special mention for the many helpful conversations on the architecture of Babybot and the many hours he spent helping me understand the software of *Babybot*.

I would also like to thank professor Schomaker for his guidance and support in writing this thesis after I returned to Groningen.

As I have had a wonderful time in Genoa, and enjoyed working on my project very much, I hope that this thesis will be a good reflection of the energy I have put into the research and in this thesis.

Groningen, 24 July 2002

Wouter Aupers

Contents

List of Figures

Chapter 1

Introduction

This thesis presents the research that has led to an extension of the developmental architecture of *Babybot* with a model for the learning of pre–grasp orientation of the hand.

Babybot is a humanoid robot being developed by LIRA–lab at the university of Genoa. At the LIRA–lab the emphasis in the research lies on implementing architectures of information processing in robotic systems that are based on physiological, neurological and psychological evidence in primates and humans. This approach is often referred to as *Biologically inspired* or *Humanoid* robotics. What sets the approach of the LIRA–lab apart is a *developmental* perspective, which is also the basis for the research that is presented in this thesis.

1.1 Research goals

In a current project at LIRA–lab called *Mirror* an architecture that shows a real-time implementation of mirror learning [1], is being de-

[1] Mirror neurons are neurons that show activity in primates during grasping actions on specific objects, not only during execution of the grasp by the primate, but also when the primate observes a similar grasping movement being executed by another person.

veloped. In the current architecture of the *Babybot* a developmental architecture of reaching behavior has been implemented. This reaching architecture basically uses three degrees of freedom in the arm of the robot, which makes *pointing*–like behavior possible. As part of the *Mirror* project an *end-effector* in the form of a functionally human–like hand will have to be added to the existing architecture of *Babybot*. Since there is no end–effector present in the *Babybot* at the moment, some knowledge will have to be developed on the control–architecture for pre–grasp orientation of the end–effector, controlling additional degrees of freedom in the arm, on top of the degrees of freedom that are already being used for the reaching.

This thesis presents *the research on the control–architecture for pre-grasp–orientation, controlling an extra degree of freedom, with a future end–effector and provides an implementation of the learning mechanism involved*. Furthermore an experiment will be used to demonstrate the workings of this architecture. For the development of this control architecture the same constraints in the sense of developmental perspective and biological inspiration, as have been used for the previous work on *Babybot*, will be used.

1.2 Approach

In the pursuit of the research goals, an overview of the developmental approach will be given. This overview will provide constraints on the type of model that can be developed. After the developmental robotics has been positioned, an explanation of the current architecture of *Babybot* will be given to illustrate the basic principles of the developmental approach in an existing system. Having introduced the *Babybot* an analysis of the pre–grasp orientation behavior in children will be presented. Since the information processing for the pre–grasp orientation behavior has visual and motor components, the next step will be to examine closer the information processing for the visual and motor–

system. Finally the insights gathered will be analysed and combined into a working model that can be implemented on *Babybot*.

1.3 Relevance

The research goals that have been laid down in the previous section fit in a trend in the field of *Artificial Intelligence* (AI) away from trying to find algorithms that mimic human problem solving capabilities, toward a different paradigm. This different paradigm tries to understand how certain cognitive capacities of humans are built–up, and are learned. In this view the interaction with the environment is essential for the capacities of a system to develop. The environment presents a rich source of sensory signals, and the fact that interaction with the world changes things in the world drives the learning. Because of the closer attention that is being paid to how human capacities are built up and learned, the research is becoming more and more interdisciplinary. Increasingly research from neurology, psychology and physiology are being used as inspiration for architectures of artificial systems.

The research that is presented here fits in this approach. It brings together evidence from developmental psychology, physiology and neurology in a robotic system that in itself forms a theory of how the human information processing works during the early stages of development of sensory–motor coordination.

1.4 About this manuscript

In the section on the approach, the basis of the structure of this manuscript has already been laid–out. In chapter 2 the paradigm of *developmental robotics* will be discussed in detail. Chapter 3 gives an overview of the developmental architecture as it is implemented in *Babybot*, followed by an analysis of the pre–grasp orientation in infants in chapter 4. Chapter 5 discusses how the visual and motor system are functioning in the brain. In chapter 6 the insights from the previous chapters will be combined into a system for learning pre–grasp orientation. This thesis concludes with some experimental results on the learning in chapter 7 and the general conclusion of this thesis in chapter 8.

Chapter 2

Developmental Robotics

2.1 Introduction

Development is a term that is not unambiguous if used in relation to robotic systems. If we take a look at the development of infants it is clear that a robotic system is not capable of the same kind of development in terms of physically growing. A robotic system starts with a given set of hardware that is fixed in size and number of components over time. In spite of these differences there are also commonalities. One of these commonalities is the term *development* that can be viewed as *change over time* in a system. Apart from the physical changes, infants change in terms of control over their motor–system and the growth of more sophisticated cognitive capabilities. The growth of control over the motor–system is related to the physical growth of the body, but the basic components of the physical makeup do not change over time. The similarity in terms of the changing control over the motor–system and the onset of cognitive abilities leaves the opportunity to model the changes that take place in infants in a robotic system. As an approach to building robots, *developmental robotics* tries, through careful examination of developmental processes in infants, to identify a set of initial behaviors, and the developmental rules that operate on these initial behaviors in order to create a robotic system that can autonomously develop toward complex human–like behaviors.

In this chapter the *developmental robotics* approach and its historical backgrounds will be described. It starts with an overview of how robotics has evolved in the field of *Artificial Intelligence* (AI) toward the developmental perspective, followed by a more detailed explanation of what this approach entails in terms of system development and how some basic changes in behavior during the development of infants are related to each other and facilitate the development of other behaviors.

2.2 Background

In the introduction to this chapter, the research paradigm of *developmental robotics* has been outlined. This paradigm though has been a relatively recent addition to the field of Artificial Intelligence (AI). The classical AI approach has been one of building complete mathematical models of the world, and having robotic–agents operate in these worlds. Brooks [9] describes in a 1997 article a move away from model–based approaches in the field of AI.

The first step in moving away from model–based approaches has been toward a behavior–based perspective. This first step has been partially motivated by a change in perspective on what intelligence actually is. Where model–based approaches view intelligence mainly as a function that takes an input and produces an output internal to an agent, the behavior–based approach emphasizes the dynamical interaction of an agent that is situated in the real world with its environment as a measure of the intelligence of that agent. Another aspect that closely relates to this point is concerned with what is actually designed into the system. In the model–based approach behaviors are emergent from the design of a planner and the goals that this planner uses, operating in the model. In the behavior–based approaches however, behavioral responses are designed into the system, which leads to the observation of *dynamically–selected* goals and emergent behavior by the outside observer.

Even though the behavior–based approaches have been important in the move away from model–based robotics, the applications of these ideas have largely been in the field of mobile–robot navigation. In mobile–robot navigation according to Brooks [9] the advances have been limited to rather simple systems in terms of degrees of freedom (DOF) that have to be controlled and the number of behaviors that are displayed at any one time. A typical robot navigation task has often two DOF (one per wheel) to control and a behavioral repertoire that is limited to obstacle avoidance and path–planning. A humanoid robot however, places different demands on the number of DOF that have to be simultaneously controlled and also the number of behaviors that are part of the behavioral repertoire of the system. For example a basic humanoid head with stereo–vision already has four DOF to control. The different behaviors employed by a head range from maintaining a zero–disparity common fixation point for the stereo–eyes, saccading toward moving objects, smooth tracking and coordinating eye, head and body motions to follow objects. In order to handle this increase in complexity of the control–problems the behavior–based approach in itself is not sufficient [9].

In order to build complete, integrated complex systems that can utilize a multitude of behaviors to survive, interact and adapt in the real world, a different methodology that allows for scaling of simple behaviors into complex adaptive behaviors is needed. An extension of the behavior–based biological perspective can provide this methodology. Human development is about scaling behavior from simple initial behaviors in infants into complex adaptive behaviors in adult humans. Taking human development as a departure point can provide insights in the set of simple behaviors that is being used in infancy, and the developmental strategies that exploit these simple behaviors to gradually grow and aggregate into complex ones.

2.3 Modeling Complex Systems

As mentioned in the previous section, human development is about scaling. The advantages though, of trying to define a set of initial behaviors and a set of developmental strategies that exploit these simple behavior to gradually develop more complex adaptive behaviors, needs some more explanation. Scassellati [1] gives a breakdown of these advantages. Development, according to this breakdown gives a structured decomposition of complex tasks, facilitates learning and allows for gradual increase in task complexity.

Dividing up a complex task in manageable component tasks greatly influences the difficulty of implementing such tasks. Unlike other approaches, where the problems are divided up according to sensory modalities or sub–parts of complex tasks, the developmental approach can suggest a decomposition that is meaningful and effective, as this decomposition is already present in a system that uses it effectively for its development, namely human infants.

Facilitation of learning is another advantage that looking at development can yield in designing complex interactive systems. One of the big problems in robotics is the so called *degree of freedom problem*. This problem, simply put, states that the problem space that has to be explored in order to generate enough examples for learning increases exponentially with each added DOF. To illustrate this problem a typical robotic arm provides a nice example. For a robotic arm to function, six DOF suffice; two DOF in the shoulder, one in the elbow, one radioulnar and two in the wrist. Let's number these DOF or joints q_1 to q_6 in the same order. If we have say 100 motor positions for each joint, the configuration–space would already be $q_1 * q_2 * \ldots * q_6 = 10^7$. The number of 100 positions has been chosen arbitrarily, and can easily be one order of magnitude bigger, which enlarges the configuration–space to close to *one billion* possible combinations that would need to be visited to build a complete map for arm–reaching. From this observation it follows that the time needed to explore this configuration space is impractical. The development of infants suggests a different approach. Infants do not use all their DOF directly from the start of development, instead some DOF are coupled. Another mechanism reducing the problem–space are reflexes that already give a basic direction for the learning to take, and avoid a complete enumeration of the configuration–space in order to learn a behavior.

As mentioned before, development is about scaling simple behaviors into more complex ones. In development this scaling is an incremental process in which earlier simple–behaviors develop knowledge that can be recruited by a more complex behavior to simplify the learning process. Another developmental phenomenon is that simpler learned behaviors serve as building blocks for more complex behaviors. By employing this incremental approach of using and re–using simple behaviors the learning problems are minimized at each stage of development.

The advantages of using development in infants as described here is that it gives some practical constraints and design guidelines that can be used in robotics research. Before we can put these insights to work we need some further exploration of the actual development in children in order to formulate practical constraints that we can use for design purposes.

2.4 Development in infants

If we take a look at the development of infants, we can see that in the early stages of development the control over behavior is quite limited. There are even though some goal–directed behaviors already present from the very beginning. Infants seem to exploit their limited capabilities to learn from the environment, and as certain behaviors mature, these provide new and richer input to other behaviors that the infant is trying to master. These limitations in the motor–control of infants seem to have several different reasons.

- They have limited postural control of the trunk, head and arms.

Appropriate head and trunk righting reactions only begin to emerge two to three months after birth.

- They have limited knowledge about the physical makeup of their bodies.
- They have only a limited movement repertoire consisting of an array of infant reflexes (i.e. grasping and sucking) and basal intra–and inter–limb synergies (coupled flexor, extensor activity, coactivation).
- They have limited visual capabilities. During the first postnatal month, the visual system provides the infant with functionally useful but unrefined vision at about 5% of adult levels.
- They have not established a finite neural control–structure. Most cortico–spinal projections are not differentiated. There might be different processes at work in the finalization of the neural control–structure, either growing– or pruning–based, which are known to occur postnatal.

If we look at these limitations from a learning point of view, they may actually be seen as an advantage for the system. From the discussion of development in complex systems in the previous section it became clear that each DOF makes an exponential contribution to the theoretical configuration–space that a system has to work with. With control having only global targets in terms of motor–actions, the number of DOF that can be controlled by an infant is dramatically reduced. The sensory immaturity can be viewed from the same perspective as reducing the input–space on the system, making mappings between sensory– and motor–events more simple. As these simple mappings improve during the development, they provide progressively richer input for other processes that can be exploited in turn, adding degrees of freedom to the system.

In-spite of the limitations of neonates mentioned, there is evidence that there is a set of coordinated motor–primitives already present in infants at birth. For example Von Hofsten [5] showed that when infants were adequately supported, they show evidence of initiating in movement towards attractive moving objects presented in their visual space as young as one week of age. Besides the point that this observation makes about the presence of basic goal directed visuo–motor capabilities, it also demonstrates that control over posture is needed for reaching behavior to be able to develop.

The behaviors that are already present consist of reflexes and intra– and inter–limb synergies. Examples of these are the palmar–reflex for the hand and the *asymmetric tonic neck reflex* (ATNR). These behaviors stem from a sub–cortical level, and will disappear in the course of development as cortical control over behaviors takes over. These reflexes though seem to serve an important function in the development of infants. As we have seen in the previous section, there is a big control–problem in terms of the size of the total configuration–space if we take all the DOF of the system together and try to learn them all at once. By coupling DOF in the early stages of development, the control problem is greatly reduced in terms of the total configuration–space that has to be explored. Reflexes play an important role in this, but seem to also serve another purpose. Metta [37] notices that these reflexes serve as a *bias* for the system as to what is actually going to be learned at the cortical level. The reflexes serve as a way of starting interaction with the world, and from that interaction learning follows.

The role of visual information in the development has been an issue of debate, but seems to serve mainly as a trigger of motor actions in the early stages of development, since the acuity of the infants vision is rather course at only 5% of the adult level. In research on infants of 6 to 25 weeks of age Alt and Trevarthen [] showed by placing a screen over the arm of the infants in reaching that none of them relied on visual information of the limb to before or during the reaching for an object.

A similar result has been presented by McCarthy et al. [35] in a study involving infants reaching for glowing objects in the dark. There is evidence though that vision of the limbs is very important in discovering what the infants own arm is, and how it reacts to motor–commands. Once a certain measure of knowledge about its own physical makeup has been learned, the visual information on the arm appears not to be used anymore in reaching.

The position put forward in this section holds that in the infant there is a reduction of the total learning problem, by reduced precision in the motor–control, coupling of DOF and by reflexes that provide an initial direction for the learning. While these mechanisms greatly reduce the total configuration–space and thus the number of states that have to be explored to learn a meaningful mapping of this space, there still is the need to explore. Metta [37] mentions that the immaturity of the neural structure of the brain causes noise in the transmission of motor–commands from the neural–controller to the muscles or motor–plant. This noise may serve an important role in the exploration of the configuration space for learning new positions in this space. By making errors, caused by the noise, new positions in the configuration space are visited and learned.

Taking a look at the development of goal–directed reaching in extra–personal space illustrates the developments mentioned. One of the pre-requisites for goal–directed reaching to take place is that the infant needs an idea of where an object is with respect to its own body. This information can be extracted from the orientation of the eyes and the head. This proprioceptic information is not yet available at birth and will have to be learned. Goal–directed arm movements are also present quickly after birth [5], but the movements that are being executed lack feedback on the success of the movements, as this is provided by the position of the object and the position of the hand, based on the proprioceptic information from the head. Movements executed in this phase show a ballistic trajectory, that is reflexive in nature. Once the information on where a target is becomes available from proprioception of the head, the reaching improves dramatically, and the motion gradually changes from ballistic coupled–joint reaches into multi–joint reaches smooth reaching toward the target. As the reaching for the target gets reliable, a new form of sensory feedback about the extra–personal space becomes available, namely touching of object. It is this contact with an object that will be used also in the orientation of the wrist, which is the next step in the development of prehension.

2.5 Summary

In this chapter the developmental perspective has been introduced, and a broad picture of the changes in the human infant from the neonatal period up until the emergence of grasping has been presented. The developmental path of prehension in infants gives some important constraints on the system that will be developed.

There appear to be important motor–primitives that help the infant interact with and explore the world. These motor primitives provide a starting point for improving and fine–tuning the motor–control and the visual–motor mappings. The imprecision in the control that is apparent in the neonatal period is not a disadvantage, instead an essential in reducing the problem of what has to be learned. As control over the motor system improves this in turn creates new possibilities for richer and more detailed interaction with the surrounding world providing a richer set of stimuli for learning, and gradually improving motor–control as new aspects of the interaction with the environment become learn-able.

Chapter 3

A Developing Robot

3.1 Introduction

The experimental setup being used in the research is called *Babybot* and is a non–mobile robot–platform that is mounted on a table on a rotating base, called the torso of the system. On this torso an industrial off–the–shelf six degree of freedom manipulator (Unimation Puma 260). It is mounted in such a way that it resembles a human arm. On top of this there is a five degree of freedom robot head with stereo camera's that have independent pan–axes, and a shared tilt axis. The head also contains two microphones that function as ears, for integrating sound–events. Figure 3.1 gives an overview of the system and its degrees of freedom. The system is controlled from a setup of standard PCs with Pentium II/III processors. One machine controls the arm and torso, another contains the framegrabbers for the camera's and does the visual processing, a third one is used for controlling the head. The PCs are connected through an ether-link connections.

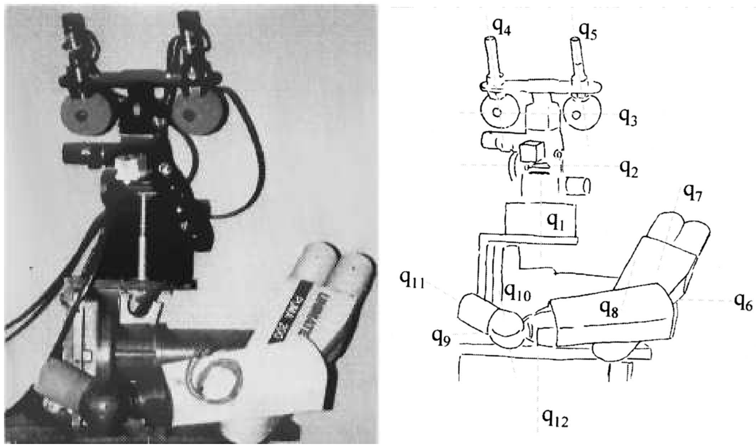

Figure 3.1: Babybot and its degrees of freedom. All rotational axes are numbered *q1* through *q12*

3.2 Development in Babybot

Development in *Babybot* is not like the usual development as we see it in biological creatures. *Babybot* being a mechanical structure, growth is only possible in terms of the use of sensory information and control over the motor apparatus. The initial system is characterized by a small number of free parameters, which can easily be estimated on–line. The concurrent controllers then learn on the basis of how the simpler loops are behaving. State–space exploration is driven by additive noise, which simulates defective command generation and muscle control.

The initial task of the control process is that of calibrating the closed loop gains. Many biological systems also have to tune delays, but in *Babybot* these delays were hard-coded into the system. From the very beginning the system has certain reflexes and couplings of joints that make it move in a goal directed way, even though the noise dominates the actual movements. Once the movements have reached a reasonable level of performance, the robot starts to use more degrees of freedom. First the head comes into play. By the time this happens, the eyes have already formed saccade maps that are relatively well formed, and can be used to help coordination of the redundant eye–head degrees of freedom.

At the same time that the system starts to explore how to use the head and eyes together, the reaching also starts to take of, and a map of the head–arm coordination starts to be built. Because reaching depends on gazing, the initial progress of the reaching is slow. Only when the gazing becomes effective, and thus the proprioceptic information of the head reliable, can the reaching be learned more effectively. Learning of the VOR is turned on all the time, and thus improves as more data is gathered. Figure 3.2 gives a schematic view of the developmental stages of the *Babybot*.

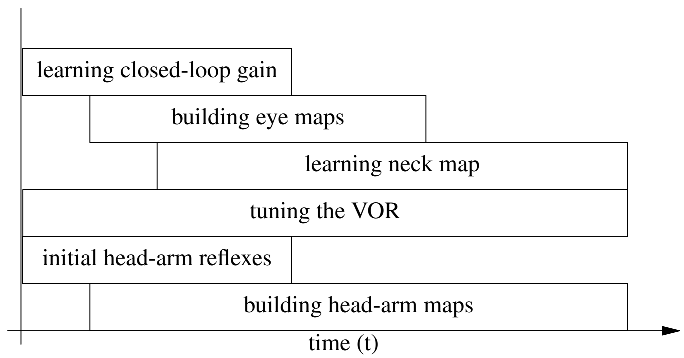

Figure 3.2: Schematic view of the developmental stages in Babybot.

3.3 Control of the Eyes and Head

After the initial vergence on a target has been learned, the control of the eyes can be learned. The control schema for the eyes is presented in figure 3.3. It features an open loop and a feed-forward loop. The secondary loop consists of an inverse model. This inverse model is activated whenever the retinal–error is bigger than a certain error, and as a consequence initiates a saccade. The block that contains the word *saccade* in the figure, contains the threshold in terms of retinal error that initiates the saccade. The goal of the network is to learn the inverse model. The big advantage in terms of performance of this setup is that control and training run in parallel, which is not only interesting from the performance point of view, but also an important feature of a developmentally plausible system.

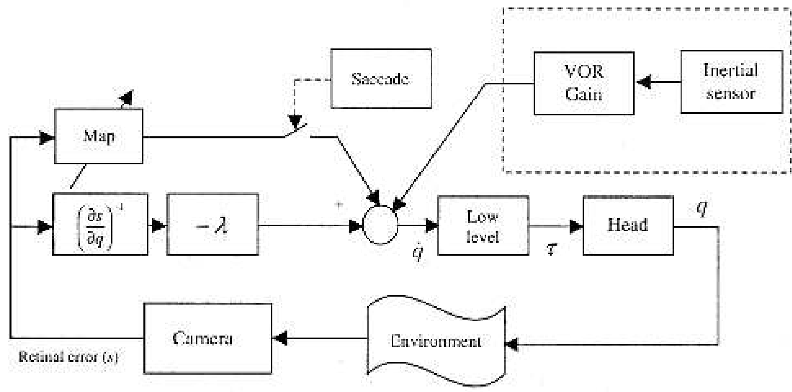

Figure 3.3: Information flow for learning eye–control. The term $\left(\frac{\partial s}{\partial q}\right)^{-1}$ is the inverse Jacobian. The λ term is a constant positive gain that is tuned to make the loop stable. q are motor commands in terms of joint–angles. τ are the motor–commands in terms of torques.

As figure 3.2 shows, the last behavior of the head that comes into play is the coordination of the neck in the tracking of objects. For this system there is also an inherent goal–directed nature. The goal of the neck–control system is to try and get the two vergence angles of the eyes the same. In terms of behavior this will mean that the neck will be oriented in such a way that the head faces the target at more or less a straight angle, giving a zero difference for the vergence angles of the two eyes.

The association of vestibular signals related to head turns and visual signals related to image motion, guide learning in biological systems.

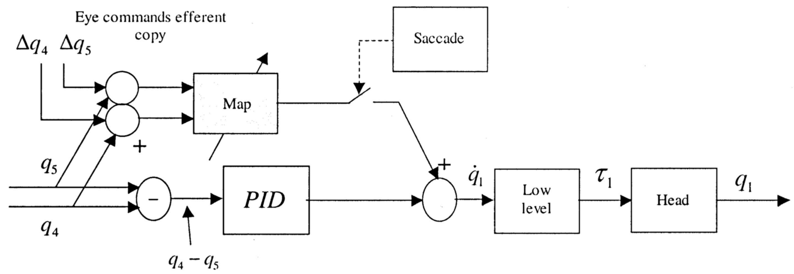

Figure 3.4: Information flow for learning neck–movements. q_1, q_4 and q_5 refer to the joint angles for the different degrees of freedom of the head as shown in figure 3.1. τ is again a command in terms of torque as it is relayed to the head.

In Babybot the learning of this relation between vestibular signals and the optic flow are learned in a neural network. The optic flow is also the teaching signal for the network itself. Because the network tries to minimize the optic flow, its success can be measured by the optic flow that remains after a correction by the network as compared to the optic flow that has been found before the compensation took place. This delayed reaction when learning is indicated by the dotted delay–line in figure 3.5. Metta [37] showed that this schema works in *Babybot*.

3.4 Control of Reaching

As figure 3.2 shows the reaching starts at the same time that the building of the eye–maps starts. As mentioned before, the reaching can only learn effectively as the orientation of head and neck reach a reasonable degree of accuracy in order to provide valuable feedback about the position of objects in the extra–personal space of the robot.

The reaching in *Babybot* uses a force–fiel approach, in which the reaching is driven towards equilibrium points in the force–field. The use of these force–fields makes compliant, safe for interaction with other

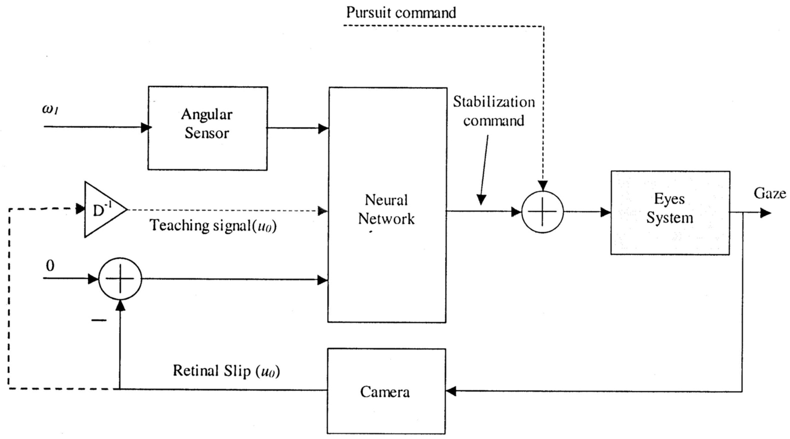

Figure 3.5: Information flow for VOR–tuning

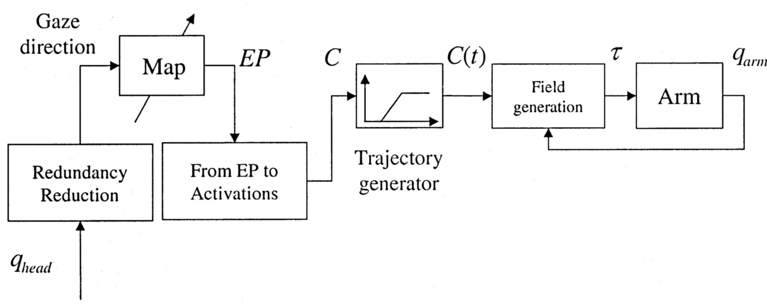

Figure 3.6: Schematic view of the reaching in Babybot. EP indicates an equilibriumpoint in the force–field, C is an activation vector, C(t) is an extrapolation of the final activation value over a number of time–steps, τ is the torque of the joints

people, motion of the arm possible. At the moment the trajectories are directed towards one point in the extra–personal space, and therefore ballistic in their nature. As I write this thesis an model for the the learning of more complicated trajectories, which are defined as a series of equilibrium points is being developed.

3.5 Summary

In this chapter the architecture of *Babybot* has been introduced. Also the control structure of the head–system has been discussed a bit more in depth. At the same time the head system is developing also the reaching starts. This reaching and its interplay with the head system has been addressed only scarcely as we will touch on this subject more in–depth later on.

Babybot shows that the developmental approach as it was introduced in chapter 2 is feasible in a real online robotic system. Even though the development of a robotic system is not directly comparable in the sense of physical growth, it shows that the developmental context can be adequately modeled into a robot, giving us the opportunity to systematically explore and test hypotheses about how certain developmental processes take place in the infant. Another important point Babybot shows is that in order to exhibit complex sensor–motor behavior, central monolithic processing is not needed, and could even be argued not wanted. It is the simultaneous development with different behaviors mutually influencing each others possibilities for gathering information from the environment that keep the system learn-able and stable.

Chapter 4

Pre-grasp Orientation in Infants

4.1 Introduction

The act of grasping an object has been extensively studied in infants. These studies range from the motor aspects of the reaching task for coordinating a movement of the hand toward an object in space, aperture of the fingers to accommodate for object–size, to the formation of precision grip with the fingers in order to manipulate a small object. Surprisingly there has been little attention for the orientation of the hand, aimed at positioning the hand conveniently for grasping. In this chapter a few of the articles about the pre-grasp orientation of the hand in anticipation of grasping an object will be discussed.

4.2 Pre-grasp Orientation

The advantage of correct orientation of the hand along the longitudinal axis of an object is easily noticed if one wants to grasp for example a drinking bottle or a toy. There are three basic preconditions to pre-grasp orientation. The first is that there needs to be an idea of orientation in the visual field, meaning that the infant needs to be able to discriminate between between different orientations of an object in its visual field. The second precondition is some measure of control over the motor apparatus to be able to orient the hand in the first place. The third precondition that is needed is some form of association between the visual information and the motor part of the orientation. The four questions that arise from these observations are: When are infants able to discriminate between different orientations of visual objects? When is the pre-grasp orientation beginning to develop in infants? Is vision of the hand and limb needed for the pre–grasp orientation to take place? And is the orientation part of the reach–for or the grasp–of an object in terms of motor–behavior?

4.2.1 Onset of pre-grasp orientation

According to [32] Essock and Siqueland already established in a 1982 experiment that infants are able to discriminate between different orientations of objects as early as two months of age. How precise this discrimination is, is not exactly known. The experiment concentrated on horizontal and vertical rods that were presented to the infants. As to the onset of observed orientation there is some early longitudinal research on pre–grasp orientation in infants by Von Hofsten [55] which indicates that infants start to contact objects that they are reaching for at the age of four months. In the beginning this contact is often with the back of the hand, and grasping if any, is slow and awkward. Also Corbetta et al. [13] point out that there are studies that report scaling of responses based on visual properties of targets in infants as early as four to five months of age. Von Hofsten [55], McCarthy [35] and others place the stable orienting of the hand as a preparation for grasping an object at about the age of eight to nine months. So in terms of the onset of the pre-gasp orientation and the period in which an important part of this behavior takes place there seems to be considerable agreement. The onset of pre–grasp orientation lies around the fourth to fifth months of age, and a stable orientation respons for objects emerges around the eighth to ninth month of development.

4.2.2 Is vision of the hand needed?

There has been some debate over the years as to the question whether visual information of the hand and limb bare needed for learning to orient the hand in the right orientation with respect to a perceived object. In a recent study by McCarthy et al. [35] it was shown that vision of the hand is not necessary for infants to coordinate a correct orientation response of the hand for an oriented object. This conjecture was tested by using glowing targets in a dark setting and having infants reach for these glowing targets.

4.2.3 The control of pre-grasp orientation

McCarthy et al. [35] observed in one condition of their experiment where they had infants grasp for a glowing object in the dark, but stooped the glowing of the object in the last part of the trajectory, that the basic ability to correctly pre–orient the hand according to the orientation of the object was not impaired by the lack of visual information of either the hand or the target. They interpret these results in terms of the character of the representation that is being built up for the control of the pre-grasp orientation. In the condition where there is a lack of visual information the infants guide their actions based on the motor representation of the previously perceived orientation of the object. Interestingly they interpret these results in the light of the theory of Milner and Goodale [38], which we will discuss in the next chapter.

At the end of this chapter I would like to mention one final interesting study that relates closely to the question of the nature of pre–grasp orientation. In a statistical study of many aspects of infants motor–

behavior [50] it was found that there was little correlation between the development of different aspects of the infants behavior, and when there actually was a correlation, the proportion of variance explained by the correlation was only 25%–30%. These findings suggest that the development of different aspects of the behavior of children develop in a rather independent way, or at least have independent brain functions responsible for their control. These findings were also extended tot the different kinds of grip that infants develop.

4.3 Summary

From the research it is clear that children as young as five months old start in orienting their hands in anticipation of grasping a rod that is oriented in the visual field of an infant. The behavior of pre–grasp orientation develops over the course of four to five months, until it becomes relatively stable at nine months of age. Anticipation of orientation for objects that rotate during the grasping movement continues to develop after this time.

The onset of the pre–grasp orientation behavior coincides with the beginning of successful reaching. Another important finding from pre–grasp orientation of infants is that this behavior does not appear to be dependent on vision of the hand during the trajectory of the arm toward the target, nor is vision of the hand necessary before the onset of the orientation. It appears that the orientation of the hand takes place in an open–loop fashion based on visual information of the object as it appears on the retina before the orientation–act is initiated.

Chapter 5

Visuo–Motor System and the Brain

5.1 Introduction

Having explored the developmental aspects of sensory–motor coordination and the gradual buildup of functions as the infants brain develops, we have left out the actual physiological aspects of the information processing within the brain. In this chapter theoretical insights into the makeup of the visual and motor–system will be discussed. Furthermore some evidence from brain–damaged patients will be presented that supports the position about the functional makeup of the visual–motor system.

5.2 The visual brain

To establish the relative importance of vision as a sensory modality in humans, we simply need to look at the percentage of the cortex that is devoted to visual processing. Evolution has provided humans with a patchwork of visual functions occupying about 50% of the cortex, mainly in the posterior areas [60]. In spite of this relative importance in terms of size, vision is in essence one of the sensory modalities that humans use to interact with their environment. A long-standing view of vision has been that it is of a higher order in sensory terms and produces complete and abstract models of the outside world, and that through reasoning about these models interaction signals are produced to initiate actions. Ungerleider and Mishkin [51] already proposed a division of the visual system into two separate "streams". The ventral stream, or *what* stream, that is used for object recognition and the dorsal stream, or *where* stream, that encodes the spatial location of an object. Milner and Goodale [38] paint a different picture, based on physiological research of the visual system in monkeys and data from patients with brain damage. In their book *The Visual Brain in Action* they propose that the dorsal and the ventral stream are better characterized as *what* and *how*. In their view the visual processing proceeds along two distinctly separate lines, vision for perception through the ventral stream and vision for action through the dorsal stream.

Figure 5.2 gives a schematic of the ventral and the dorsal streams as they were identified in the brain of macaques. As can be seen in the picture, both streams originate in the primary visual area of the cortex (V1), with the ventral stream projecting to the inferior temporal (IT) cortex and the dorsal stream projecting to the posterior parietal (PP) cortex [51]. The presence of these broad visual streams is not really debated; they have been established as anatomical features present in the brain, what is being debated is their function. As mentioned, Goodale and Milner differ in their interpretation from the classic account of Ungerleider and Mishkin. Goodale and Milner base their altered view of the function of the two visual streams on data from patients with lesions in their brain. In the case–reviews they present patients with lesions in the *ventral* as well as the *dorsal* stream, and discuss the behavioral deficits and residual capabilities that these patients have. Based in this evidence they conclude that the dorsal stream is action

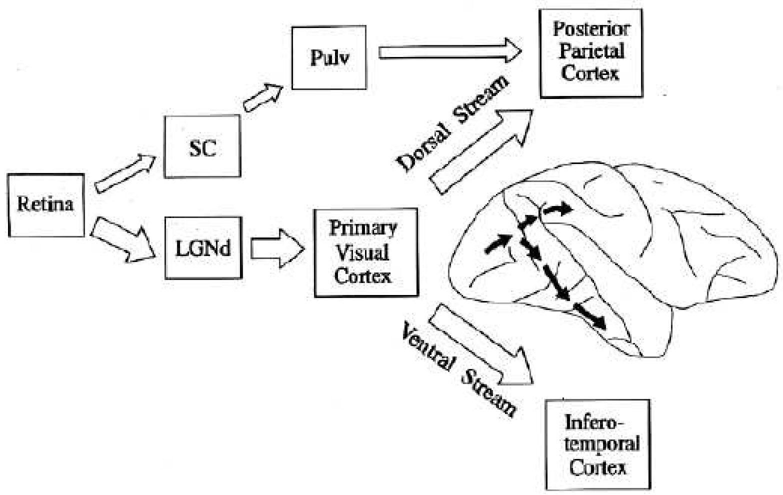

Figure 5.1: Major routes of visual input in the brain. This figure shows the dorsal and the ventral pathways, together with some of the sub–cortical connections that reach into the dorsal–area.

oriented, while the ventral stream is perceptual in character. One of the most important cases that Milner and Goodale describe is that of patient DF.

DF is a patient that has been diagnosed with visual agnosia. In short patients diagnosed with visual form agnosia, exhibit an inability to perceive objects, even more remarkably in DF there is little or no deficit in the action oriented aspects of object recognition. When presented with prototypical shapes, for example a red apple, and asked to give a verbal report of these objects, DF is incapable of telling what kind of object is being presented. The same is true if asked to draw the same object. It appears that DF is able to discriminate some objects on the basis of their color and sometimes texture, but their form or size is not being used as a cue. If on the contrary DF was asked not for a perceptual judgment of an object, but to just grasp the object presented in space, her ability to make the right grip aperture of the hand, orient the hand in the right way to match the orientation of the object and position where she needs to reach to in order to grab the object are all used successfully. DF appears to have all the information necessary to grab an object in space, while at the same time being unable to use perceptual information of that same object.

In an attempt to further explore the capabilities of DF, Milner designed an experiment that could systematically test the abilities for using orientation information about objects in DF. The design of the experiment was based on a round board, with a rectangle slot. The board could be rotated by the experimenter in all possible orientations. Figure 6.7 gives a schematic of the task that was used.

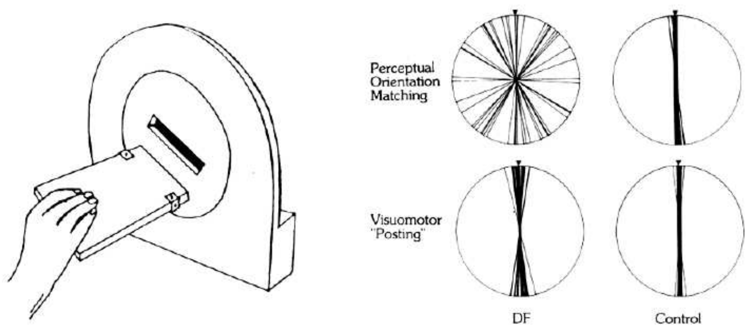

Figure 5.2: Milners experiment to test the perceptual and visuo-motor capabilities of DF

In the experiment Milner asked DF two different questions. The first question was for DF to orient a rectangular board she was holding in the same orientation as the orientation of the slot. This is a perceptual question since it involves judging the orientation of the slot and then matching the orientation of the hand to that of the slot. The second question posed to DF by Milner was to insert the board into the slot. This question being an 'action' question directed at coordinating a motor action to insert the board into the slot. Figure 6.7 also shows the results of this experiment for DF as well as for a control. The results clearly show that DF had near-normal performance on the visuo–motor coordination task, while showing almost random performance on the perceptual task [38].

5.3 Alternative views

The view of two separate visual streams has not been shared by everyone. Most of the criticism is not about interaction with the outside world as one of the driving forces in the development of the visual system, nor about the pragmatic role of visual processing in the dorsal stream for visual–motor coupling, but directed at the sharp division between perception and action.

In a review of the book of Milner and Goodale, Gallese [17] notes that there are several experimental observations that contradict a strict separation between perception and action. One of the results he mentions are the Mirror neurons which have been found in monkeys. These neurons appear to be involved in a process which is viewed as highly perceptual in nature. The mirror neurons are not only activated in goal–directed grasping for an object, but also when observing a similar goal–directed grasp in another individual. According to the author this result indicates that this matching mechanism uses the same coded actions in two ways at the same time: at the output side to act, on the input side to to analyse visual percept. One can debate the perceptual nature of such neurons, but at the least it would seem to indicate is that perception and action may be more intimately related than might be concluded from clinical data of brain–damaged patients.

Gallese also mentions two other points questioning the sharp division of perception and action. He mentions research by Craighero et al. that shows that visual presentation of a bar in a specific orientation before an object has to be grasped significantly improves the response time for the grasping action. A last finding that shows a more intimate interaction between action and perception than Milner and Goodale propose comes from Sperati and Stucchi who showed that subjects use an internal simulation of their hand to judge whether a hand that is being observed screwing is screwing or unscrewing.

5.4 Summary

There has been a lot of criticism of the position of Milner and Goodale, but even among those being critical about the sharp distinction between perception and action there is a common praise for the fact that Milner and Goodale have placed bottom–up vision–to–motor binding in the center of attention again. Previous accounts of the visual system are sometimes called the *couch–potato* model of vision, whereby observation of visual scenes would be enough to construct a visual model of the outside world. With their position about a dorsal stream dedicated to visuo–motor transformations and a ventral stream representing semantic or cognitive aspects of objects, Milner and Goodale emphasize the importance of interaction with the outside–world as the basis of the visual system.

Where does the position of Milner and Goodale help us in the subject at hand? In our search for a model of pre-grasp orientation, the views of Milner and Goodale contribute in two areas to the formulation of the constraints of our intended system. In the first place the emphasis on the direct coupling of visual sensory–information to motor commands provides a useful mechanism for the design of our system. The second constraint is derived from the first point and holds that we do not need high order understanding of visual scenes in order to be able to coordinate a motor–action toward a target. Since basic egocentric information about objects such as size, orientation and retinal–form are communicated through the dorsal stream and directly combined into motor–actions.

Chapter 6

Pre-grasp Orientation in Babybot

6.1 Introduction

In the previous chapters, the backgrounds of developmental robotics, development of pre–grasp orientation in infants and information processing in adult humans have been laid out. In order to create a model of pre–grasp orientation we will have to bring the information from these topics together and merge them into a working system that performs pre-grasp orientation for *Babybot*.

As described in the introduction, the goal of this thesis is to add a system to the developmental architecture of the *Babybot* that is able to perform a pre–grasp orientation to orient the end–effector of the Babybot to match the orientation of a stimulus presented in the visual–field. For this behavior to be added to Babybot, the following aspects are needed:

- orientation task that can be learned
- target identification

- visual orientation of the target
- reaching for a target
- orienting for a target
- learning mechanism

6.2 Task

The explanation of the pre–grasp orientation in infants in chapter 4 has shown that the orientation of objects is being learned in the extra-personal space of the infant. This learning builds on the ability of the infant to reach and touch the object, as a result of which the learning gathers in speed as the reaching becomes precise enough to consequently reach for the object. For the orientation of the end–effector to be learned, an association between visual orientation of an object and the corresponding motor commands that will direct the end–effector of the robot toward the object in the right orientation will have to be established. Even though there is no advanced grasping in the early stages of development, the palmary grasp reflex, combined with the sensory feedback of the hands touching an object give the infant a good measure whenever the hand was in the right orientation.

In the current configuration of *Babybot* there is no end–effector present that has any DOF, or sensors fitted. The end–effector that is being used is a rigid aluminum fitting with a crossection of about four centimeters and a width of about one centimeter. This limits the possibilities for the tasks that can be used for the learning of the orientation of the end–effector. The important aspects of the task that are learned though can be modeled with this system.

Since the important aspects are orientation of the object and feedback through proprioception, we can use the experiment that Milner and Goodale used to test their patient DF. As discussed in chapter 5 they used a box with a slot that could be rotated by the experimenter in an orientation that would have to be matched by the subject. What makes this task suitable for the learning of the orientation in *Babybot* is the fact that when an orientation action has successfully been performed, the end–effector will be *inside* the slot, thus giving the opportunity to use proprioceptic information of the wrist to give a measure of success for the orientation of the end–effector. If the orientation action was successively, the wrist cannot be rotated further, if it was not successful, the end–effector will not be inside the slot, and can thus be freely rotated in a different orientation.

Figure 6.1: The task that *Babybot* will learn

6.3 Vision

For the visual system we have seen in chapter 5 that in adults there is evidence that only course low–level information from the visual system is being used. Goodale and Milner place this information that is being used for reaching and orienting for manipulating objects in the dorsal stream. In view of this observation combined with the task that will be used as discussed in the previous section, the visual information that is necessary can be directly extracted from the retinal image without the need for higher order visual processing of the scene in which the object is presented to the robot.

Having established the task that will be used for learning the pre–grasp orientation, and the character of the visual information that we can use for the learning of the task, we can now focus on the visual processing. As mentioned in the introduction to this chapter, the target will have to be identified, after which the orientation of the target will have to be established. Figure 6.2 shows a block diagram of the processing steps that are needed, with a raw image from the camera going into the system and an angle of an object coming out.

Figure 6.2: Visual processing steps

In this section the object identification will be discussed first, after which a measure for the orientation of the identified object will be introduced. Since the visual system of Babybot is based on log–polar images, an explanation of this representation is given first.

6.3.1 Object identification

The identification of objects in *Babybot* is based on motion and color–segmentation. Infants display a increased attention for brightly colored objects that are moving in their central vision []. *Babybot* has a target identification procedure that parallels this same source of attention. This procedure basically takes the moving pixels in the visual field, and tries to start a color–segmentation based on these moving pixels. This approach differs from standard approaches based on color–

segmentation in that it takes into account that new objects with different colors can be presented in the visual field, and that it adapts to the different colors based on motion.

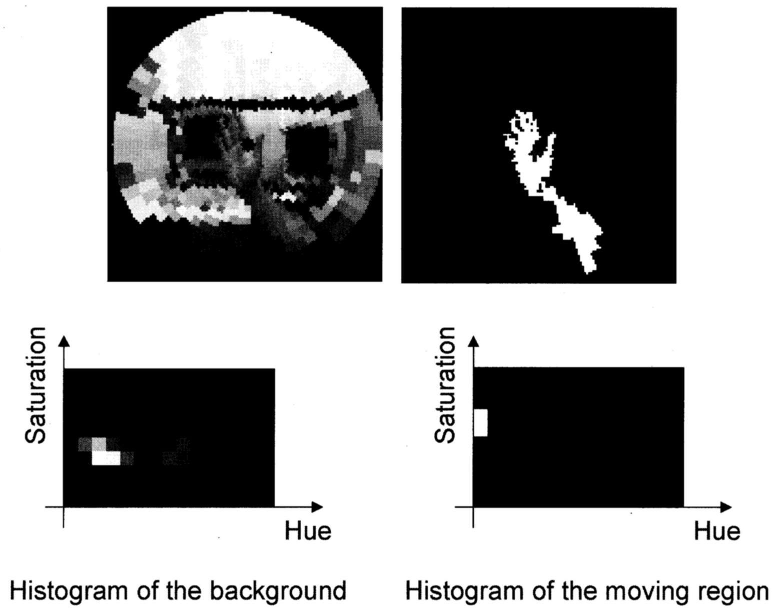

Figure 6.3: Log-polar flower

The color segmentation is based on HSV color space instead of RGB color space, because HSV gives a more efficient way of handling color information. The HS information of the HSV space give us the information about which color a pixel is, while the V gives us a measure of its value or luminance. The luminance information is not critical for judging a color and will therefore not be used for color-segmentation.

In order to detect what might be an object, motion detection is applied to the initial images. From the pixels that are moving, a histogram of the HS space is constructed. The assumption is that the dominant color in the histogram of the moving pixels represents the actual object–color, so the dominant object–color should be on or around the maximum of the histogram. A second histogram of the background is then constructed. If the object histogram and the background histogram differ enough the color segmentation itself is started. The color segmentation is done simply by checking if a pixel matches the color of the object–histogram. If no pixels of the color that is being segmented have been found for a certain period of time, the process of establishing a color for color–segmentation starts again, thus giving the robot the possibility to shift attention between objects of different colors.

6.3.2 Log-polar images

In the *Babybot* log–polar images are used as the basis of the visual system. Log-polar images resemble the physical makeup of the retina in their space–variant density of photo–receptors. Just as in the retina, there is a high density of photo–receptors in the center (or fovea) of the image, and a progressively lower density (on a log scale) toward the periphery. Figure 6.4 shows a Cartesian image and its transformation to the log-polar or *cortical* domain. The *cortical* domain is called this way because of its resemblance to the surface that the visual information of the retina maps to in the cortex. The third image is the log-polar image in a Cartesian representation.

Mathematically the transformation between a Cartesian and the *cortical* domain can be expressed as the transformation between a polar plane (the retinal plane) and a Cartesian plane (the log-polar or cortical plane):

$$\eta = q\Theta$$

,

$$\xi = K_{\xi} ln_a \frac{\rho}{\rho_0}$$

(6.1)

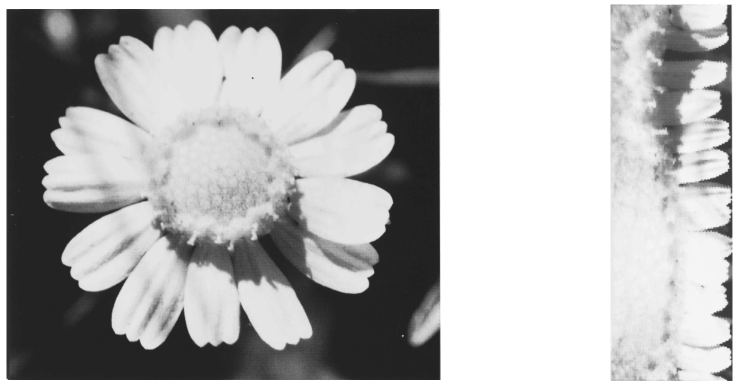

Figure 6.4: Cartesian flower and logpolar flower

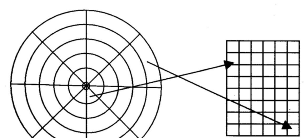

$$(\rho, \theta) \rightarrow (\eta, \xi)$$

Figure 6.5: Mapping an image to the cortical domain

In this equation ρ_0 is the radius of the innermost circle, 1/q is the minimum angular resolution of the log-polar layout, and (ρ, Θ) are the polar coordinates. K_ξ is a linear scaling parameter, which has been added to the original equation in order to fit the mapping into a fixed size squared image.

6.3.3 Rotation in Log-polar

Because the properties of *cortical* or log–polar images, a rotation in the Cartesian domain is a translation in the *cortical* domain. This property can be used to establish the orientation of an object. If we have a template of a certain object in memory the orientation of that object can be established using the cross correlation of the object with the template from memory. The point of the highest cross correlation indicates the phase–shift that is necessary to make the two images overlap exactly. Though in theory this approach works perfectly, in practice it has some drawbacks. The most important drawback is that the translation in the log-polar domain is only valid for objects that are precisely foveated. Another problem is that the template that is being used has a big influence on how the orientation is being judged. If an object is rotated in the horizontal plane with respect to the template, the matching result changes considerably. A last drawback is the rigidity of using templates for establishing the orientation. A slight change in the form of the object has a big influence of the orientation measure.

An alternative approach is to use the so called *moment of inertia* of the color–segmented (and thus binary) images. 2D images have several *moments.* Let's first look at these moments in Cartesian domain, before translating the results to the log–polar domain. The 0^{th} moment defines the *area* of a region, for a continuous image this would give:

$$A = \int\int f(x,y)dxdy \tag{6.2}$$

in which $f(x, y)$ is a function that evaluates if a pixel is part of the object or not (the function returns 0 for pixels that are not part of the object or area and 1 for pixels that are). The 1^{st} moment, defines the *center of mass* or *centroid* which is denoted by $(\bar{x}, \bar{y})$ and is defined by:

$$\bar{x} = \frac{\int\int xf(x,y)dxdy}{\int\int f(x,y)dxdy}$$
$$\bar{y} = \frac{\int\int yf(x,y)dxdy}{\int\int f(x,y)dxdy} \tag{6.3}$$

To establish the orientation of an object, the *axis of minimum inertia* is used. This is the axis of least 2^{nd} moment. For this we want to find the line through the *centroid* with an angle θ with respect to the x-axis for which the integral

$$I = \int\int ((x - \bar{x}) \sin\theta - (y - \bar{y}) \cos\theta)^2 f(x, y) dx dy$$

(6.4)

is a minimum. Expanding this equation we get

$$I = a \sin^2\theta - b \sin\theta \cos\theta + c \cos^2\theta$$

(6.5)

where the terms a, b and c are defined by

$$a = \int\int (x - \bar{x})^2 f(x, y) dx dy$$
$$b = 2 \int\int (x - \bar{x})(y - \bar{y}) f(x, y) dx dy$$
$$c = \int\int (y - \bar{y})^2 f(x, y) dx dy$$

(6.6)

differentiating this expression with respect to θ and equating it to 0 renders:

$$\tan 2\theta = \frac{b}{a - c}$$

(6.7)

which is only valid for $b \neq 0$ and $a \neq c$. This last equation can be rewritten to give us the final result for the orientation, which is also defined for $b = 0$ and $a = c$:

$$\sin 2\theta = \frac{\pm b}{\sqrt{b^2 + (a-c)^2}}$$
$$\cos 2\theta = \frac{\pm(a-c)}{\sqrt{b^2 + (a-c)^2}}$$

(6.8)

From these last two equation, the positive solution maximizes I, and the negative solution minimizes I.

Since the result given is only for Cartesian space we will have to translate this result to the log–polar domain. As we have seen in section 6.3.2 the log-polar image is defined by the coordinates (η, ξ). To use the points in the log-polar image to calculate the axis of least inertia we will have to translate these points to the Cartesian domain.

The general case for translating coordinates (x, y) to the new coordinates (u, v) in a double integral we use the Jacobian determinant. Suppose that the new coordinates u and v are defined by the relations

$$x = x(u, v)$$
$$y = y(u, v)$$

(6.9)

where there is a one-to-one correspondence between (x, y) and (u, v). Then a double integral in coordinates (x, y) can be expressed in the new coordinates by:

$$\int\int f(x,y)dxdy = \int\int f\left[x(u,v), y(u,v)\right] \left|\frac{\partial(x,y)}{\partial(u,v)}\right|$$

(6.10)

in which

$$\frac{\partial(x,y)}{\partial(\eta,\xi)}$$

(6.11)

is the Jacobian of the translation between the different coordinate systems. As we have seen in the previous section, the log–polar coordinates are given by:

$$\eta = q_r \cdot \arctan\left(\frac{y}{x}\right)$$
$$\xi = \ln_a \frac{\sqrt{x^2+y^2}}{\rho_0}$$

(6.12)

which gives us for the transformation of the Cartesian to the log–polar domain in equation 6.2:

$$\int\int f(x,y)dxdy = \int\int f\left[\eta(x,y),\xi(x,y)\right]\left|\frac{\partial(x,y)}{\partial(\eta,\xi)}\right|$$

(6.13)

with

$$\frac{\partial(x,y)}{\partial(\eta,\xi)} = \frac{1}{q_r}\rho_0^2 a^{2\xi}$$

(6.14)

This result forms the basis of the transformation that we will use. For the purpose of transforming the formula of the minimum 2^{nd} moment, or the least-axis-of-inertia, we need a definition of how to translate individual points from log–polar to the Cartesian domain. Derived from equation 6.12 this gives us for a point (x, y):

$$x = \rho_0 a^{\xi} \cos \frac{\eta}{q}$$
$$y = \rho_0 a^{\xi} \sin \frac{\eta}{q}$$

(6.15)

Now that we have the Jacobian for the double integrals and the transformation of the individual points from log–polar to Cartesian, we can substitute these results in equation 6.2 which gives us the 0^{th} moment or the *area* of the object in log–polar:

$$A = \int \int f(\eta, \xi) \left| \frac{1}{q_r} \rho_0^2 a^{2\xi} d\eta d\xi \right|$$

(6.16)

The *centroid* can be obtained by substituting equations 6.14 and 6.15 in equation 6.3:

$$\bar{x} = \frac{\rho_0 a^{\xi} \cos \frac{\eta}{q_r} f(\eta, \xi) \left| \frac{1}{q_r} \rho_0^2 a^{2\xi} d\eta d\xi \right|}{A}$$
$$\bar{y} = \frac{\rho_0 a^{\xi} \sin \frac{\eta}{q_r} f(\eta, \xi) \left| \frac{1}{q_r} \rho_0^2 a^{2\xi} d\eta d\xi \right|}{A}$$

(6.17)

using this result in the transformation of the various sub-parts of final equations 6.6 gives us:

$$a = \int\int \left[\rho_0 a^\xi \cos\frac{\eta}{q} - \bar{x}\right]^2 f(\eta,\xi) \left|\frac{1}{q_r}\rho_0^2 a^{2\xi}\right| d\eta d\xi$$

$$b = 2\int\int \left[\left[\rho_0 a^\xi \cos\frac{\eta}{q} - \bar{x}\right]\left[\rho_0 a^\xi \sin\frac{\eta}{q} - \bar{y}\right]\right] f(\eta,\xi) \left|\frac{1}{q_r}\rho_0^2 a^{2\xi}\right| d\eta d\xi$$

$$c = \int\int \left[\rho_0 a^\xi \sin\frac{\eta}{q} - \bar{y}\right]^2 f(\eta,\xi) \left|\frac{1}{q_r}\rho_0^2 a^{2\xi}\right| d\eta d\xi$$

(6.18)

Which can be used in conjunction with equation 6.8 to calculate the orientation of an object in the log–polar of cortical domain.

6.3.4 Stereo orientation

Now that the algorithm for the visual processing has been established, a further step will have to be taken in order to integrate the information from the two eyes from the stereo–head. Because both the eyes have a slightly different angle with respect to the object that they are looking at, as expressed in the vergence of the eyes, the best results for the global orientation of an object in the visual field is obtained by summing the results from both the eyes. Because there is a discontinuity at $\pm 90°$ provisions to handle this discontinuity will have to be made, or otherwise the summed orientation near the discontinuity can be zero. Figure 6.6 gives an overview of the processing steps in the visual subsystem.

As can be seen in the figure, after the summation of the separate eyes, a moving average filter of five points has been added to get more stable orientation result.

6.4 Controlling the arm

The reaching part of the behavior is already implemented. For the reaching a force–field approach is used, which guides the arm toward

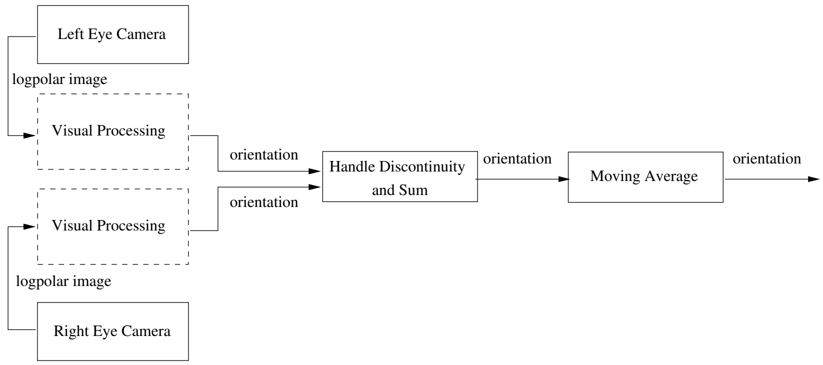

Figure 6.6: Information flow in the visual subsystem

a given equilibrium point in the force–field. For the arm to be guided toward a point in the extra–personal space of the robot, only three of the six DOF of the robot–arm are used. The orientation will be adding an extra DOF to the arm that is being learned. In order to get the orientation working the wrist also has to be positioned well for the end–effector to be able to be inserted into the slot. As was mentioned before in this chapter, the current setup for the arm has no touch–sensors, which makes the learning of the azimuth of the wrist, that is needed to insert the wrist into the slot difficult, since there is no feedback on the success of the azimuth of the wrist. For this reason, an inverse kinematic model of the azimuth of the wrist, based on the angles of the arm in the equilibrium–point and the orientation of the head, will be used. Jeannerod [] showed that the angle of the wrist has a direct relation to the position of the object. The hand is kept at an angle with respect to the body that is roughly equal to the angle of the object with respect to the body.

The proprioceptive information of the head gives us the information about the location of an object that is being reached for, and therefore the final angle of joints q_4 and q_5 of the arm. For the wrist to be oriented correctly though, also the final angles of the joints q_1 to q_3 at the end–point of the reaching trajectory have to be taken into account.

If q_{head} gives us the angle of the object with respect to the torso, then the final angle of joint q_5 can be calculated by:

$$q_5 = -(q_2 + q_3) + q_{head}$$

(6.19)

Figure ?? gives a schematic view of how the angle of an object with respect to the torso can be deduced from the proprioceptive information of the head.

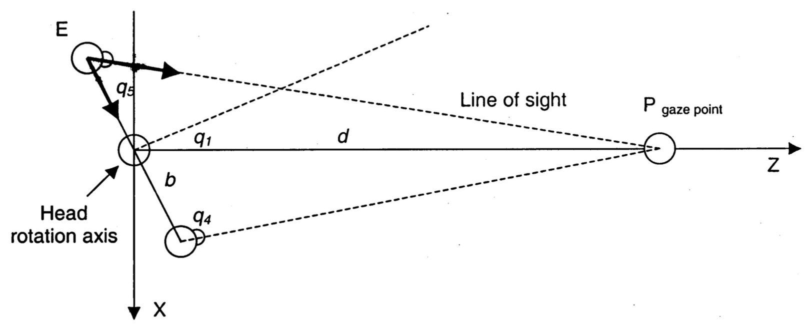

Figure 6.7: Technical layout of the head

The actual rotation of the hand is controlled by joint q_6 of the arm, the control of which will be discussed in the next section.

6.5 Learning the mapping

For the orienting behavior to be learned we have looked in the previous section for a measure of success to give a reinforcement signal to the learning process. There are more aspects to the learning though. In other parts of *Babybot* we saw initial goal–directed reflexes, which grew as more feedback for the learning process became available because of other behaviors that became more accurate. The same principle can be applied for the orienting behavior. There exist reflexes toward orienting, but as reaching becomes more accurate, more and more feedback about the success of the reaching is generated, which in turn drives the learning.

As has been shown in chapter 5 the dorsal stream description of Milner and Goodale gives the possibility of directly coupling visual information with motor commands. In our case this gives us the opportunity to use visual information about the orientation of an object in visual space to the orientation commands given to the arm/wrist. This coupling can best be learned in a simple neural network. For other systems in the *Babybot* Nearest Neighbor Look-up Tables (LT) have been used, which can learn an arbitrary mapping between an n dimensional input–space and a m dimensional output–space. For a description see [37]. These LT's behave in the same way as neural networks, but are faster when learning, which is an important factor in a real-time system like *Babybot*, and have the added advantage that the numbers that are learned have a direct meaning, which makes inspection of the learning easier.

For learning to take place, there has to be exploitation of what has been learned, and exploration to discover new positions in the configuration space. The exploitation is based on what is being learned, because the LT's can learn online, based on what is being experienced, and therefore don't need a separate training phase, that would stand in the way of exploitation. The other aspect to the learning is the exploration of the configuration–space. As we have seen in chapter 4 the neural structure of infants has not yet fully matured, causing imprecise communication between the neural controller and the *plant*, or the muscles. In the intended system, noise will function as substitute for the impress ccommunication between neural controller and plant. Adding noise to the motor commands that have been retrieved from the LT, gives random exploration of points around what already has been learned.

Given the *mailbox* task, with the error between the motor–command given and the final–position of the hand as a measure of success, and the learning of a direct association between visual orientation and the motor–commands given, the learning procedure looks as follows:

1. calculate orientation for the object in view
2. query the orientation map with the visual–orientation for a motor–command.
3. query the arm–reaching map for a reaching command based on the position of the head
4. add the orientation command to the reaching motor–command of the arm
5. execute the motor–command for the complete arm
6. test for success, learn association if successful

6.6 The complete system

Figure 6.8 gives an overview of the information–flow of the intended system. As can be seen in the figure, the wrist consists of two parts, the part that controls the azimuth of the wrist and the part that controls the pre–grasp orientation. The box that controls the azimuth takes its inputs from the head and from the output of the reaching–map, while the orientation–map that is being learned takes its input from the visual subsystem as discussed in section **??**, with the visual orientation of the object in the visual field given by θ_{object}, and gives an output to the motor–plant in terms of a motor command for the q_6 of the arm. The result of the motor–command is evaluated and gives a measure of success or the orientation, driving the learning.

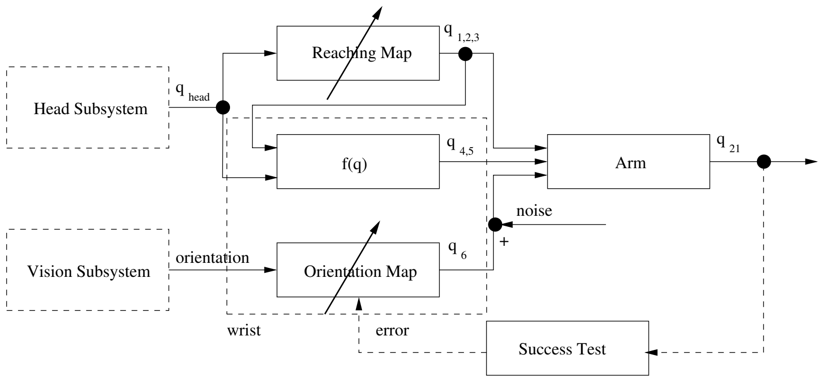

Figure 6.8: Information flow for the pre–grasp orientation

6.7 Summary

In this chapter we have seen how the different theoretical insights about development in general and the development of pre-grasp orientation especially together with the theory of Milner and Goodale about the pragmatic character of visual information processing in the dorsal stream have been combined into a single model.

From the development of pre-grasp orientation in infants we have seen that the visual capacity to distinguish between objects in different orientations is already available before the onset of the actual grasping actions. The motor system in infants was already biased for orienting, or at least for using recruiting the muscles used for orienting the hand during the development of reaching. From the general framework of the development we have seen that the learning of the pre-grasp orientation only takes off properly after the onset of more reliable grasping actions, facilitating a richer environment to learn from for the pre-grasp orientation. This point fits well in the view of simultaneous development of several relatively independent subsystems that have a bias through reflexes or motor–synergies for certain kinds of behavior, mutually influencing each others development by the new interaction possibilities that become available as the system grows.

Chapter 7

Experiments

7.1 Introduction

To test if *Babybot* system was successfull in learning, the execution of the task that has been discussed in the previous chapter was evaluated. From these experimental sessions, data was gathered and evaluated.

7.2 Procedure

The procedure to test the performance of the system was straightforward. In total five training sessions were performed and after each training session the performance of the system was evaluated.

Because of limitations in the precision of the existing reaching architecture that was used, the position of the *mailbox* has not been changed during the trials. The reaching–map as it exists is not sufficient to successfully reach in places other than the direct front of the robot. The *mailbox* was also placed under a small angle with respect to the torso to make the reaching into the slit easier.

7.3 Results

Figure 7.1 shows the orientation–map that has been learnt after 160 trials. The datapoints in the figure are the actual points that have been learnt. The line is a least mean squares fit to the data, taking into account the variance of the individual data–points. As can be seen in the figure, not all the data–points in the table fall exactly on the fitted line. This is mainly the result of the fact that the end–effector doesn't fit exactly into the slot. This leaves room for slightly different orientations of the end–effector within the same visual angle. Another factor in this variability is that the reaching does not always have the same end–position on the slot for the same orientation.

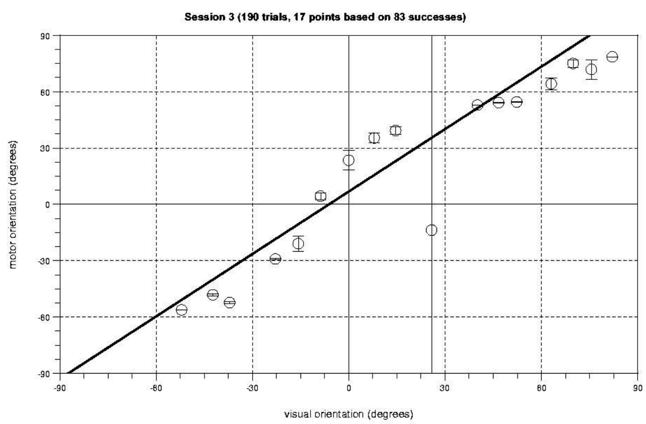

Figure 7.1: Orientation–map after session 3

Figure 7.2 shows the final orientation map that has been learnt after 351 trials. The point near the 30° of visual angle that had a very big variance in session 3, as can be seen in figure 7.1, has now a much smaller variance, and shows a closer correspondence tot the motor–orientation.

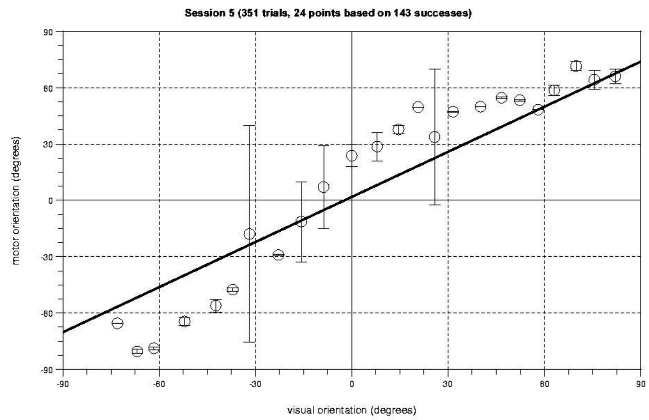

Figure 7.2: Orientation–map after session 5

When looking at the ditribution of the points in figure 7.2 around the regression line, it looks like the function that is being learnt is actually not linear one. Considering the experimental setup that is being used, this observation gains more credibility. Since the *mailbox* was placed on an angle with respect to the torso, the observed orientation of the slot is not exactly the same as the orientation that is needed for the wrist. When investigating this assumption further, a third order polynomial function has been fitted to the results of session 3. Figure 7.3 shows the same data–set that was used in figure 7.1, but now with a third order polynomial function fitted. Unfortunately this same result does not apply to the orientation–mapping after the fifth session as shown in figure 7.2. The reason behind this might be that there are too many points being learnt with respect to the precision of motor–orienation and visual–orienatation, as a result of which weaker generalisation is being achieved than might be possible. Better results may be obtained by lowering the resolution of the lookup–table.

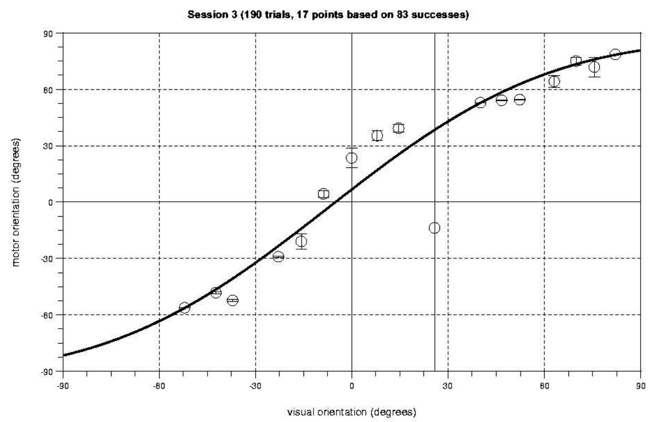

Figure 7.3: Rational function fitted to the orientation–map of the third session

After all of the five learning sessions, the success of the learning has been tested. Figure 7.4 shows the results of this learning. Since there were five separate learning sessions, a curve has been fitted that shows the interpolation over the average error of the five sessions. As can be seen in the figure, the system is improvong its performance over time, tending towards even better performance if there are more learning trials added.

7.4 Summary

The results from the experiments show that *babybot* is learning the mapping from visual orientation on motor–commands. The gradual

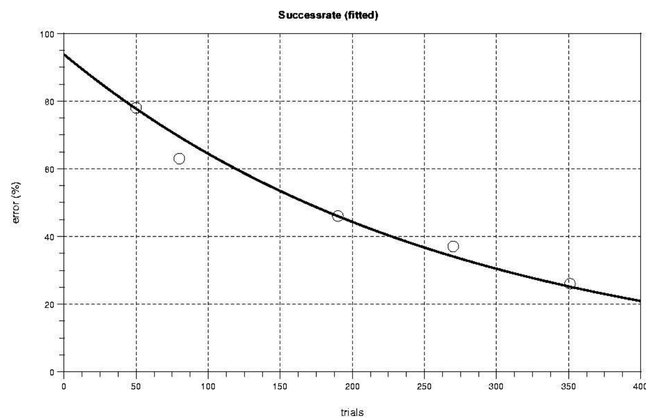

Figure 7.4: Learning results for the five sessions

diminishing of the error after 351 learning trials shows that as learning progressses, the success–rate for the mapping improves. The inspection of the orientation mappings from session three and session five show some evidence that the resolution of the orientation–map has been set too high, which warrants further investigation. Overall these results are encouraging and show that the approach that has been chosen to model the pre–grasp orientation is the right one.

Chapter 8

Conclusion

Where do we stand now at the conclusion of this thesis? In the past chapters a description has been given how a developmental perspective can help in developing complex robotic systems. A description of one such a developmental architecture namely *Babybot* has been given, followed by a presentation of how infants learn and develop pre-grasp orientation. As a point of adult level reference for the way in which sensory–motor systems are implemented and functioning we have taken a look at the theory of Milner and Goodale with respect to the buildup of the visual system, and the way in which this relates to the motor system. These chapters come together in a model of the development of pre–grasp orientation in a developmental system like the *Babybot*.

The data from chapter 7 shows that the system for pre–grasp orientation that has been developed functions as intended and is learning and improving over time as intended. The fact that the system is functioning properly provides evidence that the hypothesis of Milner and Goodale on the character of the visual system of the human brain. In our model we have shown that using only pragmatic visual information from the dorsal stream for a direct 'bottom–up' coupling to the motor system works very well, for artificial development at least. Important questions concerning the division of especially the perceptual side of the Milner and Goodale theory remain. In view of the important role the visuo–motor coupling plays in the development of an infant an interaction between the vision–for–action of the dorsal and the vision–for–perception of the ventral stream seems likely.

As this thesis has shown that the developmental approach is usable to extend the architecture of *Babybot* to pre–grasp orientation, it has only been on the fringe of the work that has to be done yet to add a working end–effector in the form of the hand. The model has been shown to be working, but there is room for improvement, especially in terms of the generalization of the model. As we have seen in chapter 7 the reaching does not have an accuracy at the moment that makes manipulation of objects and pre–orientation for objects easy. In part this is an effect of the type of task that has been chosen for the robot to learn. As was explained in chapter 6 the choice for this type of task resulted from the type of end–effector that was available at the time and the sensory information that could be gathered. As a result of these constraints the task that has been used, uses a slot in a *mailbox* for the posturing of the hand. This type of target had as an advantage that proprioceptic information could be used to drive the learning, but there are also important drawbacks to this type of target. The first important drawback is that for the robot to successfully insert the end–effector into the target, there is a far greater strain on the precision of the reaching than there would be if a *positive* object, for example a *rod* would be used. With this last object–type even sweeping motions, from many different angles in the neighborhood, could result in a contact with the object, where in the *mailbox* task there is only one final–trajectory that can be used to make a successful reach into the target possible.

The choice for a different kind of task, using *positive* objects instead of a slot also has further implications for the system. As the discussion in chapter 6 already pointed out, a measure of the success of a pre–grasp orientation action is needed in order for the system to be able to learn. Removing the constraints of the mailbox, introduces the renewed need for sensory feedback from the end–effector to be able to evaluate if an action was successful. Combining sensory feedback with a kind of palmary–reflex, grabbing a touched object as soon as contact has been made, would add developmental plausibility to the system, and would probably add more generality to the approach because of the improved exploitation of the existing reaching abilities. One final point concerning the type of task that is used for the learning concerns the learning of the azimuth of the other degrees of freedom in the wrist. Changing the task to pre–orienting for *positive* objects would put less strain on what has to be learned for these DOF in the early stages, as a simple fixation of these DOF with respect to the arms suffices in the case of a *rod* that has to be pre–oriented.

In sum, the research in this thesis has been a first step in the learning of pre–grasp orientation and the manipulation of objects, and gives an encouraging result for the model that has been chosen. Furthermore, some clear recommendations about the direction of future research have come from this thesis.

Bibliography

[1] AAAI. *Building Behaviors developmentally: A New Formalism*, Spring Symposium on Integrating Robotics Research, 1998.

[2] Christopher Assad. An hypothesis for a novel learning mechanism in the cerebellar cortex. *Autonomous Robotics*, 11:285–290, 2001.

[3] Aaron P. Batista, Christopher A. Bueno, Lawrence H. Snyder, and Richard A. Andersen. Reach plans in eye-centered coordinates. *Science magazine*, 285:257–260, july 1999.

[4] Scott A. Beardsley and Lucia M. Vaina. A laterally interconnected architecture in mst accounts for psychophysical discrimination of complex motion patterns. *Journal of Computational Neuroscience*, 10:255–280, 2001.

[5] Bennett Berenthal and Claes von Hofsten. Eye, head and trunk control: The foundation for manual development. *Neuroscience and Biobehavioral Reviews*, 22(4):515–520, 1998.

[6] Luc Berthouze and Yasou Kuniyoshi. Emergence and categorization of coordinated visual behavior throough embodied interaction. *Machine Learning*, 31:187–200, 1998.

[7] Antonio Bicchi. Hands for dexterous manipulation and powerful grasping: A difficult road towards simplicity, 1997.

[8] Aude Billard and Maja J. Mataric. Learning human arm momvements by imitation: Evaluation of a biologically inspired connectionist architecture. *Robotics and Autonomous systems*, 37:145–160, 2001.

[9] Rodney A. Brooks. From earwigs to humans. *Robotics and Autonomous Systems*, 20(2-4):291–304, 1997.

[10] David P. Carey. Eye-hand coordination: Eye to hand or hand to eye. *Current Biology*, 10:R416–R419, 2000.

[11] Adam F. Carpenter, Apostolos P. Georgopoulos, and Giuseppe Pellizer. Motor cortical encoding in a context-recall task. *Science*, 283:1752–1757, march 1999.

[12] Jefferson Coelho, Justus Piater, and Roderic Grupen. Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot. *Robotics and Autonomous Systems*, 37:195–218, 2001.

[13] Daniela Corbetta, Esther Thelen, and Kimberley Johnson. Motor constraints on the development fo perception–action matching in infant reaching. *Infant Behavior and Development*, 23:352–374, 2000.

[14] David M. Eagleman and Terrence J. Sejnowski. Motion integration and postdiction in visual awareness. *Science*, 287:2036–2038, march 2000.

[15] I. Faillenot, I. Toni, J. Decety, M. Gregoire, and M. Jeannerod. Visual pathways for object oriented action and object recognition functional anatomy with pet. *Cerebral Cortex*, 1(7):77–85, 1997.

[16] T. Flash. The control of hand equilibrium trajectories in multi–joint arm movements. *Biol Cybern*, (57):257–274, 1987.

[17] Vittorio Gallese, Laila Craighero, Luciano Fadiga, and Leonardo Fogassi. Perception through action. *Psyche*, 5(21), july 1999.

[18] Melvyn A. Goodale and G. Keith Humphrey. The objects of action and perception. *Cognition*, (67):181–207, 1998.

[19] Michael S. A. Graziano, Dylan F. Cooke, and Charlotte S.R. Taylor. Coding the location of the arm by sight. *Science*, 290:1782–1785, december 2000.

[20] Douglas Greve, Stephen Grossberg, Frank Guenther, and Daniel Bullock. Neural representations for sensory–motor control, i: Head-centered 3-d target positions from opponent eye commands. Technical report, Center for Adaptive Systems and Department of Cognitive Neural Systems Boston University, 111 Cummington Street Boston, MA 02215, 1991.

[21] Stephen Grossberg. Linking laminar circuits of visual cortex to visual perception: Development, grouping and attention. *Neuroscience and Biobehavioral Reviews*, 25:513–526, 2001.

[22] R.A. Grupen and J.A. Coelho, Jr. Structure and growth: A model of development for grasping with robot hands.

[23] G. Hager, W. Chang, and A. Morse. Robot hand-eye coordination based on stereo vision, 1995.

[24] Alexa Hauck, Georg Passog, Johanna Ruttinger, Michael Sorg, and Georg Farber. Biologically motivated hand-eye coordination for the autonomous grasping of unknow objects. *Autonome Mobile Systeme*, november 1999.

[25] Alexa Hauck, Michael Sorg, and Thomas Schenk. What can be learned from human reach-to-grasp movements for the design of robotic hand-eye systems. In *IEEE Int. Conf. on Robotics and Automation*, pages 2521–2526, 1999.

[26] S. Hutchinson, G. Hager, and P. Corke. A tutorial on visual servo control, 1996.

[27] Marco Iacoboni, Roger P Woods, Marcel Brass, Harold Bekkering, John Mazziotta, and Giacomo Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286:2526–2528, december 1999.

[28] IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'00). *On the performance of a biologically motivated visual control strategy for robotic hand-eye coordination*, november 2000.

[29] M. Jeannerod. Visuomotor channels: their integration in goal-directed prehension. *Human Movement Science*, 18:201–218, 1999.

[30] Scott H. Johnson. Thinking ahead: the case for motor imagery in prospective judgements of prehension. *Cognition*, 74:33–70, 2000.

[31] J.P. Kuhtz-Buschbeck, H. Stoltze, A. Boczek-Funcke, K. Johnk, H. Heinrichs, and M. Illert. Kinematic analysis of prehension movements in children. *Behavioural Brain Research*, 93:131–141, 1998.

[32] Jeffrey J. Lockman, Daniel H. Ashmead, and Emily W. Bashnell. The development of anticipatory hand orientation during infancy. *Journal of Experimental Child Psychology*, 37:176–186, 184.

[33] R.E.S. Manzotti. none. This is the formula for calculating the center of mass directly in log-polar, 2001.

[34] C R Mason, J E Gomez, and T J Ebner. Hand synergies during reach-to-grasp. *Journal of Neurophysiology*, 86:2896–2910, 2001.

[35] Michael E. McCarthy, Rachel H. Clifton, Daniel H. Ashmead, Philip Lee, and Nathalie Goubet. How infants use vision for grasping objects. *Child Development*, 72(4):973–987, August 2001.

[36] G. Metta, G. Sandini, and J. Konczak. A developmental approach to visually-guided reaching in artificial systems. *Neural Networks*, 12:1413 1427, 1999.

[37] Giorgio Metta. *Babybot a study on sensory-motor development.* PhD thesis, LiraLab, Department of Telecommunication, Computer and System Sciences, University of Genova, Italy, 1999.

[38] A. David Milner and Melvyn A. Goodale. *The visual brain in action.* Number 27 in Oxford psychology series. Oxford university press, 1995.

[39] Mark Mon-Williams and James R. Treselian. A simple rule of thumb for elegant prehension. *Current Biology*, 11:1058–1061, 2001.

[40] Pietro Morasso. Coordination aspects of arm trajectory formation. *Human movement science*, (2):197–210, 1983.

[41] Gerardo Pardo-Castellote, Tsai-Yen Li, Yoshihito Koga, Robert H. Cannon Jr., Jean-Claude Latombe, , and Stan Schneider. *Experimental Robotics III: The third International Symposium*, volume 200 of *Lecture Notes in Control and Information Sciences*, chapter Experimental Integration of planning in a Distributed Control system, pages 50–61. Springer-Verlag, Kyoto Japan, October 1993.

[42] Alexandre Pouget, Jean Christophe Ducom, Jeffrey Torri, and Daphne Bavelier. Multisensory spatial representations in eye centered coordinates for reaching. *Cognition*, 83:B1–B11, 2002.

[43] Daniel C. Richardson, Michael J. Spivey, and Jane Cheung. Motor representations in memory and mental models: Embodiment in cognition.

[44] Agns Roby-Brami, Nesha Bennis, Mounir Mokhtari, and Pierre Baraduc. Hand orientation for grasping depends on the direction of the reaching movement. *Brain Research*, 869:121–129, 2000.

[45] Alice C. Roy, Yves Paulignan, Alessandro Farn, Christophe Jouffrais, and Driss Boussaoud. Hand kinematics during reaching and

grasping in the macaque monkey. *Behavioural Brain Research*, 117:75–82, 2000.

[46] Dan Ryder and Oleg V. Favorov. The new associationism: a neural explanation for the predictive powers of cerebral cortex. *Brain and Mind*, 2:161–194, 2001.

[47] Ricarda I. Schubotz and D. Yves von Cramon. Functional organization of the lateral premotor cortex: fmri reveals different regions activated by anticipation of object properties, location and speed. *Cognitive Brain Research*, 11:97–112, 2001.

[48] Noel E. Sharkey. Learning from innate behaviors: a quantative evaluation of neural network controllers. *Machine Learning*, 31:115–136, februari 1998.

[49] Adrian Stoica. Robot fostering techniques for sensory-motor development of humanoid robots. *Robotics and Autonomous systems*, 37:127–143, 2001.

[50] Bert C. L. Touwen. The neurological development of prehension: a developmental neurologist's view. *International Journal of Psychophysiology*, 19:115–127, 1995.

[51] L. G. Ungerleider and M. Mishkin. *Analysis of visual behavior*, chapter Two cortical visual systems, pages 549–586. Cambridge, MA: MIT Press, 1982.

[52] J.P. Urban, J.L. Buessler, and J. Gresser. Neural networks for visual servoing in robotics. Technical report, Universite de haute Alsace, 4 rue des Freres Lumiere - 68 093 Mulhouse Cedex - France, november 1997.

[53] B. Vereijken, R.E.A. van Emmerik, H.T.A. Whitting, and K.M. Newell. Freezing degrees of freedom in skill acquisition. *Journal of Motor Behaviour*, 24:133–142, 1992.

[54] Claes von Hofsten, Peter Vishton, Elisabeth S. Spelke, Qi Feng, and Kerstin Rosander. Predictive action in infancy: tracking and reaching for moving objects. *Cognition*, 67:255–285, 1998.

[55] Cleas von Hofsten and Shirin Fazel-Zandy. Development of visually guided hand orientation in reaching. *Journal of Experimental Child Psychology*, 38:208–219, 1984.

[56] Cleas von Hofsten and Louise Rönnqvist. Preparation for grasping an object: A developmental study. *Journal of Experimental Psychology*, 14(4):610–621, 1988.

[57] W.J. Wadman, J.J. Denier van der Gon, R.H. Geuze, and C.R. Mol. Control of fast goal–directed arm movements. *Journal of Human Movement Studies*, (5):3–17, 1979.

[58] W. G. Walter. A machine that learns. *Scientific American*, 185(5):60–63, 1950.

[59] Jan D. Wolter, Richard A. Volz, and Anthony Woo. Automatic generation of gripping positions. *IEEE Transactions on systems, man and cybernetics*, SMC–15(2):204–213, 1985.

[60] S. Zeki. *A Vision of the Brain.* Oxford: Blackwell, 1993.